An Agent-based Classroom Lessons Model and Simulation

by

Fred J. Ingram

Department of Management Science
Lancaster University

A thesis submitted in part fulfilment of the degree of Doctor of Philosophy

November 2020

Abstract

In education contexts, agent-based simulations have been used at all levels, from the individual to the classroom to the state, to explore individual-level interactions and often to provide conclusions and explanations that could support decision-making. The aim of this research was to develop an empirically-grounded classroom lessons behaviour model and explore the theoretical consequences of the model design and the modelling assumptions, to understand the mechanisms of lesson interactions and to assess whether the results and explanations were plausible and realistic. The model was to be as realistic (not simplistic) as possible, addressing the limitations found in current lesson-related simulations and extending the use of agent-based simulations in classroom research to full-lessons with full-classes.

A typical empirically-driven agent-based modelling and simulation methodology was followed, incorporating an investigative case study at a UK secondary school. A comprehensive agent-based model of classroom lessons was formulated and a lesson event recording tool was developed and used to record a wide range of student and teacher activities. These data were used to calibrate and validate the simulation model. Three agent types were incorporated: students, teacher and teaching assistant. Agent decision-making was modelled using conventional production rules (one set for each agent type) that integrated the influences from, in order of significance, the lesson plan (which specifies the desired behaviour and is enacted by the teacher), the current circumstances (who is doing what and where) and, for students only, their historical empirical activity state frequencies.

The simulation model was validated with the help of experienced teachers, who considered that it embodied plausible theories of classroom behaviours. It was seen to generate plausible causal agent-level explanations for some important lesson dynamics and therefore judged to be a useful lesson analytics and decision-support tool, enabling educators to explore the consequences of a range of lesson interventions.

Declaration

I declare that the thesis is my own work and has not been submitted in substantially the same form for the award of a higher degree elsewhere. I declare also that the word count of 63,405 conforms to the permitted maximum (80,000 words excluding bibliography).

Acknowledgements

I would like to thank all those who have supported this research, particularly those who helped in the case study and model validation. I also thank my PhD supervisors, Roger Brooks and Richard Williams, for their support and advice. Particular thanks go to Helena Z. and my daughter Ruth for all their superb proof-reading, and Ruth again for her invaluable general support. Thank you.

Contents

1. Introduction and research methodology 1
1.1 The potential impact of a classroom lessons simulation 2
1.2 The research objectives and thesis structure 3
1.3 The reasons an agent-based approach was adopted 4
1.4 Existing classroom-related simulation models 8
1.5 The research methodology 10
1.5.1 ABMS development methodology 11
1.5.2 The investigative case study 16
1.6 Chapter summary 21
2. Conceptualizing and modelling classroom lessons 22
2.1 The classroom lesson system 22
2.2 The structure of classroom lessons 26
2.3 Relevant lesson activities 27
2.4 Teacher and student decision-making 30
2.4.1 The interaction between student misbehaviour and the teacher's disciplining . 33
2.5 Representing the classroom and lesson interactions 34
2.6 Assumptions and simplifications 36
2.7 Chapter summary 40
3. The Classroom Lessons Simulation Model (CLSM) 41
3.1 Agents activity states 42
3.2 How lesson plans control lessons 46
3.3 Agent behaviour and decision-making 48
3.3.1 Teacher behaviour 48
3.3.2 Student behaviour 51
3.3.3 TA behaviour 54
3.4 Agent interactions 55
3.4.1 Modelling the teacher's disciplining 58
3.4.2 Modelling the students' response to disciplining 59
3.5 The implementation in NetLogo 60
3.6 Chapter summary 64
4. Comparing simulated and empirical lessons 65
4.1 Model parameters and instantiation per lesson 65

4.2 Metrics selected for output comparison 69
4.3 The empirical data and acceptability ranges 72
4.4 Number of replications 74
4.5 Chapter summary... 76
5. Model calibration procedures... 78
5.1 Chapter summary... 83
6. Model validation 84
6.1 Conceptual model validation 85
6.2 Simulation model validation procedures... 86
6.3 Face validity testing 87
6.4 Validation results 90
6.5 Chapter summary... 95
7. Sensitivity analyses... 96
7.1 Parameter sensitivity 97
7.2 Student state attribute value sensitivity.. 99
7.3 Relative parameter importance 101
7.4 Chapter summary... 103
8. Investigation of the four research questions 104
8.1 Experiment 1: The influence of teacher support and discipline on student behaviour105
8.1.1 Analyses of the empirical data 106
8.1.2 Analyses of simulation results.. 107
8.1.2.1 Confirmation of effect of parameter value changes 109
8.1.2.2 The relationship between %Productivity and the parameters TSOL, TMT and SSRW 110
8.1.2.3 Meso-level analyses of the relationship between %Productivity and TSOL and TMT 116
8.1.3 Discussion of results and conclusions to research question 1 120
8.2 Experiment 2: The impact of a TA on productivity 122
8.2.1 The effect of removing the TA from Lesson #5 123
8.2.2 The effect of adding a TA to Lesson #1, #3 and #4 126
8.2.3 Conclusions to research question 2 131
8.3 Experiment 3: Seating arrangements and student behaviour 132
8.3.1 The chosen reseatings for each lesson 133
8.3.2 The effects of the seating rearrangements.. 137

8.3.3 Conclusions to research question 3 140
8.4 Experiment 4: Experiments with artificial classes 142
8.4.1 Experiment 4a: Investigating the effect of disruptive students 143
8.4.2 Experiment 4b: Active and passive disengagement in a full classroom 148
8.4.3 Conclusions to research question 4 150
8.5 Chapter summary 152
9. Discussion, further research and conclusions 155
9.1 Thesis summary 155
9.2 Further development to reduce model limitations 158
9.2.1 Quantifying disruption 158
9.2.2 The initiator of an interaction 158
9.2.3 The influence of student friendships 160
9.2.4 Enhanced decision-making logic 161
9.2.5 Understanding model stochasticity 161
9.2.6 Improving model validation 162
9.3 Issues with the CLSM and with ABMS in general 163
9.3.1 Agent attributes, empirical data and model parameters 164
9.3.2 The empirical data, model calibration and generality 168
9.3.3 Variability of simulation results 169
9.3.4 Using sensitivity analysis during model calibration 170
9.3.5 The validity of the conceptual and simulation models 171
9.3.6 Prediction of ‘possible futures’ 172
9.4 Further classroom research using the CLSM 173
9.5 Conclusions 174
References 177
Appendices 188
Appendix A Further details of the CLSM 188
A.1 Teacher behaviour 188
A.2 TA behaviour 193
A.3 Student behaviour 196
A.3.1 The student variables, constants, attributes and parameters 196
A.3.2 Rules for forced student behaviour 200
A.3.3 Rules for filtering out impossible student states 202
A.3.4 The students’ response to the teacher’s disciplining 206

A.3.5 How lesson parameters adjust student state scores 210
A.3.6 Modelling students leaving the classroom 212
Appendix B The derivation of model constants and functions 215
B.1 State duration modelling 215
B.2 The Time-To-Start-Helping lesson variable.. 217
B.3 TA-MIN-HELP-WAIT and TEACHER-MIN-HELP-WAIT 218
B.4 TIME-LAST-WITH-TA-THRESHOLD.. 220
B.5 The teacher's misbehaviour response delay 220
Appendix C Model calibration results 222
Appendix D Analyses of relative importance of the lesson parameters 225

List of Figures

Figure 1-1 'A typical agent' (Macal and North, 2010, p. 154) 5
Figure 1-2 'The conceptual model in the simulation project life-cycle' (Robinson, 2008a, p. 282) 12
Figure 1-3 'Real World and Simulation World Relationships with Verification and Validation' (Sargent, 2007, p. 127) 12
Figure 1-4 The LERT user interface showing student groups, activity states and participation18
Figure 1-5 An example of the lesson event data generated. 19
Figure 2-1 The influences on one student in the classroom lesson sub-system 23
Figure 2-2 Some activities of, and interactions between, lesson participants 29
Figure 2-3 A 2D model of a classroom lesson layout (Ge *et al.*, 2011, p. 96) 35
Figure 2-4 The classroom layouts available in the model from McDevitt (2017)................. 35
Figure 2-5 An example classroom lesson layout 36
Figure 2-6 The components and interactions included in the classroom lesson system 37
Figure 3-1 The principle interactions between the three agent types 41
Figure 3-2 Student agent considerations for the choice of the next activity 48
Figure 3-3 The flow of the teacher's activity choices... 49
Figure 3-4 Flowchart for student agent behaviour 52
Figure 3-5 TA agent decision-making and behaviour... 54
Figure 3-6 The main factors involved in teacher-student discipline interactions 57
Figure 3-7 How students are affected by the teacher's discipline....................................... 60
Figure 3-8 A screen capture of the simulation tool developer's interface 61

Figure 3-9 An example of the representation of a classroom lesson and participants 62
Figure 3-10 An example of the teacher and student state trajectories through a lesson 63
Figure 4-1 Comparison of distance metric results for different models showing that as the model becomes more realistic the simulations become more like the empirical lesson 72
Figure 4-2 The empirical distributions for the first six metrics and %Productivity 73
Figure 4-3 Example distributions of distance metric for 1000, 500, 200 and 80 replications 76
Figure 5-1 Cumulative metrics for Lesson #6 for one parameter set over 500 replications . 82
Figure 5-2 Cumulative overall distance metric for Lesson #6 with 95% CI (red lines) falling within the defined tolerance interval (green lines) 82
Figure 5-3 Distribution of the overall distance metric 83
Figure 6-1 Example of face validity test showing the three animations to be judged 88
Figure 6-2 Screen capture of an instant in one animation face validity test 88
Figure 6-3 Lesson #2 Distribution of the overall distance metric 93
Figure 6-4 Lesson #2 Distributions for the 7 metrics for lesson over 500 replications 93
Figure 6-5 Lesson #5 Distribution of the overall distance metric 94
Figure 6-6 Lesson #5 Distributions for the 7 metrics for lesson over 500 replications 94
Figure 7-1 The simulation model as a black box with output metric values resulting from lesson parameter values 96
Figure 7-2 Example parameter-metric relationships from Lesson #1 97
Figure 7-3 The increase in the four students' state 3 chatting time as their empirical state 3 probability was increased 99
Figure 7-4 The changes in the overall %Disruption associated with increasing state 3 (chatting) scale factors 100
Figure 7-5 The changes in the overall %Productivity associated with increasing state 3 (chatting) scale factors 100
Figure 7-6 Lesson #1 Predictor Importance (upper) for parameter rejection (lower) for closest match to empirical lesson 102
Figure 8-1 The relationship between the empirical amount of individual help and %Productive behaviour 106
Figure 8-2 The relationship between the empirical amount of disciplining and %Productive behaviour 107
Figure 8-3 The effect of adjusting the parameters in the validated lessons 110
Figure 8-4 The relationship between %Productivity and TSOL for all four lessons 111
Figure 8-5 The relationship between %Productivity and TMT for all four lessons 112

Figure 8-6 The relationship between %Productivity and SSRW for all four lessons 113
Figure 8-7 The relationship between %Productivity and TMT for different TSOL levels..... 115
Figure 8-8 Alternative consequences of increasing TSOL depending on TMT.................... 117
Figure 8-9 Examples of student lesson trajectories - with black arrows indicating when the teacher disciplined the class 118
Figure 8-10 Lesson #5 Independent %Productivity in various scenarios with the TA removed125
Figure 8-11 Teacher Help time in various scenarios with the TA removed 126
Figure 8-12 Lesson #1 metric distributions for several TASOL levels 127
Figure 8-13 Lesson #3 metric distributions for several TASOL levels 127
Figure 8-14 Lesson #4 metric distributions for several TASOL levels 128
Figure 8-15 Example Lesson #1 replication 129
Figure 8-16 Lesson #1 replay and seating arrangement showing Reseat 1 (blue arrow) and Reseat 2 (red arrow) 133
Figure 8-17 Lesson #3 replay and seating arrangement showing Reseat 1 (blue arrow) and Reseat 2 (red arrow) 134
Figure 8-18 Lesson #4 replay and seating arrangement showing Reseat 1 (blue arrow) and Reseat 2 (red arrow) 135
Figure 8-19 Lesson #5 replay and seating arrangement showing Reseat 1 (blue arrow) and Reseat 2 (red arrow) 136
Figure 8-20 Lesson #1 reseating results for %Productivity, %Distracting, %Other and TD . 137
Figure 8-21 Lesson #3 reseating results for %Productivity, %Distracting, %Other and TD . 137
Figure 8-22 Lesson #4 reseating results for %Productivity, %Distracting, %Other and TD . 139
Figure 8-23 Lesson #5 reseating results for %Productivity, %Distracting, %Other and TD . 139
Figure 8-24 Percentage of lesson time students spent productively vs disruptively 143
Figure 8-25 The students who became disruptive (labelled by seat number) 144
Figure 8-26 The effect of the disruptive students on metrics %Dist, %Other, TH, TD and TAH145
Figure 8-27 Factors affecting student misbehaviour and lesson misbehaviour metrics 147
Figure 8-28 Full class experiment with students selected for disruptive behaviour highlighted 148
Figure 8-29 The metrics %Dist, %Other, ASD, TH and TD in the five scenarios 149
Figure 9-1 Which parameter value is better: 2 or 8? 171
Figure A-1 The data required for teacher agent behaviour modelling 191
Figure A-2 The data required for TA agent behaviour modelling 194

Figure A-3 The data required for student behaviour modelling 197
Figure A-4 The empirical % off-task student time for selected teacher states 208
Figure A-5 The empirical % frequency of disengaged student states for selected teacher states 208
Figure A-6 Which parameters affect which student state attribute scores 210
Figure A-7 Fitted curve for predicting the likelihood of a student leaving the room 213
Figure B-1 An example of the type of curve fitted to the distributions in order to produce realistic random state durations 216
Figure B-2 Fitting a function for the time a teacher gives one-to-one support to a student216
Figure B-3 Finding a suitable function for generating realistic delays in helping students . 217
Figure B-4 A comparison of the empirical delays before helping and the results of the delay function 218
Figure B-5 The empirical delay between episodes of students being helped 219
Figure B-6 Empirical data on length of time between teacher and TA interactions 220
Figure D-1 SPSS analyses of the relative influence of the lesson parameters 228

List of Tables

Table 3-1 The teacher activity states modelled 43
Table 3-2 The student activity states modelled 44
Table 3-3 The TA activity states modelled 45
Table 3-4 The colours used to represent different student and teacher activities/states 63
Table 4-1 The model parameters and their use 66
Table 4-2 Summary of selected lessons 68
Table 4-3 Lesson comparison metrics (all mean values over the number of simulation runs)70
Table 4-4 The acceptability range for each metric.. 74
Table 5-1 Example of procedure for choosing a top set of parameters 80
Table 6-1 Responses of the teachers (in rows) to the 3 animation types (typical, worst, actual) for the seven selected lessons 89
Table 6-2 Simulation model validity test results (with validated lesson models highlighted)91
Table 6-3 Lesson #2 validation results 93
Table 6-4 Lesson #5 validation result 94
Table 7-1 The student state choices affected by each parameter...................................... 96

Table 7-2 Consequences of stepwise increments in parameter values................................ 98
Table 7-3 Relative parameter importance and the nature of the parameter space 103
Table 8-1 Data for the Independent Working section of the empirical lesson and the three replications (with student state times in seconds) 119
Table 8-2 Lesson #1 comparison of metrics without and with a TA (TASOL = 0) 130
Table 8-3 Lesson #1 changes in productivity after TA added, at two different TASOL levels130
Table 8-4 Lesson #1 expectations for the reseatings... 138
Table 8-5 Lesson #3 expectations for the reseatings... 138
Table 8-6 Lesson #4 expectations for the reseatings... 140
Table 8-7 Lesson #5 expectations for the reseatings... 140
Table 8-8 The relative changes between scenarios 6 and 7 for each student.................... 146
Table A-1 Teacher-related model constants... 192
Table A-2 TA-related model constants.. 195
Table A-3 Student-related model constants 198
Table A-4 Student agent variables 199
Table A-5 Conditions under which the student agent will not choose a state 202
Table A-6 The dynamic adjustments to student state scores.. 211
Table A-7 The number of occasions students were observed to leave the room 213
Table A-8 Adjusted frequency for leaving room 213
Table C-1 Parameter estimation results.. 223

1. Introduction and research methodology

The primary goal of this research was to determine how and to what extent an agent-based model could adequately represent the behaviours of students, teachers and teaching assistants in classroom lessons at a UK secondary school. This was motivated by the need for a lesson analytics tool that would enable teachers and educators to explore the mechanics and consequences of lesson interventions. The core task was to develop a model that could simulate the dynamic, asynchronous, spatial interactions between the autonomous and heterogeneous lesson participants over a lesson period. The simulation outcomes included measures of productivity and disruption, on the level of individual students and the whole class and lesson. Model adequacy was evaluated by comparing simulation outcomes against empirical lesson data and based on the opinions of experienced teachers.

This first chapter introduces the research project, explaining the motivation for the research (section 1.1) and specifying the research objectives and thesis structure (section 1.2). Chapter 1 also explains why an agent-based modelling (ABM) approach was adopted (section 1.3), considers existing classroom-related simulations (section 1.4) and describes the standard agent-based modelling and simulation (ABMS) development methodology for empirically-based models and how this was applied to developing a classroom lessons simulation model (CLSM) (section 1.5).

Concerning the terminology used:

- The initialism **ABM** can mean either agent-based model or agent-based modelling depending on context. Likewise, **ABMS** can mean either agent-based model and simulation or agent-based modelling and simulation depending on context. **ABS** means agent-based simulation.
- A **lesson** is regarded as a period of instruction/learning and a group of students who stay together for a series of lessons is called a **class**. Hence a class attends lessons. Lessons can take place anywhere (for example on a class excursion, during a one-to-one tutorial, during coaching on a sports field), but this project focused solely on lessons in classrooms, with desks, chairs, a board etc. Where **lesson** is used alone it means a **classroom lesson**.
- In UK classrooms there is often another adult besides the teacher, sometimes more than one, who assists in teaching and student support. Sometimes they are called learning assistants. In this project they have been designated **teaching assistant** or **TA** for short. A TA could be allocated to a specific student or selected students only, or as modelled in this research, be free to support any student.

1.1 The potential impact of a classroom lessons simulation

Teaching involves frequent decision-making, whether as part of the lesson planning process or on-the-spot decisions, e.g., when to use group work and for how long, when to ask questions, when and how much to help individuals, when to discipline and to what extent, and how to arrange seating (Ahoa *et al.*, 2010; Dillenbourg, 2013). Headteachers and school leaders make decisions about the duration of lessons, the range of academic abilities to place in classes, where to allocate TAs and so on. Teachers prepare lesson plans and try to follow a plan during a lesson, deviating from it to adjust to the needs of the individuals and the class as a whole. The decisions about what to plan and how to depart from the plan are typically left to the teacher's discretion (Ahoa *et al.*, 2010; Dillenbourg, 2013). Teachers are both inclined and expected to innovate, to try new methods or resources, but sometimes the innovation is not as successful as hoped. The evaluation of the consequences of decisions, in particular whether the change significantly improved the situation, is frequently left to the teacher's judgement. If the situation did not appear to improve, the teacher tries something else.

A model and simulation that enables some quantitative assessment of the consequences of a proposed change and perhaps provides a plausible chain of cause and effect, could be of significant practical benefit to a school or a teacher. For example, both are interested in the effect of an intervention on overall class productivity (perhaps *time on-task*) or the frequency of disruptions (especially *low-level disruption* - persistent, sub-critical distracting behaviour) or the amount of student participation (engagement). A simulation might help answer questions such as: 'How do you know the proposed change will likely increase student productivity?', or 'Why/How would doing X lead to situation Y?'. Experiments using such a simulation could enable teachers and others to explore alternative teaching strategies, or perhaps test a proposed or existing theory of classroom dynamics. Also, it is often quite difficult, and possibly unethical, to experiment (Gilbert, 2007; Read and Timmis, 2012) with alternative options in classes - even a small seating rearrangement can cause disputes and complaints from parents. With a simulation tool one could try out scenarios without disadvantaging anyone and without wasting resources.

In their endeavours to improve their practice, teachers also seek quantitative data and readily understandable visual information concerning what happens in their lessons (Fidalgo-Blanco *et al.*, 2015; Durall Gazulla and Leinonen, 2016; Holstein, McLaren and Aleven, 2017a; Klerkx, Verbert and Duval, 2017). Teachers are particularly interested in tools that facilitate analysis of and reflection on their lesson plans and actual teaching, especially concerning student engagement (Sergis and Sampson, 2016; Xhakaj, Aleven and McLaren, 2016). When a teacher or school wants to try out a new idea, it would be very useful to be able to measure the current situation, then simulate the new idea and measure the results. A realistic lesson simulation would require data about individual students, the teacher, the planned lesson and the teacher's and students' past behaviour in previous lessons. With these data and a validated model, the simulation could provide an idea about the consequences of an initial scenario or intervention, perhaps highlighting why anticipated outcomes may not occur and suggesting causal mechanisms for different outcomes. The research presented in this thesis was motivated by the benefits that a classroom lessons ABS could provide for decision-making in schools and colleges.

1.2 The research objectives and thesis structure

The primary objective was to answer this question:

How and to what extent can an agent-based model adequately represent the behaviours of, and interactions between, students, teacher and teaching assistant in classroom lessons at a UK secondary school?

To answer this, four experiments were designed. The answer would depend on how well the simulation model enabled the following four questions to be investigated (the first three of which are of on-going interest to educators).

1. Does the simulation model provide realistic explanations of the effects on lesson behaviours and lesson outcomes of alterations to the teacher's inclination to offer one-to-one support and to take disciplinary action?
2. Does the simulation model provide realistic explanations of the effects on overall student productivity of providing or withdrawing a TA who gives individual support to any student?
3. Does the simulation model provide realistic explanations of the effects on lesson behaviours and lesson outcomes of different student seating arrangements?
4. Does the simulation model provide realistic results in experiments with a class of artificial students?

These objectives were achieved by adopting an agent-based modelling and simulation methodology (described in section 1.5.1). The bulk of this thesis explains how this methodology was applied, how classroom lessons were conceptualized and represented, how empirical data were used, how lessons were compared and how the plausibility of simulations was established. The aim was to develop an empirically-grounded theoretical classroom lessons behaviour model that could be used to explore experiments on empirical lessons and provide plausible explanations. The conceptual and simulation models were intended to be as realistic (not simplistic) as possible, to address the limitations found in current simulations (described in section 1.4) and extend the use of ABS in classroom research to full-lessons with full-classes.

As the questions indicate, the goal was to explore the theoretical consequences of the model design and the modelling assumptions, to understand the mechanisms of interactions, and to assess whether the results and explanations were plausible and realistic. If the results and explanations were considered satisfactory, then this would show the extent to which the model could be useful. It would also indicate that the model embodies plausible theories of classroom behaviours, which could then be investigated further.

To facilitate explanation of this research, this document is structured around the stages in the ABMS methodology, summarized below and explained in section 1.5.1. Note that relevant background research is embedded in the appropriate sections (e.g., the next two sub-sections review previous applications of ABMS in educational contexts, section 2.3 discusses classroom activities investigated by other researchers, and section 2.4 discusses agent decision-making literature).

Chapter	Content
1	Introduction and research methodology.
2	Stage 1: Clarification of the problem domain, the classroom lesson system. Stage 2: Development of a conceptual model of classroom lessons.
3	Stage 3: Implementation of the conceptual model as a simulation model.
4	An explanation of how simulated and empirical lessons were compared.
5	Stage 4: Simulation model calibration using empirical data and suitable criteria.
6	Stage 5: Validation of the conceptual model and simulation model.
7	Stage 6: Sensitivity analyses.
8	Stage 7: The research experiments and results.
9	Stage 8: Evaluation and conclusions.

1.3 The reasons an agent-based approach was adopted

While there are many techniques for modelling systems, when it comes to modelling and simulating social systems, popular choices are discrete event simulation, system dynamics and agent-based modelling and simulation (ABMS) (Borshchev and Filippov, 2004; Katsaliaki and

Mustafee, 2011). The distinction between these approaches is not clear-cut however and several authors have shown or attempted to show that the three are equivalent to some extent (Bonabeau, 2002; Onggo, 2010). Others view ABMS as an add-on to discrete event simulation and system dynamics (Borshchev and Filippov, 2004). In addition, numerous hybrid combinations of these approaches exist (Martin and Schlüter, 2015).

The foundation of all simulations is a conceptual model of the system of interest. In agent-based models, the participants in the system – the agents or actors – are explicitly modelled as separate, autonomous, interacting entities, with their own characteristics or attributes, engaged in various activity states (Bonabeau, 2002). Agents have attributes (i.e., properties or characteristics) and methods (i.e., actions they can do or have done to them) – see Figure 1-1 from Macal and North (2010). Agents can have a mixture of heterogeneity (individuality) and homogeneity (uniformity) in their attributes and behaviours. An agent could be a single, indivisible entity (an actor or participant) or it could be an entire system itself (e.g., a company, an industry, a state), conceived as having characteristics and objectives and interacting with other individual agents or systems. Agents interact with other agents, and all reside within an environment which has some structure and properties (constants and variables) and with which the agents interact.

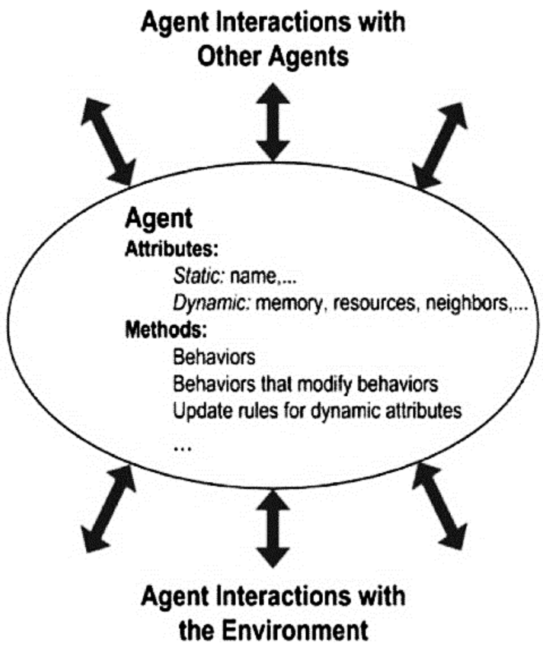

Figure 1-1 'A typical agent' (Macal and North, 2010, p. 154)

The decision-making and principal interactions between agents and their environment are often abstracted and simplified as formulae and/or decision rules. This is particularly useful when modelling social situations because it enables explicit modelling of the behaviour of the individual actors involved. An agent's decision-making can be based on their attribute values (past and present), what other agents have done or are doing, the current environment and the agent's goal(s). The formulae or rules often incorporate some stochasticity (using a pseudo-random number generator (PRNG)) to emulate real life systems where events are subject to random fluctuations (e.g., in choice thresholds). The appeal of rules is that they can offer suggestions for causality, explaining chains of interactions between agents and their environment - a micro-level explanation of macro-level results (Axtell, 2000; Macal, 2016; Edmonds *et al.*, 2019). This is something that purely stochastic or equation-based models generally do not provide, but an ABMS can.

If agents affect other agents and are affected by the other agents, then an agent effectively affects itself (van Geert, 1994). ABMs have often been used to investigate the emergent behaviour of an overall system as it evolves through the interactions of the agents (Macal, 2016). Such emergent behaviour may not be easily predicted from the formulae/rules used to describe the behaviour of individual agents. In complex systems there is often no way to determine the state of a system at a future time other than calculating the system state at every intermediate time step (Axtell, 2000). This is known as a generative approach to explanation (Epstein, 2008; Manzo and Matthews, 2014).

To put all this in the research context: It is natural to view a teacher, TA and students as heterogenous agents because they are each unique actors who relate to each other both spatially (in a classroom) and dynamically (over the lesson duration). Their characteristics and activities would be the attributes and states of agents, respectively. It is also natural to envisage a 2D or 3D representation of the classroom as the participants exist, act and interact from specific locations and have a physical range of influence. In addition, the goal was to provide causal explanations of the simulation outputs. These considerations made an ABMS approach the most appropriate for modelling classroom lessons. (For more information about ABMS one could consult Manzo and Matthews (2014) and Macal (2016)).

There are also many examples of where ABSs have been used to model and investigate agent-level behaviours in social systems, often to provide evidence that could support decision-making. For example: the analysis of smart energy grids (Ringler, Keles and Fichtner, 2016); forest management practices (Sotirov, Sallnäs and Eriksson, 2017); predicting refugee destinations (Suleimenova, Bell and Groen, 2017); simulating COVID-19 pandemic dynamics (Silva *et al.*, 2020). (See Macal and North (2010, 2014) and Gómez-Cruz, Loaiza Saa and Ortega Hurtado (2017) for many more.) In education contexts, ABMs have been used at all levels, from the individual to the classroom to the state. For example, they have been used to investigate:

- the advantages and disadvantages of team-work in problem-solving activities (Abrahamson, Blikstein and Wilensky, 2007);
- the spread of an infectious disease in a university classroom, with a teacher and a group of students interacting in a 2D representation of the classroom, complete with seats, desks and aisles (Ge *et al.*, 2011);
- the college sorting process[1] in the US (Maroulis *et al.*, 2014; Reardon *et al.*, 2016) and school choice in the UK (Harland and Heppenstall, 2012);
- how a network of social influences can give rise to the unequal distribution of educational choices across social groups in France (Manzo, 2013);
- differential school effectiveness[2] in London schools (Salgado, Marchione and Gilbert, 2014);
- whether a class should be judged on test results or cooperation and participation (Raca and Dillenbourg, 2014);
- the factors that affect the success or failure of proposed school projects (Mital, Moore and Llewellyn, 2014);
- the link between what students feel about their tuition, their teacher, their lessons and their actual attainment (Gamboa-Brooks-Gray, 2015);
- justifying training for classroom evacuation procedures (Liu, Jiang and Shi, 2016);
- energy expenditure and beverage consumption of school children after physical exercise (Chen *et al.*, 2017);
- the effect of seating arrangements on prejudice and variation in academic performance across ethnic groups (Radó and Takács, 2019);
- the impact on educational achievement of the same-race effect (a slight academic boost when students and teachers are of the same race) (Montes, 2012);
- a model of cheating behaviour when completing homework assignments (Paul *et al.*, 2020).

Although these are quite diverse applications, the above authors generally proposed that agent-based modelling could be used to run experiments *in silico* (i.e., on a computer) instead of in real life, to investigate causal chains of behaviour, simulate alternatives, test a theory and therefore aid decision-making. Over half of the above examples were empirically-based

[1] The process in which students choose which colleges to apply to and colleges choose which students to target during promotional activities and which students to ultimately accept.
[2] The comparison of schools on their students' performance (normally academic grades) for the purpose of identifying how differences might be explained.

simulations, built using observational data. The research presented in this thesis extends educational policy and strategy simulation from school-region level to classroom level.

As mentioned above, a common reason for building a model is to be able to explain (to some extent) what happens in the actual system (Epstein, 2008). As Macal wrote: 'What ABMS brings to social simulation is a framework for explicitly specifying causal mechanisms' (Macal, 2016, p. 146). And: 'We desire to produce a model whose results are more or less in one-to-one correspondence with the real world so that the jump in explanation from model to real world referent is minimal and convincing' (Macal, 2016, p. 145). Edmonds *et al.* (2019, p. 6) considered explanation to mean: 'establishing a possible causal chain from a set-up to its consequences in terms of the mechanisms in a simulation'. For example, in the case of lessons, if a lesson without a TA had one added, then experts (teachers) might anticipate certain results, depending on what other factors are also involved. One could investigate whether the anticipated results were indeed evident in simulation outputs and hence judge whether the sequences of interactions that led to those results (as a consequence of the model rules) were a plausible explanation. This form of investigation motivated the choice of ABMS to address the research questions.

Even though a model is intrinsically wrong[3] (in that it is an approximation to, or a simplification of, a system), it might be useful. Just the act of developing a model can lead to a better understanding of a system. Also, even when a model is unable to explain a phenomenon, at least it might suggest hypotheses that could be investigated. As Robinson (2014) put it:

> the purpose of a simulation can be described as obtaining a better understanding of and/or identifying improvements to a system. Improved understanding of a system, as well as the identification of improvements, is important since it informs future decision-making in the real system. (Robinson, 2014, p. 3)

1.4 Existing classroom-related simulation models

Simulations have previously been used to explore the consequences of what happens in lessons. For example, in 1982, Clauset and Gaynor published the results of their School Effectiveness Model. They used a system dynamics approach to formulate mathematical relationships between student achievement (particularly reading in junior school) and teachers' expectations and instruction. Their goal was 'to assess the likely consequences of various school improvement policies' (Clauset and Gaynor, 1982).

[3] '*All models are wrong but some are useful*' (Box, 1979) is often quoted by modellers when explaining their work.

Koster *et al.* (2016) reported on a proof-of-concept Classroom Simulator that modelled a classroom in which students were following the teacher during a lecture and their interactions depended on their engagement level, the different educational contents and teacher characteristics. However, this simulation did not have the teacher being influenced by the students, nor were agents spatially arranged in a classroom – both of which were addressed in this thesis research.

One key application of lesson simulations that has been explored to date is in teacher training to help teachers make better decisions (Strang and Loper, 1984). At the start of this research project, several teacher training simulation packages were commercially available, all based on custom-built conceptual models of teaching (Gibson, 2009; Gibson and Baek, 2009; Kim and Cheong, 2009). The purpose of these packages was to provide scenarios in which the trainee teacher chooses a course of action, and then to demonstrate to the trainee the consequences of their actions and provide feedback on their reactions to different types of student responses. The packages had sophisticated models for student behaviour, for example *ClassSim* (Ferry, Kervin and Carrington, 2010) and *simSchool* (Deale and Pastore, 2014). In *simSchool*, the simulated student behaviour was driven by empirical student profiles derived from the characteristics of thousands of real US students. The attributes comprising the student profiles were based on 'established theories of cognition, emotion, social behavior' and incorporated students' 'emotional intelligence' (simEd LLC, 2020). Unfortunately, although the *simSchool* software appeared to be highly relevant, it was not available for customization. The *aha! ProcessClassroom SIMs* (Piccolo and Oskorus, 2009) also appeared relevant despite being limited to trainees practising classroom management strategies from Payne (2007). This entailed responding to events (disruptive and non-disruptive) by selecting from a list of alternative actions. Lugrin *et al.* (2016) described a sophisticated immersive virtual reality system (named *3B*) for classroom management training that generated appropriately stressful situations for teachers in front of classes. Another simulation that modelled student learning used a machine-learning agent (*SimStudent*) and allowed the trainee teacher to practise one-to-one teaching (Matsuda *et al.*, 2010; Matsuda, Cohen and Koedinger, 2015). There were also intelligent tutoring systems which incorporated one or more (typically anthropomorphized) teaching agents (Kim and Baylor, 2016), but no students – because the human student is using the system.

However, in general, in the above simulations:

- there was no autonomous teacher agent (because that role was being played by the trainee teacher);

- the focus was whole-class teaching and did not appear to include other types of lesson activities such as students working independently or in groups with the teacher providing individual assistance;
- the simulated students (often only a few) were stationary and appeared to interact only with the teacher, not each other;
- there was no autonomous teaching assistant;
- the packages provided only segments of a lesson, not a full simulation of an entire lesson.

As mentioned in section 1.2, the aim of this research was to develop a realistic conceptual model and simulation that would emulate the dynamic and spatial behaviour of autonomous students interacting with each other and with an autonomous teacher and TA for whole lessons. This was achieved by addressing the limitations listed above.

Intriguingly, although discovered only towards the end of this project, it seems another researcher had developed a NetLogo-based simulation model '*Classroom Model with Learning Assistants*' and published it on the NetLogo Modeling Commons website (McDevitt, 2017). Although the simulation does not run under the latest version of NetLogo and there is no documentation, looking at the code, it appears that the author had various room configurations (see Figure 2-4), numbers of students interacting with each other (forming networks) and interacting with their teacher and learning assistants (by raising their hands). Stochasticity was incorporated (using a pseudo-random number generator) in several places, e.g., to form random student networks and in various probabilistic rules that decided whether students had learnt or not.

1.5 The research methodology

As explained at the start of Chapter 1, the central goal of the research was to construct and experiment with an agent-based simulation model of classroom lessons. Hence the standard ABMS development methodology for empirically-based models was applied. This is described in section 1.5.1. Further, the intention was to base the development more on primary than secondary research. The investigative case study (described in section 1.5.2) was thus central to the project.

As a brief overview, the first stage in building a model of classroom lessons was to clarify and understand the classroom lesson system (as viewed from the perspective of the research objectives). The resulting conceptualization was then simplified and restricted (scoped) in order to formulate a conceptual model of classroom lessons. Despite the simplifications and limitations, the conceptual model was designed to be as realistic, comprehensive and generic as possible, capable of representing a wide range of lesson behaviours. Finally, an agent-based simulation model (implemented in NetLogo) was built, incorporating as much of the conceptual model as was practicable.

Although the overall project approach was quantitative, the way in which a conceptualization of lessons was constructed could be termed a critical realist informed grounded theory approach (Thornberg, 2012; Hoddy, 2019). It was assumed that there is no universal concept of a classroom lesson to be revealed, but that through interactions with the participants in the social activities of school lessons an understanding of the phenomenon would be constructed. And, instead of attempting to start with a completely blank piece of paper, the researcher's own background was acknowledged to have an influence on the conceptualization and the viewpoints of other researchers and teachers were considered during conceptualization. The author is a qualified schoolteacher with over 14 years' teaching experience and can be considered a *domain expert*, a classroom lessons expert. The author is also an ex IT project manager with 25 years' experience in software development so had the appropriate experience for developing the simulation software.

1.5.1 ABMS development methodology

Modelling actors as autonomous agents is not a new approach, having its origins in the 1980s, when it was known as individual-level or micro-level representation (Reynolds, 1987; Troitzsch, 2018). The ABMS development methodology is quite well established (Balci, 1995; Sargent, 2007, 2011; Robinson, 2008b, 2014; Macal and North, 2010, 2014; Salgado and Gilbert, 2013; Siebers and Klügl, 2018) and is generally acknowledged to be an iterative process. Figure 1-2 (from Robinson (2008a)) shows that the core process is the development of a conceptual model of the system of interest and that this is transformed into a computer model but that the processes flow both ways (double-headed arrows). Sargent (2007) went into greater detail in Figure 1-3, also indicating the iterative nature of ABMS development and highlighting that the development of the conceptual model is based on hypotheses and abstractions of the system of interest, and that the conceptual model is transformed into a simulation model specification which is then transformed into a simulation model.

The iterative nature of ABMS development is particularly apparent when a novel application is being investigated, as discoveries at any point may lead to earlier stages being revisited and handled differently (Siebers and Klügl, 2018). In other words, a model can grow organically and take several cycles. As Norling, Edmonds and Meyer (2018) wrote: ‘Often, when modelling some complex phenomena (and especially social phenomena), one simply does not know beforehand which parts of the system will turn out to be important to the outcomes and which can be safely omitted’ (Norling, Edmonds and Meyer, 2018, p. 62).

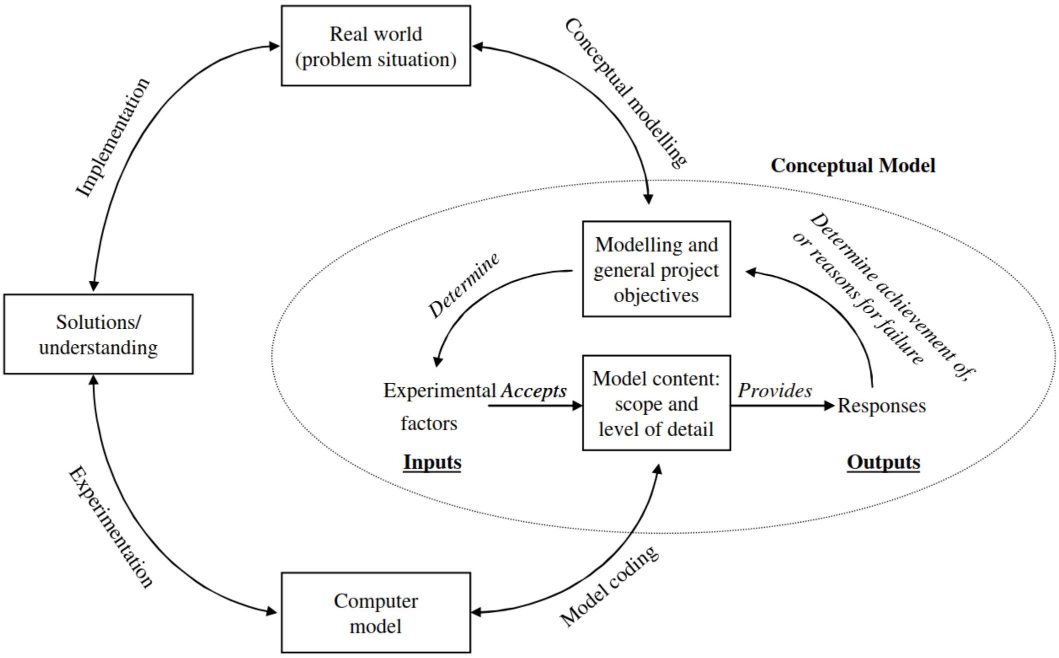

Figure 1-2 ‘The conceptual model in the simulation project life-cycle’ (Robinson, 2008a, p. 282)

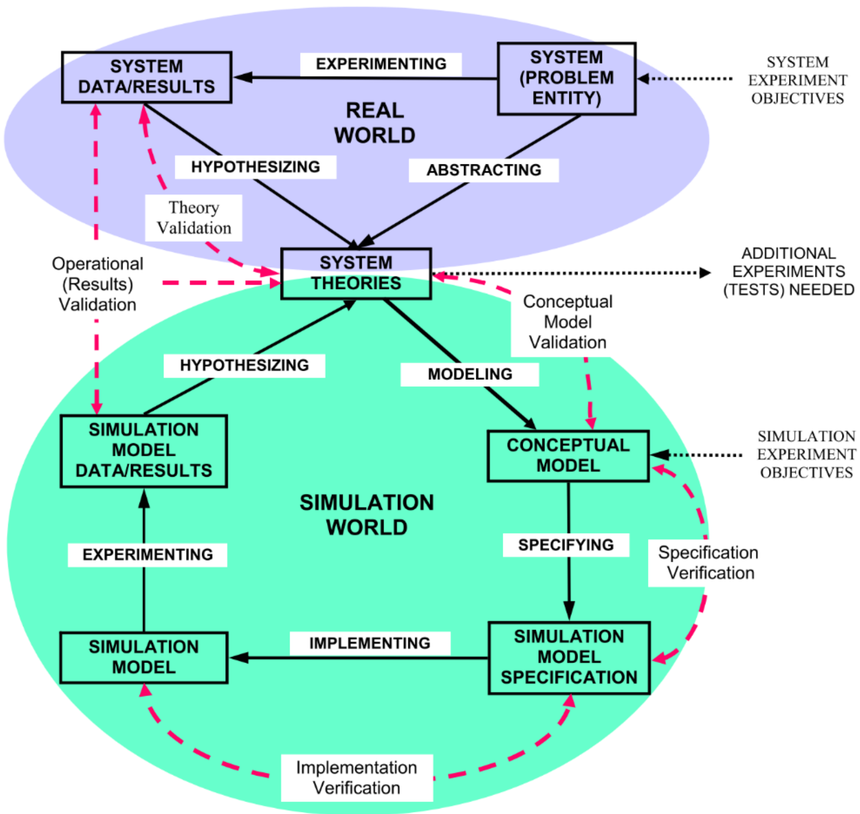

Figure 1-3 ‘Real World and Simulation World Relationships with Verification and Validation’ (Sargent, 2007, p. 127)

It is also important to appreciate that, despite the diagrams seeming to indicate that one stage follows another, typically multiple activities are being undertaken in parallel (e.g., data collection can be underway while a conceptual model is being developed). It is not uncommon for an iterative prototyping approach to be adopted in software development. A partial model can be partially implemented long before the model has been fully specified and system testing occurs. One of the purposes of a partial prototype is to confirm that the development tools and target platform (software and hardware) have the required features and capacities.

Figure 1-3 also refers to *validation*, that is confirmation that the correct model (the one the experts have specified) is built and that it delivers accurate and reliable results (Schlesinger *et al.*, 1979; Sargent, 2007, 2011; Tsioptsias, Tako and Robinson, 2016). Balci (1995) referred to validation as confirming that the '*right* model' is being built, that the right simulation has been built. There are several aspects to validating simulation models, the most pertinent being conceptual model validation, white-box validation (looking at the detailed micro-level workings of a system to establish internal validity) and black-box validation (looking at the macro- and meso-level [4] outputs of a system to establish external validity) (Robinson, 2014). These topics are covered in Chapter 6 in relation to the ABM developed as part of this thesis.

The exact steps followed during development depend on the purpose for which the simulation is being developed (Edmonds *et al.*, 2019). For example, when developing a simulation for predictive purposes and where empirical data is available, ABMS developers often split the data into two subsets. One subset (often called the '*training*' or '*in-sample*' dataset) is used to 'calibrate' the simulation model (explained in Chapter 5) while the other test ('*out-of-sample*') subset is used to attempt to validate [5] the model. If the simulation produces acceptable results for the test scenarios, then the model will be considered validated. As declared in section 1.2, the goal in this research was to develop an empirically-grounded theoretical classroom lessons behaviour model that could be used to explore experiments on empirical lessons and provided explanations – not predictions. For this type of project, it was more appropriate to use all the empirical data during model development so that the model is based on as much information as possible. However, this does not mean that no 'predicting' is attempted: experiments are conducted to test the simulation model, to see if the model can provide plausible causal

[4] In this research the labels micro, meso and macro refer to the level of agent-to-agent interactions, the level of agent statistics summaries, the level of overall lesson statistics (metrics), respectively.

[5] Perhaps a more rigorous, 'Popperian', modeller would be attempting to invalidate the simulation outputs in order to reject the hypothesis that it is a sufficiently realistic representation of the real-world system. See Epstein (2008).

relationships in new scenarios. But this is not quite the same as using the simulation to obtain only the predicted results of a previously unseen scenario.

Figure 1-3 also refers to *verification*, a quality control and assurance process to confirm that the development is being done *right* (Balci, 1995). Those familiar with software development will be well aware that it is generally impossible to prove that an entire computer system has no errors or bugs: ABMS development is no different (Galán *et al.*, 2009; Norling, Edmonds and Meyer, 2018). Standard software development practices were employed to reduce the chances of errors and confirm that the simulation was operating correctly. These practices included:

- following a defensive programming approach, so, for example, a significant amount of the code is dedicated to error-trapping;
- using preventative techniques, e.g., giving procedures, functions, variables and constants highly descriptive (and therefore long) names;
- unit testing (procedure/functions) via code walkthroughs and confirmation of output against known results (test data) and extreme and invalid value testing (destructive testing);
- black-box testing of output metric values, including confirming that lesson replay statistics matched lesson totals obtained from raw event data files, in order to establish external validity;
- visual checking of the lesson animations looking for inconsistencies or unrealistic situations;
- white-box testing of the simulation logic: detailed study of sequences of agent interactions.

Adopting procedures from the various authors mentioned above, the development of the classroom lessons ABMS broadly followed the stages listed below. As explained in 1.2, the structure of the thesis reflects these stages.

<u>Stage 1</u>: Understand the domain, clarifying what the classroom lesson system is. Identify the components of and participants in the system, the interactions between participants and the system inputs and outputs. These are necessary in order to have a clear understanding of the goals of the proposed system and would form part of a system specification (Robinson, 2013).

<u>Stage 2</u>: Develop a conceptual model of the classroom lesson system, specifying the system components included in the model, the selected participant characteristics and activities (states)[6] and interactions, and the model inputs and outputs. Formulate the rules and formulae that

[6] In this thesis the terms activity, state and activity state are used interchangeably: the state of an agent means the activity the agent is engaged in.

describe how the participants interact with each other and their environment. List all assumptions and modelling simplifications.

Although not emphasised in the diagrams above, because this was to be an empirically-based ABS, empirical data on lesson behaviours were needed to calibrate and validate the model. In addition, the goal was to develop a novel conceptualisation of the classroom lesson system (the domain) and hence a comprehensive conceptual model. For these reasons, stages 1 and 2 were developed during an investigative case study conducted at a UK secondary school.

Stage 3: Implement the ABM in NetLogo (a suitable agent-based simulation software package), with rigorous testing to minimize 'bugs'.

Stage 4: Calibrate the simulation model using the empirical data and suitable criteria. The goal of this stage is to produce a model that generates realistic results. This is accomplished by using the empirical data and finding model parameters that generate simulation outcomes within the general ranges of the empirical data. Chapter 5 explains this process.

Stage 5: Validate the conceptual model and simulation model.[7] The goal of this stage is to check whether a lesson simulation model generates outcomes that, although already considered realistic after Stage 4, are considered an adequate match to the specific empirical lesson being modelled. Chapter 6 explains with this process.

Stage 6: Conduct sensitivity analyses.

Although this was not included in the diagrams above, it is normal practice to conduct some sensitivity analyses to establish how the model outputs respond to changes in model parameter values and initial conditions (Balci, 1994; Thiele, Kurth and Grimm, 2014).

Stage 7: Conduct experiments and analyse the results. The purpose of the research experiments was introduced in section 1.2.

Stage 8: Evaluate the results and the model.

The simulation model was implemented in NetLogo (Wilensky, 1999), a popular, free, open-source agent-based simulation development package as it provided the technologies required

[7] Normally one would validate the conceptual model before developing the simulation model, but, for the convenience of the teachers involved in the evaluation, this was conducted immediately prior to face validity testing.

(such as autonomous agents interacting, agent attributes, provision for if-then-else rules to emulate agent-level decision-making and built-in animation facilities (so representing a classroom of interacting people needed minimal coding)). NetLogo is a popular choice amongst modellers for these and many other reasons and there are thousands of published NetLogo models (many on the NetLogo Modeling Commons website[8]). Building the simulation in NetLogo meant that interested parties could obtain the NetLogo package and download and run the lessons model. The NetLogo code plus data files and instructions are available at Ingram (2020a).

As explained at the start of this chapter, an investigative case study was central to the ABMS development. This is described next.

1.5.2 The investigative case study

The investigative case study took place at a very small UK secondary school at which the author was teaching and included informal discussions with individual teachers and focus group discussions with classes, plus lesson observations and data collection. It was undertaken to provide primary insights for the development of the simulation model. The goals were to:

- contribute to the conceptualization of classroom lessons;
- contribute to the design of a comprehensive conceptual model;
- identify and collect empirical data on lesson behaviours, and use these data to further contribute to the conceptualization and to conceptual model design;
- use that data for calibration and validation of the simulation model.

The simulation model needed specific data concerning what students and teachers do, and for how long: for example, how frequently and for how long students work alone or together, chat or make remarks, or receive individual help from a teacher or TA. There were some published data, including videos, that had some overlap with the data required for the proposed ABS, but not sufficient to be useful. For example, the videos did not provide the coverage of the whole class as they often followed the teacher around the class and thus lost sight of what most of the students were doing (TIMSSVIDEO, 1999). There were also data on discourses between a student and a teacher, but these data were not accompanied by sufficient other data for the timings to be used in the proposed ABMS (e.g., Kovalainen and Kumpulainen (2009)). In the absence of suitable secondary data, it was necessary that primary data be collected.

[8] Available at http://modelingcommons.org/account/login (last accessed on 26th October 2020).

Much classroom data collection is still conducted using forms (whether paper or digital), often recording the situation at regular intervals (often referred to as interval sampling or momentary time sampling), e.g., 20 seconds, and thus missing intervening events (Christle and Schuster, 2003; Brühwiler and Blatchford, 2011; Zohrabi, 2013; Goldring *et al.*, 2015). Forms (paper or digital) are still employed even when researchers capture video or audio recordings of lesson interactions because the recordings are first transcribed before coding (Hmelo-Silver, Liu and Jordan, 2009; Kovalainen and Kumpulainen, 2009). There were also several commercial lesson capture tools that enabled easy video and audio recording of lessons for later viewing (e.g., *IRIS Connect* (IRIS Connect UK, 2020), *Lessonbox* (Lessonbox Ltd, 2020) and *Lessonvu* (ONVU Technologies Group, 2020)), but, at the time of the start of this research project in 2016, these seldom provided the ability to classify events.

The use of software and hardware tools to *mark-up*, i.e., capture, categorize and code, event data directly (without the need for transcription), either live or from recorded material, has become increasingly popular, e.g., the Human Affect Recording Tool (HART) (Ocumpaugh, Baker and Rodrigo, 2014). Collecting data on classroom events is similar to collecting sports analytics data during sporting events. Commercial software, such as *Vosaic Connect* (Vosaic, 2020), enable observers to *mark-up* events (such as player interactions) by clicking user-defined buttons and entering data into standard or user-defined fields (on standard screens or on user-created forms). Most significantly though, the *Vosaic Connect* system has been adapted to recording and coding how teachers and students interact, improving on standard lesson observations by providing detailed performance feedback to the teachers (Ramos, Esslinger and Pyle, 2015).

Several other research teams have investigated the use of automated tools for collecting learning and teaching analytics data. These technologies include facial recognition, facial expression recognition and gaze tracking to determine student engagement and comprehension (Bidwell and Fuchs, 2011; Wei and Yang, 2012; Raca and Dillenbourg, 2013, 2014; Prieto *et al.*, 2016; Bosch *et al.*, 2018; Shvarts and Abrahamson, 2019).

Collecting the data necessary for a classroom lesson ABS required a lesson event recording tool that would quickly and efficiently capture the timings of more than a dozen student, teacher and TA activities, with a facility to undo or tag mistakes, and without the need to video record lessons. Due to project constraints it also needed to be free or inexpensive (ruling out *Vosaic Connect*) and customizable. It was also considered important that the tool be available to and customizable by others (schools, teachers, researchers etc.).[9] None of the available software were found to meet these requirements. For example, the Android app HART used by Holstein, McLaren and Aleven (2017b) was not considered fast enough to record multiple events near simultaneously. Hence a purpose-built Lesson Event Recording Tool (LERT) was developed and used. It is a Microsoft® Excel (2016) spreadsheet with added VBA code which can log the activities of any practical number of students. Paired and group/team working can be recorded, as can seating rearrangements. This is all accomplished by specific mouse button clicking (left, right and double) with or without Shift, Ctrl and Alt keys. The LERT spreadsheet is very efficient and quick: in one mouse operation the observer can assign an activity to a number of students. Additional information about the LERT and the program itself, along with documentation, is available online at Ingram (2018). The top part of the LERT user interface is shown in Figure 1-4.

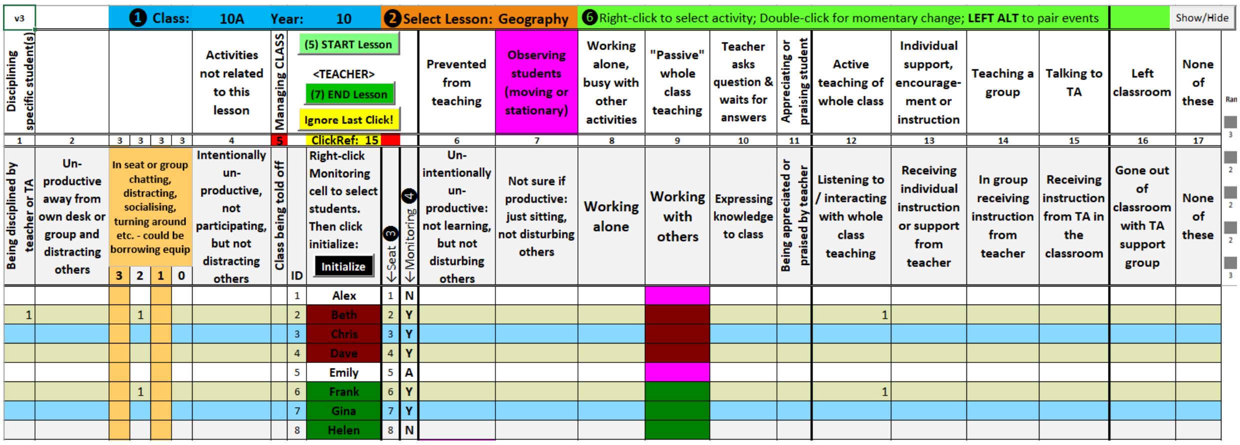

Figure 1-4 The LERT user interface showing student groups, activity states and participation

The wide top row contains the teacher activities; the lower wide row contains the student activities. Seventeen distinct student activities or states were recorded and fifteen teacher activities or states. (These are explained in section 3.1.) The state numbers can be assigned any meaning and are just labels internally. During operation, the LERT creates a log of time-stamped student and teacher activity changes (see example in Figure 1-5). Each lesson produces a CSV format file which can be 'replayed' and analysed using the ABS, which is described in Chapter 3.

Collecting classroom data involves well-known challenges, including common method bias (e.g., students further away from the observer might be monitored in less detail and less accurately, or some students may constantly initiate interactions unnoticed, but the observer always captures the respondent), observer reliability, observer effect distortion (the change in

[9] Besides providing data for the simulation model, the idea was that this tool would also be useful in its own right as a stand-alone tool to support data-driven lesson analytics and decision-making by teachers and management.

behaviour of the observed) and observer bias (e.g., if the observer knows the student and expects misbehaviour) (Christle and Schuster, 2003; Allday *et al.*, 2011; Imler and Eichelberger, 2011). Video-recording does not eliminate these problems (Savola, 2008; Dalland, Klette and Svenkerud, 2020). In addition, the practicalities of school routine (e.g., one class always unavailable) and classrooms (e.g., the availability of a suitable seat for an observer), could lead to sampling bias: only some lessons may be able to be observed and only some students.

EventNu	Year	SID	Name	Monit	Activity	Time	Date	LessonSub	Range	ClickC	RightCli	Last	Paired	Curren	WhoWith	LessonID
73	10	7		Y	8	15:42:22	08/01/2018	Geography		7	TRUE		FALSE			18010810GEO
74	10	8		N	8	15:42:22	08/01/2018	Geography		7	TRUE		FALSE			18010810GEO
75	10	0		Y	7	15:42:23	08/01/2018	Geography		7	TRUE		FALSE			18010810GEO
76	10	2		Y	3	15:42:25	08/01/2018	Geography	2	8	FALSE		TRUE		"1006"	18010810GEO
77	10	6		Y	3	15:42:25	08/01/2018	Geography	2	8	FALSE		TRUE		"1002"	18010810GEO
78	10	2		Y	8	15:42:27	08/01/2018	Geography		8	FALSE		TRUE			18010810GEO
79	10	6		Y	8	15:42:27	08/01/2018	Geography		8	FALSE		TRUE			18010810GEO
80	10	0		Y	1	15:42:27	08/01/2018	Geography		9	FALSE		FALSE			18010810GEO
81	10	2		Y	1	15:42:27	08/01/2018	Geography		9	FALSE		FALSE			18010810GEO
82	10	0		Y	7	15:42:29	08/01/2018	Geography		9	FALSE		FALSE			18010810GEO
83	10	2		Y	8	15:42:29	08/01/2018	Geography		9	FALSE		FALSE			18010810GEO
84	10	1		N	9	15:42:29	08/01/2018	Geography		10	TRUE		FALSE			18010810GEO
85	10	2		Y	9	15:42:29	08/01/2018	Geography		10	TRUE		FALSE		"1003,1004"	18010810GEO
86	10	3		Y	9	15:42:29	08/01/2018	Geography		10	TRUE		FALSE		"1002,1004"	18010810GEO
87	10	4		Y	9	15:42:29	08/01/2018	Geography		10	TRUE		FALSE		"1002,1003"	18010810GEO
88	10	5		A	9	15:42:29	08/01/2018	Geography		10	TRUE		FALSE			18010810GEO
89	10	6		Y	9	15:42:29	08/01/2018	Geography		10	TRUE		FALSE		"1007,1008"	18010810GEO
90	10	7		Y	9	15:42:29	08/01/2018	Geography		10	TRUE		FALSE		"1006,1008"	18010810GEO
91	10	8		N	9	15:42:29	08/01/2018	Geography		10	TRUE		FALSE		"1006,1007"	18010810GEO
92	10	0		Y	13	15:42:32	08/01/2018	Geography		11	TRUE		FALSE			18010810GEO
93	10	1		N	13	15:42:32	08/01/2018	Geography		11	TRUE		FALSE			18010810GEO
93	10	1		N	-1	15:42:34	08/01/2018	Geography		11	TRUE		FALSE			18010810GEO
94	10	1		N	9	15:42:34	08/01/2018	Geography		12	TRUE		FALSE			18010810GEO
95	10	0		Y	7	15:42:34	08/01/2018	Geography		12	TRUE		FALSE			18010810GEO
96	10	0		Y	13	15:42:36	08/01/2018	Geography		13	TRUE		FALSE			18010810GEO
97	10	3		Y	13	15:42:36	08/01/2018	Geography		13	TRUE		FALSE			18010810GEO
98	10	3		Y	9	15:42:38	08/01/2018	Geography		14	TRUE		FALSE		"1002,1004"	18010810GEO
99	10	0		Y	7	15:42:38	08/01/2018	Geography		14	TRUE		FALSE			18010810GEO

Figure 1-5 An example of the lesson event data generated.

Also, one can observe an interaction but be unsure who started it. This situation is well-known to teachers when they challenge a student only to be told: 'But I didn't start it!'. (Not that this is necessarily true.) Similarly, during the case study, when a teacher was moving around a class providing one-to-one support, it was often not clear whether the student asked for help as the teacher approached or whether the teacher offered assistance first. Hence, for the purposes of this project, interactions were considered directionless.

To minimize the consequences of these issues, systematic procedures for classroom observations and for collecting data on students' interactions and behaviour in classes were adapted from various researchers (Kaplan, Gheen and Midgley, 2002; Blatchford, Bassett and Brown, 2005; Bidwell and Fuchs, 2011; Ocumpaugh, Baker and Rodrigo, 2014). These procedures included giving observers standardised and comprehensive training in activity classification and the operation of the LERT data capture program. But even with perfect training the type of errors listed above could occur (Burns and Knox, 2011). There seems to be little that can be done about these issues without more intrusive observation – which would probably distort behaviour. There were three observers, the author plus two others trained by the author: a teacher at the school (who knew all the students) and a trainee teacher (and ex-pupil). They were given:

- a 4-page instruction manual that explained the principles of the data collection and how to record all the different events, including many examples of the types of events that could be encountered;
- an initial training session to go through the instructions for documenting a lesson observation, practise using the data collection tool and ask questions;
- a practice run, recording the teacher and just a few students, with a debriefing afterwards to ask questions – the data collected were not included in the analyses.

For practical reasons, lessons were not video recorded. This meant that often not all the students could be monitored – there was too much going on for the observer to record everything. As a consequence, many times the teacher, the TA or a student interacted with someone who was not being monitored. Although this resulted in an incomplete picture of lessons, simulation data for the students who were monitored could be compared to their empirical data and used to calibrate and validate the model.

Another caveat is that lesson observations are limited samples of the whole range of a student's behaviour. For example, one student may have been observed being highly disruptive. But if this was an extremely rare occurrence and this student was observed in only one or two lessons, then the relative frequency recorded for this behaviour would be disproportionately high – and the simulated behaviour for that student would be correspondingly distorted. Some students have data that shows they spent almost no time in certain activities (such as being helped by the teacher or TA), even though they often did those activities in other lessons – it was just that, in the lessons they were observed, they did not, either by choice or because they did not have the opportunity (e.g., no TA that lesson, or the teacher helped them instead). The empirical data revealed that one student can behave in a very different manner in different lessons: effectively they can become a different student. This is something that teachers had described and students had acknowledged in class discussions.

To summarize, the case study involved informal discussions with individual teachers and focus group discussions with classes and the observation of 42 1-hour lessons (comprising over 20,000 lesson events). The lesson observations involved 7 subjects, 7 teachers and 67 students

(33 females; ages 11 to 15). During the first 21 lessons, conceptualization and model design were in their infancy and the lesson event recording tool (LERT) was being developed. As has been emphasised, conceptualizing the classroom lesson system and forming a comprehensive conceptual model of classroom lessons was an iterative process (as described in section 1.5.1): every lesson observation or discussion with a teacher or focus group discussion with a class led to new insights and previous insights being re-evaluated. In parallel, the tool for recording lesson events (the LERT) was being developed and modified as different types of events were observed. Hence only the second half of the data (lesson 22 to 42) were considered complete and reflective of the final conceptualization and conceptual model. These 21 lessons involved 6 subjects, 6 teachers, 52 students (27 females; ages 11 to 15) and comprised approximately 9,000 lesson events (roughly 420 events per 1-hour lesson).

1.6 Chapter summary

This chapter explained the motivation for the research and the research objectives. It explained the reasons an ABMS approach was adopted, giving examples of other ABMS applications. This was followed by a review of existing classroom-related models and their limitations, as these are to be addressed in this research. The stages in the ABMS development methodology were then outlined. Finally, the role of the investigative case study was explained, particularly the collection of lesson event data concerning the behaviour of lesson participants.

The following chapter explains the results of stages 1 and 2 of the ABMS development - the conceptualization of the classroom lesson system and the design of a conceptual model.

2. Conceptualizing and modelling classroom lessons

As explained in section 1.5.1, the first stage in an ABMS project is to clarify one's understanding of the domain of interest, classroom lessons in this case. The second stage is to take that conceptualization and develop a simplified and formalized conceptual model of the system (Robinson, Arbez and Birta, 2015). For example, in real lessons there are chairs and books, but these may not necessarily be represented in a conceptual model.

The purpose of this chapter is to clarify what classroom lessons were considered to be, what is considered to happen in them, how participants interact, and hence justify what aspects were included in the conceptual model. As has been emphasised, the conceptualization and modelling were guided initially by the published literature and then by insights from the case study and the empirical data collected during the study, and by the authors 14 years of teaching experience.

This chapter covers:

- how the classroom lesson system was conceptualized;
- the structure of classroom lessons;
- what people do in lessons;
- teacher and student decision-making, particularly concerning misbehaving and disciplining;
- how the physical classroom was modelled;
- modelling assumptions and simplifications.

2.1 The classroom lesson system

The notion of a *classroom lesson* will mean different things to different people and will depend on the purpose for the concept. As Robinson noted, one needs to clarify which *real world* is being conceptualized (Robinson, 1997, 2014). Different researchers with different objectives would construct different description of classroom lessons, perhaps in terms of types of discourse or use of resources. The focus in this project was the dynamic behaviour of people during lessons. Hence it was important to consider what might influence their behaviour. The diagram in Figure 2-1 summarizes some of these influences, focusing on one student. Each element in the diagram (apart from the classroom resources) could be considered as a system in itself, although the people in the lessons are being considered as individual participants, not systems. Each person and system has an effect on and is affected by other people and other systems (e.g., the school and the community).

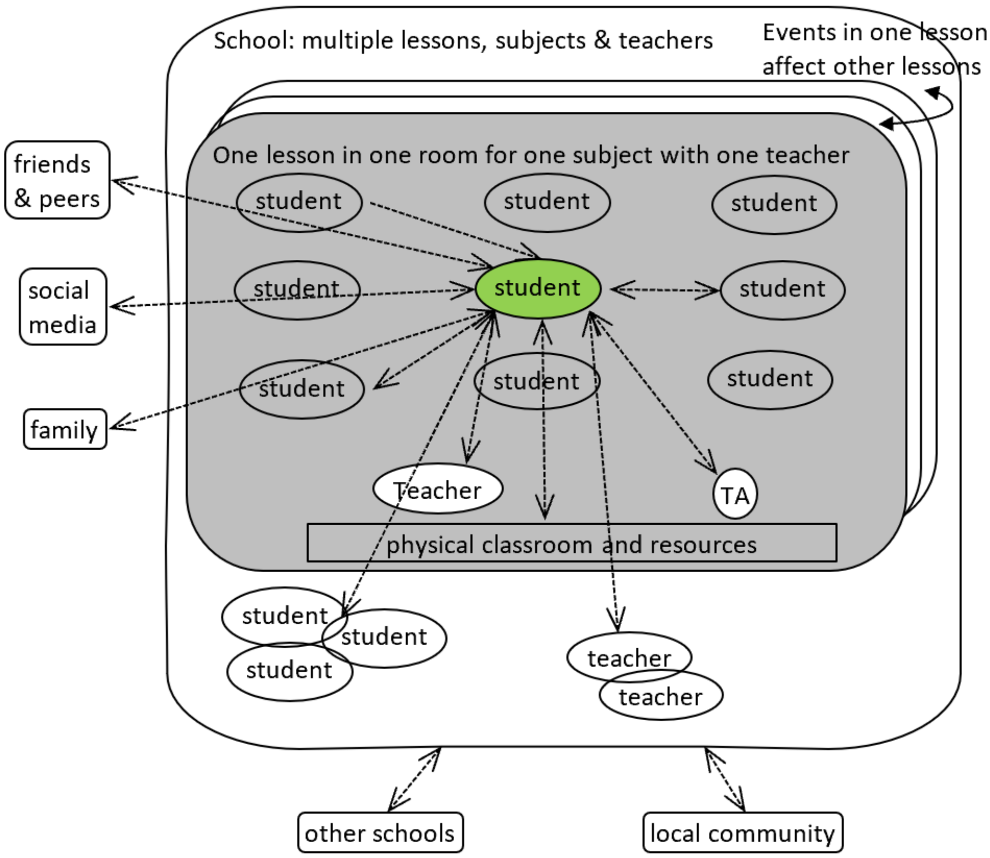

Figure 2-1 The influences on one student in the classroom lesson sub-system

The following list contains some of the influences on lesson participants:

- The school, with all the staff and all the other students is one eco-system in the local community of eco-systems. There are other schools, clubs, organisations etc. that all influence the school, the school staff, the students and their families.
- The school is an eco-system in which classroom lessons are situated. There are other classes, students, teachers, several subjects, clubs and sports, after-school activities, meals, other buildings, uniforms, etc.
- Classes are subsystems, collections of individuals, each with his/her own characteristics. The students, most of whom know each other and have formed into friendship groups, have often been together for several terms.
- Classes are entities. Teachers speak of classes having an individual ethos or quality ('an environment with its own ecology' (Parsonson, 2012, p. 16)). For example, one class may be extremely difficult to get a response from, while another may be so bubbly that it requires constant management.

- The participants may come from a wide variety of social backgrounds, or they may all be very similar. Some of the students may have special educational needs that should be catered for in every lesson.
- The participants need to adapt to internal influences (e.g., tiredness, emotions) and external influences, those that they are experiencing in the present (e.g., the teacher asking a question or a student chatting to them) and those from the past (e.g., what happened during the previous break or lesson, or what happened at home yesterday).
- Students decide how to behave, whether to follow the teacher's instructions or chat with their neighbours or play on their phones, for example. They may be driven by their own values, or by pressure from peers, parents, or other adults.
- The nature and extent of student misbehaviour often seems unrelated to the activities of a lesson and is often triggered by external factors. For example, a student might behave poorly after spending a weekend with the other parent in a separated family, or after spending social time with other, rowdier, students, or having been reprimanded and warned by the teacher in the lesson before.
- The students are in a classroom, possibly fixed for all lessons for one year, or possibly just for one subject. The classroom is probably one of many in the school, each with its own arrangement of fixtures and fittings, desks and chairs, posters etc. Perhaps one subject is always taught in one particular room, or one teacher always holds lessons in one room.
- The classroom will have physical properties, e.g., tidy/messy, warm/cold, quiet/noisy, spacious/cramped, bare/covered in displays, etc. In other words, the room itself has a character that influences its occupants.
- The seating arrangements have an enormous influence on behaviour in lessons, because they increase or decrease certain student interactions (e.g., working productively together or chatting) and possibly student-teacher interactions (such as making one-to-one support difficult because there is not enough room for the teacher and the student might feel awkward about coming to a separate desk with the teacher).
- Students reported that they felt different and interacted differently according to their seat position in a lesson. Sociable students, when seated by a wall, often faced into the class and socialized. More introvert students sometimes felt less exposed when next to a wall, as it cuts out half the goings-on of others - so many may feel more secure and able to focus productively.

- The external environment can also affect lessons, e.g., noisy building work outside or another class making lots of noise.
- The lesson is one of many the students attend, day after day, term after term, year after year. There are routines on the level of a lesson, a day, week, term and year. Students and teachers will have formed habits of behaviour.
- There are also infrequent or unexpected events that might alter lessons, e.g., fire drill, visitors, celebrations, leaving day for the final year students, a special assembly or presentation, a birthday.
- Teachers differ and students have different relationships with each one, ranging from receptive and cooperative to antagonistic and defiant. Students see different teachers as having different attitudes towards and tolerances for chatting, writing tests, asking questions, amount of help, etc.
- Lessons in different subjects differ, e.g., science lab work vs art classes vs historical role play vs mathematics investigations. Teachers typically plan a variety of activities, some of which a student might enjoy more than others, or be more competent and successful in, for example whole class teaching, working alone, working in pairs, group work, competitive or cooperative activities, exercises or investigations, more quiet work or much discussion.

The points above are just some of the influences on lesson participants. Many of the characteristics associated with complex systems (Abrahamson, Blikstein and Wilensky, 2007; Blikstein, Abrahamson and Wilensky, 2008; Burns and Knox, 2011) are evident in classroom lessons (Larsen-Freeman, 2016), such as the sudden transitions or tipping points that can occur (e.g., in the students' attitudes to the teacher after some disciplining), or feedback loops such as a student-teacher-student interaction or when one student affects other students and is in turn influenced by them. But classroom lessons are also examples of 'organized complexity' (Weaver, 1948). Since it is likely that the participants in lessons adjust their behaviour (adapt) in response to the actions of the other participants and past interactions, we might also consider classroom lessons as examples of Complex Adaptive Systems, systems in which multiple components interact (hence complex) in such a way that the behaviour of the components changes in response to events (hence adaptive).

2.2 The structure of classroom lessons

Most people are familiar with the traditional school education system: sitting in lessons with a group of peers in a classroom with a teacher purposefully guiding the lesson. Often a TA is present, to provide extra assistance to one or more students. A common lesson in a UK secondary school might have 20 to 30 students, a teacher and TA, desks, chairs, a whiteboard where the teacher is projecting material that is being discussed, the walls have various posters and displays of students' work, etc. Lessons are dynamic processes and what happens in them varies enormously. The details depend on what is being taught, the skill levels of the students, the activities planned, the teacher's teaching style, the resources available, etc. A typical lesson may have the following structure and activities.

After taking the register sometimes there will be announcements. Homework may be collected and possibly resources distributed. Once everyone is settled – and this will often require some *behaviour management* by the teacher – the teacher begins. Most schools have procedures for managing and monitoring student behaviour. The teacher will be expected to follow the procedures (it causes problems when teachers are inconsistent) although what is acceptable does vary from subject to subject and teacher to teacher and activity to activity.

It is commonplace for teachers to have a lesson plan. A central part of lesson planning is providing multiple activities and learning objectives to satisfy all the ability levels and academic levels that are present in the class. The teacher will also construct or obtain resources to aid learning and for assessment. Each lesson, the teacher tries to follow the lesson plan while responding to the students and circumstances. The timing of events is not rigid and teachers make frequent, mental, on-the-spot adjustments to the lesson plan.

During lessons, the teacher expects students to be productive, either listening or following instructions. Students either comply with the teacher's instructions or behave in ways that the teacher considers disruptive or disengaged. The teacher intervenes as necessary and according to their nature and the circumstances. Intervention can also mean spending more time revising presumed knowledge or more time helping individuals or handling an outside interruption or technical problems with equipment.

Depending on the lesson plan (which often reflects the teacher's teaching style) the teacher may spell out exactly what will happen and the specific learning objectives, or may encourage the students to investigate and discover principles. The lesson could begin with a recap or some assessment.

Teachers instruct students to do something specific, such as to work on an activity alone (without interacting with others, e.g., writing an essay), or to work in pairs (which might be their neighbour, but could mean swapping seats to move further away and will normally involve talking), or work in larger groups (again, this may involve seating changes or even furniture reconfiguring, also talking, moving around, etc.). The learning activities could be cooperative, competitive, or neither.

During student activities, the teacher and TA (if present) discreetly observe the students to ensure everyone is engaged/participating. This may involve moving between the desks/tables. If there appears to be a need, the teacher or TA will offer assistance, either to an individual or to a pair or a group. This could be declined or last just a few seconds or many minutes.

At the end of an activity (section of the lesson plan), the teacher may get the attention of the whole class and instigate discussion to bring out students' experiences. There could also be a minor assessment (verbal, as a group, or a brief, individual, written question) to monitor progress.

At the end of a lesson the teacher will often conduct a plenary discussion to summarize what has been learnt and to provide students with an opportunity to explain/teach some point and reflect on their experiences, etc. There could also be an end-of-lesson assessment. The teacher may well describe what they will be doing next lesson, perhaps to stimulate interest.

Finally, the teacher could set homework, which might involve personal interactions with specific students. The teacher will then dismiss the class.

The above description gives an overview of a common lesson scenario and indicates some of the facets that a model would need to capture.

2.3 Relevant lesson activities

Part of the design of an ABMS is the specification of the simulation outputs and the agent states or activities. Some of these outputs and states had been partially pre-selected by the objectives of the research and the research questions. For example, some of the objectives required simulating the students' and the teacher's productive and other behaviour and measuring lesson productivity and misbehaviour.

These choices reflected the constant and pressing interest in classroom behaviours evident in educational research and expressed by teachers (Marzano, Marzano and Pickering, 2003; Petty, 2006). Disruptive disengaged behaviour, particularly low-level disruption, is a concern across UK schools as teachers feel that these behaviours are very detrimental to student learning and to the well-being of teachers (Brouwers and Tomic, 1999; Lugrin *et al.*, 2016). Brouwers and Tomic (1999) defined behaviour as disruptive when 'the student in question is not engaged in a task structured for him or her by the teacher and when this behaviour is noticed by and/or interferes with the efforts of fellow students' (Brouwers and Tomic, 1999, p. 9). The teachers in the case study believed that the effect on a lesson of a single student misbehaving was far greater than that of exemplary student behaviour. Hence it was decided that a variety of engaged and disengaged behaviours would be modelled and that data would be needed to quantify the frequencies and durations of these behaviours.

From the literature various pertinent lesson activities were identified in advance of the case study. For example, regarding the teacher it was thought important to capture how often, how quickly and for how long they disciplined students (Poplin *et al.*, 2011; Baumann and Krskova, 2016). For students, potentially relevant activities included:

- being on-task: listening to instructions, doing what the teacher instructed, seeking help in the proper manner (e.g., raising hand), being in the proper location (e.g., being in own seat when 'seatwork' is required); being off-task: not doing what the teacher instructed, not in the proper location (Allday *et al.*, 2011);
- behaviours that could be classified as: engaged, passively attending, transition (e.g., preparation), non-productive (e.g., fidgeting), inappropriate, attention seeking, resistive (non-compliant), aggressive (Bidwell and Fuchs, 2011);
- various types of on-task behaviour depending on whether the students were in their seats, talking, raising their hands or using their response cards (Christle and Schuster, 2003);
- the time consumed by copying from the board, writing tests, random unpredictable distractions, reorganizing the classroom, managing books and handouts (Dillenbourg, 2013);
- 'engaged time', also known as 'time-on-task' (Carroll, 1963; Cotton, 1989).

The California Beginning Teacher Evaluation Study (BTES) (Fisher, 1978) and other studies reported that the amount of time students spent on-task depended on many factors, including the relative position of the teacher to the students, how easy it was for the teacher to monitor students, the norms for lesson behaviour, the teacher's behaviour management, the teacher's organization, how much student-teacher interaction occurs during whole-class teaching and the amount of time students spent working alone in their seats ('seatwork') or with others (Stallings, 1980; Karweit, 1984; Godwin *et al.*, 2015). Whilst acknowledging that the time a student spends on-task as determined by an external observer may be a fairly crude measure of academic learning – the presumed focus for the teacher – it seemed that teachers nevertheless strived to maximize this time (and minimize time off-task) (Gettinger and Seibert, 2002).

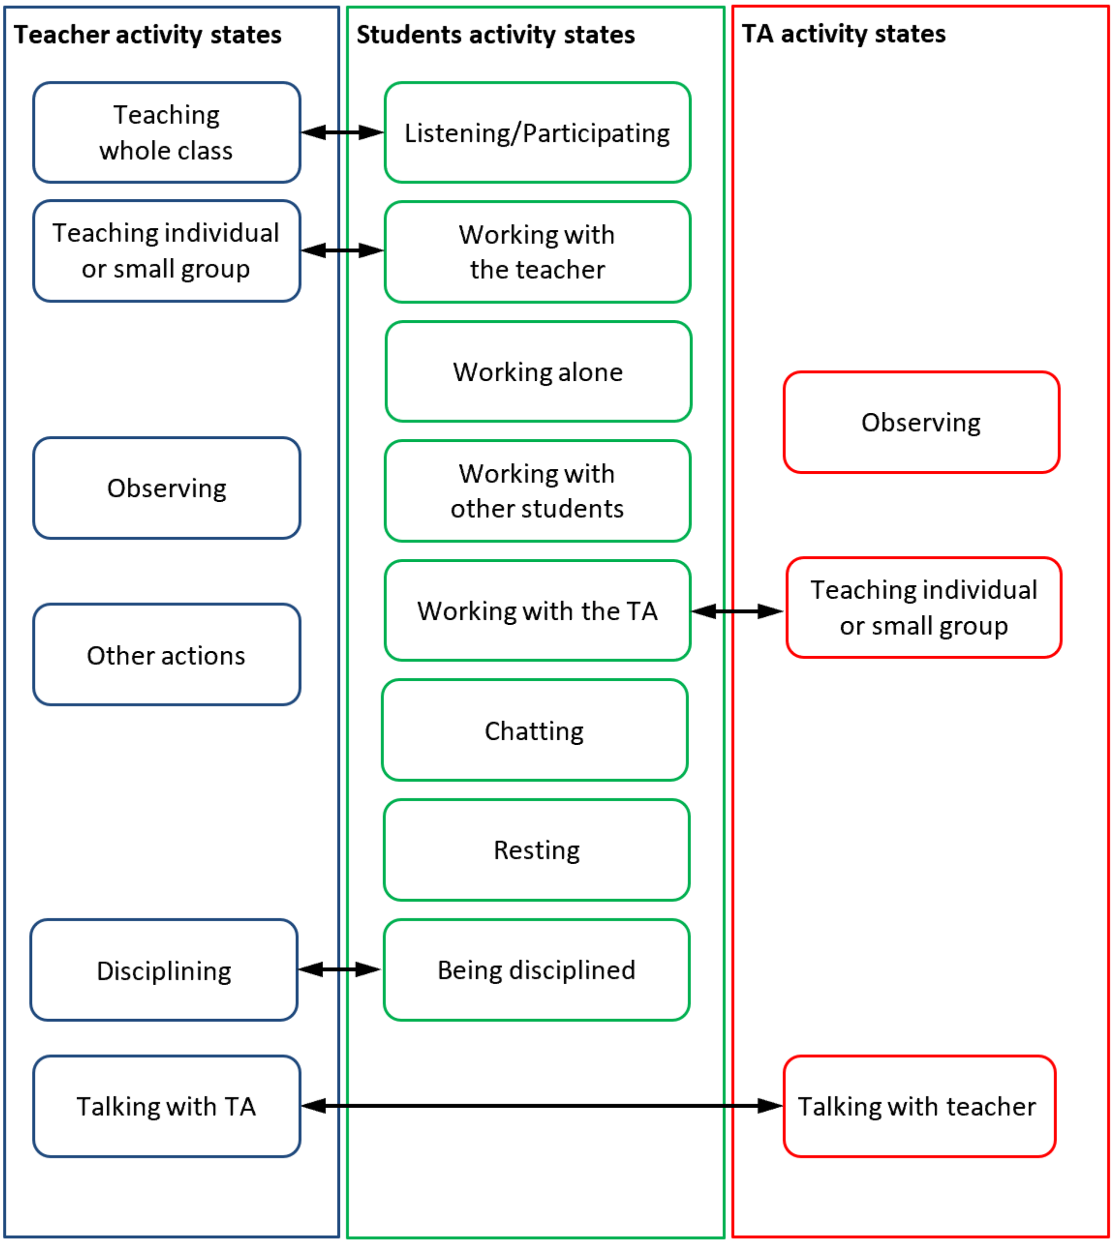

Figure 2-2 Some activities of, and interactions between, lesson participants [10]

All of the above activities were considered important and were also observed during the case study. Some states could be entered by choice (e.g., a student could choose to chat). Some states were forced on a participant and could not be entered by choice (e.g., a student could not be being reprimanded by choice – the teacher forces that state on a student). Some states could be either chosen or forced. For example, a teacher offering help might force a student into the

[10] Note that this is not a 'use case' diagram and it is not meant to convey any sequence of interactions or events.

corresponding state, but a student asking for help might force the teacher into the corresponding state.

The lesson observations led to an understanding of the variety of interactions that occur in lessons. Figure 2-2 shows a few of the observed interactions between the teacher, TA and students, with student-student interactions omitted for clarity. The arrows indicate situations where two people are interacting and where one person could not be doing this activity unless the other person was involved. The diagram does not show transitions from one activity to another. The complete list of selected activity states is given in Chapter 3.

2.4 Teacher and student decision-making

The core of an ABS is the logic that describes how an agent chooses what to do at each time instant. The goal of the classroom lessons ABMS was to model the dynamic, spatial, asynchronous, autonomous, interactive decision-making that happens in lessons. But, because teachers, TAs and students play such different roles, with different choices and consequences, it was expected that each would need a different sub-model.

There were many existing models of decision-making available (Balke and Gilbert, 2014). For example, the Belief-Desire-Intention (BDI) model has been very successful (Norling, Sonenberg and Rönnquist, 2000; Gibson, 2009; Norling, 2009; Belhaj, Kebair and Ben Said, 2014; Padgham, Singh and Zambetta, 2015). In this model, agents are considered to be intelligent decision-makers who plan for the accomplishment of a complex task based on their beliefs, desires and goals (intentions), and then carry out the planned actions. The teachers did have a goal, yet this seemed to be very simple: follow the lesson plan, handling any interruptions. (One could also claim that a teacher's goal is to maximize student learning or increase academic grades, but that would be the objective that the lesson plan was designed to achieve.) There seemed to be no need for a sophisticated goal-driven model for decision-making in lessons. However, the manner in which a teacher responds to an interruption (such as class or individual misbehaviour or disengagement, or technical problems, or deciding to modify the lesson plan to give more time for revision) did seem more complex, and, most importantly, needed to be rapid and on-the-spot. To cope with these common-place scenarios, a 'recognition-primed decision' (RPD) model (Klein, 2008) was appropriate as it was developed to model how expert leaders make a 'situation assessment' (Norling, Sonenberg and Rönnquist, 2000), taking quick decisions in complex situations under time pressure (such as when firefighting or during military engagements). The main principle is that an expert would rapidly consider alternative courses of action, choosing the first alternative that matched the current situation. This is effectively how teacher decision-making was modelled, but using the most basic RPD approach:

- look through a fixed list of rules that are ordered in decreasing order of gravity (i.e., their disruptive impact on the lesson);
- apply the first rule whose conditions are satisfied.

This is a common ABMS approach to modelling agent decision-making. A primary assumption in any ABMS is that the behaviour of the actors can be emulated formally using mathematical symbolism and/or a computer programming language, with the inclusion of some pseudo-random (stochastic) fluctuations to make behaviour more life-like (slightly random and unpredictable). Agents look through a set of *production rules* (also known as a *ruleset*) in order to choose an appropriate course of action (Urban and Schmidt, 2001; Balke and Gilbert, 2014). During the case study, observations about what happened in lessons were noted and plausible rules or principles that might be at work were formulated (c.f. 'causal rules' from 'causal stories' (Bex, 2015)). These could be called *stylized rules* (c.f. *stylized facts*) that could be configured as a multitude of single *if A then Z* rules, or a set of more complex *if A and B but not C then Z_1 and Z_2* type rules. The choice was a major design decision between a collection of simple cause → effect rules, several of which could be triggered individually at each simulation step, versus highly complex rules that more fully describe mini scenarios and only one of which is activated. The latter more RPD-like option was selected. This *production rules* approach was adopted for all three lesson participant types, teachers, TAs and students and the observations from the case study were translated into plausible rules and logic for behaviours. However, the way student agents process the student ruleset is different from the simple method just described for teachers (and TAs) and is explained in section 3.3.2.

During discussions in the case study, students did explain a variety of goals for a lesson, for example: have fun with friends; do well, be productive and learn; please the teacher; survive the lesson without getting into trouble (particularly when a student had been disciplined and given an ultimatum). However, these all seemed to play a secondary role compared to reacting appropriately to the current situation in the lesson, especially to what the teacher and the other students were doing. The students seemed to respond to the current situation and behave according to their habits. To model such decision-making, ideas from the well-established Theory of Interpersonal Behaviour (TIB), formulated by Triandis in 1971 (Jackson, 2005; Li and Lee, 2013; Moody and Siponen, 2013), were incorporated. Behaviour is considered to be driven by the intention and the habits of the person but moderated by the prevailing conditions. It explicitly acknowledges the influence of past behaviour, or habits, on the decision process. This feature was thought to be highly appropriate in lessons, where habit and routine play a major role in behaviour choices. To implement this feature for students, habits were represented by the historical relative frequencies of activity states. This essentially meant that students were given attributes (characteristics or properties) to represent the activity states they could engage in. For example, there is an attribute representing the tendency for a student to listen to the teacher during whole-class teaching. The value of an attribute at any particular instant represents the relative likelihood that the student would choose that activity compared to other activities. It was felt that this use of activity attributes, rather than generic attributes such as personality traits, made the model more specifically a lesson behaviour model and more likely to be implemented successfully.

Below are some examples of case study observations that were incorporated in student rules:

- it appeared that students tended to copy the behaviour of their peers;
- somehow students needed to cope with contradictory pressures, such as wanting to follow the teacher's instructions but also wanting to do what their friend is doing (e.g., chatting or doing nothing);
- when a student was reprimanded this appeared to have the effect of reducing misbehaviour, although the effect seemed to wear off and students may again disengage;
- students appeared less inclined to misbehave when the teacher was close by - hence the students' behaviour was considered to be affected by the proximity of the teacher;
- it also seemed that students were less inclined to misbehave in smaller classes, when the teacher-student ratio was greater.

To summarize, all agents interact autonomously, adapting (changing) their behaviour according to their current situation. In terms of the definitions offered in Macal (2016, p. 150), the ABS is an 'interactive' ABS, with the most basic 'adaptive' abilities due to 'memory' of when events last occurred. Teacher behaviour is characterized by the logic: follow the lesson plan, handling any interruptions. Student behaviour is considered to be driven by the current situation and their past habits, with the current situation comprising mainly:

- what the student is supposed to be doing, i.e., what the teacher has instructed - based on the lesson plan;
- what the student is currently doing (e.g., attempting to chat to another student) and for how long they have been doing this;

- what the teacher and TA are currently doing (possibly with that student);
- what other students may be doing to the student (e.g., attempting to chat);
- recent interactions with the teacher (e.g., being disciplined, this lesson or previously).

Some examples of the rules are given in section 3.2, with complete rulesets included in Appendix A. Specific details of where stochasticity was added are also given in the appendices.

As highlighted in section 1.1, the issue of classroom behaviour management is particularly relevant. For this reason, this aspect of the model is explained in more detail in the following subsection.

2.4.1 The interaction between student misbehaviour and the teacher's disciplining

It was apparent during the case study and is generally well acknowledged that students respond to the disciplinary actions of the teacher. Two types of disciplining were observed: individual and whole class. The teacher intervened in both cases, but the mechanisms appeared different. For example, it might have been that at some point in the lesson some students were not engaged in the learning activity (e.g., listening to the teacher). Typically, when a sufficient number of students were involved, the teacher would switch into a whole class disciplining mode, apparently addressing the whole class but actually just the subset who needed the message.

Besides catering for individual and class disciplining, it was considered important that the model incorporate low-level disruption (persistent, sub-critical, distracting behaviour). This required noting misbehaviour over a period so that it gradually built up and could be responded to. This required some sort of misbehaviour count and time period.

It also seemed that different students responded differently: some seemed more compliant and followed the teacher's instructions while others can become more argumentative or disengaged further and participated even less. It was not clear, therefore, whether disciplining always increases productivity and/or decreases disengagement and in what situations and to what extent.

Nevertheless, based on the lesson observations and discussions with the teachers, it appeared that the following behaviour was common:

1. Some teachers were stricter than others, addressing misbehaviour earlier than their colleagues would; other teachers used more light-hearted responses to get students to participate. Also, this depended on the subject, so, for example, more chatting was acceptable in art lessons than in mathematics lessons.

2. Teachers seemed to discipline proportionately, the more disruptive a student's behaviour (e.g., shouting across the whole classroom), the sterner the discipline. But the teacher would also discipline students who disengaged passively (e.g., not being productive by fiddling excessively with things, or breaking rules such as eating or using their mobile phone).
3. When the teacher was observing the class or whole class teaching, students misbehaved less frequently than when the teacher was helping someone or talking to the TA.
4. The students were even less likely to misbehave when the teacher was currently telling someone or the whole class off.
5. Students appeared to misbehave less when the teacher was nearer them than when the teacher was further away.
6. After disciplining, behaviour was more settled and possibly more productive. However, this seemed to wear off and some misbehaviour can reappear, earlier in some students than in others: this response is assumed to depend on some characteristics of the student.

Details of how these observations were incorporated into the simulation model are provided in section 3.4, with further details in Appendix A.3.4.

2.5 Representing the classroom and lesson interactions

The events that occur in lessons are affected by many factors associated with the physical classroom, for example, the shape of the room, the position of cupboards, windows and doors, desk size and arrangement. In particular, student seating arrangement is known to have a significant impact on the amount of time students are 'on task' and 'off task' or disruptive (Wheldall and Lam, 1987; Schwieso, 1995; Bicard *et al.*, 2012; Kaplan, Gheen and Midgley, 2002). This was also observed during the case study. It was therefore considered essential that the model incorporate spatial aspects of classrooms to some extent. This meant that every classroom lesson model needed a floorplan, with desks correctly arranged and the students correctly located, to ensure realistic visualization and to enable the simulations to take spatial distances into account in agent interactions. Several other classroom research projects have used 2D representations of classrooms, complete with seats, desks and aisles and the locations of students and the teacher. Figure 2-3 shows one example; other examples can be seen in Holstein, McLaren and Aleven (2017b, p. 3) and Raca and Dillenbourg (2013, p. 266). Figure 2-4 shows the classroom layouts available in the model from McDevitt (2017).

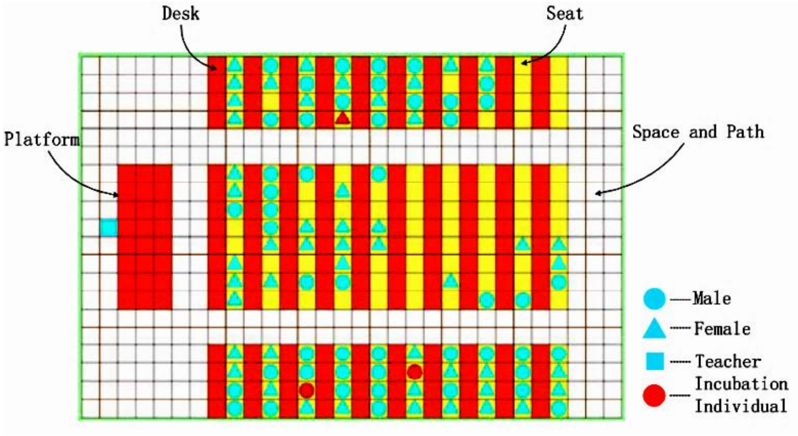

Figure 2-3 A 2D model of a classroom lesson layout (Ge *et al.*, 2011, p. 96)

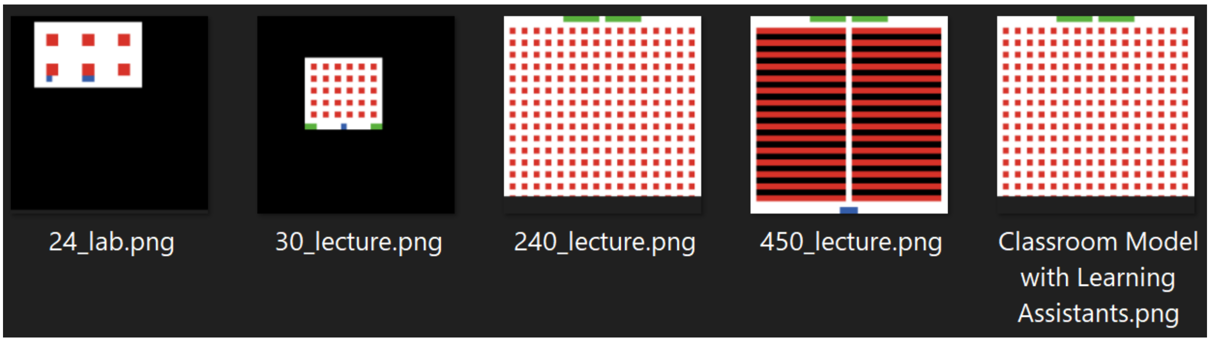

Figure 2-4 The classroom layouts available in the model from McDevitt (2017)

It was also considered essential to have visual representations that showed what was happening in dynamic detail during lesson simulations whilst also providing a general overview (Hmelo-Silver, Liu and Jordan, 2009; Dorin and Geard, 2014; García-Magariño and Plaza, 2015; García-Magariño *et al.*, 2017). Visualizations are especially useful when experts evaluate model outputs, providing face validity assessment (explained in section 6.3) (Wilensky and Rand, 2007; Klügl, 2008). The chosen representation is described in section 3.5, but an example lesson layout is shown in Figure 2-5. The blacked-out areas indicate either walls or cupboards, anywhere not accessible; a door is indicated by a white bar; desks are brown blocks; the whiteboard is a magenta line. A classroom floor plan is converted into a full lesson layout by adding the lesson participants in their correct places. The teacher (dog-like icon), the TA (cat-like icon), the observer (ant-like icon) are shown in their primary positions. Some of the students (the bear cub icons) are slightly transparent indicating that their activities were not fully monitored. The labels on students show the unique student reference/current state number.

The lesson animations also proved an important tool for 'visual debugging' (Grimm, 2002) of the simulation code.

Figure 2-5 An example classroom lesson layout
showing students at their desks, with the teacher at the front and the TA to the side

2.6 Assumptions and simplifications

From the comprehensive conceptualization of classroom lessons, a full conceptual model was formulated, which was itself transformed into the final simulation model. As is usual when formulating an abstraction of a system, some details of the real-world system were omitted and simplifying assumptions were made (Hillston, 2003, 2017; Jackson, 2005). However, in order to make the simulation as realistic as possible, the model included as much of the real-world complexity as was possible. This approach is in keeping with the 'keep it descriptive stupid' (KIDS) approach described by Edmonds and Moss (2005). They pointed out that it might be more effective to start with a relatively complex model that more accurately describes the real system, rather than the *keep it simple stupid* (KISS) approach which advocates starting as simply as possible.

The simplified component diagram in Figure 2-6 shows the components of the classroom lesson system model. As was indicated in Figure 2-1, many factors that probably affect participants' behaviour have been excluded, such as the influence of parents and what happened in earlier lessons.

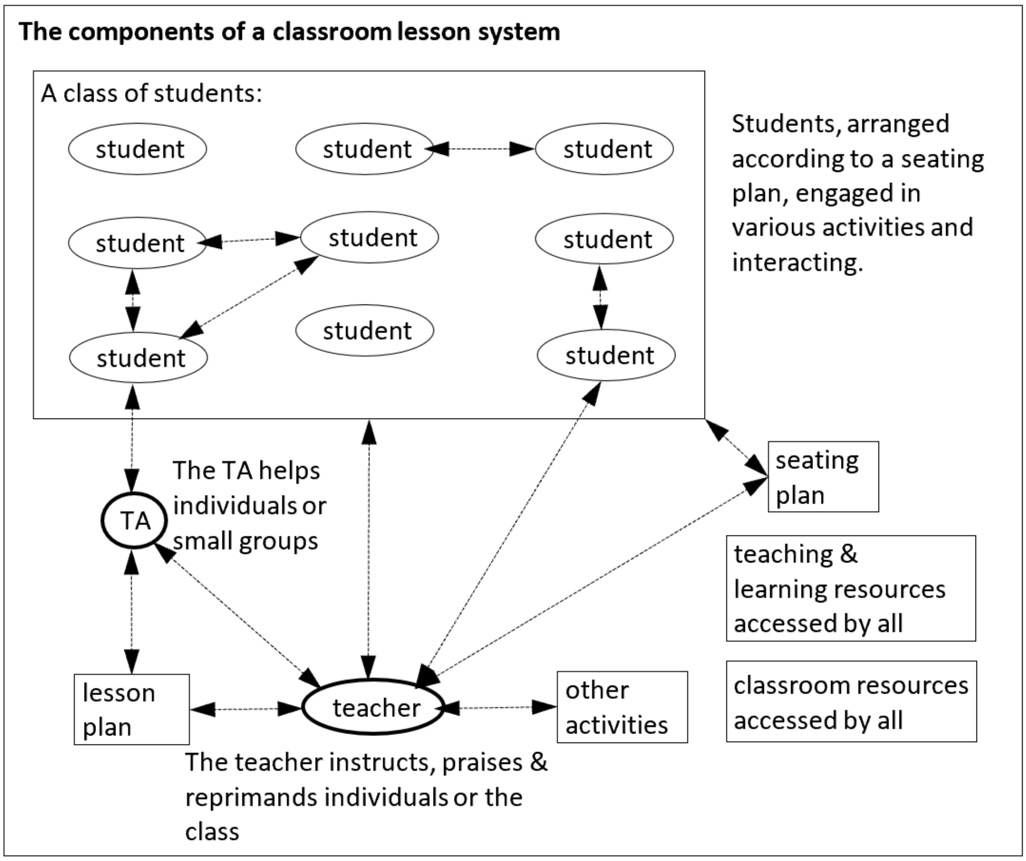

Figure 2-6 The components and interactions included in the classroom lesson system

The following decisions and assumptions were made:

1. The system being modelled is one classroom lesson, not a series of lessons.
2. The goal of a classroom lesson is assumed to be to increase the academic achievement of each student and the class overall. This is achieved by maximizing the amount of time spent productively and minimizing the amount of time spent disruptively.
3. The lesson plan is assumed to be the teacher's solution to maximizing productive time, hence the teacher's objective is to adhere to the lesson plan as closely as appropriate.
4. The teacher does not alter the lesson plan during the lesson.
5. A classroom lesson is considered to be the collection of activities performed in various locations in a classroom by a teacher, a TA (if present) and the students of a particular class while the teacher follows a lesson plan, handling any interruptions.
6. The components of a classroom lesson are the classroom (its contents and layout), the student seating arrangements, the human participants/actors and the lesson plan. All other objects (e.g., teaching and learning materials, equipment) are considered part of the environment of lessons.
7. A class is a collection of individual students identified by the school as a separate group.

8. A class is not modelled as an entity in its own right.
9. A class attends a lesson, which is run by one teacher teaching one subject in that lesson.
10. Each person has characteristics, some internal (private), some external (that other people can observe).
11. All these people interact autonomously and are always aware of each other's external characteristics, including location.
12. All lesson participants know the current activity state of all others at all times.
13. All students are always sensitive to the teacher's proximity and strictness and remember when they were last disciplined or praised.
14. The participants occupy a typical classroom, with walls, door(s), whiteboard(s), desks, chairs, a bin, etc. in any arrangement. Desk and aisle width are also important as people's behaviour can be affected by the location of and distances from other people.
15. Students have a seating position at a desk (of a specified seating size) but can be reseated during a lesson.
16. The students are free to interact with each other, with the teacher, with the TA, and with the classroom (e.g., the bin, leaving the room) at any time.
17. The classroom is small enough that any student can choose to interact with any other student, not just their immediate neighbours.
18. Students can leave the classroom, but there are restrictions, such as one student at a time and maximum number of departures.
19. Students are in the 11 to 16 age range (UK secondary school), meaning that a narrow range of teacher and student behaviour is being considered.
20. The location of lesson participants is explicitly represented, but their motion is not: for simplicity participants can change positions, but they move from one location to another instantly in one simulation time step. For example, the teacher can be helping a student on one side of the classroom in one time interval and in the next time interval be immediately at the front of the classroom.
21. There is only one TA in a lesson.
22. Although the TAs were seen to discipline students, this was not modelled.
23. Where a TA is present, the TA is considered to either respond to requests for assistance from any student or to offer assistance to any student, i.e., the TA is not allocated to just one designated student.
24. The TA first responds to the teacher, then the students and only then initiates independent action.

25. The TA always accepts an interaction proposed by a teacher or student.
26. The TA does not choose to interact with the teacher.
27. The only factors involved in the TA's decision to offer help are the state of the student and the time since they last had help from the TA - the TA does not consider factors such as which student has had the least help or a student's academic ability.
28. After giving the students instructions (e.g., to work on an exercise), the teacher and TA wait a few minutes before offering help, to give the students time to get organized and start on the activity. (This assumption is not always realistic as in many instances a teacher or TA will go to specific students immediately to confirm that they understand what to do.)
29. If the teacher is interrupted while helping a student, instead of handling the interruption then resuming helping, the help is terminated by the interruption and the teacher enters the default state specified in the lesson plan.
30. The teacher never suspends a lesson completely.
31. While the student can seek 1-to-1 assistance from the teacher or the teacher can offer it, only the teacher can conclude that assistance.
32. Agents have simple 'memory' of recent events and when they occurred.
33. Agent reasoning is basic condition-response logic: there is no reasoning about what is known, or what other agents might be considering (e.g., is the teacher looking at me?), or meta-logic that changes rule logic.
34. Agents consider only the current state of the other agents, not their past states. For example, before proposing to chat, a student agent does not consider whether another student has just been disciplined.
35. The model does not include well-known factors that affect student behaviour, such as students disengaging due to fatigue or lack of success, or being invigorated by some activities.
36. Because the empirical data collection did not enable the instigator of an interaction to be recorded, interactions were considered to be directionless, as explained in section 1.5.2.

2.7 Chapter summary

This chapter described the results of the first two stages of the ABMS development process. By combining understandings obtained from published research with the first-hand observations during the case study (and the author's personal experiences of 14 years of teaching), a comprehensive conceptualization of the classroom lesson system was set out, including descriptions of the typical structure of lessons and the activities in them. After consideration of the types of decisions teacher and students make, it was decided that decision-making would be modelled using 'production rules'. From this analysis a comprehensive conceptual model of behaviour in classroom lessons was formulated. The emphasis was on producing a realistic rather than simplistic model of lesson behaviours, despite this being complex and complicated. The focus of the model is lesson behaviours, highlighting particularly the interaction between student misbehaviour and the teacher's disciplining. The chapter included an explanation of how classrooms and lessons would be represented on a computer, and listed some of the assumptions and simplifications made.

The following chapter explains the results of stage 3 of the ABMS development - the simulation model.

3. The Classroom Lessons Simulation Model (CLSM)

The next stage in the ABMS development was to take the conceptual model of classroom lessons and transform it into a simulation model. This chapter is a summary description of the simulation model - additional details are provided in the appendices.

In the model there are three types of agents: student, teacher and TA, each having its own algorithm formalizing its decision-making. One can view the classroom lessons behaviour model as comprising three interacting sub-models. Figure 3-1 shows some of the ways the agent sub-models interact. The rectangles represent factors that either influence the agents or are affected by agents. The hollow arrows indicate bidirectional influences.

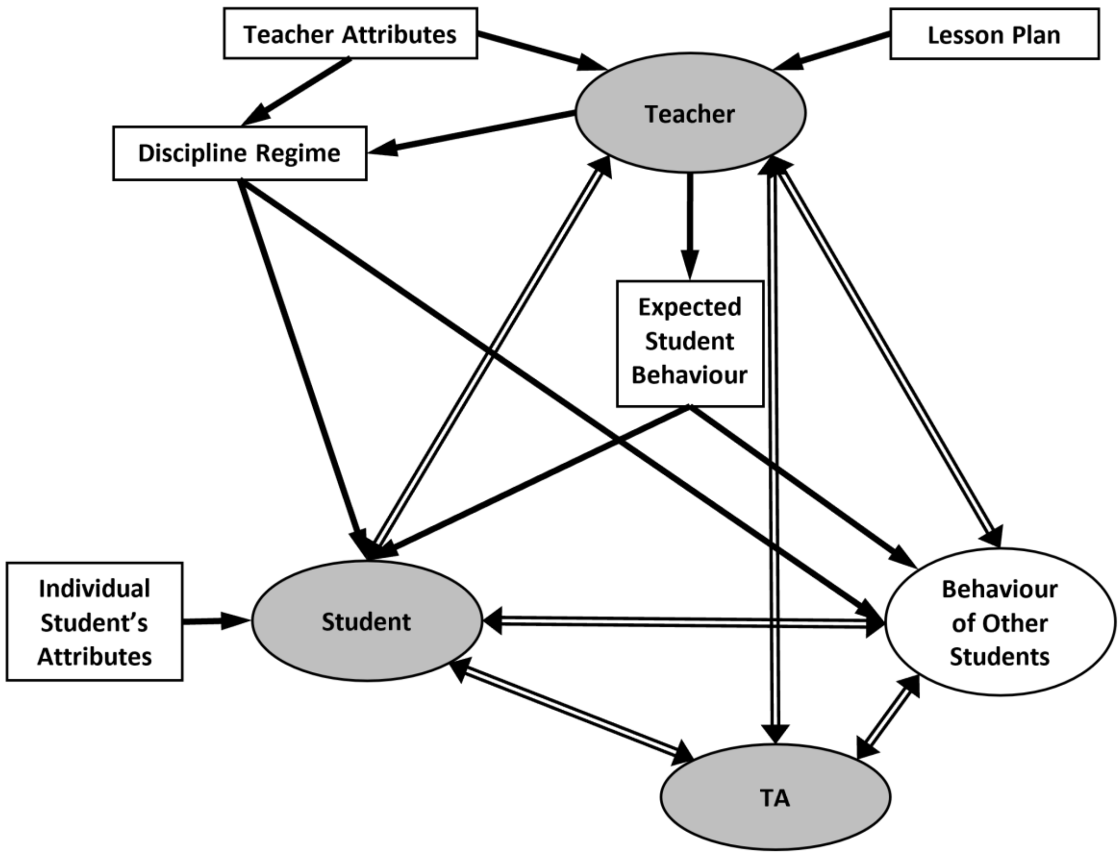

Figure 3-1 The principle interactions between the three agent types

In order to run simulations of lessons, the following inputs are needed:

- a classroom layout and student seating layout;
- a class of students each with their historical behaviour data;
- a lesson plan (which divides the lesson into sections of specific student and teacher activities, as explained in section 3.2);

There are two types of simulation outputs:

- animated visualizations of the dynamics of a lesson, including animated charts showing individual agent information and overall lesson information;
- for each agent a time series of activity states, plus summary statistics for each agent and for the lesson overall (such as the percentage of productive time and disruptive time).

This chapter explains:

- the lesson activities that were selected for the agent states;
- how lesson plans control what people do at different times in lessons;
- how agent behaviour and decision-making is represented using 'production rules', and introduces the model parameters;
- how the intricacies of agent interactions are managed, with particular emphasis on how the students and the teacher interact when it comes to misbehaviour and discipline;
- the implementation of the model in NetLogo.

3.1 Agents activity states

Each agent type has a set of mutually exclusive activity states, shown in the tables below. The colour-coding shows which states were considered Productive, Disruptive/Disciplinary or Other (other non-productive/non-teaching activities). For students, 'disruptive' behaviour is behaviour that distracts another student from the expected behaviour (e.g., chatting), whereas 'other' behaviour is non-disruptive disengagement, where the student does not (directly) interact with another student (e.g., fiddling with something). As explained in section 2.3, some states are not entered at the agent's volition but are 'forced' by other agents or circumstances. These are indicated via a solid underline. Some states can be entered either voluntarily or by 'force'. These are indicated via a dashed underline.

As explained in section 2.4, the student state attribute values at the start of a lesson are the historical relative frequency of the activity states over all lessons. These attribute values are adjusted during the lesson so that each state has a new relative likelihood. This is explained further in section 3.3.2.

The following tables summarize the agent states modelled.

Table 3-1 The teacher activity states modelled

Abbr.	State	Teacher activities
SDIS	1	Disciplining specific student(s). This is not restricted to telling off as it may involve reasoning, exhorting, encouraging, referring to student's strengths, responsibilities etc.
	2	Unused
	3	Unused
IRRE	4	Deliberate activities clearly unrelated to this lesson, e.g., creating display for another class while students are working
CDIS	5	Managing the whole class, e.g., 'Take your books out', shushing, homework telling off, disciplining (but not individual disciplining)
PREV	6	Prevented from teaching, e.g., by announcements, school matters, external interruptions, technical problems
COBS	7	Observing students (while teacher is moving around class or stationary); default mode for teacher if not in any other mode
SOLO	8	Working alone, busy with other lesson-related activities
PWCT	9	'Passive' whole class teaching, e.g., showing a video clip
Q&A	10	Asking questions and waits for answers; listening to a student's answer
SPRA	11	Appreciating or praising a student in front of the class
AWCT	12	Active teaching of whole class, e.g., class discussion, active demonstrations (e.g., software or interactive video), handing out and collecting in materials
SSUP	13	Offering or providing individual support, encouragement, or instruction
GT	14	Offering or providing support to a group
TTA	15	Talking to the teaching assistant

Abbr.	State	Teacher activities
OUT	16	Left classroom
NONE	17	None of these, e.g., attending to personal matter

Table 3-2 The student activity states modelled

Abbrev.	State	Student activities
DISC	1	Being disciplined by teacher or teaching assistant
MESS	2	Unproductive, away from own desk and distracting others, e.g., chatting, fiddling with things
CHAT	3	In own seat chatting, distracting, socialising, turning around etc.
NOT	4	Intentionally unproductive, not participating, but not distracting others, at or away from own desk (e.g., messing about at a bin or fiddling in bag or doodling, breaking rules such as eating or using mobile)
CLOFF	5	Listening to teacher in response to class being told off; this response by the student is considered good, but non-productive behaviour
PREV	6	Unintentionally unproductive: not learning, but not disturbing others; e.g., getting out books before start of lesson or packing away at the end, doing a job for teacher, technical problems
REST	7	Not sure if productive: just sitting, not disturbing others - maybe thinking or waiting for instructions or help?
ALONE	8	Working alone, incl. ignoring teacher or video, reading ahead, using laptop, away from desk getting resources
OTHER	9	Working with others, incl. at whiteboard or others' desks or elsewhere
EXPR	10	Q&A participation including expressing knowledge to the class
APPR	11	Being appreciated or praised by the teacher

Abbrev.	State	Student activities
ATTEN	12	Listening to, watching, interacting with teaching, incl. watching video, asking or answering questions, coming to the whiteboard
TSUPP	13	Receiving individual instruction or support from the teacher
GSUPP	14	In a group receiving instruction from the teacher
TASUPP	15	Receiving support from the TA in the classroom
OUTTA	16	Out of classroom with the TA support group
NONE	17	None of these, e.g., left room alone or attending to personal matter

Note that student state 3 covers a student chatting to another student or just speaking out aloud to no-one in particular (because no-one happened to respond). The distance between two students chatting is considered an important factor: the greater the distance the more disruptive to the lesson (and the stronger the teacher's discipline response).

Table 3-3 The TA activity states modelled

Abbr.	State	TA activities
COBS	7	Observing students (while moving around or stationary); default state for the TA if not in any other state
TTA	15A	Talking to the teacher
GTA	15B	Assisting a group of students
SSUP	15C	Assisting one student
OUT	16	Out of the classroom with one or more students

3.2 How lesson plans control lessons

As indicated in section 2.4, one of the main observations from the case study was that the factor that most influenced behaviour in lessons was the teacher's lesson plan. All three agents take their cue from the lesson plan. It specifies what the teacher should tell the students to do and when, and thus directly influences student behaviour. Both empirical and simulated lesson events are driven by the lesson plan. One basic assumption, built into the simulation model, is that whenever the lesson plan specifies a new activity (i.e., the lesson section changes – see below), all agents are forced to reconsider their state.

Below is an example of just the bare essentials of a lesson plan:

Section	Actions	Minutes
1	Lesson Overview/Introduction + Q&A	5-10
2	Activity 1 – in pairs	10-15
3	Whole class discussion	5
4	Activity 2 – in groups	10-15
5	Whole class discussion	5
6	Activity 3 – in groups	10-15
7	Whole class discussion + Plenary	5-10
8	Homework, next lesson etc.	5
	Total time	60

For the purposes of modelling and simulation, a lesson plan was considered to be a sequence of lesson sections, represented as follows: Lesson ID, [section1 section2 ...]

Each section has the structure (where section# is just a sequential integer):

[section# duration [acceptable student states] [planned teacher states]]

The [acceptable student states] list is a list of the student states that the teacher expects to be predominant during that section. It begins with the most desirable state first, then others in decreasing order of probability/desirability. The most desirable student state is called the *expected state* of the students. The [planned teacher states] list is similar. The most desirable teacher state is called the *primary state* of the teacher. Apart from misbehaviour states, excluding a state from a list effectively prevents that state being entered during that lesson section. One can see from this that the lesson plan effectively specifies the ideal behaviour desired.

Lesson sections are of different types. This was based on the observation that students spent most of their time in one of the following types of activities (similar to the 'social planes' (individual, team and class) identified by Dillenbourg, Prieto and Olsen (2018)):

- working alone,
- working in pairs,

- working in small groups/teams, or
- working as a whole class when the teacher is whole-class teaching.

As an example, the lesson plan:

MAT08190423, [1 20 [12;10;11] [12;9;10;11]] [2 15 [8;9;13;14;15;12] [7;13;14;12;15]] …

describes a lesson starting with 20 minutes of the teacher whole-class teaching (state 12 or 9) or some possible Q&A (state 10) or some praise (state 11) and the students listening (state 12) plus some possible 10 (Q&A) and 11 (praise). This is followed by 15 minutes during which the students are expected to either work alone (8) or with others (9) or get one-to-one help from the teacher (state 13) or TA (state 15) or maybe listen to the teacher giving some more instructions (state 12) and the teacher either observes (7), helps individuals (13) or groups of students (14) or give some whole class instruction (state 12) or interacts with the TA (state 15).

When empirical lesson event data were replayed it became apparent that, in all lessons, the lessons did not unfold minute-by-minute as specified in the teacher's lesson plan. (In fact, teacher's lesson plans typically had entries such as *10-15 mins* or *20-25 mins* reflecting the fact that teachers adjust the lesson dynamically.) For example, one lesson start was delayed for 5 minutes by several students who wanted to explain why they had not brought their homework. If the actual lesson were compared to the lesson plan, every transition to the next lesson plan section would be out by 5 minutes. A simulation using the lesson plan gave a poor match with the actual lesson. To manage this phenomenon, the lesson plan was inferred from the lesson events themselves. A small program was developed that took the raw lesson event data for the teacher and extracted a rough lesson plan based on the timings of the teacher's activities. By ignoring minor interactions and using the actual lesson plan to see what was intended, these rough plans were manually adjusted to create realistic realigned lesson plans. For example, for the lesson with a 5-minute delay at the start, the adjusted lesson plan had a separate 5-minute section of non-teaching, administrative time added, as though it were planned.

The choice of the primary state for the teacher is critical for the simulation to operate properly. When the teacher is 'passive' whole-class teaching (e.g., playing a video clip) the primary teacher state is 9; when the teacher is active whole-class teaching the primary state should be state 12; and when the teacher expects the students to work either with others or alone, the primary teacher state should be state 7 (observing).

3.3 Agent behaviour and decision-making

As explained above, on the basis of the case study observations, the lesson plan was considered to be the most important factor influencing lesson behaviours. The second most important determinant of behaviour appeared to be the current circumstances, who is doing what and where, and what has happened recently. The third level of influence was considered to be the person's habits. These latter were summarized as state duration distributions and, for students only, their historical empirical activity state frequencies. The three levels are summarised in Figure 3-2.

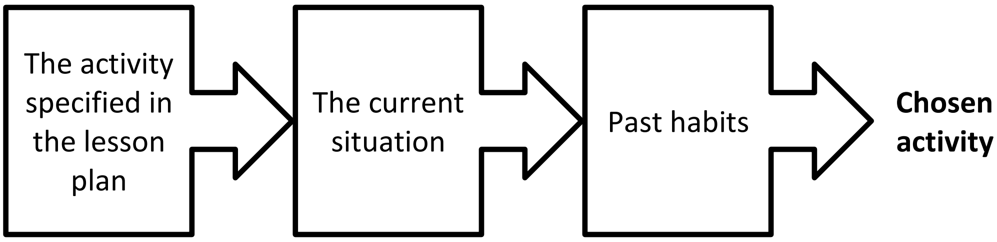

Figure 3-2 Student agent considerations for the choice of the next activity

As explained in section 2.4, agent behaviour is described using rulesets, with the choice of the next activity state being modelled by a set of rules. Most of the rules concern how an agent is to respond to what others are doing. The order in which agent types are processed is teacher, students, TA. This order was chosen because, during the case study, it appeared that predominantly the students respond to the teacher, and TA's appeared to be more reactive, responding to student requests before offering help on their own initiative. The student agents are processed in a random order. The following sub-sections provide a brief overview of how agent behaviour was modelled. Further details are provided in Appendix A.

3.3.1 Teacher behaviour

The teacher sub-model applies to all teachers in all lessons. There was too little empirical data to profile each teacher sufficiently, yet it was clear from the case study that teachers did not behave alike. Not only that, one teacher could behave quite differently:

- with different classes (e.g., the extent to which they tended to discipline);
- in different subjects (e.g., when using IT equipment vs during maths exercises);
- in different rooms (e.g., where it may be difficult to be sure what the back row are up to);
- with different lesson plans (e.g., no group work vs extensive group work).

In order to differentiate the behaviour of teachers, two teacher-lesson parameters were incorporated. These adjust, for each specific lesson, the teacher's tolerance of misbehaviour and inclination to offer one-to-one support.

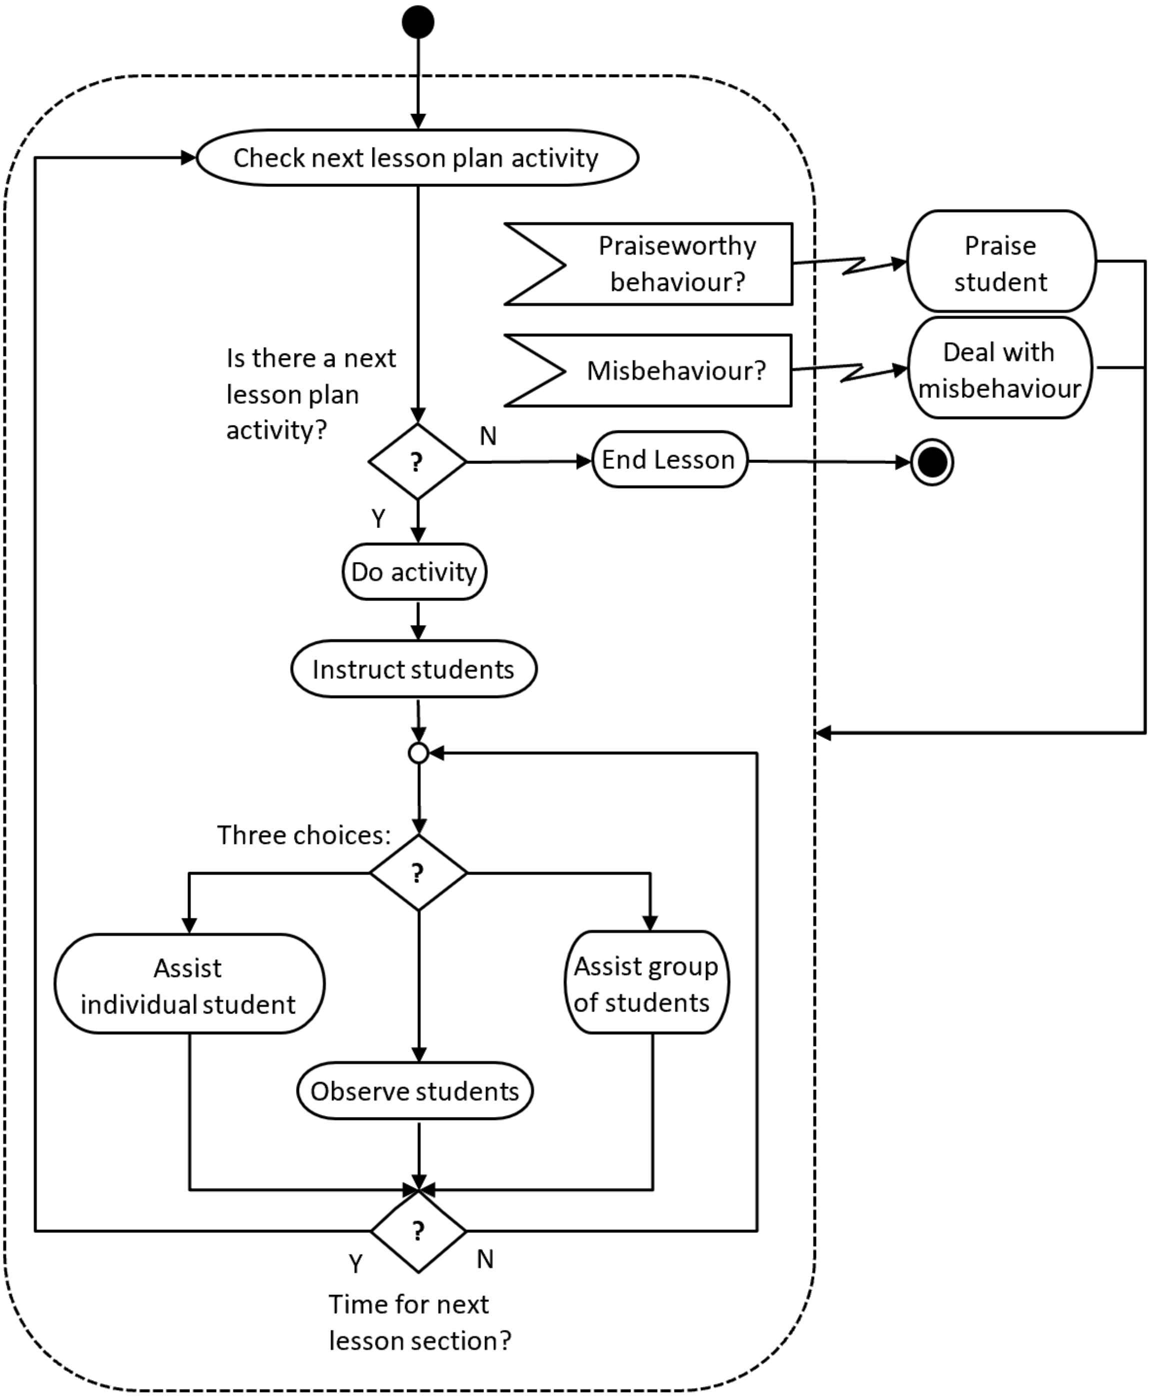

Figure 3-3 The flow of the teacher's activity choices
including providing ad hoc praise or discipline in response to student behaviour

The teacher follows their lesson plan but responds to what the students are doing. The teacher behaviour algorithm is summarized in Figure 3-3. Behaviour management and praise are considered to be 'interrupt-driven' subsystems in that the teacher temporarily deviates from the lesson plan in order to respond. After dealing with such situations the teacher will return to whatever activity the lesson plan specifies.

The following examples demonstrate the structure and nature of the teacher behaviour rules. The overall teacher rule is: Do what the lesson plan says (stay in or enter the planned primary state) but handle exceptions, including dealing with misbehaviour and praising. The first example models whole class disciplining in response to some persistent, low-level disruption:

IF there is generally too much chatting or disengagement (student state 3 or 4), according to how long the messing about has persisted and according to my misbehaviour tolerance
THEN for a suitable amount of time remind the class to focus on the activity (state 5) (treat state 3 as more serious than state 4 so intervene earlier)

IF I have been disciplining the class for a suitable amount of time
THEN stop and do what the lesson plan specifies
ELSE continue

The next example models the teacher offering one-to-one support to a student:

IF I am not helping anyone
AND I am available to help someone (which means states 4, 7, 8, 15)
AND it has been a while since I helped anyone (varies according to inclination to help and according to how long since helped anyone)
THEN
    choose a student who is not being helped by the TA and appears to be doing nothing (student state 4 or 7) or is chatting (student state 2 or 3), but not someone who has recently received help
    IF such a student was found
    THEN offer support to him/her
    ELSE (so as to allow students time to think and work)
        IF the TA is not busy and it is not too recent that I spoke to the TA
        THEN confer with the TA (state 15)
        ELSE start/resume the main activity specified in the lesson plan

3.3.2 Student behaviour

Figure 3-4 shows the logic that a student agent follows in order to decide what state to occupy during the next simulation time interval. This logic reflects the type of situations students find themselves in:

- sometimes students have no choice about what to do next – it is forced upon them (e.g., the teacher disciplines them or offers help);
- sometimes they cannot continue doing what they were doing (e.g., the person they were chatting to is now being helped by the teacher) and they are forced to reconsider what to do;
- when reconsidering a state, a choice may be made that turns out to be impossible (e.g., the student wanted to chat or work with someone but there is no-one available).

As mentioned in section 2.4, each student has an attribute for each activity state and the value of each attribute at any particular instant represents the relative likelihood that the student would choose that activity compared to other activities. For example, the attribute 'state 3' means the students tendency to be 'in own seat chatting, distracting, socialising, turning around etc.' The value of this attribute at the start of a lesson is the student's historical relative frequency of this state over all lessons. These attribute values are used to restrict and rank the choices of activity states. The state attributes are considered together as a probability mass function (PMF). To be precise, each student has three PMFs, one for each lesson section type (whole-class teaching, working alone and working with others – a combination of the working in pairs and working in small groups sections described in section 3.2). At the start of each lesson section (including the start of the lesson), the appropriate set of empirically-derived relative state frequencies are loaded into the PMFs. These state frequencies – representing the students' historical behaviour habits – are averages, formed by the students' behaviours over several different lessons. The PMF values are adjusted according to the current circumstances a student is facing, with some situations increasing the score for a state, others decreasing it. (The term 'score' is preferred for the adjusted values as they are not strictly probabilities.) The student rules make these dynamic adjustments. Where a state is 'forbidden' or impossible, its score is set to zero. From this explanation one can see that the lesson plan is highly influential, specifying the desired ideal activity states and assigning the appropriate student 'habit' data.

Figure 3-4 shows the overall student behaviour logic. Several aspects of the logic involve the management of agent interactions, such as a student waiting for another student to respond. These interactions are elaborated further in section 3.4.

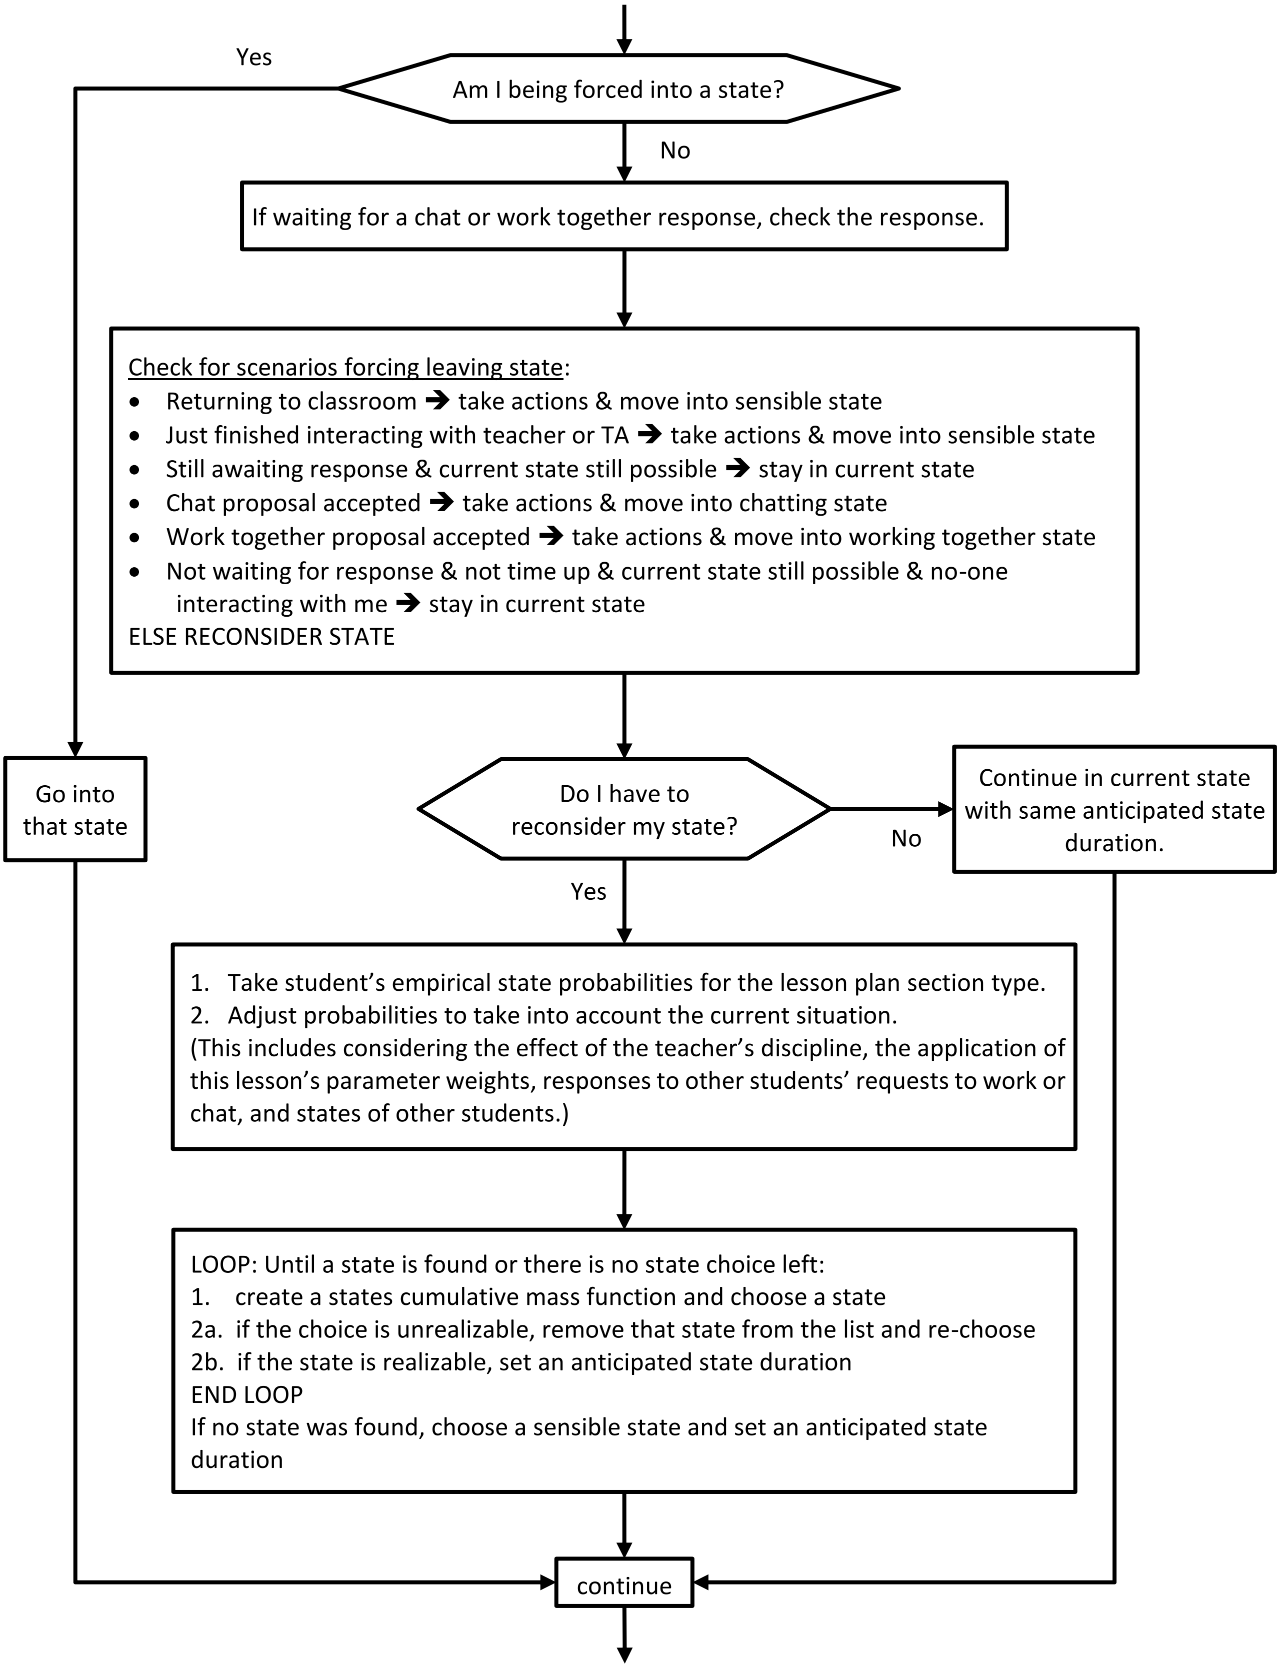

Figure 3-4 Flowchart for student agent behaviour

As an example of the type of rules preventing a state, the following conditions cause the score for state 13 (being helped by the teacher) to be set to zero (the 'I' and 'me' is the student agent doing the deliberation):

> the teacher is in any of these states: 1 5 9 10 11 12 14 16 17
> OR the teacher is already helping someone else
> OR I am not expected to be working alone or with others
> OR I am out of the room
> OR the teacher helped me too recently
> OR the teacher has said they are not giving help

As explained above, each student has a set of PMF attributes that, initially, contain their <u>average</u> historical activity state frequencies, formed by the students' behaviours over several different lessons. But students do not generally behave in any lesson in the same way as their average behaviour. Students behave differently:

- with different teachers (e.g., according to their relationship with the teacher);
- in different subjects (e.g., according to their interest in the subject);
- in different rooms with different seating arrangements (e.g., whether they are beside a friend);
- in different activities (e.g., preference for solo or group work).

Students and classes can behave quite differently from one lesson to the next. Activity state averages were used because there were insufficient empirical data to characterise each student in one type of lesson (e.g., mathematics) with one specific teacher and one specific TA in one specific classroom. Practically this meant that lesson simulations needed individual adjustment. For example, disruptive behaviour or other behaviour or the incidence of one-to-one support needed turning up or down relative to their average values.

To accomplish these adjustments, several student-lesson parameters were introduced. These adjust the averaged historical state frequencies so that behaviour more closely matches that of the specific empirical lesson. They affect every student in the class and are more accurately termed class-lesson parameters. They are:

- Student-Support-Request-Weight (SSRW) – to alter the students' inclination to ask for assistance;
- Student-Interaction-Weight (SIW) – to alter the students' inclination to initiate an interaction;

- Interaction-Response-Weight (IRW) – to alter the students' inclination to respond to an interaction;
- Peer-Weight (PW) – to alter the extent to which students copy their peers' behaviour;
- Relative-Lesson-ES-Weight (RLESW) – to boost or reduce the expected and acceptable states;
- Relative-Lesson-Disruption-Weight (RLDW) – to boost or reduce disruptive states;
- Relative-Lesson-Other-Weight (RLOW) – to boost or reduce other disengaged states.

Appendix A.3.5 contains details about how the student activity state scores are adjusted by these parameters.

3.3.3 TA behaviour

The TA sub-model applies to all TAs in all lessons. As in the case of teachers, there was too little empirical data to profile each TA adequately, yet it was clear from the case study that TAs did not all behave alike. Not only that, one TA could behave quite differently with different classes, even with the same class in another lesson. In order to differentiate the behaviour of TAs, a TA-lesson parameter was incorporated. This adjusts, for each specific lesson, the TA's inclination to offer one-to-one support. The TA's decision-making is summarised in Figure 3-5.

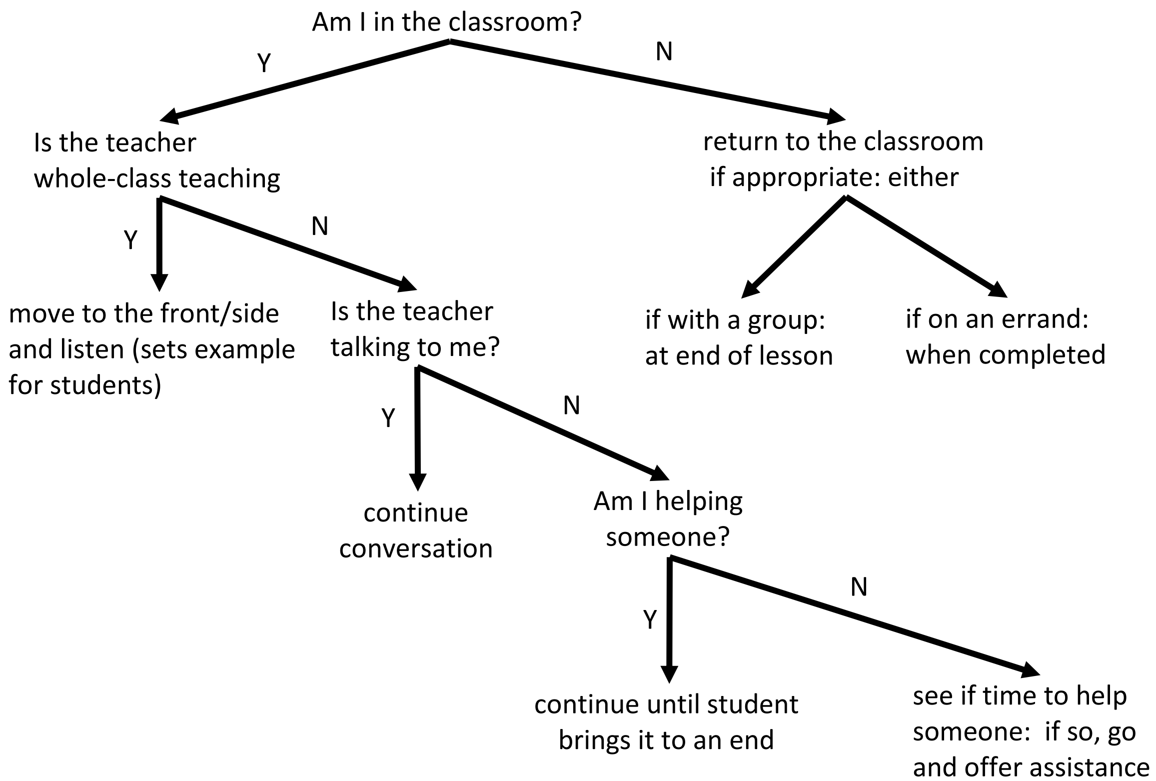

Figure 3-5 TA agent decision-making and behaviour

The following example demonstrates the structure and characteristics of the TA behaviour rules.

```
IF I am free (not helping anyone, state 7)
  AND    the teacher is not talking to me
  AND    the lesson plan allows TA support for students (student state 15)
  AND    I feel like offering support
  THEN   look for a student to help (someone in state 2, 3, 4, 6, 7, 8 or 9
             who has not been helped for a while)
      IF      a student was found
      THEN offer help to him/her
             go into state 15
(the student will enter state 15 when they next consider what to do)
```

3.4 Agent interactions

One of the reasons an agent-based modelling approach was adopted was so that causal chains of interactions between the autonomous agents could be investigated. The following example provides an insight into the mechanisms used in the model to manage the complexities of agent interactions.

There are two main student-student interactions, working together (state 9) and chatting (state 3). When one student proposes interacting with another student and waits for a response, several possible scenarios could unfold, depending on the order of student processing and the actual response, accept or reject. When an agent proposes an interaction, the agent remains in their current state until that proposal is accepted or rejected. The presence of an interaction proposal increases the probability of the proposed state for the respondent.

Consider the scenario where student P proposes working with student Q and notifies Q (and effectively the rest of the class). Notification is implemented by recording the proposal in a globally-visible list.[11] When student Q has a turn, during the same time-interval, student Q rejects the proposal (by deleting the entry in the list). A rejection could be for any reason, including choosing to interact with someone else or being forced into a state by the teacher, or just

[11] In computer programming, variables or constants that are accessible to all procedures or objects are sometimes called global variables or constants, globals for short.

choosing to work alone. It is entirely possible that several students proposed partnering with student Q: the algorithm takes the first proposal.

In the next time interval, P notices the rejection (the proposal is no longer in the list), but, when choosing a next state, again chooses state 9. What to do? Asking Q again would be odd – and possibly cause unrealistic looping. To stop this, a *rejected-by* list was implemented as a student attribute, along with a note of the time of the rejection. Student P looks for any potential partner excluding anyone on his/her *rejected-by* list.

This is fine for a while until you have situations where two students are working at a table and everyone around is either too far away or busy or has rejected P. If that rejection is permanently remembered, then student P will never ask Q again. And if P had rejected Q, then Q would never ask again. Hence, even if the lesson plan (and thus the teacher) had instructed students to work together, they never would. To avoid this, rejections have to be forgotten after a while. This required yet another model variable whose value could not easily be established empirically and, even if such data were available, it would likely be extremely variable. The logic is enhanced: if student P does not find a partner this time step, then if it has been some time since the last rejection (with some plausible amount of randomness), his/her *rejected-by* list is cleared so that future state 9 decisions have a fresh start.

Numerous rules and mechanisms were created to cope with such interaction intricacies. Agents, each with a unique reference (ref), remember certain past events and communicate asynchronously using either inter-agent messaging or conventional *blackboard* variables (aka global variables or flags). These are used for recording (amongst other things):

- who the teacher is currently helping or about to help (Teacher-Allocated-To list);
- which student is engaging in Q&A or being praised;
- current and proposed chatting (state 3) partnerships (Proposed-3-Partnerships list);
- current and proposed working together (state 9) partnerships (Proposed-9-Partnerships list).

For example, when the teacher is available, a student requests help from the teacher by adding his/her ref to the Teacher-Allocated-To list; the teacher consults this list on the next simulation step. Each student also has an attribute *paired-with*, which is a list of the refs of other students that the student is currently partnered with (working or chatting).

In section 2.4, two particularly important interactions were mentioned:

- the students' responses to the behaviours of their peers;
- the students' responses to the teacher's discipline regime.

Peer response was implemented by simply counting the number of students in each activity state and increasing or decreasing a student's own state scores proportionately – so students tend to copy others. For states where only one student can be in that state (or two or three for state 15 (being helped by the TA)), the fraction of the class in any productive state was used. The impact of this peer influence factor in relation to other factors was unknown, hence the lesson parameter Peer-Weight was introduced into the procedures that adjust student state scores. An underlying assumption was that all students responded to the activities of other students to the same degree.

The implementation in the simulation model of the teacher's disciplining and the students' responses to this is more complicated than the above interactions and is explained in the following subsections, with further details provided in Appendix A.3.4. Figure 3-6 shows the main factors involved in teacher-student discipline interactions.

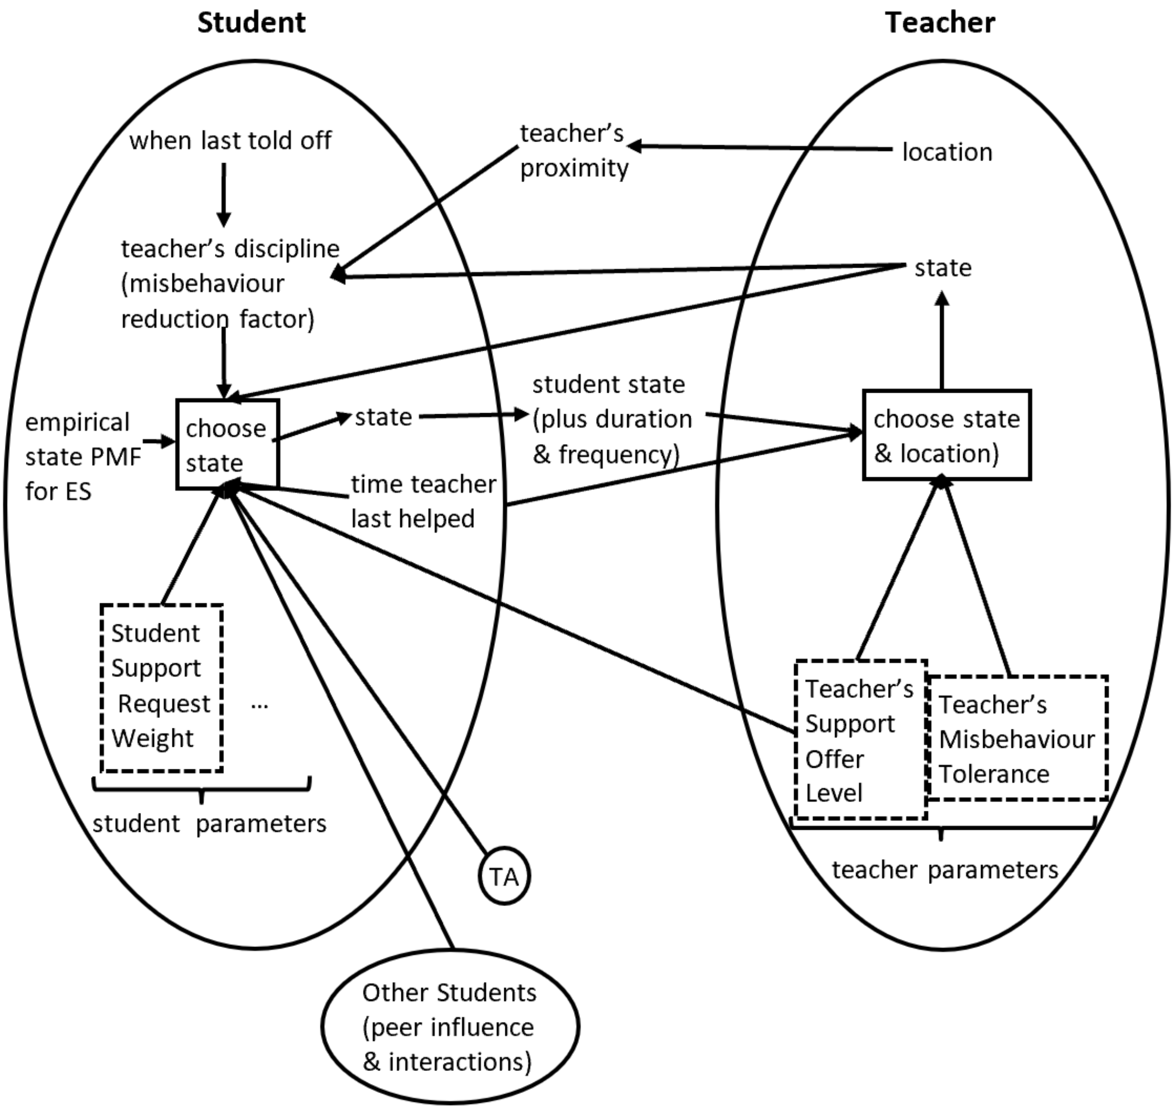

Figure 3-6 The main factors involved in teacher-student discipline interactions

3.4.1 Modelling the teacher's disciplining

This subsection describes how the interactions concerning misbehaviour described in subsection 2.4.1 were modelled. To recap, during the case study it had been observed that teachers disciplined individuals or a whole class, but that in some lessons teachers seemed stricter than in others. Also, students misbehaved more when the teacher was helping someone and when the teacher was further away, but after being disciplined, behaviour was more settled, although the effect seemed to wear off over time. The model needed to accommodate those aspects plus the following points:

1. UK Secondary school classes might have roughly 5 to 35 students: hence the approach needed to suit this range.
2. In a larger class, the teacher might tolerate more students disengaging before reprimanding the class: hence calculations needed to raise the threshold with class size.
3. The teacher's response seemed to depend on what the teacher was doing, whole-class teaching, helping an individual, etc.: hence calculations needed to change with the state of the teacher.
4. The teacher's response seemed to depend on the type of student disengagement, whether it distracted or disrupted others or was just passive disengagement, and the magnitude of the disruption (for example whether a student shouted out across the classroom versus just made comments quietly to their neighbour): hence calculations should depend on the students' states and locations.
5. As with individual misbehaviour, there needed to be some mechanism for capturing ongoing but intermittent distractions (e.g., ongoing chatting-listening-chatting-listening): hence calculations should incorporate students' recent misbehaviour.
6. Teachers were not consistent in their responses but varied from lesson to lesson and within a lesson: hence, calculations needed to add some random effects.
7. In the literature, there appeared to be no standardised scale to rate misbehaviour tolerance and to compare the differences between teachers and lessons, however, the empirical data could be used to determine distributions for the time teachers did wait.
8. The empirical data could also be used to set durations for the teacher to be disciplining, both individual and whole-class disciplining.

All teachers were assumed to respond to student misbehaviour in the same way, but randomness was added to the timing of their responses. This seemed plausible because there were times when teachers appeared to delay disciplining a persistent misbehaver, for example to first finish off an explanation to another student.

This led to the following mechanism for the teacher's response to student misbehaviour:

- The teachers 'remember' each students' misbehaviour for a period (currently set at 10-minutes).
- This record is kept on a rolling cycle, the most recent time interval pushing out the oldest record (from 10 minutes ago).
- If the amount of individual misbehaviour in the past 10 minutes crosses a threshold (which includes some random fluctuations), the teacher will discipline the individual.
- If the amount of class misbehaviour in the past 10 minutes crosses a threshold (which includes some random fluctuations), the teacher will discipline the class.
- Different misbehaviours have different thresholds (each with some random fluctuation) (see Appendix B.5); the threshold may be 0 s for a serious misdemeanour.
- The record of a student's misbehaviour is cleared whenever the teacher tells off that individual or the whole class.

3.4.2 Modelling the students' response to disciplining

Following the common viewpoint that disciplining reduces misbehaviour and improves productivity, it was assumed that being disciplined increased the likelihood of a productive state being chosen and decreased that of all disengaged states (active distracting and passive disengagement). To cause the students to respond to discipline in this way, their state scores are adjusted to reflect the following factors:

- the distance between the teacher and the student;
- a distance falloff function that mimics how, as distance increases, the students feel freer to misbehave;
- the elapsed time since the student or the class were last disciplined;
- a time decay function used to mimic the effect of discipline fading over time;
- the current state of the teacher (e.g., the effect is stronger if the teacher is observing the class, and even stronger if the teacher is currently telling someone, or the class, off).

These factors are combined (see Figure 3-7) to produce an EffectOfDiscipline factor. The students' misbehaviour state scores are reduced by this factor. Another assumption was that all students respond to the teacher's discipline to the same degree.

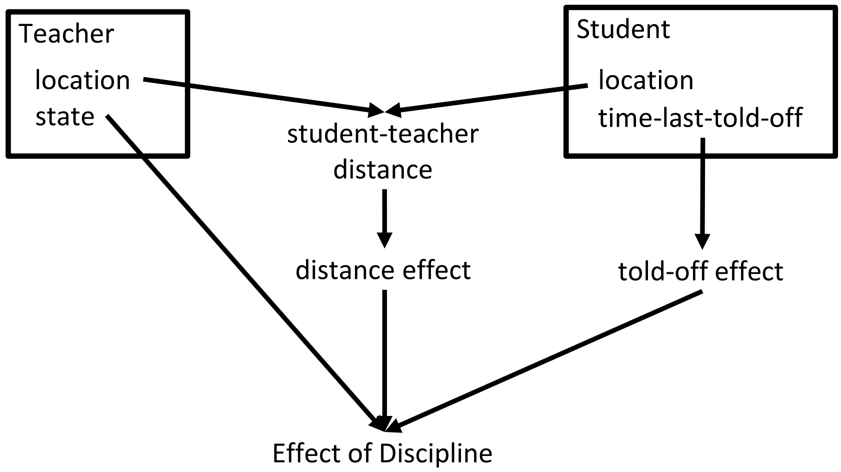

Figure 3-7 How students are affected by the teacher's discipline

3.5 The implementation in NetLogo

As mentioned in section 1.5.1, the simulation model was implemented in NetLogo (Wilensky, 1999). This took the form of several code modules comprising in total around 9,700 lines of (commented) code, plus several data files. To run the simulation code, one needs the NetLogo software installed, with all the required extensions. The simulation tool developer's interface is shown in Figure 3-8. It comprises four vertical panels. The second panel contains many statistics summarizing the lesson dynamically and includes the 2D representation of the lesson (also shown in Figure 2-5).

A larger image of a lesson is shown in Figure 3-9, an instant in an empirical lesson replay. The teacher (dog icon) is helping one student, two students are working together (bottom-left), the TA (cat icon) is helping two students, one student is out of the class (#1001 by the door) and two students are chatting – with one trying to interact with the student in the top-left corner. The others are working alone (brown, heads down). The agent icons move and dynamic links form and disappear during interactions. These links thicken over time.

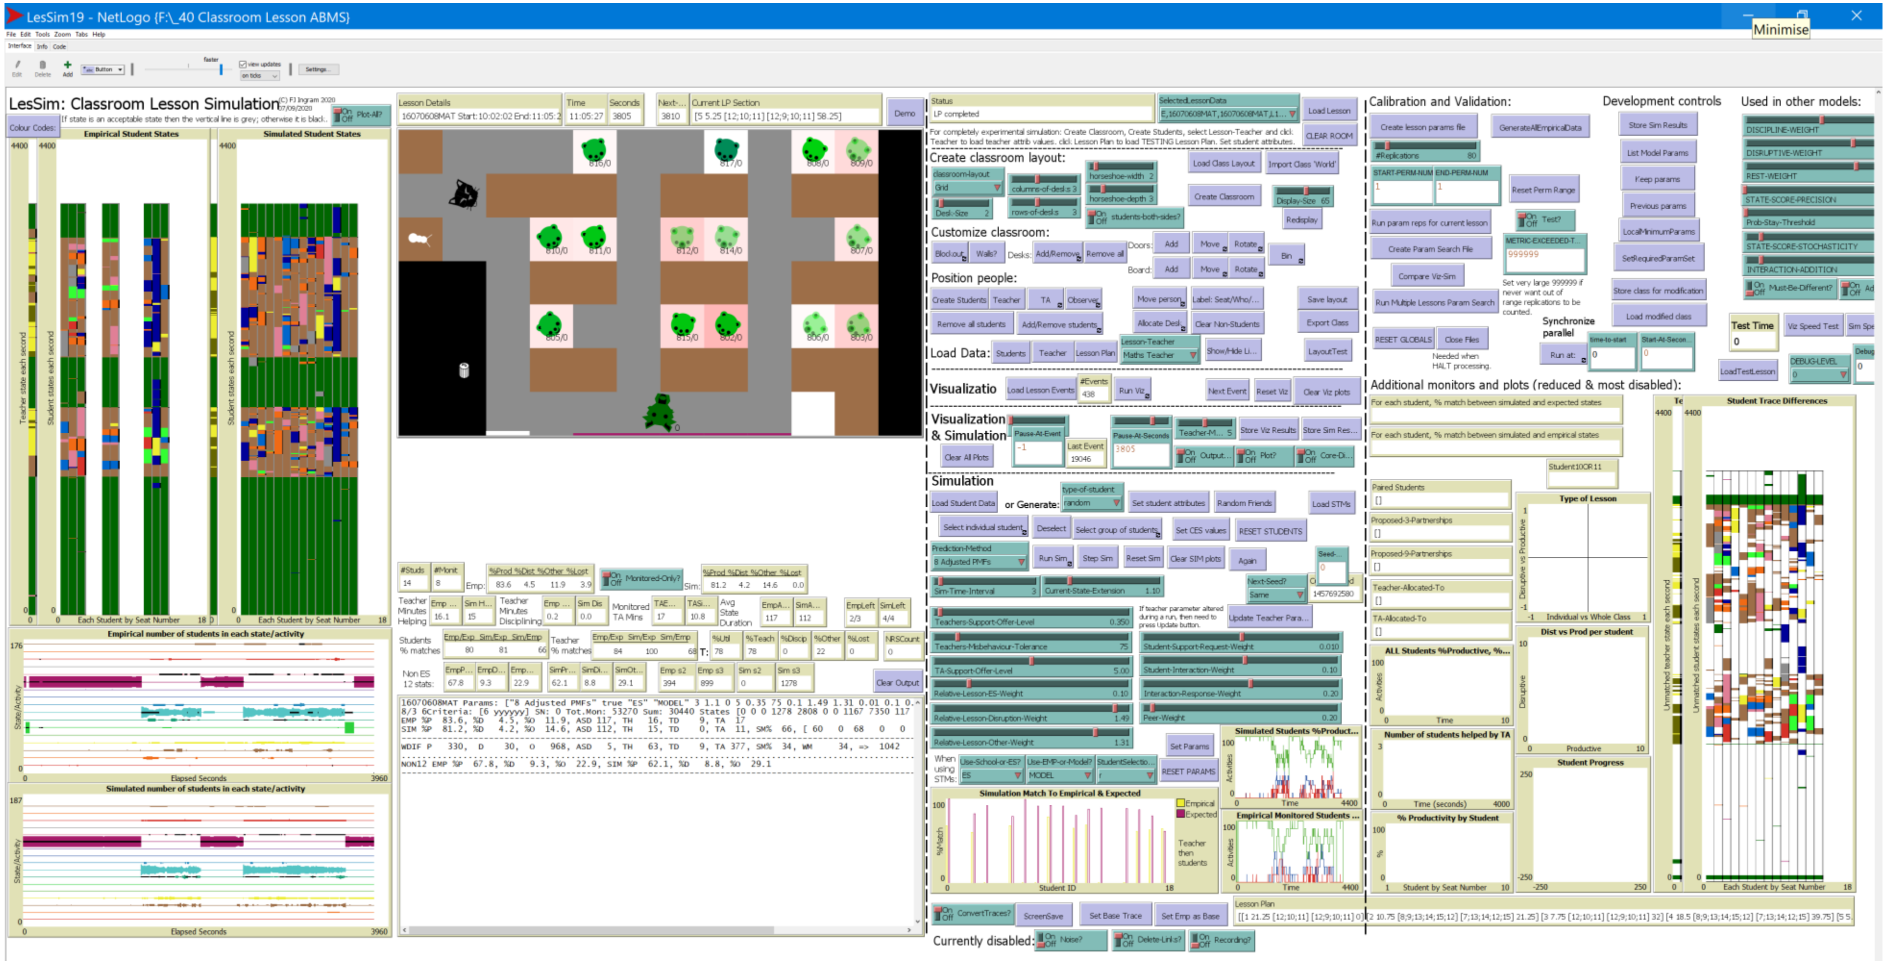

Figure 3-8 A screen capture of the simulation tool developer's interface [12]

[12] It is appreciated that the current user-interface needs to be significantly simplified before it can be used by others. Further research is needed to establish what teachers and school management would require (Ali et al., 2013).

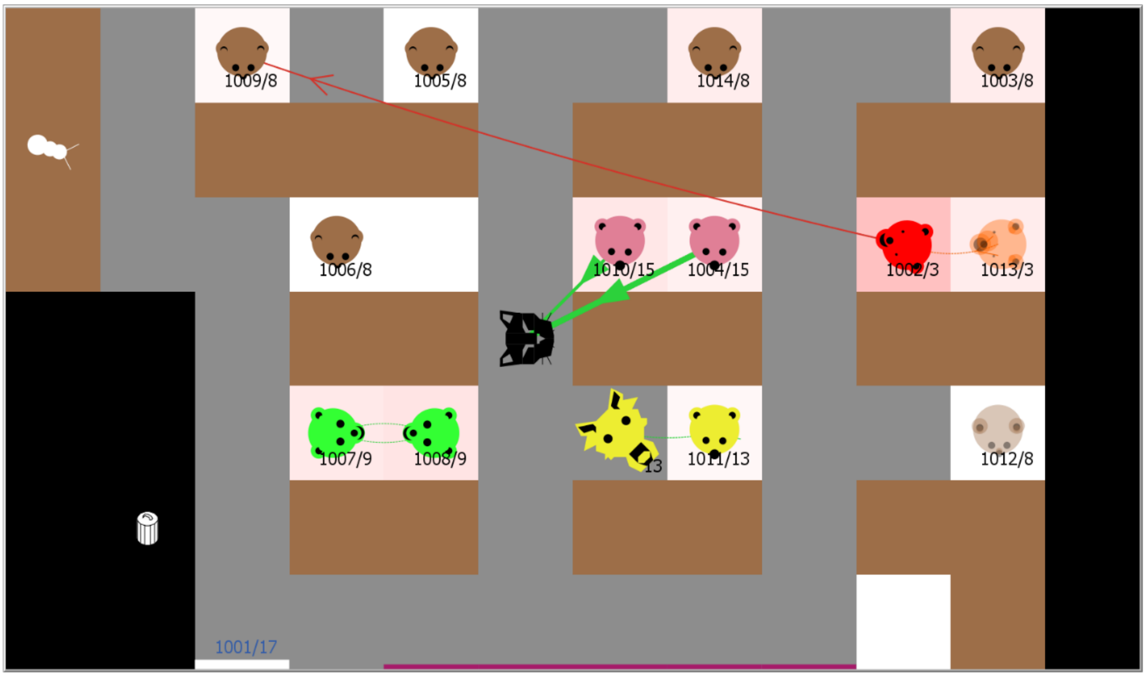

Figure 3-9 An example of the representation of a classroom lesson and participants

The colours of the icons represent different activities/states and are explained in Table 3-4. For example, a student icon turns bluer the more the student disengages and the background of each student becomes pinker then redder as their record of disruptiveness worsens.

A classroom floorplan is a specification of a NetLogo *world*. A significant portion of the code is explicitly for the purpose of designing classroom lesson layouts. This is particularly relevant during simulations where, for example, a different desk arrangement is being investigated. NetLogo worlds are built from a grid of patches, the units of the world. In the physical classroom model, a patch is approximately a square of 0.6 m side. This is based on the dimensions of the 2-seater 0.6 m x 1.2 m desks used in the school in the case study.

During whole-class teaching, the teacher is simulated to move pseudo-randomly from side-to-side at the front of the class, avoiding the corners of the room and spending less time directly in front of the middle of the whiteboard. The reason for this mechanism, apart from it being quite realistic, is that the students respond to the proximity of the teacher, which would otherwise be constant.

The left-hand panel in Figure 3-8 contains the charts that were considered most useful for visually determining what was happening in a lesson. (The right-hand panels contain other charts that were used although many were disabled to increase simulation speed.)

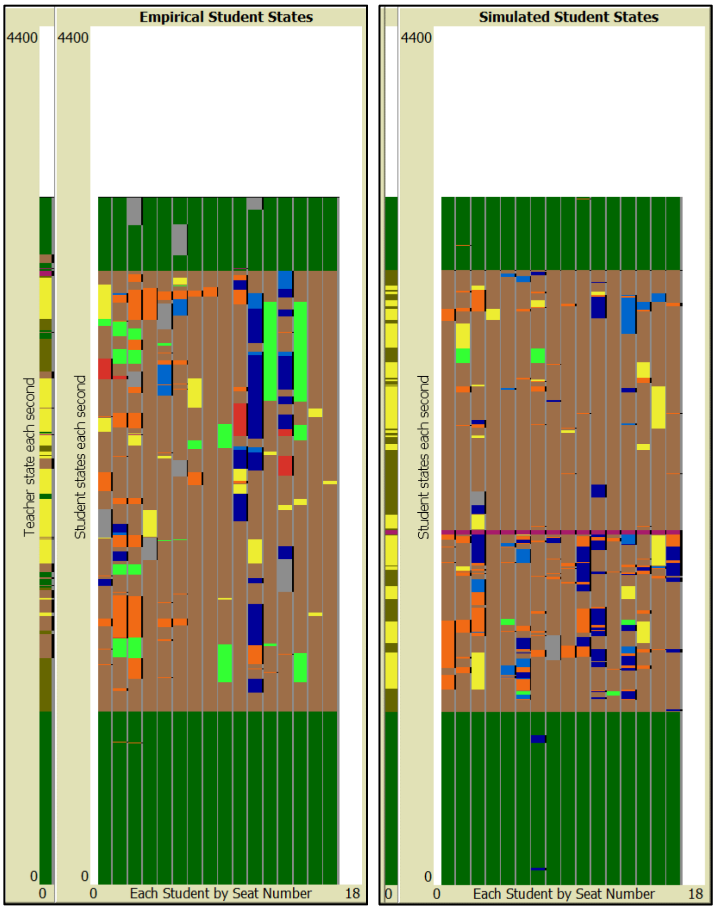

(a) (b)

Figure 3-10 An example of the teacher and student state trajectories through a lesson (a) the empirical lesson replay and (b) one simulation run.

Table 3-4 The colours used to represent different student and teacher activities/states

State#	Colour	Student state type	State#	Colour	Teacher state type
1		Being disciplined	1		Disciplining individual
2		Serious misbehaving	2		Unused
3		Chatting	3		Unused
4		Passively disengaged	4		Other, non-lesson
5		Class being disciplined	5		Managing/disciplining class
6		Prevented from working	6		Prevented from teaching
7		Resting?	7		Observing
8		Working alone	8		Working alone
9		Working with Others	9		Passive whole-class teaching
10		Q&A participation	10		Q&A
11		Being praised	11		Praising student
12		Whole-class learning	12		Whole-class teaching
13		1-1 teacher support	13		1-1 student support
14		Teacher support in group	14		Group support
15		1-1 TA support	15		Interacting with TA
16		Out of class with TA	16		Out of class
17		None of these e.g., out	17		None of these

A larger image of this type of lesson tracing is shown in Figure 3-10. This shows the teacher and student traces from an empirical lesson replay next to one simulation run. The separate thin strips on the left are the teacher's state traces. Time flows upwards and the colours indicate the different activity states (as per Table 3-4). These traces were used extensively during the analyses of experiments (in Chapter 8).

The simulation time interval (essentially another model parameter) was set at 3s, but internally every agent and statistic is updated every one second. The reason for this was that to replay empirical lesson data the software must be able to process an event at any time, not just at the start of a simulation interval.

3.6 Chapter summary

This chapter provided a summary of the classroom lessons simulation model (CLSM). It explained the range of factors that need to be taken into account to construct a more generic behavioural model. The activity states of the three agent types, teacher, TA and students were described. The central role of lesson plans in determining behaviour was explained, followed by an overview of how each agent type behaves. Several examples of the behavioural rules for each agent type were given. The teacher follows a simple algorithm (basically follow the lesson plan handling any interruptions along the way), the TA follows a few simple rules, but students follow more complicated logic. The intricacies of managing agent interactions were described, with emphasis on how the students and the teacher interact when it comes to misbehaviour and discipline. The model parameters that adjust behaviour were introduced. The chapter ended with a brief description of how the simulation model was implemented in NetLogo, in particular the developers interface.

Together, Chapters 2 and 3 have covered stages 1 to 3 of the ABMS development. The following chapter explains topics that are relevant to the following stages 4 and 5 in the ABMS development: model calibration and validation.

4. Comparing simulated and empirical lessons

As outlined in section 1.5.1, a crucial activity in ABMS development and use is the evaluation of simulation output to decide whether it is acceptably realistic. In empirically-based ABMS development methodology, after a conceptual model has been transformed into a simulation (computational) model, the next two stages are model calibration then validation. Both stages require the comparison of outputs. One of the goals in building an empirically-based classroom lessons simulation model (CLSM) was to obtain a model that, given the same conditions (input data) that were present in an empirical lesson, generated outputs that matched (to an extent) the empirical lesson (Balci, 1994, 1995; Lee *et al.*, 2015).

This chapter explains the decisions that applied to both the calibration and validation stages. Section 4.1 explains why the model needed instantiation per empirical lesson and how this was accomplished. Section 4.2 describes the lesson metrics that were chosen and explains how they were used. This addresses one of the research objectives, to determine suitable metrics and methods for comparing lessons. Section 4.3 presents the empirical data and acceptability ranges. Section 4.4 justifies the number of replications used during simulation runs.

4.1 Model parameters and instantiation per lesson

The CLSM, like most models, includes variables and constants for which ranges or values were initially unknown. Some values were derived directly from empirical data (see Appendix B). Others, such as realistic state durations (e.g., for a teacher providing one-to-one assistance to a student), were obtained by fitting functions to the empirical data and sampling randomly (see Appendix B.1). Where values for model parameters could not be determined directly from the empirical data, their values were inferred after running multiple simulation replications. Table 4-1 summarizes the model parameters, which, apart from the first, Current-State-Extension (CSE), have all been introduced earlier. In the simulation animations one could often see that a simulated lesson was too frenetic or too stable compared to the empirical lesson: the state durations were too short or long. The CSE parameter was introduced to increase or reduce state durations. The ‘class’ parameters are also termed ‘student’ parameters although they are applied to the class as a whole, not specific students. More detail about the effect of the parameters on agent behaviour and lesson outputs is presented in Chapter 6.5 in the context of sensitivity analysis. (Section 9.3.1 considers reduction of the number of parameters.)

Table 4-1 The model parameters and their use

Parameter name	Category	Explanation
Teachers-Support-Offer-Level (TSOL)	Teacher	A weighting in [0,10] used to adjust the likelihood that the teacher will offer 1-to-1 support; 0=never offers, 10=always offers if appropriate.
Teachers-Misbehaviour-Tolerance (TMT)	Teacher	A weighting in [0,600] used to adjust the likelihood that the teacher will ignore student misbehaviour; 0=completely intolerant (with small random fluctuations).
TA-Support-Offer-Level (TASOL)	TA	A weighting in [0,10] used to adjust the likelihood that the TA will offer 1-to-1 support; 0=never offers, 10=always offers if appropriate.
Relative-Lesson-ES-Weight (RLESW)	Class	This parameter adjusts the extent to which students comply with the teacher's instructions (engage in the expected state, ES) relative to the aggregate of all students in all lessons.
Relative-Lesson-Disruption-Weight (RLDW)	Class	This parameter adjusts the extent of student disruptive behaviour relative to the aggregate of all students in all lessons.
Relative-Lesson-Other-Weight (RLOW)	Class	This parameter adjusts the extent of student passive disengagement behaviour relative to the aggregate of all students in all lessons.
Student-Support-Request-Weight (SSRW)	Class	A weighting in [-1,1] used to adjust the likelihood that students will request 1-to-1 support.
Peer-Weight (PW)	Class	This implements a modelling assumption that students are influenced by what their peers are doing; this parameter increases the likelihood of the state that other students are in.

Parameter name	Category	Explanation
Student-Interaction-Weight (SIW)	Class	A weighting in [-2,2] used to adjust the students' inclinations to propose interactions (states 3 & 9); this was considered useful to adjust students to be more interactive than their average behaviour; this value represents the magnitude of that effect.
Interaction-Response-Weight (IRW)	Class	This implements a modelling assumption that a student who has been approached for interaction is more likely to choose interaction; this value represents the increase in probability.
Current-State-Extension (CSE)	Lesson	The models use the empirically-derived average state duration (ASD) distributions; this parameter is a state duration adjustment that affects all agents; ASD is reduced by incidents of student disengagement.

As explained in section 3.3.2, the students are driven by their historical average activity state probabilities (aggregated over all lessons), and aggregated teacher and TA data was used to model teacher and TA behaviour. But, as also explained, students do not generally behave in any lesson in the same way as their average behaviour, and teachers and TAs do not generally behave in the manner of some 'average' teacher or TA. This was the reason for introducing lesson parameters, to:

- adjust the teacher agent from its representation of a stylized teacher to the specific teacher-lesson in question;
- adjust the TA agent from its representation of a stylized TA to the specific TA-lesson in question;
- adjust the class of 'averaged' students to their behaviour in the specific lesson in question.

For example, in one specific lesson, disruptive behaviour or the incident of one-to-one support might need turning up or down relative to the average lesson behaviour. While this could have been accomplished at the individual student level, this was considered to be over-fitting. Instead, a lesson as a whole was adjusted with each student being adjusted in the same degree.

There was not one set of parameter values that suited all the lessons (see Appendix C for an investigation into this). Hence, although a generic classroom lesson behaviour model had been developed, model calibration (parameter estimation) and validation were performed per lesson. Note that there was no guarantee that the best parameterization found would enable the lesson model to pass the validation tests.

Time constraints meant that only seven of the twenty-one lessons available were taken through calibration and validation procedures:

Lesson#1	16062207GEO	Year 7	Geography
Lesson#2	16063010SCI	Year 10	Science
Lesson#3	16070508MAT	Year 8	Maths
Lesson#4	16070510MAT	Year 10	Maths
Lesson#5	16070607MAT	Year 7	Maths
Lesson#6	16070608MAT	Year 8	Maths
Lesson#7	16070809SCI	Year 9	Science

Table 4-2 is a break-down of these lessons by year group and subject showing which lessons had a TA.

Table 4-2 Summary of selected lessons

Subject	Year 7	Year 8	Year 9	Year 10
Geography	No TA			
Maths	TA	TA & No TA		TA
Science			No TA	TA

There is one aspect of the empirical data collection that had particular impact on lesson comparisons. As mentioned in section 1.5.2, practical restrictions meant that in many of the observed lessons only a selection of the students were monitored, i.e., had their behaviour recorded. As a consequence, sometimes the teacher, the TA and students interacted with unmonitored students. For example, suppose the teacher or TA helped several unmonitored students. The events of the teacher and TA were always recorded, but that data would not be well correlated with the data from the monitored students: it was unmonitored students that had been the triggers for the behaviour of the teacher and/or TA. Simulation runs in which the teacher agent happened to choose students in the appropriate monitored or unmonitored category would then better match an empirical lesson.

4.2 Metrics selected for output comparison

The empirical data collected during the case study provided a continuous event history of agent states for the teacher, the TA and each student. In principle, lessons could be compared by comparing these event histories, using a variety of techniques such as 'Hamming distance', 'confusion' or 'error' matrices, or event history modelling (Bennett, Klimas Blanc and Bloom, 2018; Moore and Hayward, 2019). However, a completely plausible simulated sequence of agent states might not resemble the empirical lesson at all. For example, consider a simulation where, at the start of a lesson section, the teacher helps one student (who was not helped in the real lesson and misbehaved), and, because the teacher is now in a different position in the classroom, nearby students also do not misbehave (which they did in the real lesson). One ends up with a realistic lesson that does not match the empirical lesson for entirely plausible reasons: a simulated event early on triggered quite different trajectories for the agents. For this reason, agent state sequence comparisons, although taken into account, were not the primary way that lesson outputs were compared - although a simulation model that gave a better empirical-simulation state sequence match would be considered superior to one that did not.

Instead, overall lesson metrics were used to compare lessons - see Table 4-3. Each is a key indicator of lesson outcomes and shows how two lessons differ in some specific way. Also, each has a visual aspect that can be seen clearly in the lesson output charts. These visual aspects enabled subjective judgements as to whether a simulated lesson was realistic and a sufficiently close match to an empirical lesson over time. To ensure these macro-level statistics were credible, the behaviour rules were carefully designed to maintain internal, micro-level validity, so that the simulation never generates impossible scenarios.

In addition, %Productivity (or %Prod) was often used to compare lessons (e.g., in experiments in Chapter 8) but not as part of the calibration and validation procedures. Productive time is just total student-lesson time minus the time spent by students on unproductive behaviour, i.e., distracting/disruptive or other, passively disengaged activities. Section 3.1 explained the categorization of student activity states into Productive, Disruptive and Other.

To compare simulated and empirical lesson metrics only monitored students' data were used, but when comparing simulations or just evaluating a single simulation, all students, monitored or not, were taken into account.

Table 4-3 Lesson comparison metrics (all mean values over the number of simulation runs)

Name	Description (all are means)	Participants	Visual Counterpart
%Dist D	total seconds of student disruptive/distracting behaviour; converted to % of lesson time	students	amount of red & orange in the charts
%Other O	total seconds of student other behaviour (passive disengagement); converted to % of lesson time	students	amount of blue in the charts
ASD	student average state duration (seconds)	students	apparent volatility of states (length & frequency of change)
TH	time teacher spent helping any student one-to-one (seconds)	teacher	amount of yellow in the charts
TD	time teacher spent disciplining an individual or the class (seconds)	teacher	amount of red & magenta in the charts
TAH	total time students were helped by the TA (seconds)	TA	amount of pink in the charts
ESM	%match between student simulated and empirical states	students	N/A

Two sets of values were obtained, one for the empirical lesson (subscript E) and one for the simulated lesson (subscript S). To compare outputs, the absolute difference between the means of the simulation values and empirical values was used (the variables are defined in Table 4-3):

$$\Delta \mathrm{D} = |\mathrm{D_S} - \mathrm{D_E}|$$
$$\Delta \mathrm{O} = |\mathrm{O_S} - \mathrm{O_E}|$$
$$\Delta \mathrm{ASD} = |\mathrm{ASD_S} - \mathrm{ASD_E}|$$
$$\Delta \mathrm{TH} = |\mathrm{TH_S} - \mathrm{TH_E}|$$
$$\Delta \mathrm{TD} = |\mathrm{TD_S} - \mathrm{TD_E}|$$
$$\Delta \mathrm{TA} = |\mathrm{TA_S} - \mathrm{TA_E}|$$

To compare student simulated and empirical state traces, the mismatch between student simulated and empirical states was calculated, as follows:

1. For each student calculate the percentage match between the student's simulated and empirical states (ESM).

2. Take the average of these percentages over all students, giving a measure of the match between the simulated and empirical lesson (average ESM) for the class.
3. Subtract from 100% to get a percentage mismatch (100% - average ESM).

In addition, a single, overall distance metric was constructed. It is the Euclidean distance from the origin (the empirical data point) of the point representing the set of 7 difference metrics, where each of the metrics has an equal weighting:

$$\text{distance} = \sqrt{(\Delta \mathrm{D})^2 + (\Delta \mathrm{O})^2 + (\Delta \mathrm{ASD})^2 + (\Delta \mathrm{TH})^2 + (\Delta \mathrm{TD})^2 + (\Delta \mathrm{TA})^2 + (100 - \mathrm{ESM})^2}$$

This metric is open to criticism:

- it is unclear what distance physically represents;
- each factor is given the same weight;
- the units are mixed (seconds and a percentage).

These points were not addressed directly, but to provide some degree of construct validity, the practical utility of this distance metric was evaluated by comparing the results from several alternative, simpler models. The reasoning was that if the overall distance metric value improves when the model improves, then one can have greater faith in the usefulness of the metric. Five student behaviour models were built, each intended to be increasingly realistic, with the following mechanisms for generating student activity states:

1. random (for the teacher also) – so generates impossible state sequences;
2. the expected state as per the lesson plan (for the teacher, the primary state) – the 'perfect' lesson;
3. each student selects a state at random using the empirically-derived state transition matrix that summarised the relative frequencies of all the state-transitions observed – so also generates impossible situations;
4. each student selects a state at random using their own empirically-derived individual state PMFs (explained earlier in section 3.3.2) – so also generates impossible situations;
5. as in 4 plus the students respond to the teacher (forced states) and impossible states are filtered out (but there are no adjustments to the probabilities to take into account the current situation and interactions as described section 3.3.2).

In models 3, 4 & 5 the final teacher and TA behaviour sub-models were used: it was the student behaviour sub-model that was being changed.

As an example, Figure 4-1 shows the results for one lesson (L#6 16070608MAT), but the pattern was the same for all lessons. Taking a range of twenty-seven parameter set values (so that the results were not specific to one parameterization), 150 replications were run for each. The left-hand chart shows the distance metric scores for each replication while the right-hand chart shows the cumulative average score. (Models 4 and 5 had two alternatives but those details are not relevant here.) The distance metric scores clearly decrease from model 1 (black) to 2 (red) to 3 (green) to 4/5, showing that the models were becoming better matched to the empirical data as the models become more complex and realistic. This showed that the metrics responded to the increased realism of the model. It also indicated that the increased complexity was worthwhile incorporating.

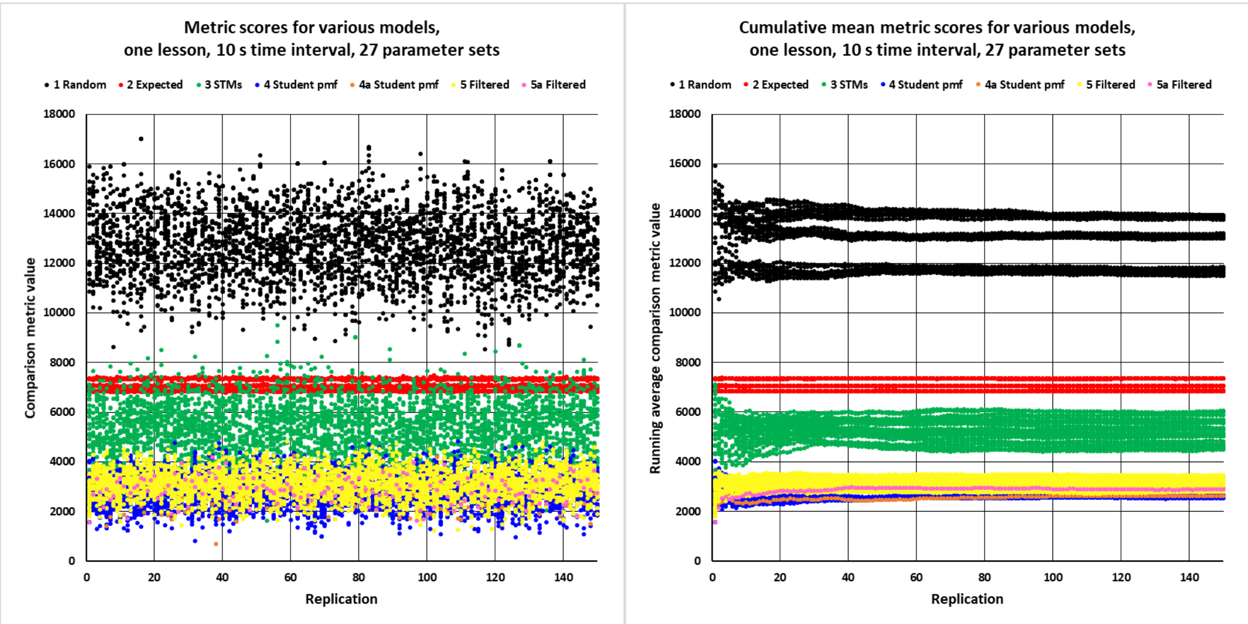

Figure 4-1 Comparison of distance metric results for different models showing that as the model becomes more realistic the simulations become more like the empirical lesson

4.3 The empirical data and acceptability ranges

Figure 4-2 shows the empirical distributions for the first six lesson comparison metrics (i.e., all except ESM) and %Productivity over all the lessons. The data showed that students were productive 90% of the time (modal class 90-95%), which meant that they were in one of the activity states prescribed by the lesson plan (as instructed by the teacher). The data also showed that for 60% to 70% of the lesson time students were in the expected state for the lesson section (not just productive in one of the acceptable states).

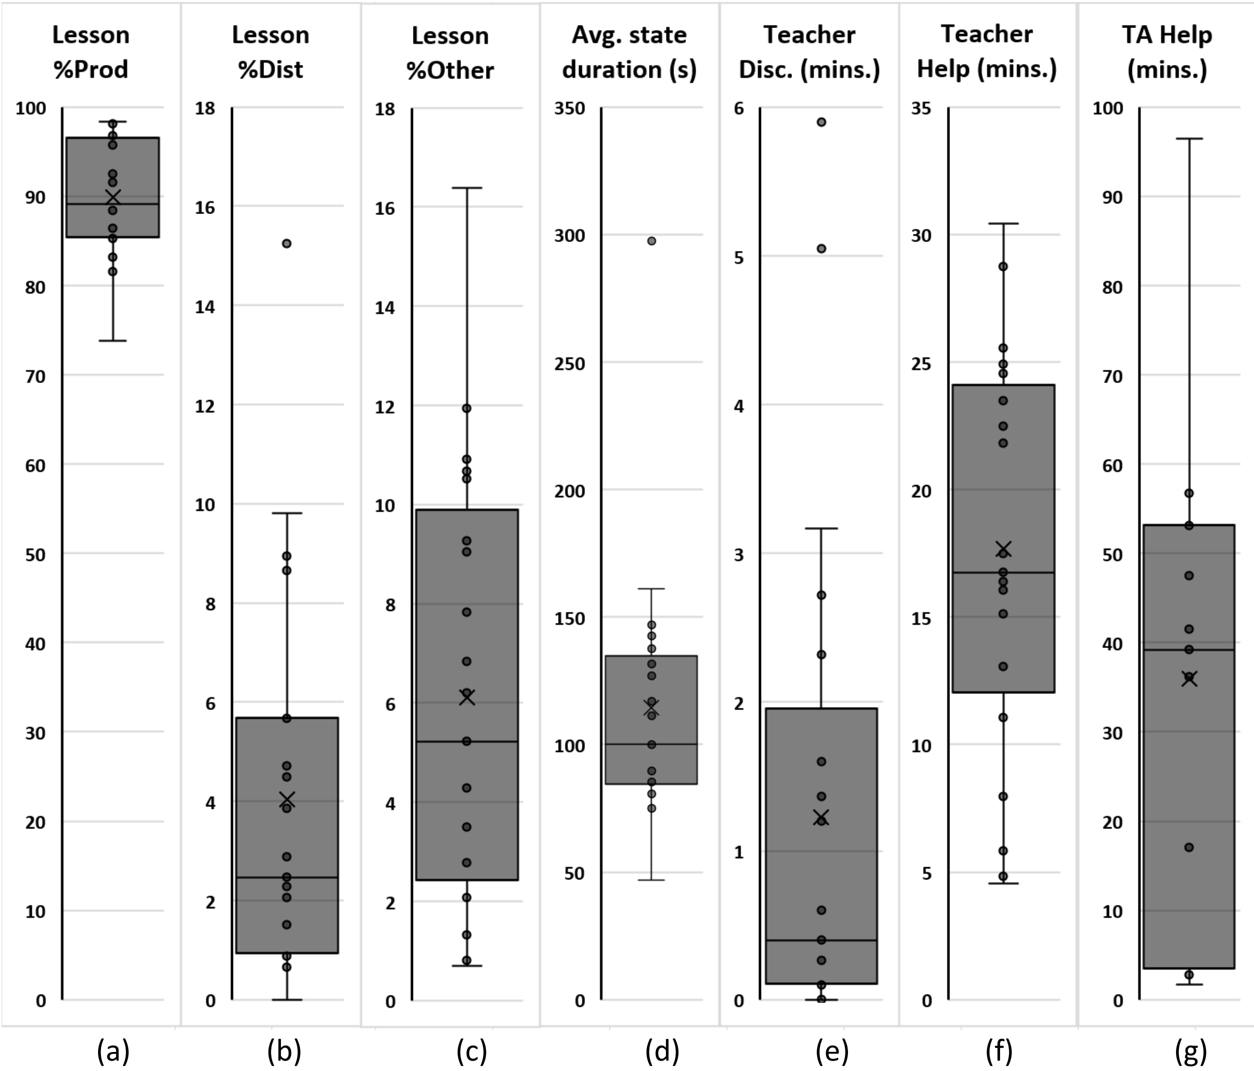

Figure 4-2 The empirical distributions for the first six metrics and %Productivity

The simulation outputs were not required or even expected to generate the exact empirical metric values – that would be highly suspicious and raise concerns about over-fitting. One would expect there to be some variability, some randomness, in both empirical and simulated lessons. A range of acceptability was defined for each metric by extending the empirical range by ±25%. The empirical values were thus viewed as intervals: (minimum value - 25%, maximum value + 25%) with the lower bound cut off at 0 because all the metrics are non-negative. These acceptability ranges are shown in Table 4-4.

Table 4-4 The acceptability range for each metric

Name	Description	Participants	Min	Max
%D	% of lesson student disruptive behaviour	students	0	20
%O	% of lesson student other behaviour	students	0	21
ASD	student average state duration (seconds)	students	35	375
TH	time teacher spent helping any student one-to-one (seconds)	teacher	204	2283
TD	time teacher spent disciplining an individual or the class (seconds)	teacher	0	443
TA	total time students were helped by the TA (minutes)	TA	0	121
ESM	the %match between student simulated and empirical states	students	n/a	n/a

Note that since students were in the expected state 60% to 70% of the time and in a productive state 90% of the time, this meant that a very simple, purely stochastic, student model that merely implemented these probabilities produced satisfactory %Productivity values. However, the other metrics were not as well matched and in addition this approach did not provide causal chains of interactions.

4.4 Number of replications

Running simulations entails running multiple replications of the same scenario. To emulate some of the aleatory variability apparent in the choices and durations in lessons, the simulation uses a pseudo-random number generator (PRNG) to generate streams of numbers. Each replication starts from a different seed. (A fixed sequence of seeds was used to enable simulations to be reproduced easily.) The results from many replications are aggregated to generate distributions and means for each output metric.

Although the goal was to ascertain whether each simulation metric mean fell within an associated empirical interval, it was vital to confirm that the estimate of the simulation mean was precise enough, i.e., close enough to the true simulation mean that would be found after infinitely many replications. The procedures recommended by Robinson (2014) were followed:

> ... more replications (samples) are performed until the interval becomes sufficiently narrow to satisfy the model user … (p184)
>
> The model user must determine what constitutes a sufficiently narrow interval. The number of replications is selected at the point where the interval reaches and remains below the desired level of deviation. (p185)
>
> ... it is important to obtain output data from more replications than are required in order to be sure that the cumulative mean line has flattened and that the confidence interval remains narrow. (p186)
>
> If there is more than one key response … the number of replications should be selected on the basis of the response that requires the most replications. (p186)

One can see from these instructions that the procedures involve some subjectivity.

The condition for a metric's cumulative mean to be an acceptable estimate of the simulation model's true mean was defined as follows:

The final cumulative mean has a 95% confidence interval that lies within the interval:

(max{0, final cumulative mean – 7.5%}, final cumulative mean + 7.5%).

As Robinson advised, because there were several metrics and each one would stabilize at its own rate, the number of replications chosen had to accommodate the slowest converger. Furthermore, this number needed to suit all lessons as these were separate instantiations of the CLSM and needed individual calibration and validation.

Another important consideration was the shape of the distributions obtained after various numbers of replications. Taking 1000 replications as representative of the 'actual' distribution, the number of replications was reduced to find when the shape of the resulting distribution was too dissimilar to the actual distribution. The plots in Figure 4-3 are an example. These show the distributions obtained for the distance metric from four different replication numbers. These were from just one lesson model and parameter set but the same result was seen across all lessons with many different parameter sets. One can see that the shape degrades as the number of replications decreases.

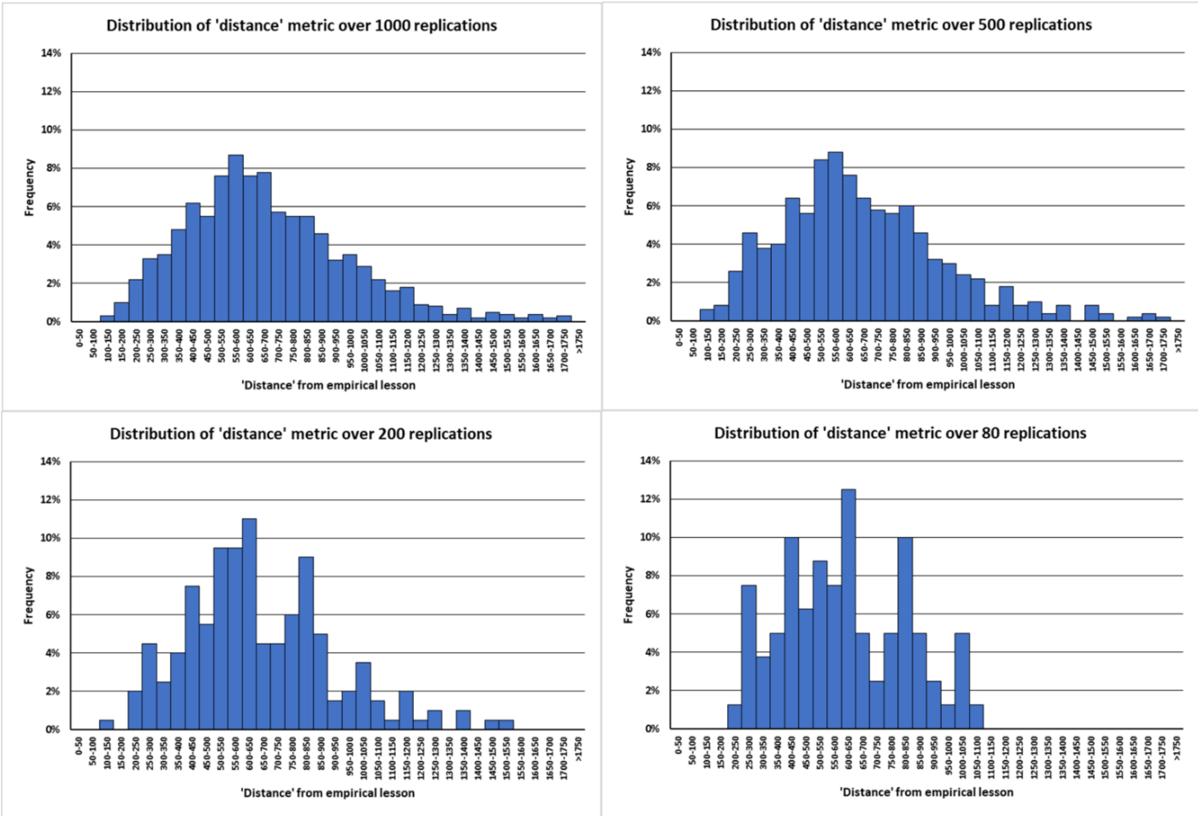

Figure 4-3 Example distributions of distance metric for 1000, 500, 200 and 80 replications

This process was quantified by running independent 2-sample Kolmogorov-Smirnov tests [13] to see if the overall distance metric was similarly distributed in the various situations. In all cases tested, the 500-replication distributions were found to be similar to the 1000-replication distributions, whereas some 200-replication distributions were found to be significantly different.

Putting all the evidence together, the large, conservative value of 500 was chosen for the number of replications.

4.5 Chapter summary

This chapter explained some topics that are common to both the model calibration and model validation stages in the ABMS development and hence useful to outline before discussing these stages in full. The main topic is that, although a generic behavioural CLSM had been developed, this model needed instantiating for each empirical lesson. This was the primary use of the lesson parameters: to tune the model to a specific lesson. The empirical lesson event data had been aggregated so that each student had their 'average' behaviour as their profile. The

[13] The Kolmogorov-Smirnov test is non-parametric so does not require a specific (e.g., normal) distribution for the data sets being compared. Also, it does not require datasets of equal size. It reports the maximum difference between the two cumulative distributions and an estimate of its significance.

lesson parameters adjusted the class's behaviour so that it more closely matched the specific lesson in question. Similarly, parameters adjusted the 'average' teacher and TA behaviour to the specific lesson. The outputs from a simulation are the state sequences over time of each agent, plus summaries of these sequences over the lesson. Several overall lesson metrics were identified, such as the proportion of the lesson that was disrupted or the amount of time the teacher spent helping individuals. The empirical data was discussed and the metric acceptability ranges defined. These ranges are one of the criteria used to decide whether simulation outputs are realistic. Finally, the process for determining the number of simulation replications was presented.

The following chapter explains the next stage in the ABMS development: how the lesson models were calibrated ('tuned') to the specific lessons using the metrics described above to compare simulation outputs and empirical lesson data.

5. Model calibration procedures

The next stage in the ABMS development was to calibrate the simulation model. The goal of model calibration (also known as parameter estimation) was to find, for each lesson, a set of model parameter values that generated output metric values that were realistic (i.e., satisfied the acceptability ranges explained in section 4.3), that matched (as closely as possible) the empirical lesson's values and did so stably (i.e., the model generated a smooth distribution of metric results). The acceptability criteria provided one check of the realism of the simulation. There may be outliers, but these should be plausible rare cases. This required exploring the parameter space, hunting for the best parameter settings. Techniques for parameter space exploration range from manual exploration through heuristic-guided to exhaustive search (Lee *et al.*, 2015). However, because the number of parameters was large and the appropriate range and resolution of the parameter values was unknown, the parameter space (the number of possible model configurations) was enormous. From the many search techniques that were available (Galán *et al.*, 2009; Salgado and Gilbert, 2013; Buwaya and Cleophas, 2015; Lee *et al.*, 2015) a pragmatic approach was adopted, using 'categorical calibration' (Thiele, Kurth and Grimm, 2014) to filter out parameter sets that failed the acceptable criteria (specified in section 4.3). Time constraints and resource limitations meant that the best parameter set may not have been found, but at least an acceptable one was.

The first stage in the search of the parameter space was to conduct a quick coarse grid-search and supplement this with finer grid-searches at manually-selected promising points. The procedure was:

A. Conduct a coarse grid-search of the parameter space, using 80 replications[14] per parameter set, but immediately eliminating a parameter set that caused any metric to exceed the acceptability ranges; from those that passed this test, rank the results and select a top set; for this top set run 500 replications.

B. Manually find a point (parameter set) in the parameter space that appears promising; conduct a fine grid-search around this point, using 80 replications per parameter set, but immediately eliminating a parameter set that caused any outputs to exceed the acceptability ranges; from

[14] After 80 replications most parameter sets in most lessons had metrics that seemed to stabilize, possibly converge, but at least remain tightly bounded.

those that passed this test, rank the results and select a 'top set' (see explanation below); for this top set run 500 replications.

C. Combine the 500-replication top set results from the coarse grid-search and manual search; rank the parameter sets and choose a final winning parameter set.

The procedure for choosing a top set of parameters involved the parameter sets being compared using the 7 metrics and the overall distance metric. As explained in section 4.2, simulation results were evaluated by calculating the absolute difference between the empirical metric values and the means of the simulation metric values (6 metrics), comparing student simulated and empirical state traces (1 metric) and comparing the overall distance metric. The procedure is described in the box below and then via an example.

Procedure for choosing a top set of parameters:

for each parameter set
    run simulations (for required number of replications) aggregating the difference metric values
    calculate the 7 difference metric means over all replications
    assign a rank on each of the 7 difference metric means (so each parameter set has 7 rankings)
    identify the worst ranking

sort the parameter sets on the worst ranking
select a top set of parameters based on their least worst ranking AND their overall distance metric
    - this choice is slightly subjective as it involved weighing up the two scores (see example)

The final step in the whole calibration procedure is to choose a single, winning parameter set. This is accomplished using the same, slightly subjective, logic of weighing up the least worst ranking and the overall distance metric.

The example in Table 5-1 is provided to help clarify the process. Suppose four parameter sets were to be compared on three difference metrics and that a top set of two was required. Each parameter set receives a ranking for each metric. The worst ranking for each parameter set is then identified. In this example, Param Set 4 has the least worst rank so it is selected for the top set. However, the next best parameter set, Param Set 2, has a better distance score (22 vs

33) and is only ranked 3 (the next in line), so it too is included in the top set. Note that Param Set 3 was discounted even though it had the same distance score as Param Set 4. The reason is it that its performance was unbalanced, 1 to 4. If this were the final stage of selecting a winning parameter set, Param Set 4 would be chosen.

Table 5-1 Example of procedure for choosing a top set of parameters

Param Set	Difference Metric A	Difference Metric B	Difference Metric C	Rank Metric A	Rank Metric B	Rank Metric C	Worst Rank	Distance Metric	Top Set
1	100	200	300	1	2	4	4	55	No
2	200	100	250	3	1	3	3	22	Yes
3	250	240	220	4	4	1	4	33	No
4	150	200	180	2	2	2	**2**	33	Yes

The reason this method was adopted was that a parameter set may have had a rank of 1, i.e., be the best, on one metric but perform poorly on all the other metrics (as in the case of Param Set 3 in the above example). Candidates for the top parameter sets were those with a low worst ranking, i.e., they gave a more balanced match with the empirical lesson on all metrics. The top set typically comprised 10 to 100 parameter sets, depending on how rapidly they differed on their worst ranking and their distance metric.

The final selection always ended with several equally viable parameter sets, i.e., parameter sets that resulted in overall lesson metrics being almost identical, so that there was no one distinct best parameter set. Even if one could find the absolute optimum parameter set (possibly set of sets) the output metrics it generates may differ so little from other parameter sets that the exact choice seems irrelevant.[15] The fact that a model can generate almost the same summary output metric results using different parameter values is often a concern (although it is also confirmation that over-fitting[16] has not taken place). However, this equifinality of simulation results was not found at the level of agent state sequences. Never was one simulation found to be identical to another given the same inputs (excluding pseudo-random number generator seed).

[15] Attempts to use cluster analyses to partition the parameter space based on metric outputs did not yield any significant grouping.
[16] The selection of highly specific parameter values in order to generate acceptable simulation results.

For some simulation models the inability to identify unique parameter values (parameter identifiability) is a serious issue (Evans, Heppenstall and Birkin, 2018), while for others there is a concern that the simulation produces acceptable results only for highly specific parameter values (Edmonds *et al.*, 2019). But Evans, Heppenstall and Birkin (2018) pointed out that if one is merely interested in consistent, accurate predictions one may not be worried about this, but if the model is being used to explain why specific results were obtained then the correctness of the causal explanations is important. The concern is that if more than one set of conditions can give rise to an effect, then inference from an effect back to an initial cause is dubious. However, in the context of lesson modelling where, for example, two parameters can increase the same output metric, this lack of identifiability was to be expected. For example, raising the teacher's inclination to offer one-to-one assistance and raising the students' tendency to ask for one-to-one assistance both have the desired effect of increasing the time spent on one-to-one assistance. Another example is that increasing the students' tendency to behave as expected is similar to, but not the same as, reducing the students' inclination to chat or disengage passively. These mechanisms were designed based on what was observed in lessons and were intended to make simulations more realistic. It is a strength of the model that it is able to substantiate alternative explanations for the same macro-level simulation outcomes.

Another relevant point is that in certain situations one parameter may be less effectual than another. For example, in a lesson where the students are constantly seeking help, the teacher's own tendency to offer help may be almost irrelevant: any value given to the TSOL parameter would yield almost the same results because the students' attributes make the teacher's inclination to offer support almost superfluous.

The simulation outputs were judged graphically and numerically as per Robinson (2014). All seven lessons were subjected to the same calibration procedures and all generated very similar results. As an example, using the winning parameter set found for Lesson #6, Figure 5-1 shows plots of the cumulative metrics, which highlight visually that the cumulative mean metric values tended to stabilize well over the replications. The scale of the graph makes the teacher's disciplining time appear to vary significantly, but closer inspection shows that it is fluctuating between 16 and 19 seconds, not of practical significance. Plotting the overall distance metric described in section 4.2 (the Euclidean distance of the point representing the vector of 7 difference metrics from the origin, the empirical data point), the overall stability of the cumulative metric mean over the replications can be seen graphically in Figure 5-2. The cumulative mean after 500 replications was 670 and the associated 95% CI was (645,696) which lies within the acceptable ±7.5% tolerance interval of [619,720]. These results show that the selected parameter set satisfied the calibration criteria.

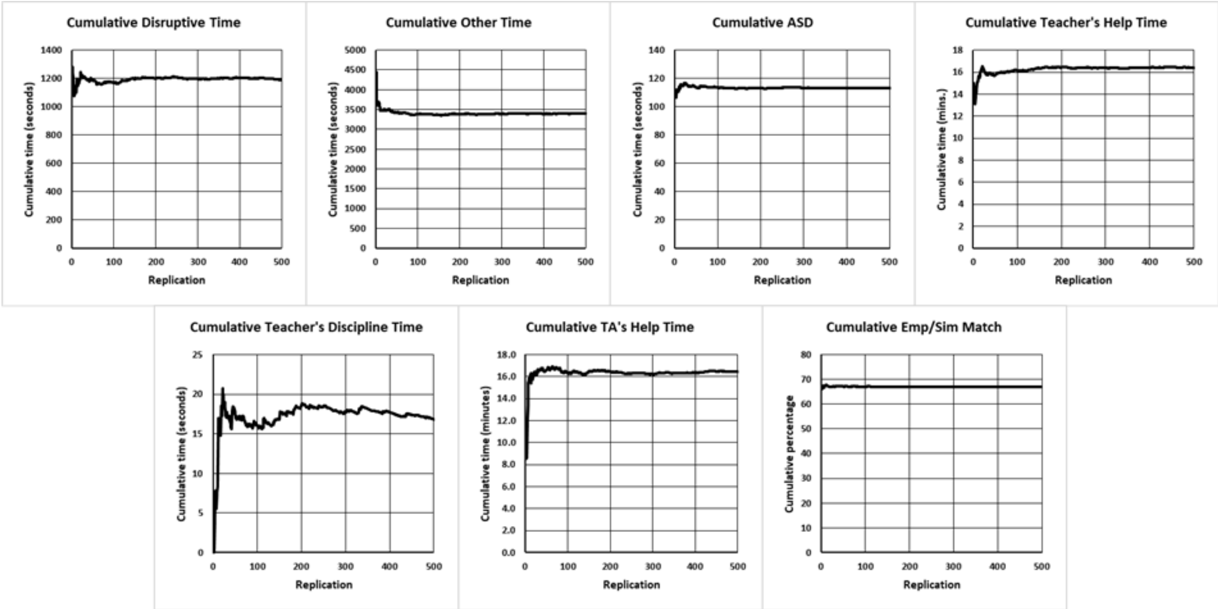

Figure 5-1 Cumulative metrics for Lesson #6 for one parameter set over 500 replications showing stability of convergence or at least boundedness

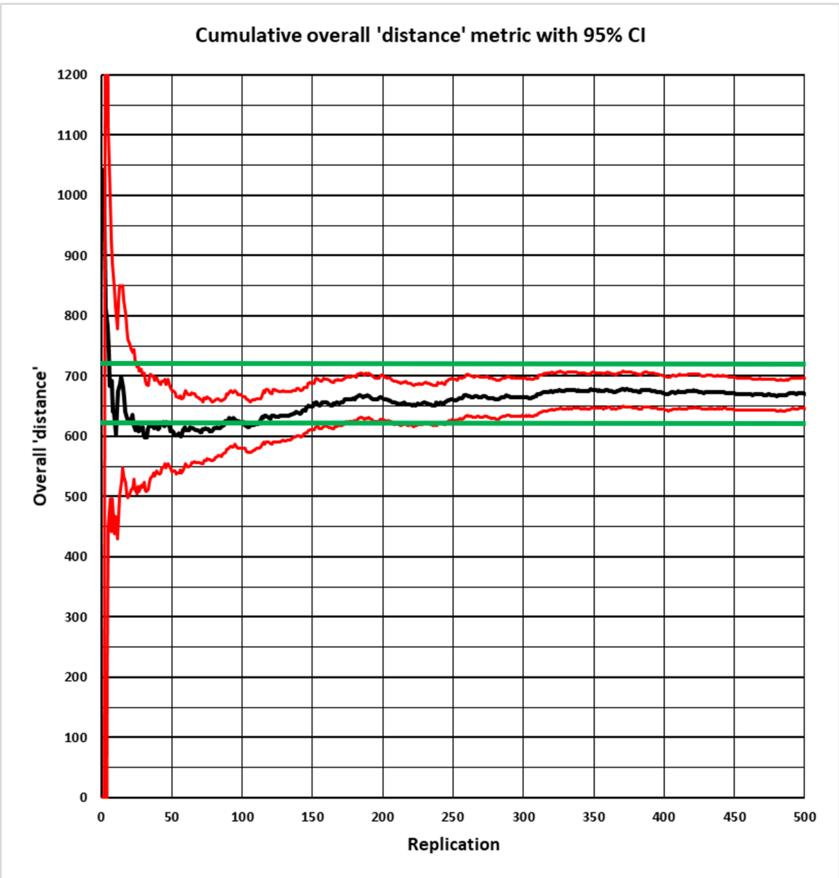

Figure 5-2 Cumulative overall distance metric for Lesson #6 with 95% CI (red lines) falling within the defined tolerance interval (green lines)

Figure 5-3 shows that the model generated a smooth, unimodal distribution of the overall distance metric over 500 replications (with bins of 50 units).

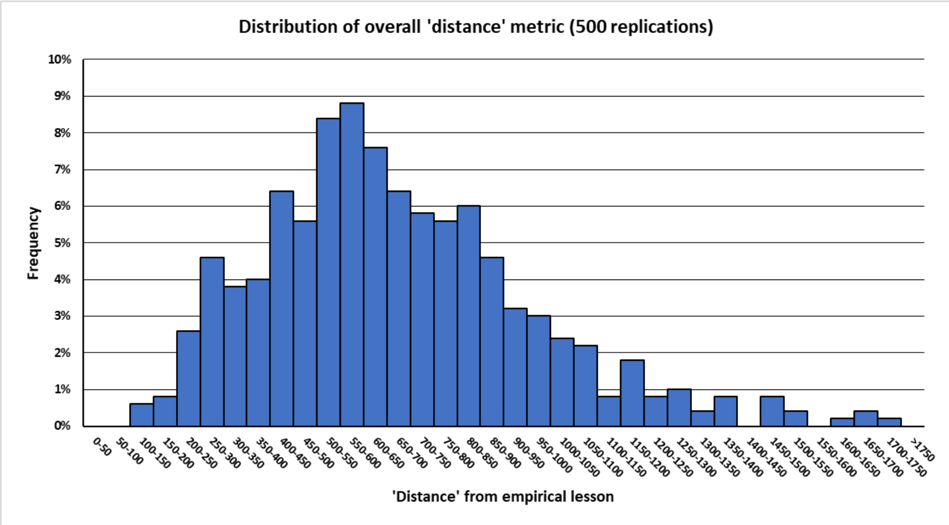

Figure 5-3 Distribution of the overall distance metric

Each of the seven selected lessons was successfully parameterized. A summary of the winning parameter sets for each is shown in Appendix C. Having calibrated the lesson models, the next stage was to validate the lesson models using the calibration parameter sets.

5.1 Chapter summary

This chapter explained the process of finding a parameter set for each lesson model. This process involved two steps. The first step was to run 80 replications and discard parameter sets that lead to a replication that produced unrealistic results (using the 'categorical calibration' approach), i.e., results that fall outside the acceptability ranges (explained in section 4.3). The second step was to explore the remaining parameters in finer detail over more replications (500). A parameter set still had to generate output metric values that were realistic in all replications, but it must also match (as closely as possible) the empirical lesson's values and do so stably, i.e., the model generated a smooth distribution of metric results. Of the many candidate parameter sets, the one that was least worst on all the metrics was chosen. Each of the selected seven lesson models was successfully calibrated.

The following chapter explains the next stage in the ABMS development process: how the calibrated lesson models were validated.

6. Model validation

Before the outputs from a simulation model can be trusted and before the simulation is used to run experiments, one needs to have established that the simulation outputs are good enough for the purposes required. This is the next stage in the ABMS development: to validate the simulation model. The lesson simulation results had been judged realistic overall, but the goal of model validation is to increase confidence in the simulation model and its results by demonstrating that it is sufficiently realistic and reliable (Balci, 1995; Robinson, 1997, 2008a; Siebers and Onggo, 2014). Typically for empirically-based ABMSs this means that the simulation outputs are compared to empirical data to confirm there is an adequate match, and, in addition, the plausibility of the micro-level agent interactions is confirmed, including testing face validity. If the lesson model passed all these tests it would be considered a plausible simulation of that specific empirical lesson. As mentioned before, a lesson model could be successfully calibrated but fail validation.

But to *match* or be *realistic* are often quite subjective concepts, even though they might be unambiguously mathematically defined for the process – as was done in section 4. This will depend on many factors, including the accuracy and precision of the empirical data. 'Despite its apparently scientific nature, modelling is a matter of judgement' (Salgado and Gilbert, 2013, p. 254). While verification and validation increase confidence in a model, they do not prove absolute accuracy (Robinson, 1997). Despite a lesson model passing all the tests it was given, each model was not tested in all possible scenarios, hence it was not absolutely validated (Balci, 1995).

As mentioned in section 1.5.1, validating a simulation model involves validating the conceptual model, white-box validation (looking at the micro-level of agent interactions to establish internal validity) and black-box validation (looking at the outputs of the system at the macro- and meso-levels to establish external validity). In all simulation modelling one needs to confirm the structural and behavioural validity of the model (Qudrat-Ullah, 2005, 2008; Edmonds *et al.*, 2019). There is a risk that the explanation provided by the simulation model is given in terms of model constructs/rules that are not actually real, perhaps based on actual mechanisms but too simplified, so that the outcomes, while realistic, are realistic for the wrong reasons (c.f. a true conclusion from false premises). It might also be the case that the model is just one of many models that could generate the same outcomes. These are concerns about the 'structural validity' of the model: a model may generate excellent results, but it is vital that the model exhibits the 'right behavior for the right reasons' (Qudrat-Ullah, 2005, 2008).

As Norling *et al*. (2018) wrote:

> With an explanatory model, if one has demonstrated that a certain set of assumptions can result in a set of outcomes (e.g. by exhibiting an acceptable fit to some outcome data), this shows that the modelled process is a possible explanation for those outcomes. Thus, the model generates an explanation, but only in terms of the assumptions in the setup of the simulation. If these assumptions are severe ones, i.e. the model is very far from the target phenomena, the explanation it suggests in terms of the modelled process will not correspond to a real explanation in terms of observed processes. The chosen assumptions in an explanatory model are crucial to its purpose in contrast to the case of a predictive model—this is an example of how the purpose of a model might greatly influence its construction. (Norling, Edmonds and Meyer, 2018, p. 65)

Structural validation of the conceptual model (described in section 6.1) is the first step in mitigating for these concerns. Behavioural validity of an ABMS developed to explore a theory and/or generate explanations can be tested using experiments. Similar to the type of tests that a predictive simulation would undergo, the simulation can be given scenarios like the ones it was developed from and the behaviour at all levels (overall system output (macro-level), overall individual agent behaviour (meso-level), individual agent behaviour and agent-interactions (micro-level)) should be scrutinised. For example, if a lesson without a TA had one added, the conceptual model anticipates certain consequences. One could investigate whether the anticipated results were evident in all levels of the simulation results and whether the simulated interactions seemed plausible. Part of this process involves face validity testing by experts – discussed in section 6.3. Successful results in these activities would increase confidence in the realism of the behaviours generated by the simulation model.

Three schoolteachers participated in conceptual model validation and simulation model face validity tests. Teacher1 and Teacher2 knew the school, the classrooms, the teachers and the students, and had participated in the case study, collecting lesson event data (section 1.5.2). Teacher1 had more experience at the school than Teacher2, who was a trainee teacher and ex-pupil. Teacher3, a highly-experienced teacher, had no knowledge of the school or people.

6.1 Conceptual model validation

A simulation model is not validated merely by checking that the outputs seem realistic (black-box validation): the core conceptual model must be validated too. It is important to confirm that the model rules adequately represent real-world behaviour. One way to accomplish this is to have experts confirm the assumptions upon which the model was built and endorse the proposed rules specifying agent behaviour. This attends to the requirement for structural validity mentioned above. The goal is to reduce the chances that realistic simulation outputs are generated by constructs/rules that are not real, perhaps being too simplistic (Edmonds *et al.*, 2019), or that crucial facets have not been omitted from the model (Norling, Edmonds and Meyer, 2018). For example, it was possible to calibrate a lesson model that omitted the rules that manage students asking for help (so only the teacher offering help was modelled), but such a model was then missing a key aspect of real classroom behaviour.

The core conceptual model - including the overall simulation algorithm - was informally validated individually by the three teachers who took part in the face validity tests (described below). Each teacher was shown the lists of activity states that had been identified for students, teachers and TAs. The teachers all independently agreed that this list was comprehensive and matched their experiences of teaching (and being taught). Each teacher was shown the behavioural rules in narrative form and was asked whether the rules seemed plausible and whether there should be other rules. The teachers all considered that the rules did reflect the reasons and mechanisms for specific behaviour in lessons, for students, teachers and TAs. They sometimes suggested additional rules but then acknowledged that the simplified rules they had been shown covered the most common situations in lessons.

6.2 Simulation model validation procedures

For a lesson model to be considered valid, 500 replications of the lesson were aggregated and the macro- and micro-level validity criteria below were applied to the outputs. Note that the model calibration procedures had already ensured that:

- all simulation output metrics fell within their acceptability ranges (specified in section 4.3);
- the means of all the metrics (including distance) converged or stabilized within the specified 95% confidence interval.

<u>Macro-level validity</u>

- Simulation metric distributions encompass the empirical metric values;
- Simulation metrics means are within ±25% of the empirical value (although an exception was made for TD, the time in seconds that the teacher spent disciplining, as it was felt that, for example, comparing means of 17 s for simulations and 9 s for the empirical to within 25% was artificially precise).

Micro-level and macro-level validity

- In face validity tests (discussed below), a lesson model was considered validated only if all the teachers considered both the simulations and the empirical lesson replay to be realistic; the two simulations were the replication with the worst overall distance score and one from the modal class of distance scores.

6.3 Face validity testing

The face validity tests effectively enabled some degree of both macro-level and micro-level checking as one could observe individual interaction sequences as well as see the overall picture. The basic objective of the face validity testing was to find evidence to refute the claim that the lesson animations were plausible. This could be done by observing an event or interaction that was inconsistent or impossible or implausible.[17] It was essential that the animations had no glitches, ensuring that there was nothing that would give away to the viewer that a lesson was simulated vs an empirical lesson replay.

To establish face validity, for each selected lesson, three animations were presented (in a random order) on a webpage [18] (see Figure 6-1):

- the empirical lesson replay;
- a typical replication, one from the modal class of the overall distance metric;
- the worst replication, the one furthest from the empirical lesson (largest overall distance metric).

Three teachers participated in the face validity tests, two of whom had also assisted in data collection during the case study (described in section 1.5.2). They were asked to decide whether an animation represented a realistic and plausible lesson, or not, and whether an animation was an actual lesson replay, or not. They were told that all three animations could be simulations. The tests were thus a form of *Turing test* to see if the teachers could discriminate real lessons from simulations (Xiang, Kennedy and Madey, 2005). They were also asked to explain what strategies they had used to make their decisions. There were no clues as to what the 'correct' answers were, or what outcomes the author might be hoping for. So, even if they had a motive (e.g., to please the author) the three had no way of compromising the tests (other than to spoil

[17] The animations also made it possible to detect and debug irregularities in agent behaviours visually.
[18] The website files and all the animation video clips are available at Ingram (2020b).

them by taking decisions randomly – and they would still have to explain their decisions). They were given an initial training session to explain what was required and in which they were shown how to control the animation replays and how to interpret the colours and interaction links, etc.

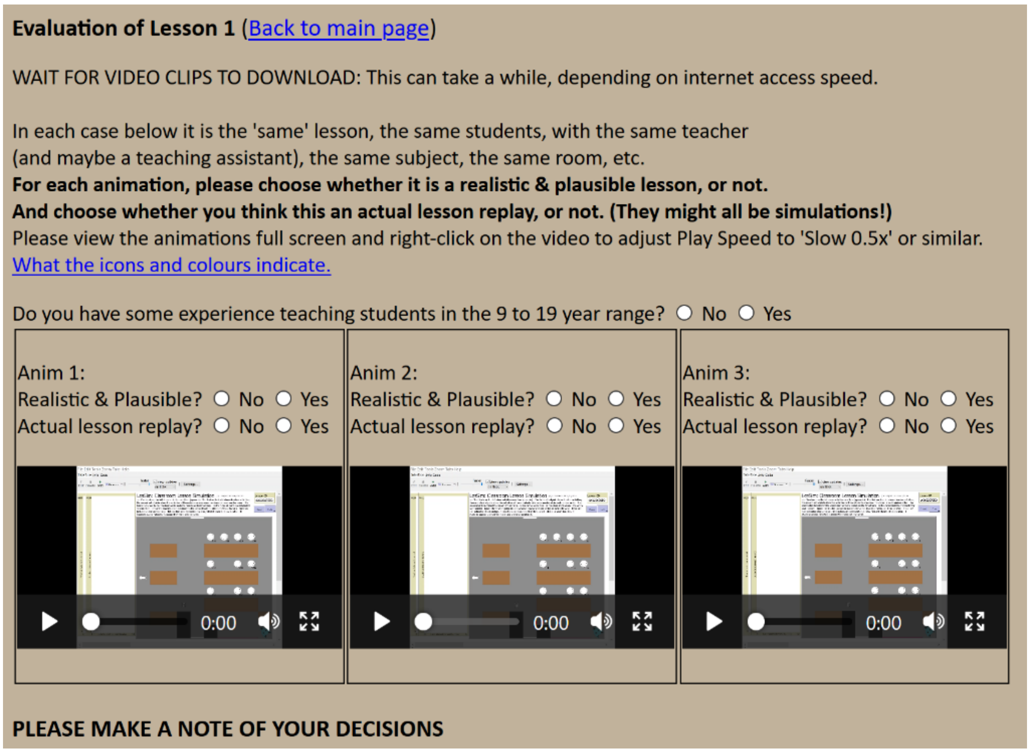

Figure 6-1 Example of face validity test showing the three animations to be judged
Note that the animations are to be viewed full screen

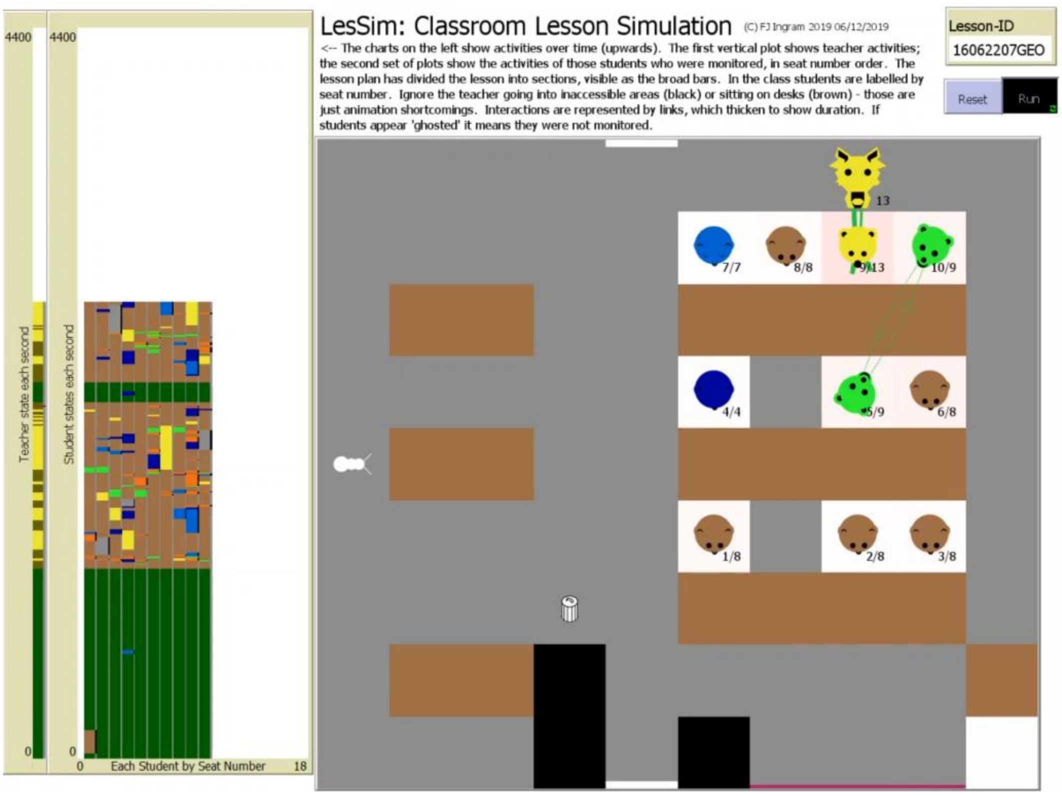

Figure 6-2 Screen capture of an instant in one animation face validity test

Figure 6-2 shows one of the animations mid viewing. Each of the lessons was accessed from a central webpage which contained some explanatory context. They were able to slow down and replay the animations in order to examine interactions more closely. The results of the face validity tests are listed in Table 6-1.

Table 6-1 Responses of the teachers (in rows) to the 3 animation types (typical, worst, actual) for the seven selected lessons

Lesson	Teacher	Typical simulation plausible/realistic?	Worst simulation plausible/realistic?	Actual lesson realistic?	Actual lesson identified?	Face validity verdict?
#1	1	Yes	Yes	Yes	Yes	Yes
	2	Yes	Yes	Yes	No (chose worst)	
	3	Yes	No	Yes	No	
#2	1	Yes	Yes	Yes	Yes	Yes
	2	Yes	Yes	Yes	No	
	3	Yes	No	Yes	Yes	
#3	1	Yes	Yes	Yes	No (chose worst)	Yes
	2	Yes	Yes	Yes	No	
	3	Yes	Yes	Yes	Yes	
#4	1	Yes	Yes	Yes	Yes	Yes
	2	No	No	Yes	Yes	
	3	Yes	Yes	Yes	Yes	
#5	1	Yes	Yes	Yes	No	Yes
	2	Yes	Yes	Yes	No	
	3	Yes	Yes	Yes	No	
#6	1	No	?	Yes	Yes	No
	2	No	Yes	Yes	Yes	
	3	Yes	Yes	Yes	Yes	
#7	1	Yes	Yes	Yes	Yes	No
	2	No	No	Yes	Yes	
	3	No	Yes	Yes	No (chose worst)	

Teacher1 identified 5 of the 7 actual lesson replays and considered all the animations (apart from Lesson #6) to be realistic. Teacher1 and Teacher2 both explained that their decision-making was influenced by their knowledge of the students and the teacher. They recognised the classroom layout and speculated where particular students sat. In effect they were comparing the animations against their own mental model of that class's real lessons with that teacher. They sometimes explained that they thought that the animations showed plausible behaviour in general, but that the behaviour was not plausible for that class in that lesson. Their decision-making about being realistic or the empirical lesson included observations such as:

- that is just like Teacher X, disciplining the class at the start of the lesson;
- that is just like Teacher X, talking for so long (or maybe they are showing a video?);
- that is far too long for the students in that class to sit without disengaging or chatting.

Teacher3, who was unfamiliar with the school and students, came to some different conclusions based on her experiences. Some actual student behaviour was considered unrealistic (e.g., students would not sit so long without chatting), some was considered realistic even though it was the worst match with the empirical lesson.

It seemed that a teacher may judge the animations to be plausible - because they show what could happen in general - but nevertheless reject the simulation on the grounds that it does not fit with their experiences of that specific class, teacher, TA and classroom.

6.4 Validation results

On completion of the validation procedures, four of the seven lesson models were considered validated:

Lesson#1	16062207GEO	Year 7	Geography
Lesson#3	16070508MAT	Year 8	Maths
Lesson#4	16070510MAT	Year 10	Maths
Lesson#5	16070607MAT	Year 7	Maths

Lessons 2, 6 and 7 failed validation. Table 6-2 shows a summary of the results of validation.

Table 6-2 Simulation model validity test results (with validated lesson models highlighted)

Lesson#➔	1	2	3	4	5	6	7
Lesson ID	16062207GEO	16063010SCI	16070508MAT	16070510MAT	16070607MAT	16070608MAT	16070809SCI
Proportion of students monitored	10/10	10/10	16/16	8/12	8/10	8/14	7/7
TA present?	No	Yes	No	Yes	Yes	Yes	No
Distributions visually encompass the empirical metric values?	Yes	No	Yes	Yes	Yes	Yes	No
Mean metric values visually sufficiently close to empirical values, i.e., within ±25%?	All except TD (7s vs 24s)	No	Yes	Yes	Yes	(Yes)	No
Conclusion: the lesson simulation model has macro-level validity (i.e., closely matches empirical lesson outcomes)?	Yes	No	Yes	Yes	Yes	Yes	No
Face validity testing results: plausible and realistic?	Yes	Yes	Yes	Yes	Yes	No	No
Is the explanation of the worst replication plausible?	Yes	-	Yes	Yes	Yes	-	-
Conclusion: simulation model valid? The lesson simulation model is producing plausible and realistic lessons AND these are sufficiently similar to the empirical lesson.	Yes	No	Yes	Yes	Yes	No	No

Two examples of validation procedure results summaries are shown below, one for a lesson model that failed (Lesson #2) and one for a lesson model that passed (Lesson #5). Although Lesson #2 passed the face validity tests, it failed validation because, apart from %Dist (%D) and TD, the metric distributions did not adequately encompass the empirical metric values (marked with a red x in the boxplots in Figure 6-4) and the simulation mean metric values (black x) were not sufficiently close to the empirical metric values.

There are many reasons that might explain why a lesson model would fail validation. For example:

- the model may be generally adequate but it had not been sufficiently well calibrated – a better parameter set is needed;
- the metric acceptability criteria were too broad, enabling a model with inappropriate parameter values to pass calibration tests;
- the model rules as simplifications and abstractions of reality might just be too inaccurate;
- the empirical student behaviour in that lesson might have been too unusual;
- there were too few empirical data to assign reliable student state probabilities to those students.

Lessons #2, #6 and #7 could have failed validation for any of these reasons.

Table 6-3 Lesson #2 validation results

<table>
<tr><td>Lesson Validation: Lesson #2</td><td>LID 16063010SCI</td></tr>
<tr><td colspan="2">

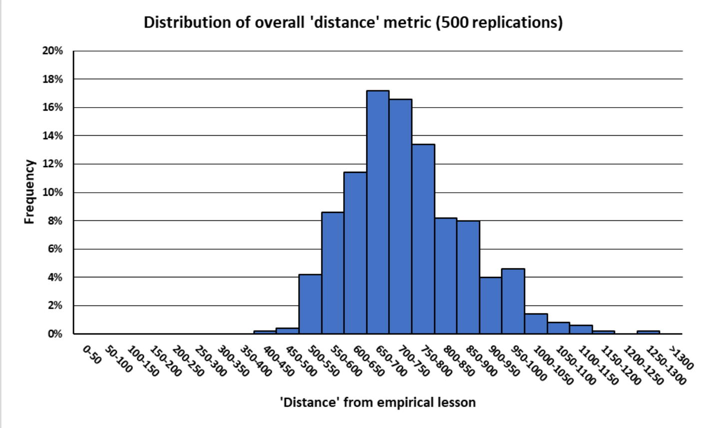

Figure 6-3 Lesson #2 Distribution of the overall distance metric

</td></tr>
<tr><td colspan="2">

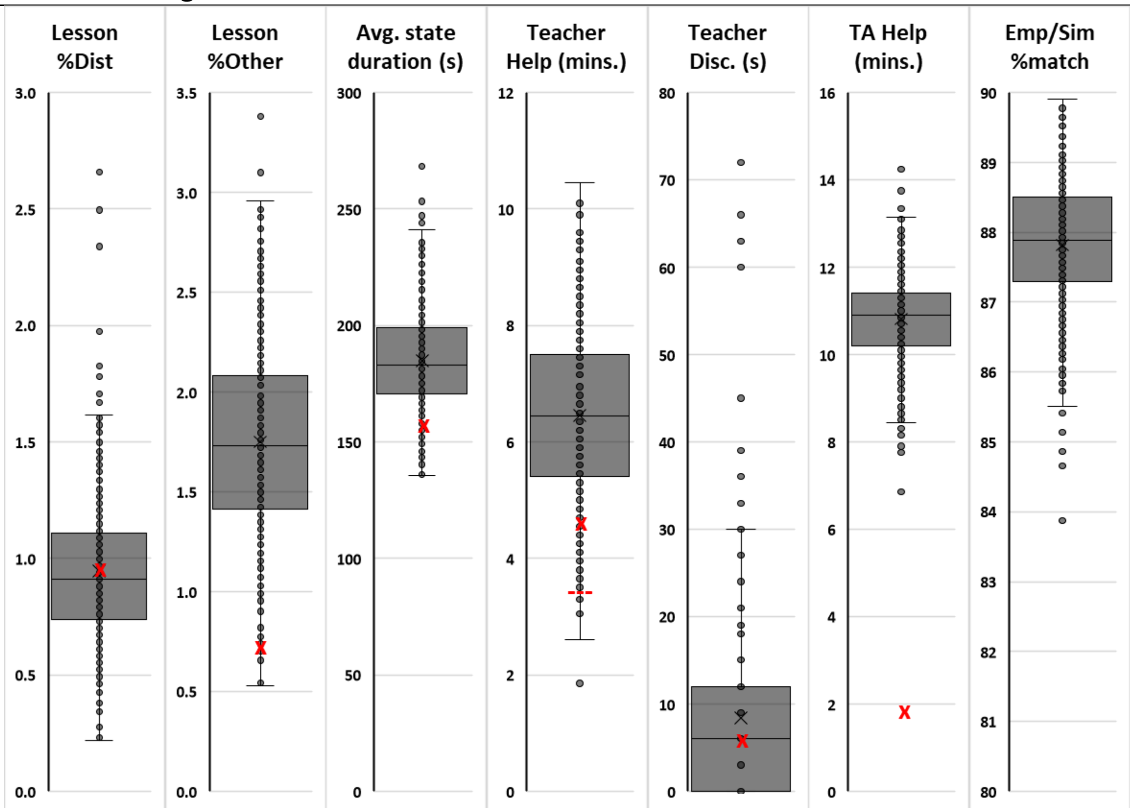

Figure 6-4 Lesson #2 Distributions for the 7 metrics for lesson over 500 replications (with the empirical lesson value shown as x)

</td></tr>
<tr><td>Do distributions encompass the empirical metric values?</td><td>Not adequately</td></tr>
<tr><td>Are mean metric values sufficiently close to empirical metric values, i.e., within ±25%?</td><td>No
(only %Dist and TD)</td></tr>
<tr><td>Face validity test of worst overall distance replication?</td><td>Pass</td></tr>
<tr><td>Face validity test of example replication from modal class (650,700) of overall distance distribution?</td><td>Pass</td></tr>
<tr><td>Consider lesson model validated?</td><td>No</td></tr>
</table>

Table 6-4 Lesson #5 validation result

Lesson Validation: Lesson #5	LID 16070607MAT

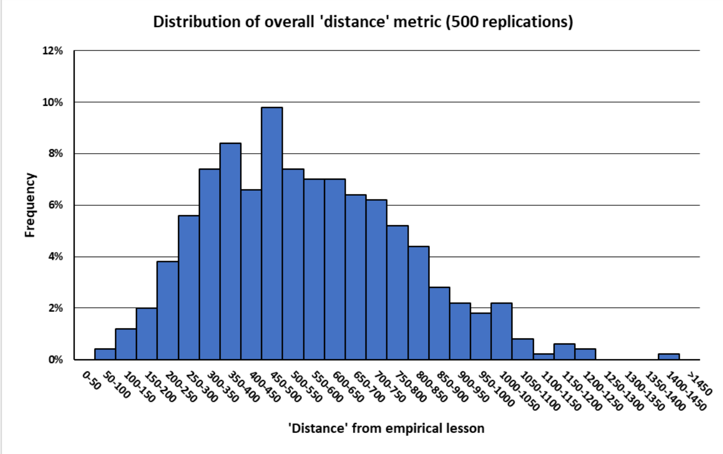

Figure 6-5 Lesson #5 Distribution of the overall distance metric

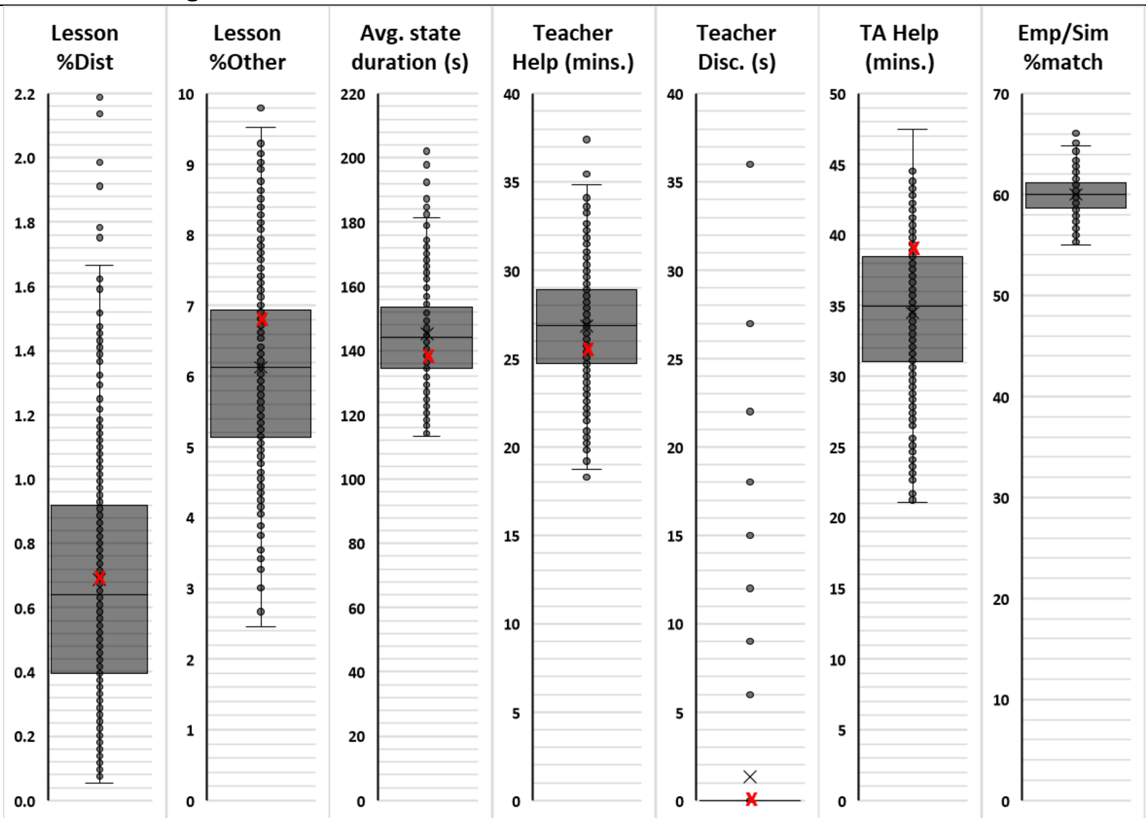

Figure 6-6 Lesson #5 Distributions for the 7 metrics for lesson over 500 replications (with the empirical lesson value shown as x)

Do distributions encompass the empirical metric values?	Yes (although TD skew)
Are mean metric values sufficiently close to empirical metric values, i.e., within ±25%?	Yes
Face validity test of worst overall distance replication?	Pass
Face validity test of example replication from modal class (450,500) of overall distance distribution?	Pass
Consider lesson model validated?	**Yes**

6.5 Chapter summary

This chapter explained the model validation procedures undertaken to extend confidence in the lesson models. This is necessary in order to have enough evidence that a simulation can be trusted for use in experiments. This included consulting three other teachers to confirm the plausibility of the core conceptual model, the modelling assumptions and the behavioural rules. Following calibration, the lesson simulation results had been judged realistic at the overall lesson level (macro-level), now they were checked to see if they were close enough to the specific empirical lesson results to be considered a plausible simulation model of that specific empirical lesson. Plausibility was investigated at the micro-level of agent interactions and the meso-level of overall behaviour for individual students. The three teachers also took part in 'face validity' testing, visually inspecting the simulation and empirical lesson animations and results, looking for anything that was inconsistent or impossible or implausible. Note that a lesson model could be successfully calibrated but fail validation: in fact, only four of the seven calibrated lessons were judged to be sufficiently like their empirical lessons that they could be used in the research experiments. But just because a lesson model passed the validation tests does not mean that the model is correct.

The following chapter explains the next stage in ABMS development for the four validated lesson models: sensitivity analysis - the analysis of the models' responses to changes in parameter values and student base state probabilities (e.g., inclination to chat) to check that the model does not exhibit any undesirable behaviour.

7. Sensitivity analyses

As described in section 1.5.1, sensitivity analyses are conducted as part of the ABMS development methodology. Some basic sensitivity analyses (Robinson, 2013) were conducted to assess the impact on simulation outputs of changing parameter values and student base state probabilities (e.g., inclination to chat). The latter would also mimic the effect of errors in the empirical data. In addition to describing the results of these sensitivity analyses, the relative importance of the parameters is discussed. Figure 7-1 summarizes the 11 parameters that can be adjusted and the 6 output metrics they affect.

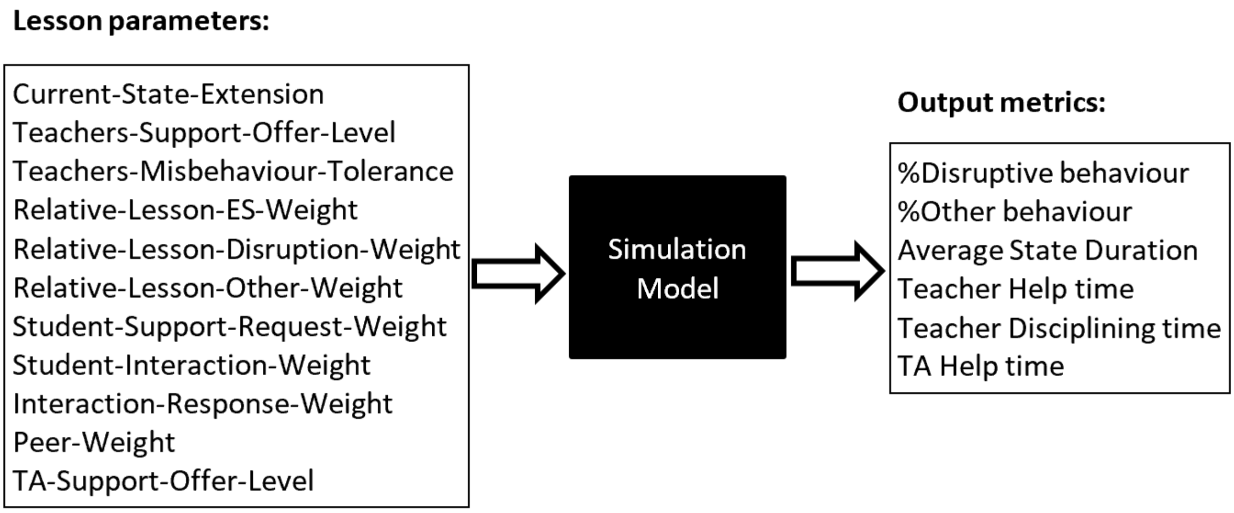

Figure 7-1 The simulation model as a black box with output metric values resulting from lesson parameter values

Table 7-1 The student state choices affected by each parameter

Parameter	Abbrev.	Student State Choices
Teachers-Support-Offer-Level	TSOL	13
Teachers-Misbehaviour-Tolerance	TMT	2, 3, 4, 7
Relative-Lesson-ES-Weight	RLESW	current expected state (either 8, 9, 12 or 14)
Relative-Lesson-Disruption-Weight	RLDW	2, 3
Relative-Lesson-Other-Weight	RLOW	4, 7
Student-Support-Request-Weight	SSRW	13
Student-Interaction-Weight	SIW	3, 9
Interaction-Response-Weight	IRW	3
Peer-Weight	PW	all except 1, 5, 6, 11, 16, 17
TA-Support-Offer-Level	TASOL	15

Table 7-1 summarizes which parameters affect which student state choices. Current-State-Extension affects all state durations but not choices, whereas Peer-Weight affects all student state choices (except those that are 'forced').

The following section summarizes the sensitivity analyses on these parameters.

7.1 Parameter sensitivity

Full factorial investigation was infeasible due to the number of parameters and output metrics, so 'one-factor-at-a-time' (OAT) analyses (Thiele, Kurth and Grimm, 2014) were conducted, followed by bivariate analyses during research Experiment 1 (section 8.1). It was recognized that this would limit the ability to investigate interaction effects between parameters and that this would limit conclusions (Lorscheid, Heine and Meyer, 2012).

Investigating the effect of each of the 11 parameters on each of the 6 output metrics for all four validated lessons resulted in 264 data series (with each datapoint being the mean of 80 replications). The charts in Figure 7-2 show example results from Lesson #1. The results are summarized in Table 7-2. Overall, the output metrics all changed smoothly as parameter values were incremented by small amounts – with the exception of TD which often showed non-linear fluctuations. Some parameters had little or no impact at all, e.g., PW. Others (intentionally) had a very specific effect, e.g., TASOL affected TAH only. The impact of parameter values on ASD is relatively simple to understand: those parameters that increased the entropy in behaviour (e.g., chatting) reduced ASD, whereas those parameters that increased the choice of the expected state (e.g., RLESW) increased ASD.

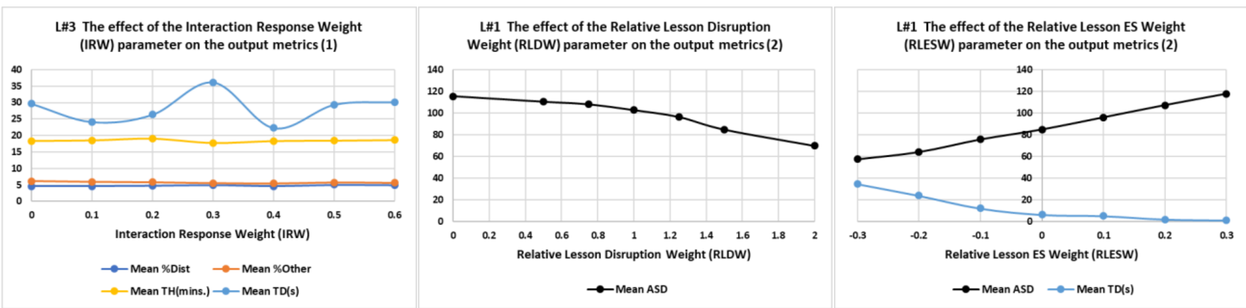

(a) (b) (c)

Figure 7-2 Example parameter-metric relationships from Lesson #1
(a) the effect of IRW on four metrics; the effect of (b) RLDW on ASD;
(c) the effect of RLESW on ASD and TD

Table 7-2 Consequences of stepwise increments in parameter values

Parameter	Consequences of stepwise increments in parameter values
CSE	little or no effect on most metrics except ASD; smooth increase in TH; TD showed fluctuations (especially Lesson #3)
TSOL	affected all metrics smoothly except in Lesson #1 ASD had one fluctuation and TD showed fluctuations
TMT	affected all metrics smoothly
RLESW	affected all metrics smoothly; little or no effect on TAH
RLDW	affected all metrics smoothly; little or no effect on TH
RLOW	affected all metrics smoothly except in Lesson #1 TD showed fluctuations
SSRW	little or no effect on most metrics except TH smooth increase; in Lesson #3 TD showed fluctuations
SIW	little or no effect on most metrics except smooth effect on ASD, TAH and %Other; TD showed some fluctuations
IRW	little or no effect on all metrics except smooth effect on ASD and TD showed some fluctuations (especially Lesson #3)
PW	little or no effect on all metrics except smooth effect on ASD; TD showed some fluctuations
TASOL	little or no effect on all metrics except smooth effect on TAH and ASD; TD showed some fluctuations

The fluctuations in the TD metric were expected and desired consequences of the nature of the rules that control the amount of teacher disciplining. Many have thresholds that must be crossed before the teacher disciplines. But these thresholds are not fixed - they have a random component. The purpose of these random fluctuations was to make the teacher agent more realistic, e.g., waiting different amounts of time before intervening or needing different numbers of students to be involved (as discussed in section 3.4.1). But the use of one PRNG seed will cause a different sequence of pseudo-random numbers, and one sequence might trigger a chain of agent interactions whereas another sequence might not: it will depend on whether a variable threshold has been crossed. The effect is that a very small change early in a lesson can significantly alter the remainder of the lesson.

In conclusion, the lesson models seemed appropriately well-behaved in their response to parameter value changes.

7.2 Student state attribute value sensitivity

To investigate the consequences of adjusting students' base state attribute values (their empirical state probabilities), a (relatively) chatty student was chosen in each of the four validated lessons. Their base state 3 (chatting) probability was scaled by multiplying by a factor ranging from 0 to 3 (so 1 corresponds to their empirical probability). Note that even when that probability is adjusted to 0, some chatting happens in response to a chat request from other students. In all the charts in this subsection, the four plots presented are for four different students: (a) Lesson #1 student 710 (the chattiest of the four), (b) Lesson #3 student 802, (c) Lesson #4 student 1008 and (d) Lesson #5 student 702 (the least chatty of the four). Each coloured boxplot shows the distribution obtained from 80 replications, for each scaling factor (0, 0.2, 0.4 ... 2.5, 3).

The first step in this sensitivity analysis procedure was to confirm that turning up the state 3 chat probabilities has the desired effect of increasing a student's chatting time. This is shown for all four students in Figure 7-3 by the boxplots shifting upwards. The charts also show that the changes are smooth. The slight fluctuations are due primarily to the teacher's response, whether more disciplining results from the increased disruption or not.

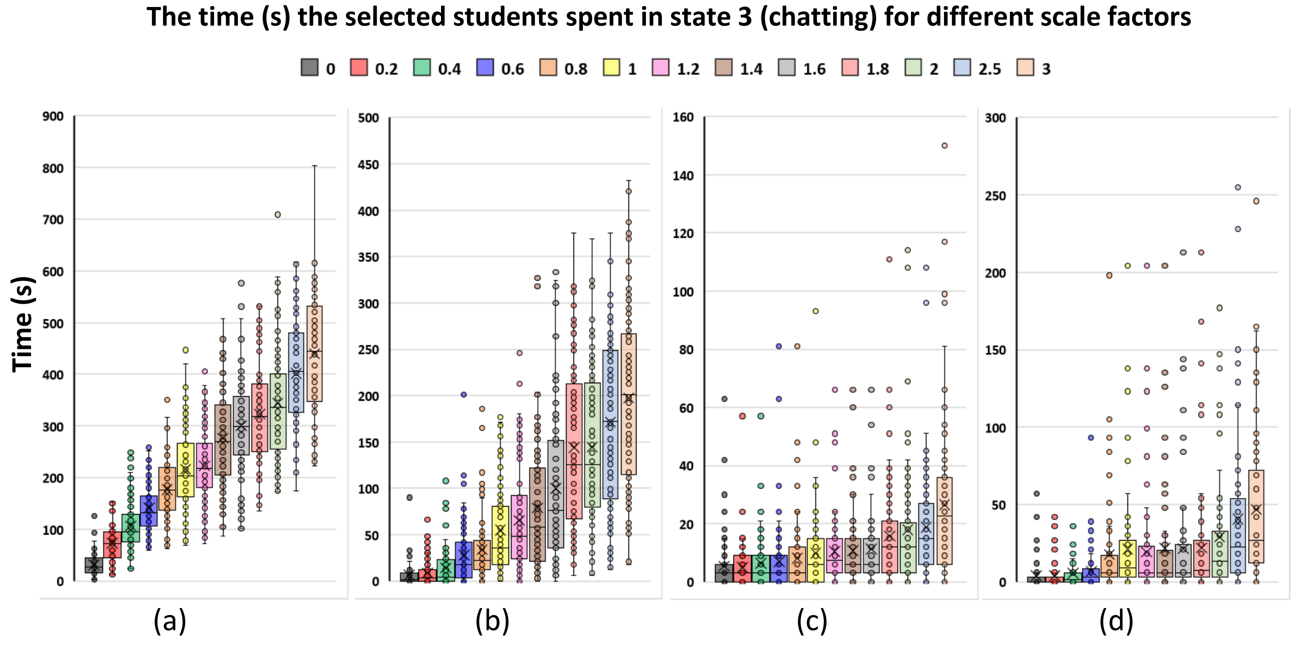

Figure 7-3 The increase in the four students' state 3 chatting time as their empirical state 3 probability was increased

The effects of changes to each student's chattiness on the overall lesson are shown in Figure 7-4 and Figure 7-5. The vertical axis in Figure 7-4 is %Disruption (%D) and in Figure 7-5 it is %Productivity. From the (a) plots one can see that this rather chatty student has affected the whole class in terms of %Disruption and %Productivity by 1 or 2 %. The effect of less chatty students (b, c and d) is much less and appears mixed.

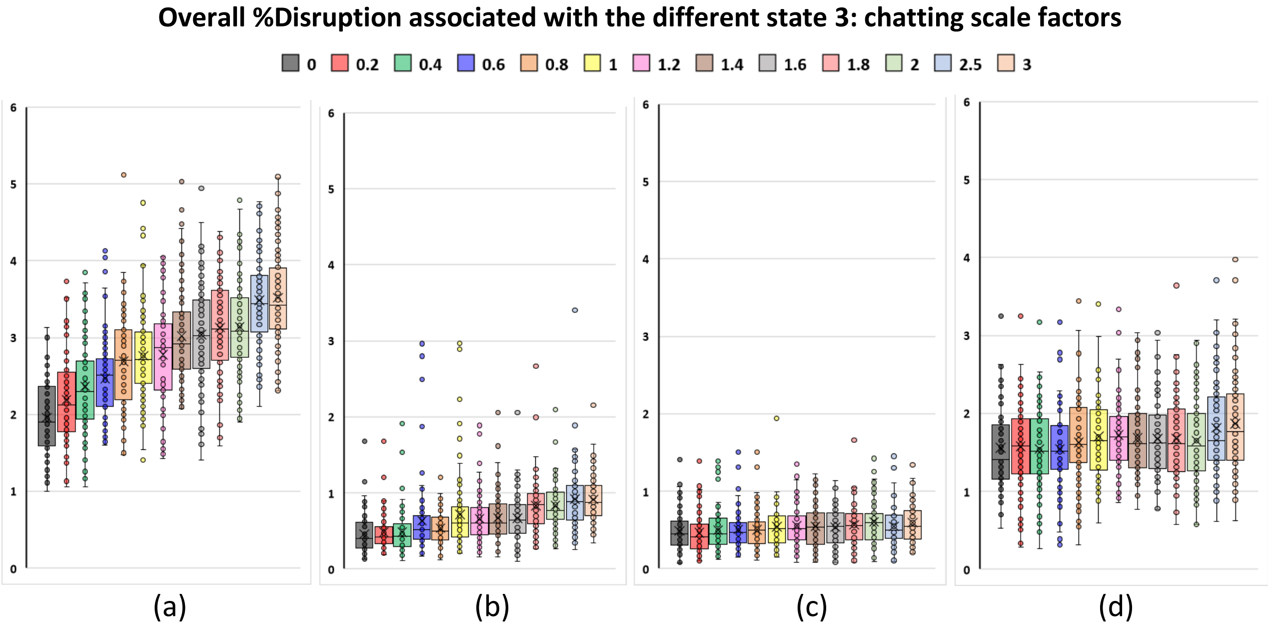

Figure 7-4 The changes in the overall %Disruption associated with increasing state 3 (chatting) scale factors

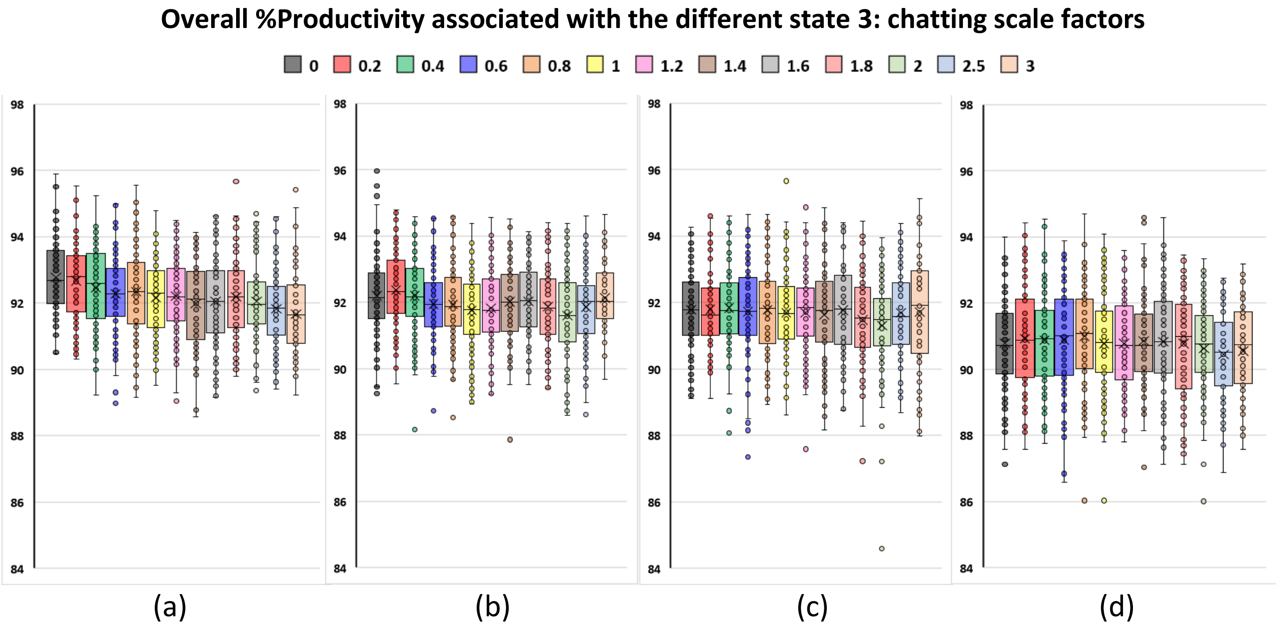

Figure 7-5 The changes in the overall %Productivity associated with increasing state 3 (chatting) scale factors

The objective of the analysis was to check that the model behaved plausibly to the changes in state 3 base probabilities. It is not sensible to compare the actual base state 3 probabilities of the four students at the start of the four different lessons as there were several other factors at work that would influence the simulation results, such as the position of the student in the classroom and the proximity of other students and their chattiness and inclination to work productively. From these results, the conclusion is that the lesson simulations seem to respond reasonably to changes in student state 3 base state probability attributes, as a relevant test-case.

7.3 Relative parameter importance

The lesson models contain 11 parameters. This section describes the analyses conducted to investigate the relative importance of the parameters. 'Importance' was evaluated from two perspectives:

1. What parameters were the most important predictors of whether a parameter set would pass all the metric acceptability checks (i.e., not be rejected) – part of the model calibration process;
2. What parameters were the most important predictors of whether a parameter set would cause a good match with the empirical lesson – part of the model validation process.

As part of the model calibration procedures, for each lesson a coarse grid-search of the parameter space had been conducted. Each search looked at 177,147 points (parameter sets) and ran 80 replications for each point. This provided the data for the importance analyses, which was carried out using the well-known SPSS software package from IBM.

For the first analysis of importance, for each parameter set, a count was made of the number of replications where all output metrics were within the acceptability ranges. The count ranges from 0 (no replication stayed within all the acceptability ranges) to 80 (every replication stayed within the acceptability ranges). If a parameter set had all 80 replications acceptable it was considered viable, otherwise it was rejected.

For the second analysis of importance, linear modelling was used to calculate normalized predictor importance scores in each of the four validated lessons. In SPSS, the *leave-one-out* method was used. In this method one predictor at a time is removed from the final full model and the result is ranked on the residual sum of squares. The value obtained is the normalized, relative importance of each parameter. This method enabled interactions and correlations to be taken into consideration.

Figure 7-6 shows an example of the results, for Lesson #1 (full details for all calibrated lessons are presented in Appendix D). The upper plot shows the parameters most influential in the rejection stage: in this case TSOL was the most important predictor (3x more important than

RLESW, for example); the lower plot shows the most influential parameters in terms of finding values that match the empirical lesson values: in this case RLOW is the most important parameter (more than 2x the importance of RLESW, for example).

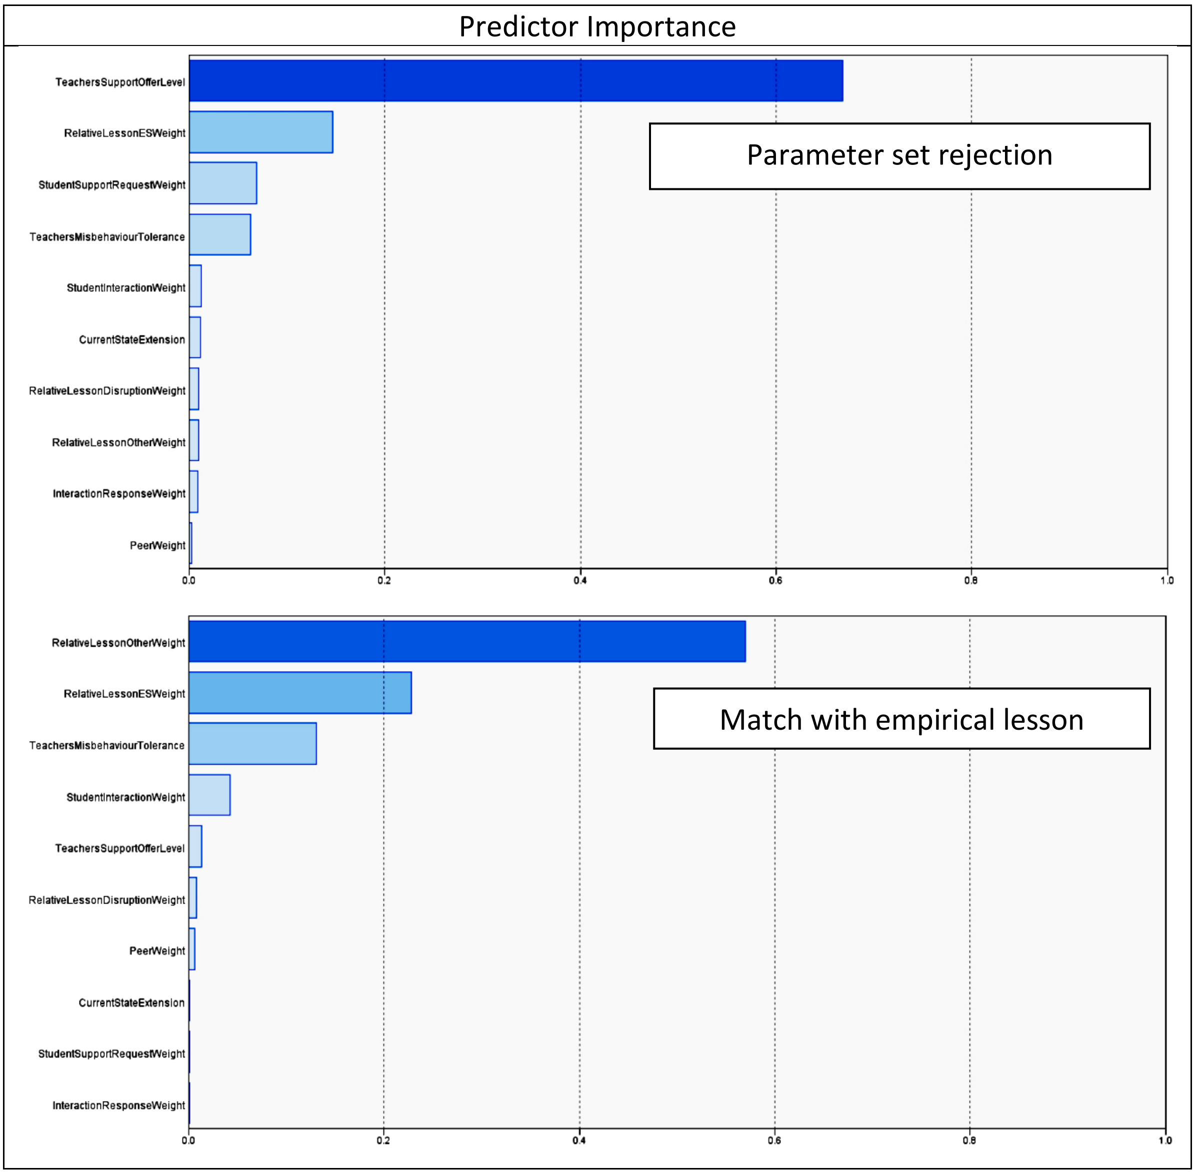

Figure 7-6 Lesson #1 Predictor Importance (upper) for parameter rejection (lower) for closest match to empirical lesson

The overall results of the sensitivity analyses are summarized in Table 7-3, which lists, for each lesson, the topmost influential parameters (in order) for both measures of importance. The most important predictor, relatively (across all lessons), for the first stage of calibration was TSOL, followed by RLESW and RLOW. For predicting the match with the empirical lesson, the most important predictors were RLOW and RLESW, followed by TMT – with TSOL being found irrelevant. Overall, informally combining the two measures of importance, RLESW and RLOW appeared to be the most influential parameters.

Table 7-3 Relative parameter importance and the nature of the parameter space

Lesson#	Lesson ID	TA present?	Calibration: most influential parameters	Validation: most influential parameters
1	16062207GEO	N	TSOL, RLESW	RLOW, RLESW, TMT
2	16063010SCI	Y	SIW, TSOL, RLDW, RLESW	RLOW, RLESW
3	16070508MAT	N	RLDW, RLOW, SIR, TSOL	TMT, RLESW, RLOW
4	16070510MAT	(N)	TSOL (RLOW)	RLOW, RLESW, TMT
5	16070607MAT	Y	TSOL, SSRW (RLOW)	RLOW
6	16070608MAT	Y	TSOL, SSRW (RLOW)	RLOW
7	16070809SCI	N	RLESW, SIW, TSOL, RLOW	RLESW, RLOW, SIW, TMT

These results are consequences of the rules, formulae and constants built into the CLSM, and also the empirical lesson data collected during the case study. The important question is whether these results have an interpretation for real lessons. RLESW adjusts the importance of the state expected for the students and has a natural interpretation as a magnifier of student compliance to the teacher's instructions. RLOW adjusts student passive disengagement. Passive disengagement could be a measure of the extent to which students find the lesson materials and activities engaging (whereas active disengagement – chatting – could be more a measure of socialness) and could be affected by how well the materials and activities suit the students' academic levels, interests, etc. Compliance and ease of disengagement could be pertinent latent constructs for student attributes and these parameters are effectively mediating for them.

7.4 Chapter summary

As part of the ABMS development methodology, some basic sensitivity analyses were conducted. This chapter explained the procedures involved and the results. 'One-factor-at-a-time' (OAT) analyses were conducted to explore how changes in the 11 model parameters affected the 6 simulation output metrics. The consequences of changes to the student base state probabilities for inclination to chat (as a relevant test-case) were also examined. The conclusions were that the lesson simulation models seem to respond reasonably to both types of changes.

Having established that the four validated lesson models (as well as the other three lesson models) were appropriately well-behaved, the following chapter explains how the research questions (described in section 1.2) were investigated and answered using these lesson models.

8. Investigation of the four research questions

This part of the thesis describes the experiments conducted to answer the research questions listed in section 1.2. Each experiment was also a contribution to answering the primary research question:

> **How and to what extent can an agent-based model adequately represent the behaviours of, and interactions between, students, teacher and teaching assistant in classroom lessons at a UK secondary school?**

The following quote explains the justification for conducting experiments with a validated simulation model:

> If we design the micro level of an ABM making the best use of available data that we can (calibration) and that ABM proves capable of producing simulated aggregate data which resembles real aggregate data (validation), then we have reason to believe that our ABM is not arbitrary ... and doesn't simply match the real aggregate data ... but that it might actually explain observed patterns because of the similarity between the real and simulated social processes in key respects.
>
> (Chattoe-Brown, 2014, p. 14)

Each experiment takes each of the validated lesson models and investigates the consequences of a change to the initial conditions, for example a change in the teacher's tendencies or a change in the student seating arrangements. For each experiment, the aim was to explore the theoretical consequences of the model design and the modelling assumptions, to understand the mechanisms of interactions, and to assess whether the results and explanations were plausible and realistic. As stated in section 1.2, if the results and explanations were considered satisfactory, then this would show the extent to which the model could be useful and would also indicate that the model embodies plausible theories of classroom behaviours, which could then be investigated further.

The ABMS had been validated for several lessons and passed sensitivity checks. However, validation (and calibration) had involved comparing statistics for monitored students only. It was now assumed that if the monitored students were adequately modelled then the unmonitored students would also be adequately modelled.

The framework used for the investigation of each research question was:

1. Consider what the empirical data indicated and what factors might influence results.
2. Explain how the scenario was modelled and controlled.

3. Confirm that the model responded sensibly to changes in the relevant lesson parameters or variables.
4. Take each of the four validated lesson models and run multiple replications of the research scenario.
5. Combine the data for all the validated lessons and interpret the results.

In many cases results were summarised in boxplots and compared visually, comparing means and distribution overlap.

8.1 Experiment 1: The influence of teacher support and discipline on student behaviour

This experiment addressed research question 1:

Does the simulation model provide realistic explanations of the effects on lesson behaviours and lesson outcomes of alterations to the teacher's inclination to offer one-to-one support and to take disciplinary action?

Teachers try to increase student learning/productivity/time-on-task by helping individuals and by managing behaviour. Hence teachers will often move around a classroom offering support to individuals, pairs or small groups. In terms of lesson metrics, helping and being helped are productive states. Teachers try to keep students on-task by maintaining a structured and disciplined environment, but misbehaviour is often troublesome. As teachers will relate, sometimes a mild form of discipline early on (e.g., a reminder of the importance of an activity) is all that is needed for everyone to be focused and productive. Sometimes though, just a few seconds of disciplining can create a negative atmosphere in a lesson, causing students to 'work to rule', minimizing interactions, spoiling learning and making it hard work for the teacher. It should be remembered that the opposite of students being productive is their being disengaged, either passively or actively distracting others. Consider the overall lesson metric %Disruption. Suppose in a small class of 12 students one of them does something highly disruptive and receives a serious (and extended) telling off for 10 s, which affects the rest of the 1-hour lesson. These 10 s are considered disruption. As a percentage of student-lesson time this %Disruption is $^{10}/_{3600x12} \approx$ 0.023%. So even a very small percentage of misbehaviour and disciplining can lead to an unpleasant and less productive lesson.

This research question was answered by investigating whether:

- class productivity increases or decreases as the amount of help from the teacher increases;
- there is a relationship between class productivity and teacher's help, and, if so, what the mechanism is;
- class productivity is affected by the teacher's discipline regime;
- individual student behaviour alters when the teacher behaves differently, even if overall lesson statistics do not alter significantly.

8.1.1 Analyses of the empirical data

From the empirical data there appeared to be a negative relationship between the amount of individual help and overall lesson productive behaviour (Figure 8-1a). The more time the teacher spent helping individuals the less the class was productive. But this was a comparison of entire lessons so included a lot of whole-class teaching time, when students would not be given individual assistance. When only the lesson time where the students could be helped was considered, the relationship appeared weak (Figure 8-1b). This time is labelled *independent* because the students are working alone or in pairs. It means that the expected state for students is not state 12: listening to the teacher whole class teaching.

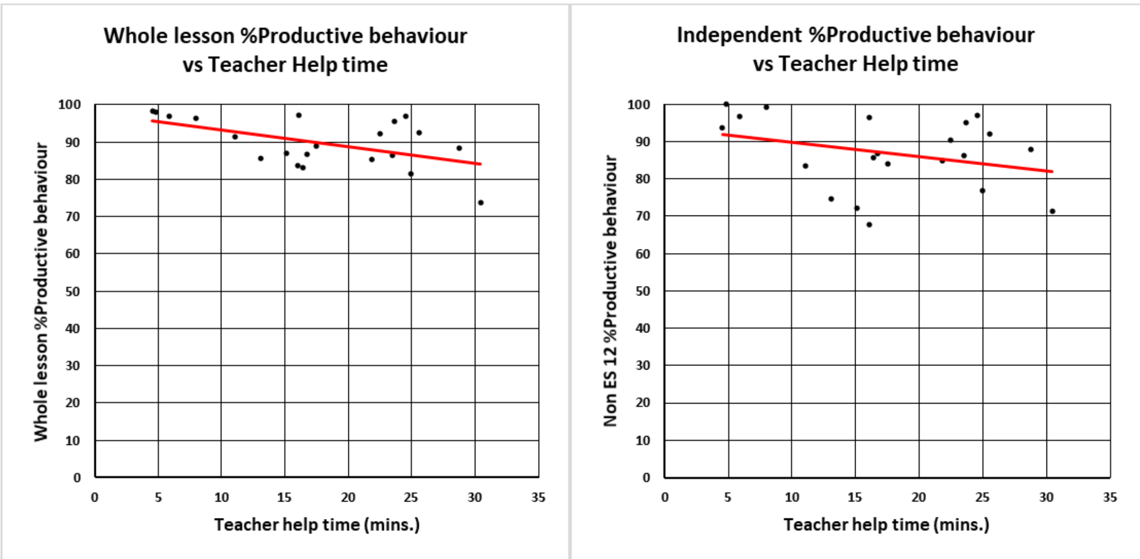

(a) (b)

Figure 8-1 The relationship between the empirical amount of individual help and %Productive behaviour
(a) whole lesson (b) when students were working alone or with others

To assess the significance of the relationship, regression analyses were performed (in MS Excel). For the data in Figure 8-1a, $r^2 = 0.30$ and the significance was $p < .05$. This meant that the amount of time the teacher spent helping appeared to explain 30% of the variance in the %Productivity, in lessons overall. For the data in Figure 8-1b, $r^2 = 0.10$ and the results were not significant ($p = .17$). The implication was that the amount of time the teacher spends helping individuals did not appear to be related to a class's amount of productive behaviour during the independent working periods. This prompted the following questions: If there was no effect, or even a negative effect, on productivity, why would teachers offer one-to-one assistance? Were other variables involved?

Further analyses of the empirical data yielded no significant relationships, in particular there seemed to be no significant relationship between %Productive behaviour and teacher discipline time – see Figure 8-2. For plot (a) $r^2 = 0.01$ and $p = .77$, and for (b) $r^2 = 0.02$ and $p = .53$.

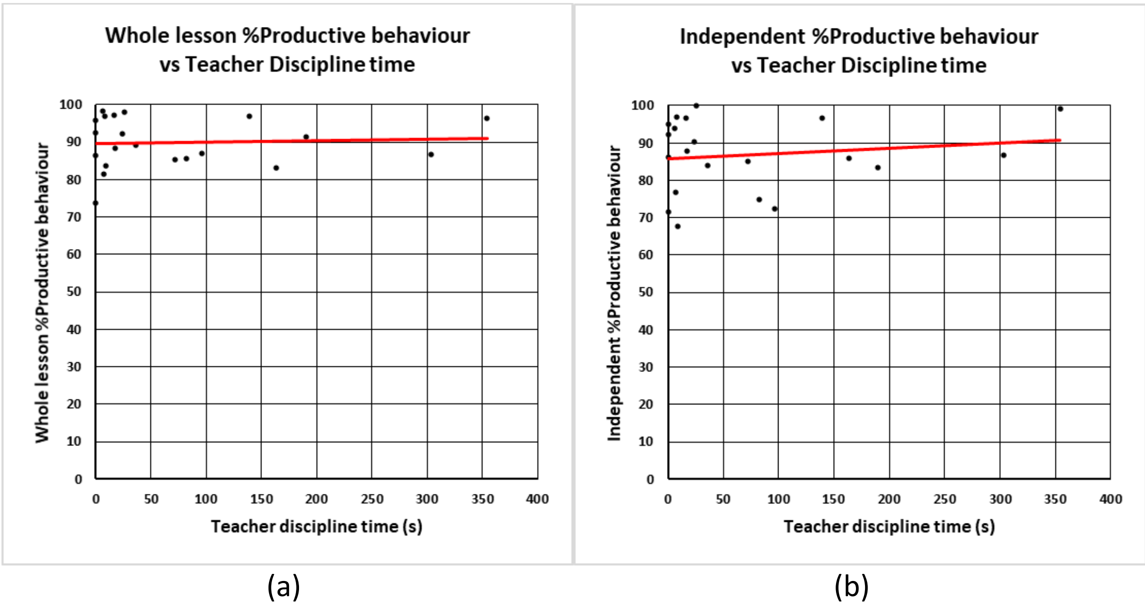

Figure 8-2 The relationship between the empirical amount of disciplining and %Productive behaviour
(a) whole lesson (b) when students were working alone or with others.

8.1.2 Analyses of simulation results

The simulation was used to investigate the impact on productive behaviour of increasing and decreasing the teacher's inclination to offer individual help and tolerance of misbehaviour. Below is a summary of the model aspects most pertinent to this investigation:

- The model assumes that the teacher moves to the student – it does not model the situation where a teacher calls students to the teacher's desk for discussions.

- For both the teacher and the students the choice to enter state 13 (being helped by the teacher) is a stochastic one, involving historical probabilities and the current situation, as well as the use of pseudo-random numbers.
- The teacher parameter Teachers-Support-Offer-Level (TSOL), with values from 0 to 10, can be adjusted to either increase or decrease the likelihood that the teacher will offer help. $^{TSOL}/_{10}$ is used as a threshold against which a pseudo-random number is tested. TSOL=10 results in the teacher always offering help if that is possible and appropriate. TSOL=0 results in the teacher never offering help.
- A student parameter, Student-Support-Request-Weight (SSRW), can be adjusted to either increase or decrease the likelihood that students will ask for help. $^{SSRW}/_{10}$ is added to the score for state 13.
- During the case study it was observed that, when students see the teacher offering help to others, they appeared more inclined to ask for help themselves. During focus group discussions this topic came up and students explained that they might hold back a little until it was clear that the teacher was available for giving assistance and that they would not look stupid asking the teacher. To mimic this relationship, that students will more likely approach the teacher for help when the teacher is more likely to be giving help, $^{TSOL}/_{10}$ is added to a student's score for state 13.
- The teacher parameter Teachers-Misbehaviour-Tolerance (TMT) can be adjusted to either increase or decrease the teacher's tolerance of misbehaviour – a mixture of the severity and amount of student misbehaviour tolerated and the tolerated duration of the misbehaviour before intervening. This is important for modelling realistic variations in teachers' approaches to disciplining. TMT=0 means that the teacher is completely intolerant and will intervene immediately for the slightest misbehaviour – with some random fluctuations introduced by the model.

Note also that the model does not allow the teacher to incessantly help students. The current model constants have the teacher waiting at least 4 minutes (with a random addition of up to 2 minutes) before offering the same student further help. There is also a random delay of a couple of minutes at the start of an independent working lesson section before the teacher offers help.

Each of the validated lessons underwent the following investigation. Each validated lesson model had a set of selected parameter values resulting from validation. For the three relevant parameters (TSOL, TMT and SSRW) a set of alternative values was selected, on either side of the validated values. For example, for Lesson #1, the parameter values investigated were (with the validated value highlighted):

TSOL [0 0.1 0.3 0.5 0.55 0.6 0.65 0.7 1 2 3 6 8]

TMT [5 10 20 30 40 45 50 55 60 70 80 90]

SSRW [-0.4 -0.3 -0.2 -0.1 -0.06 -0.04 -0.02 0 0.02 0.04 0.06 0.1 0.2 0.3 0.4]

The other parameters were fixed at their validated values. 350 replications were run and aggregated for each of the several thousand combinations of parameter values. Means and standard deviations were calculated to compare results.

8.1.2.1 Confirmation of effect of parameter value changes

The first step was to confirm that each lesson simulation model responded appropriately to changes in the TSOL, TMT and SSRW parameters, namely:

- when TSOL or SSRW increase, TH (the amount of time the teacher spent helping students (state 13)) increased;
- when TMT increased, TD (the time the teacher spent disciplining) decreased.

The results for each of the four lessons are shown in the rows of Figure 8-3. In column (a) one can see that increasing TSOL increased mean TH. The plots also show that TH saturated: the teacher could not help for longer than the time available. When the teacher was helping maximally, turning up TSOL had little effect - the model introduces random fluctuations that might cause a few second's rest. The ±1 SD error bars show that variance decreases as TH saturates. In column (b) one can see that as TMT increased mean TD decreased. Similar to (a), the effect levels out at zero as eventually increasing the value causes no further reduction: the teacher does not discipline. In column (c) one can see that increasing SSRW increased TH, but that this saturated. With a decreasing inclination of the students to ask for help, TH decreased, but also levelled out.

These results established that the parameters TSOL and SSRW influence TH appropriately and that TMT affects TD appropriately.

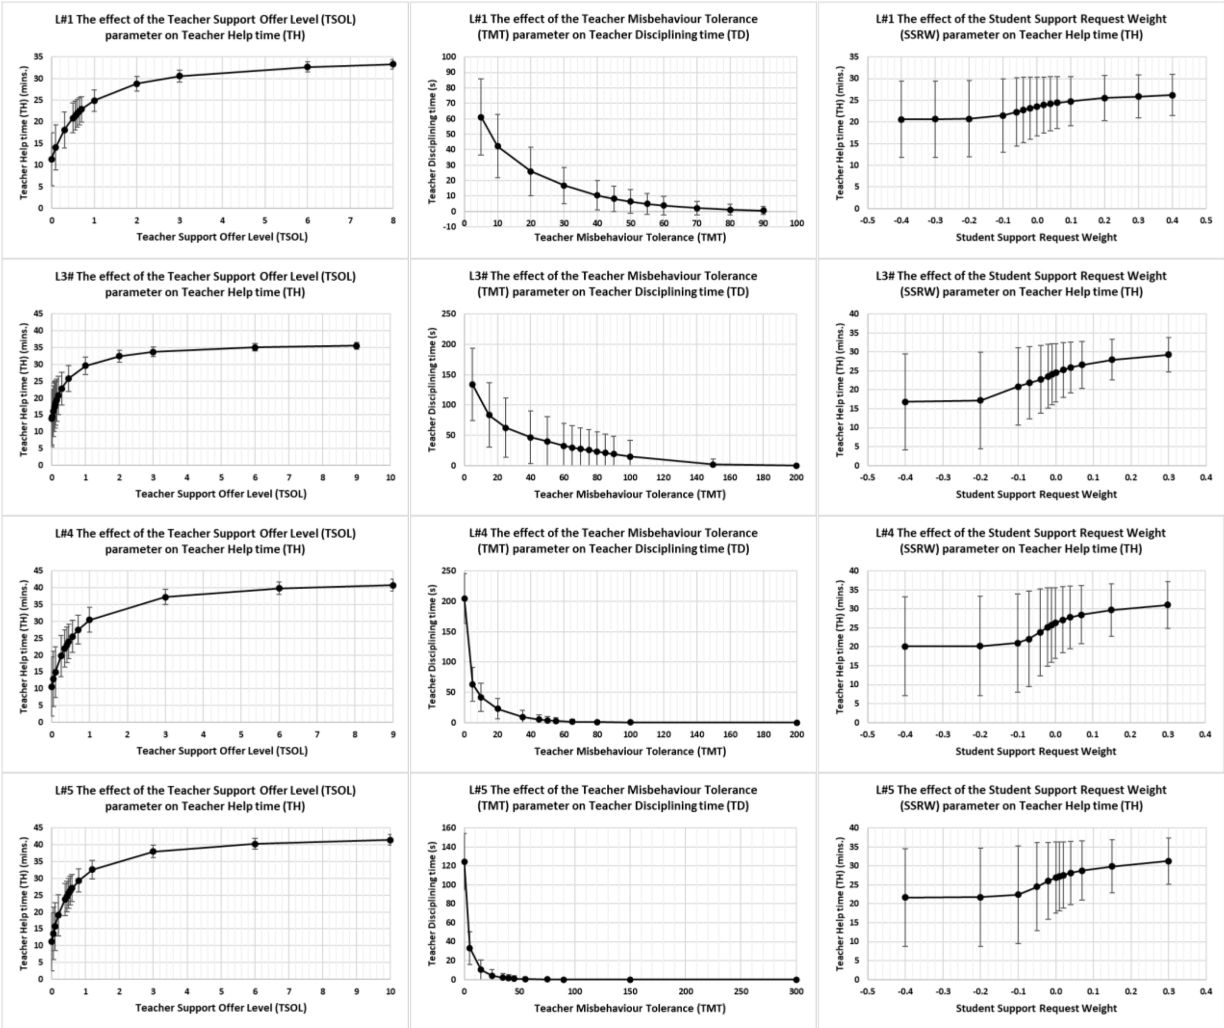

(a) (b) (c)

Figure 8-3 The effect of adjusting the parameters in the validated lessons
(a) TSOL (b) TMT (c) SSRW

8.1.2.2 The relationship between %Productivity and the parameters TSOL, TMT and SSRW

The next step was to observe the effect of the parameters on the student productivity, disruption and other disengagement metrics. The set of graphs in Figure 8-4 show that in all four lessons, increasing TSOL reduced %Productivity (whether measured over a whole lesson or just during independent working), until it levelled out. The set of graphs in Figure 8-5 show that in all four lessons, increasing TMT (meaning reducing response to misbehaviour) reduced %Productivity, until it levelled out. Both of these results were expected for the following reasons. Firstly, the model rules were explicitly designed so that students are more likely to disengage (actively and passively) when the teacher is helping another student – and this help increases as TSOL increases. The model rules were also explicitly designed so that the more tolerant the teacher is of disengagement (active and passive), the more likely students are to disengage. Hence %Productivity declines in both cases.

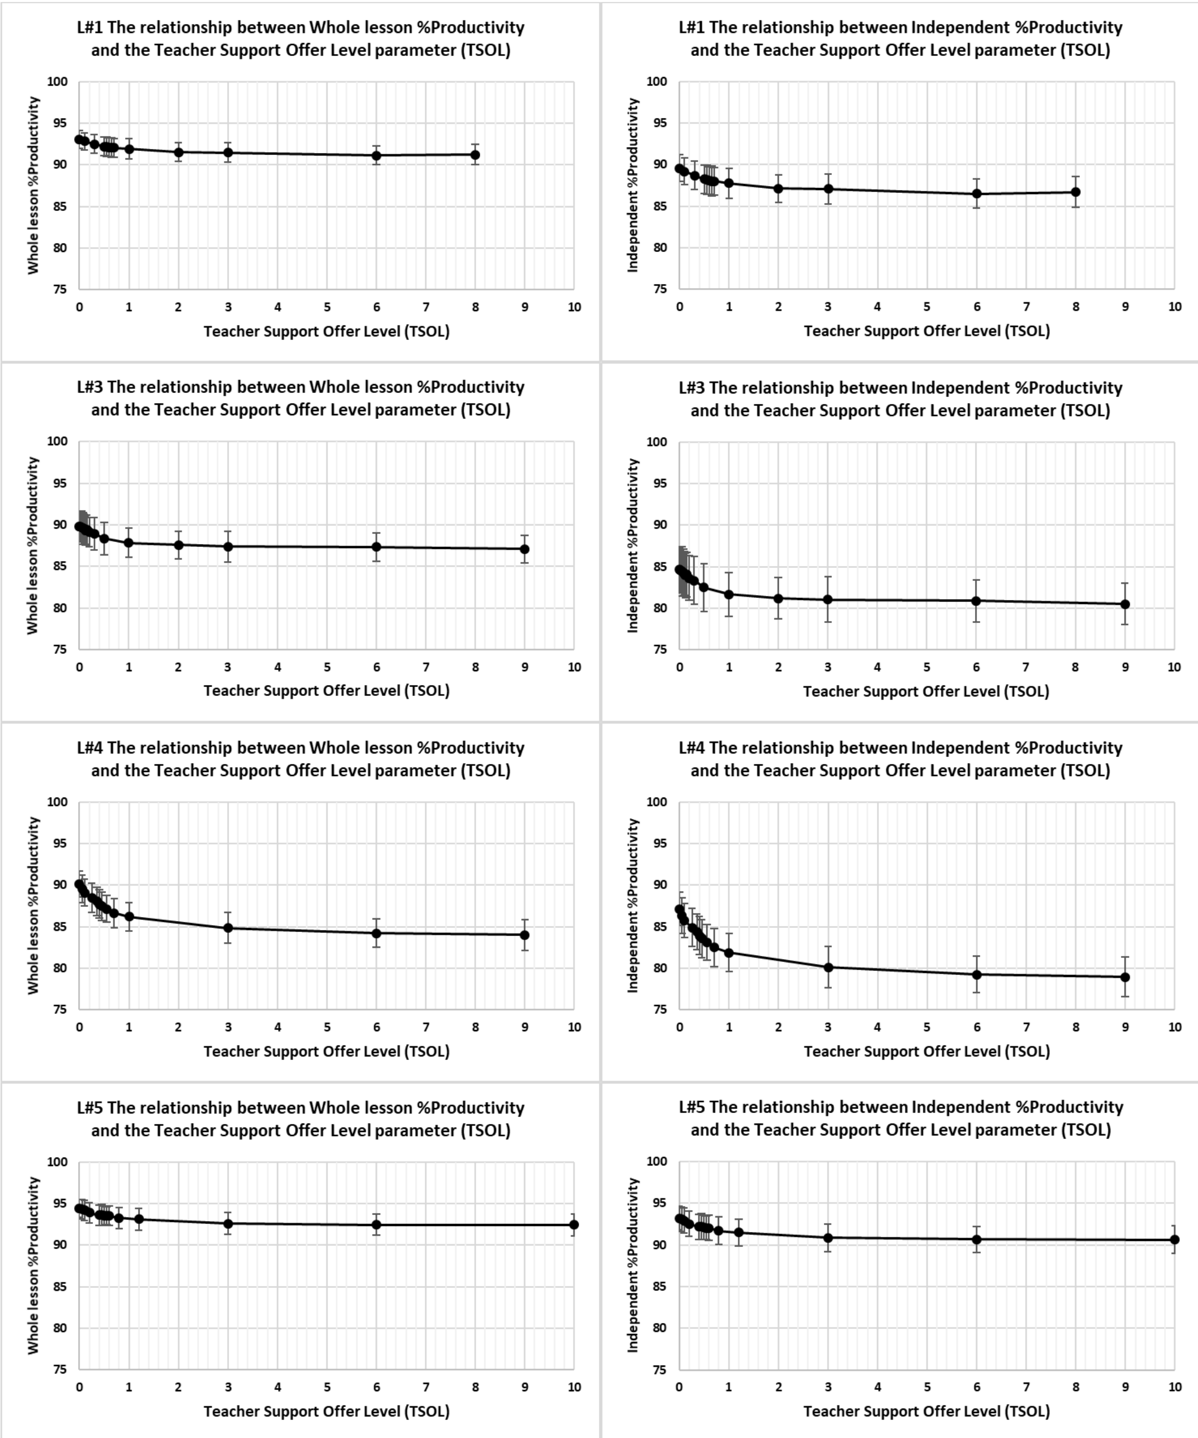

(a) (b)

Figure 8-4 The relationship between %Productivity and TSOL for all four lessons (a) over the whole lesson (b) during student independent working time.

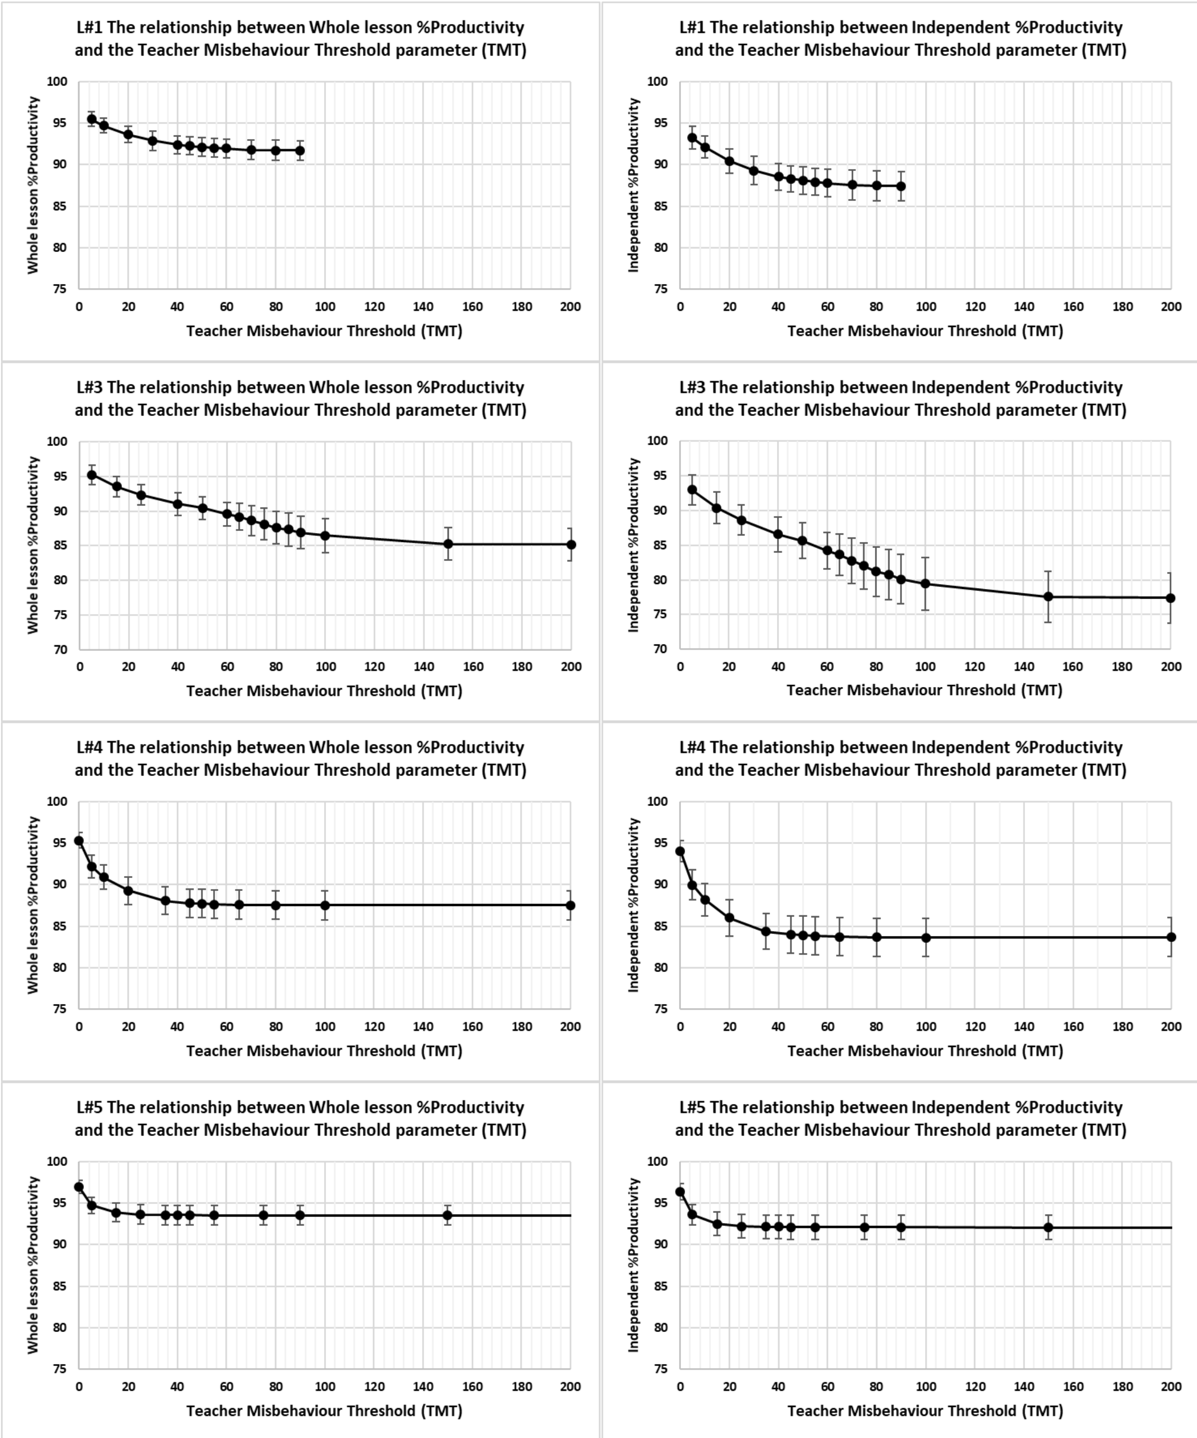

(a) (b)

Figure 8-5 The relationship between %Productivity and TMT for all four lessons (a) over the whole lesson (b) during student independent working time.

The set of charts in Figure 8-6 show that in all four lessons, increasing SSRW can cause some reduction in %Productivity, but that the effect was relatively minor. Note though that this was a consequence of the range of parameter values investigated: SSRW values of greater magnitude do have a strong impact.

In all cases, the whole lesson characteristics seem to follow the pattern of the independent work sections.

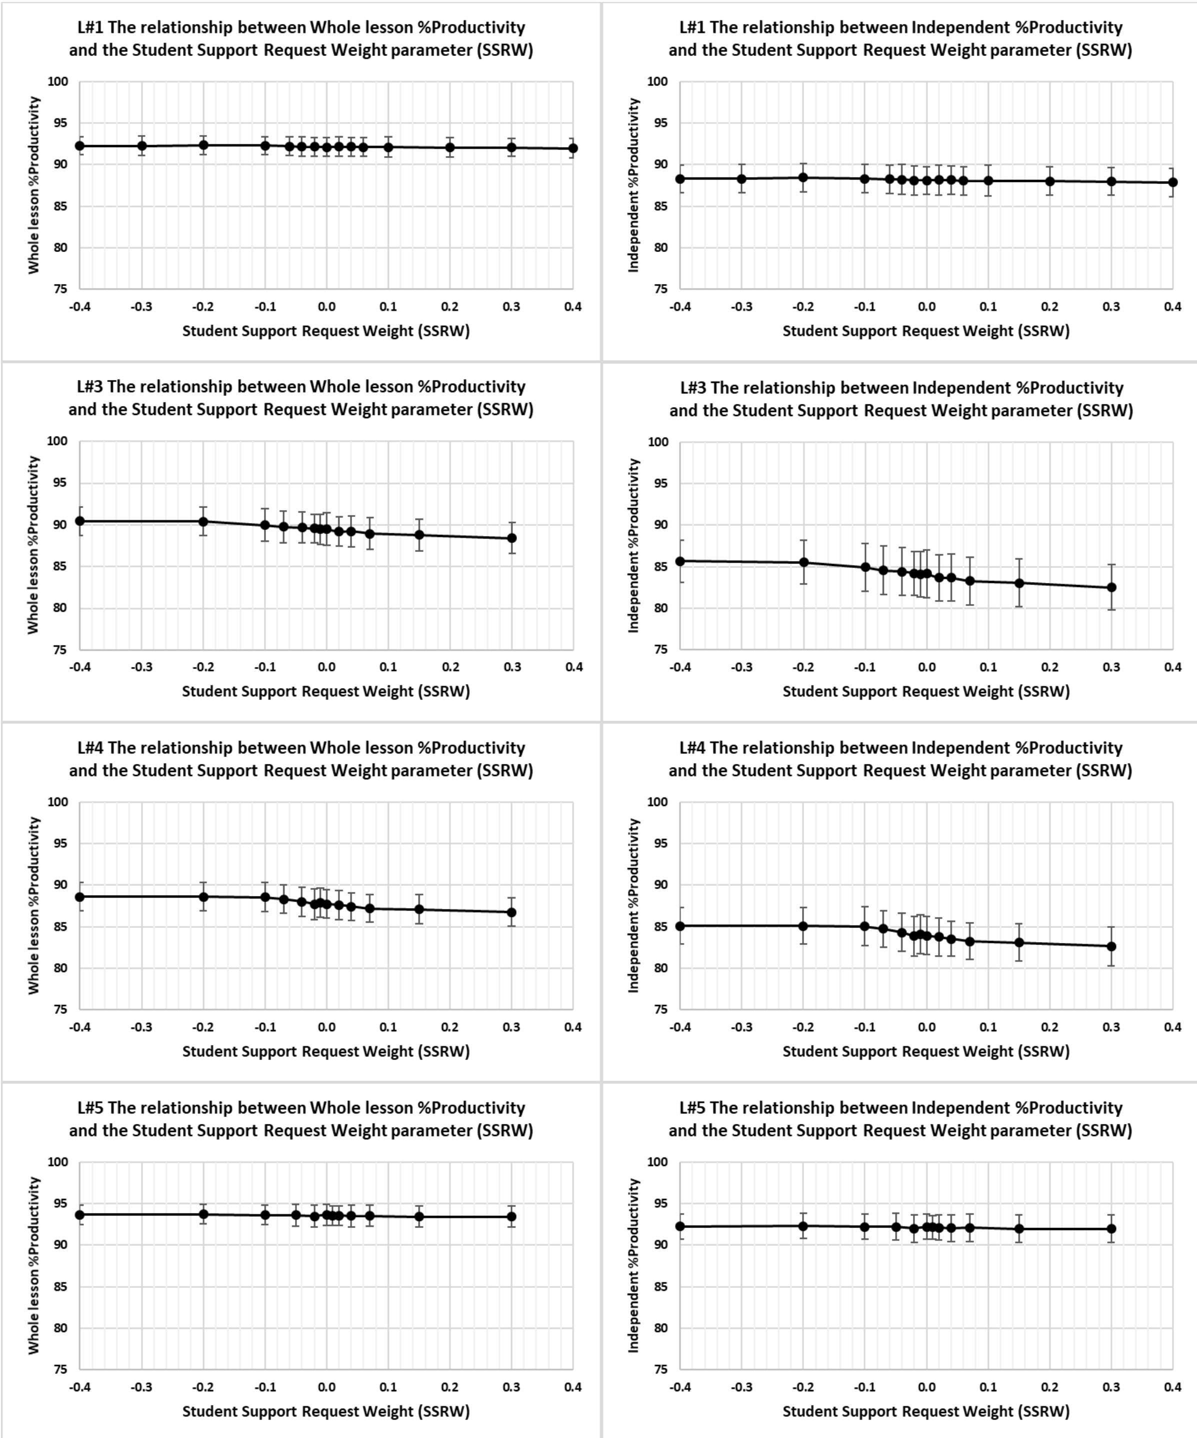

(a) (b)

Figure 8-6 The relationship between %Productivity and SSRW for all four lessons (a) over the whole lesson (b) during student independent working time.

The set of graphs in Figure 8-7 show %Productivity as a function of TMT for the selected set of TSOL values. In all four lessons, %Productivity (both whole lesson and independent working time) decreased or flattened out as both TSOL and TMT were increased, that is as the teacher helped more and was made more tolerant of misbehaviour. In all the plots, moving to the right along the x-axis represents increasing TMT, i.e., being more tolerant of misbehaviour, and moving from one TSOL line down to another corresponds to increasing TSOL, i.e., increasing the amount of individual help. The wider dispersion and lower productivity in Lesson #3 simulations suggests that, when the teacher gave more individual help, this class increased misbehaviour at a greater rate other classes. (Although not shown here, similar plots were produced for %Productivity as a function of TSOL for different TMT values. They also showed downward curves, one below the other.)

One could interpret the graphs in Figure 8-4 to Figure 8-7 in two ways. They suggest that if the teacher wants to increase the time spent providing individual attention (TSOL) but maintain the same level of %Productivity (specifically time on-task), the teacher needs to reduce their tolerance of misbehaviour (i.e., intervene more quickly to smaller incidents). Alternatively, to maintain the same level of productivity without extra disciplining, the teacher needs to be seen to be observing the class more, so giving less individual support. The fact is, in UK schools during periods of independent or paired or group work, teachers do often provide individual support – because they believe that this benefits the learning of the individual – and they knowingly take the risk that this may trigger some chatting or passive disengagement behind their backs.

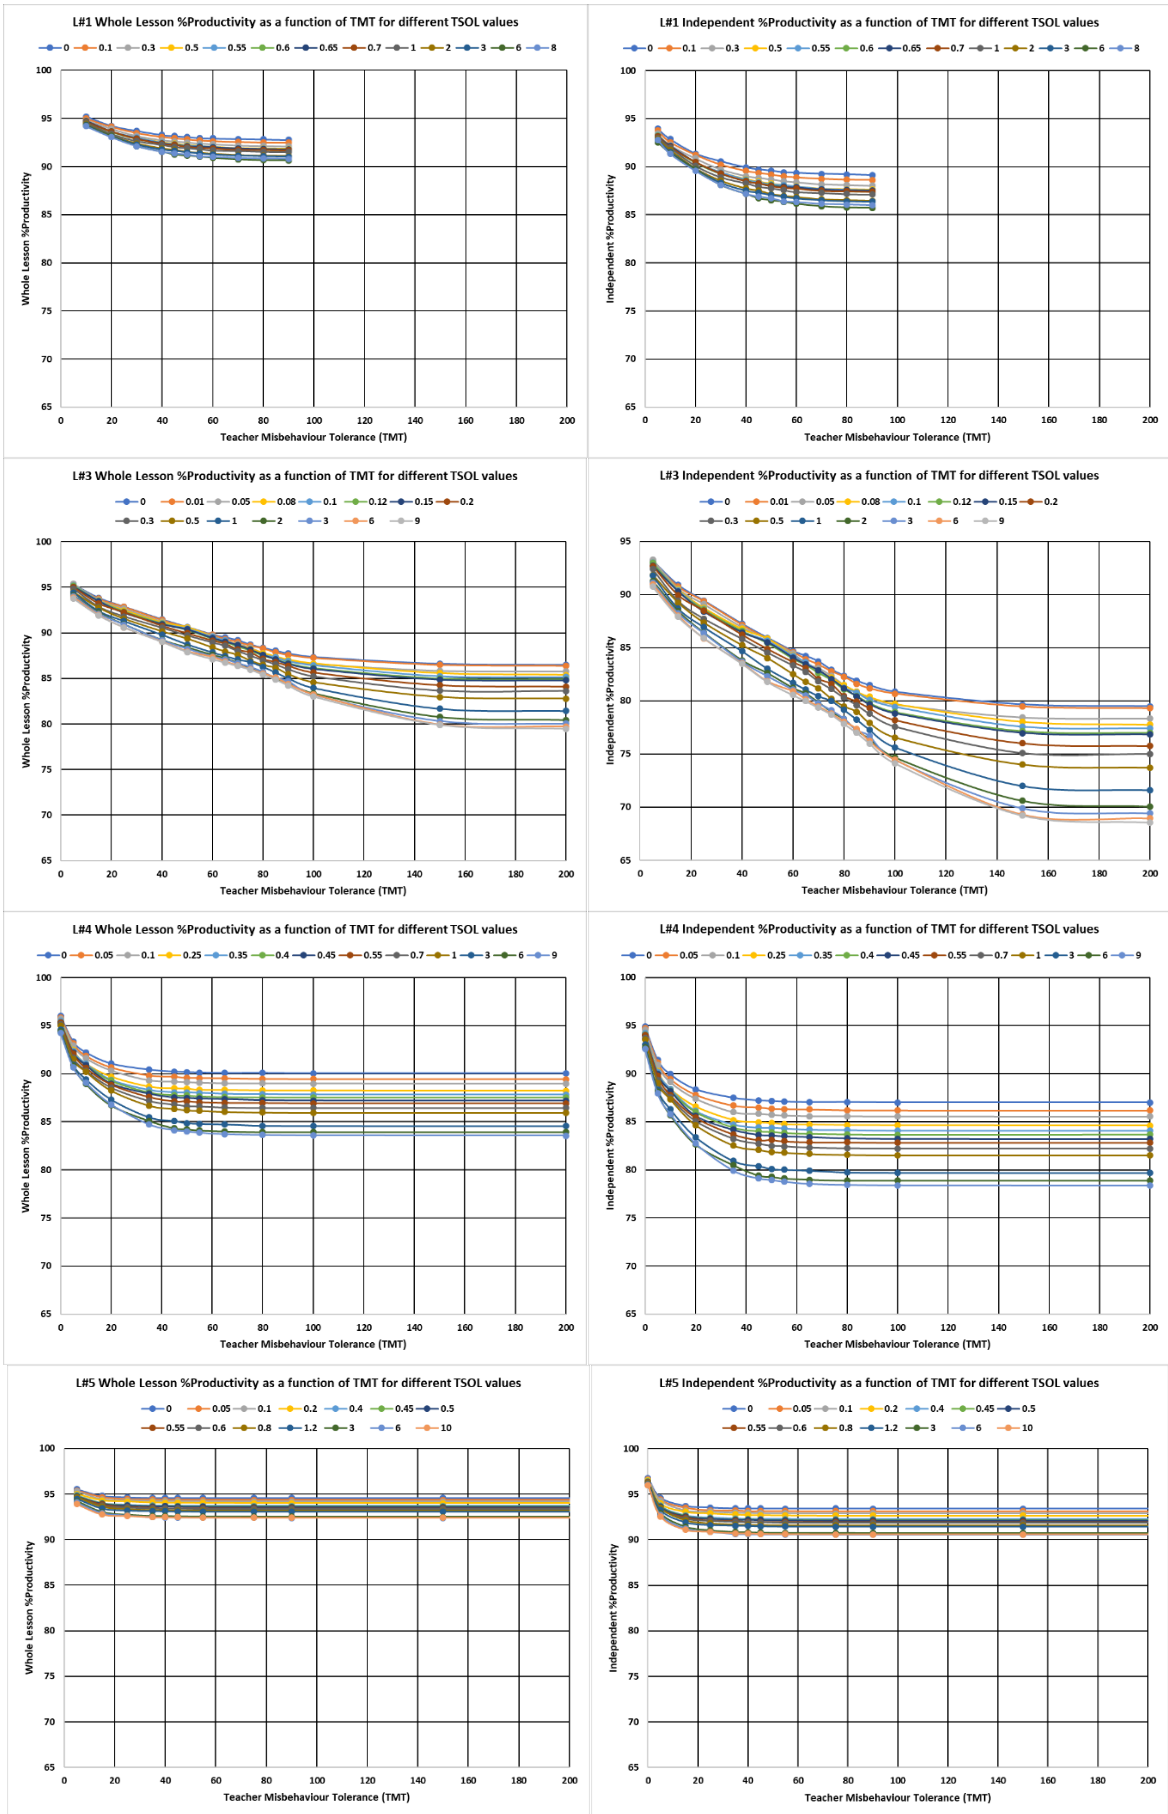

Figure 8-7 The relationship between %Productivity and TMT for different TSOL levels

8.1.2.3 Meso-level analyses of the relationship between %Productivity and TSOL and TMT

The analyses so far have shown that the parameters influence the output metrics in consistent ways and that the model appears to be generating plausible overall results. The understanding of the model so far is that:

- increasing the TSOL parameter value (the teacher's inclination to offer individual assistance) increases the TH metric and reduces the %Productivity metric;
- increasing the TMT parameter value (the teacher's tolerance of misbehaviour) decreases the TD metric and decreases the %Productivity metric.

However, it is not clear how the TSOL and TMT parameters affect %Productivity. %Productivity is determined by the amount of time students spend in productive activity states. %Disruption and %Other metrics are defined similarly. The amount of time involved depends on many other factors and relationships besides the parameter values.

The purpose of this subsection is to show how the agent-based simulation can provide plausible sequences of agent interactions that lead to the overall results or explain the variations in overall results. Figure 3-6 shows the main factors involved in modelling these interactions, specifically the students' and the teacher's choice of their next states. The expectation was that if one increased the likelihood that the teacher offered support (TSOL), then overall student productivity would increase – that is why teachers offer help. It was also expected that if the teacher was more tolerant of misbehaviour (increased TMT) then productivity would decrease. The overall metric results supported the latter expectation but not the former.

One factor behind this could be that the metrics are means of distributions. In many of the replications for a fixed value of TSOL the %Productivity values were well above or below the mean. The same happened with TMT. There were obviously other factors involved. To investigate what lesson dynamics resulted in %Productivity being significantly higher or lower than the mean, the discussion focuses on the hypothesised mechanism shown in Figure 8-8. This shows one possible explanation: student behaviour in response to the teacher helping others is moderated by the teacher's TMT but the triggering of disciplining is partly stochastic. If, due to the stochasticity, the teacher does not intervene then overall productivity drops; if the teacher intervenes, productivity increases.

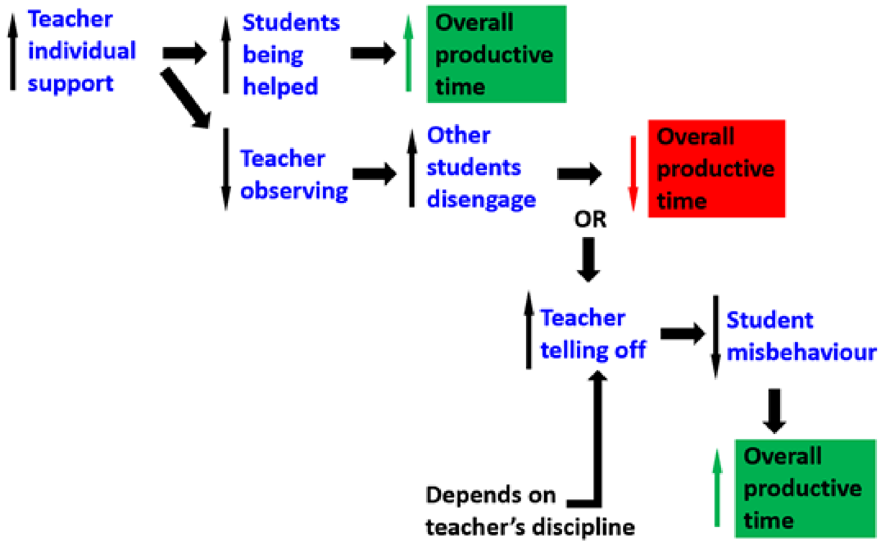

Figure 8-8 Alternative consequences of increasing TSOL depending on TMT

Using Lesson #3 as an example, consider the three replications that gave the lowest overall lesson %Productivity, the mean %Productivity (84.1%) and the highest %Productivity. Figure 8-9 shows the state trajectories for these replications (with the empirical lesson for comparison). In each panel, time flows up and the teacher trace is the thin plot on the left-hand side. One of the clearest differences is the time in the lesson when the teacher disciplined the whole class (indicated by the magenta line across the whole trace, highlighted by the black arrows). In the empirical lesson this did not happen. In the least productive replication this happened at the very end of the independent working period (the brown section of the traces). Prior to that the teacher had been providing extensive individual support and the students had misbehaved extensively (indicated by the orange and red). In the average replication, the whole-class disciplining happened less than halfway through the independent working section. The result was that the remainder of that lesson section was more on-task. In the most productive lesson, the teacher intervenes much earlier, after finishing helping one student and because several students had been misbehaving. One student had even been out of their seat and distracting others – red block. It is also true that the teacher did much less individual support overall.

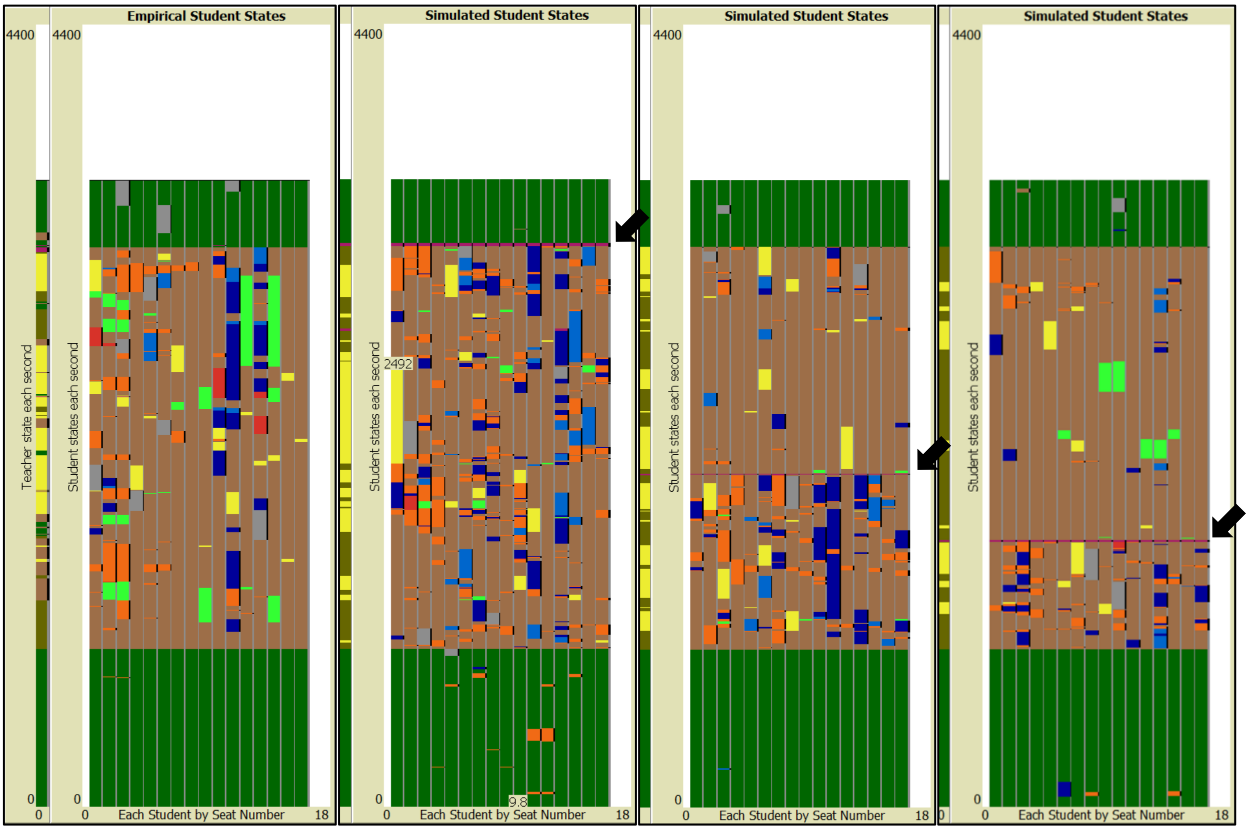

(a) (b) (c) (d)

Figure 8-9 Examples of student lesson trajectories - with black arrows indicating when the teacher disciplined the class
(a) Empirical (b) Lowest %Productivity (c) Average %Productivity (d) Highest %Productivity

The simulation results are naturally a consequence of the model rules plus the built-in stochasticity. There are rules which alter a factor that adjusts scores for students' misbehaviour states according to the state of the teacher. (This was explained in section 3.4.2 with further details provided in Appendix A.3.4.) When the teacher is possibly observing the class (especially when actively teaching) the misbehaviour reduction factor is larger and the students are less likely to choose to misbehave. When the teacher is helping a student, the teacher is considered to be less likely to detect or react to misbehaviour and the students respond by only slightly reducing their misbehaviour score. And, after being disciplined, students are much less inclined to misbehave.

Table 8-1 shows the metrics for the Independent Working section of the lessons shown in Figure 8-9. The data includes the total time (in seconds) that all students spent in each of the student states. The colour-coding shows which states are considered Productive, Disruptive/Disciplinary or Other (other non-productive activities). Note that in the least productive lesson the teacher did the most disciplining (30 s compared to 9 s) and students spent the most time in state 1 (being disciplined individually) and state 5 (being disciplined as a class): 12 s and 288 s versus 3 s and 96 s, respectively. This shows that disciplining by itself may not be very effective: the timing of the disciplining seems to be important too. In the most productive replication, the teacher addressed misbehaviour earlier, disciplining the class before misbehaviour got out of hand – see arrow in Figure 8-9(d). But this was triggered by the amount of serious disruption (36 s state 2, the red block) and chatting (state 3, the orange blocks). Note also that the amount of time the teacher spent helping (Teacher Help and student state 13) in the most productive lesson was half that of the least productive lesson. This is discussed in the next section.

Table 8-1 Data for the Independent Working section of the empirical lesson and the three replications (with student state times in seconds)

Metric	Empirical	Lowest %Productivity	Average %Productivity	Highest %Productivity
%Productivity	84.0	76.0	84.1	91.5
%Disruption	7.3	11.3	6.5	4.2
%Other	8.7	12.7	9.4	4.4
Teacher Help (mins.)	18	21	19	10
Teacher Discipline (s)	36	**30**	**9**	**9**
Student state 1: individual disciplining	n/a	**12**	**3**	**3**
Student state 2: very disruptive	n/a	84	0	**36**
Student state 3: chatting	n/a	4266	2253	1419
Student state 4: intentionally unproductive	n/a	2514	2298	1095
Student state 5: whole-class disciplining	n/a	**288**	**96**	**96**
Student state 7: not sure if productive	n/a	1611	759	252
Student state 8: working alone	n/a	26028	29217	31908
Student state 9: working with others	n/a	267	72	660
Student state 13: being helped by teacher	n/a	1251	1182	600
Student state 17: out of room	n/a	513	417	411

8.1.3 Discussion of results and conclusions to research question 1

To interpret these results, it is important to understand that there are several feedback loops in the lesson system. For example, as explained earlier, the model uses the state of the teacher to adjust the probability of student disengagement and this is considered more likely when the teacher is busy with a student (state 13) than when the teacher is observing (state 7) the class. The mechanism for the adjustments (explained in section 3.4.2, with further details provided in Appendix A.3.4) uses a misbehaviour reduction factor to reduce the influence of the teacher. Teacher state 13 does not affect the factor, while state 7 reduces the factor by 0.9, leading to a greater reduction in misbehaviour. Hence student disruption can increase when a teacher stops observing and starts helping an individual, and this disruptive behaviour can lead to the teacher spending more time disciplining – which then affects student behaviour. This is one feedback mechanism but there are others.

Simulations with one setting of TMT will lead to a distribution of student misbehaviour and some teacher disciplining. One might expect that increasing TMT would always lead to less disciplining by the teacher (less TD). However, it can happen that this greater tolerance leads to students misbehaving slightly more, pushing the teacher over some threshold and leading to the teacher doing more disciplining – which reduces further student misbehaviour.

There is another agent-level factor to consider. Since both working alone and being helped are productive states (so contribute to %Productivity), if in a simulation replication the teacher happens to always choose to help students who would have been productive anyway, then there will be no increase in overall productivity. Increased teacher help would not increase overall productivity. To increase productivity, non-productive (i.e., disengaged) time has to be replaced by productive time. Suppose on the contrary, that there is always one student disengaged and that in the simulation replication the teacher always chooses to help a disengaged student, thus converting them to a productive state. If the class had, say, 10 students and 40 minutes of a 1-hour lesson were allocated to independent working, the equivalent would be one student disengaged for the whole 40 minutes. Without any teacher help the overall %Productivity would be $\frac{(10\times 20)+(9\times 40)}{10\times 60} \approx 93.3\%$, but this would rise to 100% if that equivalent student were helped for the full 40 minutes. In summary, in the most extreme case, where there is always someone disengaged and the teacher always chooses that student, then, in a class of 10 students the maximum impact of the help over a 40-minute period would be to increase %Productivity by 6.7%.

This analysis indicates that the range in %Productivity between the lowest and highest replications (see Table 8-1) can be explained by which individual students were provided one-to-one support.

Taking this result further, if the teacher were to focus their assistance on those students who tended to misbehave or looked as though they were about to disengage, then this may well increase productivity – by converting non-productive time to productive time. Even just going over to a student and offering assistance would reduce misbehaviour by and around that student. But that strategy would be unfair on well-behaved students who would benefit from support.

There are other factors influencing the simulations, for example:

- If there is always another student asking for help as soon as the teacher has finished helping one student, then the teacher will never need to offer help, irrespective of the TSOL setting.
- If the teacher is helping as much as physically possible then no more help can be offered and turning up TSOL further would have no effect.
- Even when the teacher is helping maximally, this does not mean that the teacher is providing help all the time: the model rules require that students have a gap between episodes of one-to-one support so that the teacher does not incessantly provide help. If there are no more students eligible for help then no more help can be given.
- Because the students are less likely to misbehave the closer the teacher is, if the teacher helps someone in the centre of the classroom, this reduces misbehaviour in more students than if helping someone in a corner.

The goal of the above investigations was to answer research question 1: Does the simulation model provide realistic explanations of the effects on lesson behaviours and lesson outcomes of alterations to the teacher's inclination to offer one-to-one support and to take disciplinary action? The investigations provided comprehensive explanations for the complex relationships between teacher help, teacher discipline and student productive behaviour. The simulation model appeared to mimic observed behaviour and give sensible results for sensible reasons. In fact, the simulation results highlighted that if the teacher wants to increase the time spent providing individual attention, then, to maintain the same level of time on-task, the teacher needs to become less tolerant of misbehaviour, intervening more quickly to smaller incidents in order to prevent misbehaviour spreading. Alternatively, a teacher could manage student behaviour without additional disciplining by providing less individual support and being seen to observe more. Another possibility would be for the teacher to ask for a TA to provide individual support – as investigated in Experiment 2.

The analysis also highlighted that if the teacher were to focus their assistance on those students who tend to misbehave or looked as though they were about to disengage then this may well increase productivity – by converting non-productive time to productive time. Even just going over to a student and offering assistance would reduce misbehaviour by and around that student. But, as noted above, that strategy would be unfair on well-behaved students who would benefit from support.

In conclusion, the simulation model provided realistic explanations of the effects on lesson behaviours and lesson outcomes of alterations to the teacher's inclination to offer one-to-one support and to take disciplinary action.

8.2 Experiment 2: The impact of a TA on productivity

This experiment addressed research question 2:

Does the simulation model provide realistic explanations of the effects on overall student productivity of providing or withdrawing a teaching assistant who gives individual support to any student?

In general, the expectation is that having a TA in a class can be a great boost for the students and a significant help to the teacher, perhaps by enabling the teacher to provide assistance to more students. But many factors affect the potential benefits, such as:

- the extent of the assistance the TA offers;
- the extent to which students ask the TA for assistance;
- whether students are influenced by the behaviour of their peers in asking for TA help;
- whether the TA disciplines students or not;
- the proportion of a lesson available for one-to-one support - there is little point having a TA if the teacher is mainly whole-class teaching;
- whether the teacher spends time interacting with the TA (either for instructions, feedback on students, chatting).

The following assumptions and decisions had been built into the model:

- the teacher must instruct the TA as to when they should or should not offer help – this is implemented via the lesson plan;
- the TA is expected to assist any student - the lesson plan implementation does not allow specific students to be specified (which is sometimes the case when a TA is more of a one-to-one support for a designated student);
- the TA does not discipline students.

The lesson outputs investigated were Independent %Productivity (student productivity during lesson sections where the students can be helped by the TA), TAH (the amount of student time spent receiving TA help) and TH (the amount of time the teacher provided help). The two specific student states being discussed are state 13 (being helped by the teacher) and state 15 (being helped by the TA). Note that where the TA helps more than one student at a time this counts as additional student-TA help time (TAH). For this investigation, apart from adjusting TSOL and TASOL, all parameters were held at the lessons' validated values.

Of the four lessons that passed the validation tests, only one (Lesson #5) had an active TA (Lesson #4 had a TA, but she was inactive (at the teacher's instructions)). This meant that there could be only one experiment to remove a TA, from Lesson #5.

8.2.1 The effect of removing the TA from Lesson #5

Suppose a class is used to having help from a TA: what will the students do if there is suddenly no TA in their lesson? Would the students just not ask for help because it was the TA they felt comfortable with, or would they be inclined to ask the teacher instead? How might lesson productivity be affected? The following scenarios were compared:

(a) the original validated lesson with all the validated parameter values (TSOL=0.5);

(b) the teacher and the students do not change their behaviour; the teacher's TSOL stays at the validated value 0.5 and the students' base state probabilities stay the same; but because there is no TA, student state 15 (being helped by the TA) never occurs;

(c) the students adapt by instead asking the teacher (who behaves just as before, TSOL=0.5); the students' state 15 base scores are moved to the state 13 scores (being helped by the teacher);

(d) the students adapt by instead asking the teacher (who behaves just as before) and the teacher adapts by offering more assistance (TSOL=5);

(e) the original validated lesson with all the validated parameter values except TSOL=5, to see the effect of this on the original lesson (that had a TA).

All the other lesson parameters were held at their calibrated values.

Taking the validated parameter set for Lesson #5, 500 replications were run for all scenarios. The aggregated results of the TA-less replications were, of course, quite different from the empirical lesson: TAH was now 0. However, the goal of this experiment was not to compare the simulations to the empirical lesson, but to compare the different simulated scenarios. The plots in Figure 8-10 show the effects on Independent %Productivity (**x** marks the empirical lesson value) of the various scenarios:

(a) the original distribution;
(b) with the TA removed but nothing else changed there was a very slight drop (0.5%) in Independent %Productivity;
(c) with the TA removed and the students base state 15 scores moved to their state 13 scores (meaning that they would ask the teacher instead), there was negligible impact;
(d) with the TA removed and the students base state 15 scores were moved to their state 13 scores (meaning that they would ask the teacher instead) and the teacher took over and offered more help (TSOL=5 instead of 0.5), there was now a noticeable drop in Independent %Productivity, likely due to the increased TH – as discussed in section 8.1;
(e) the effect of TSOL=5 on the original lesson (that had a TA) was similar to scenario (d) in that it showed a drop in Independent %Productivity, again probably due to the increased TH.

In all cases there was a decline in Independent %Productivity, particularly in (d) and (e). However, the amount of TH in scenarios (d) and (e) is much greater.

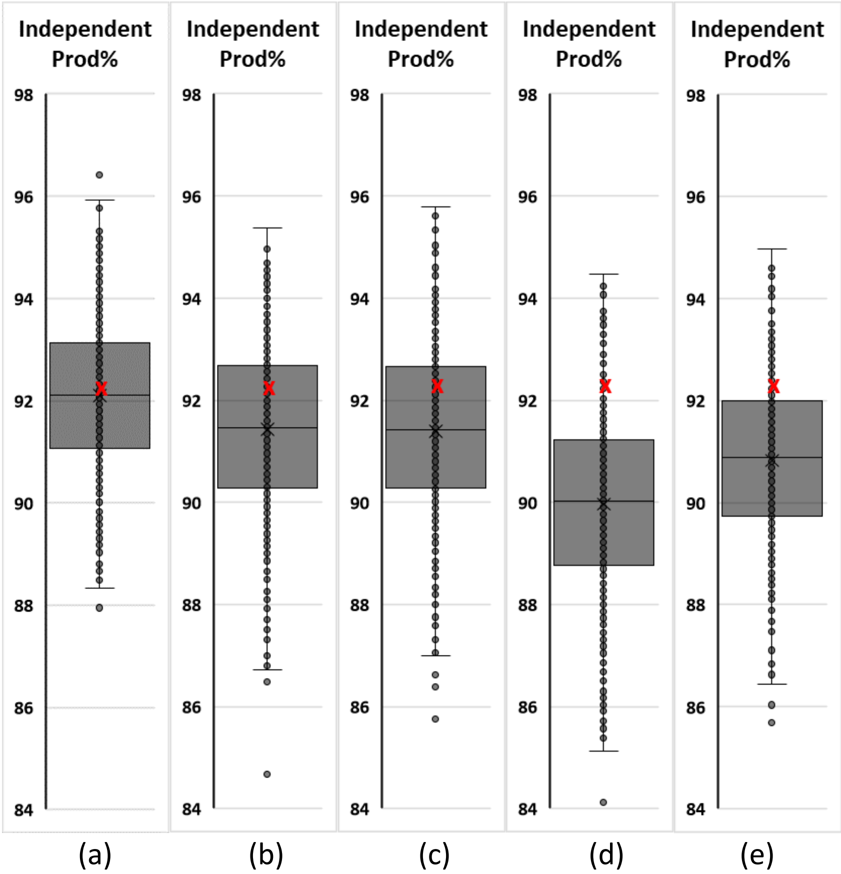

Figure 8-10 Lesson #5 Independent %Productivity in various scenarios with the TA removed (a) original validated lesson results with TA present and low TSOL=0.5 (b) TA removed but no other change (c) TA removed and students are instead inclined to ask the teacher (still low TSOL=0.5) (d) TA removed and students are instead inclined to ask the teacher but TSOL=5 (e) original validated lesson but TSOL=5 (Empirical lesson value **x**)

The plots in Figure 8-11 show the amount of TH associated with each of the scenarios. With TSOL=0.5, TH is approximately 25 minutes in (a), (b) and (c), but in (d) and (e), with TSOL=5, TH increased noticeably. This suggests that the changes in Independent %Productivity are due mainly to the increased TSOL=5. It seems that an increase in TH is associated with lower Independent %Productivity.

These results for Lesson #5 indicate that the removal of a TA did reduce Independent %Productivity slightly. However, the response of the teacher, particularly whether to compensate for the absence of a TA by increasing one-to-one support, may be a more significant factor, as discussed in Experiment 1.

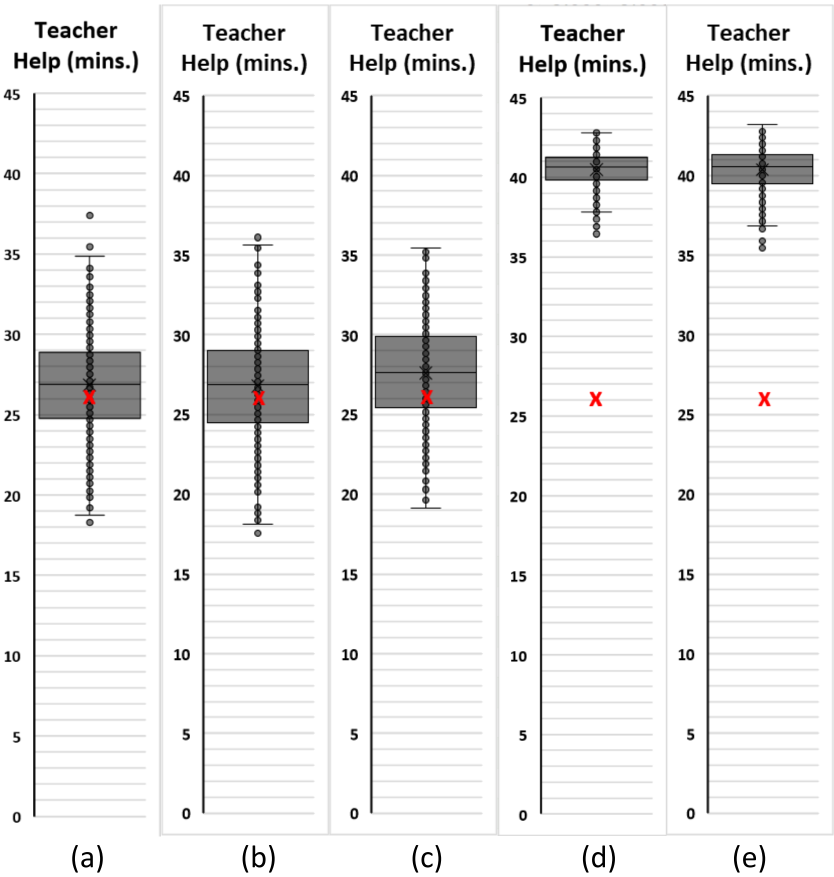

Figure 8-11 Teacher Help time in various scenarios with the TA removed
(a) original validated lesson results with TA present and low TSOL=0.5 (b) TA removed but no other change (c) TA removed and students are instead inclined to ask the teacher (still low TSOL=0.5) (d) TA removed and students are instead inclined to ask the teacher but TSOL=5
(e) original validated lesson but TSOL=5 (Empirical lesson value **x**)

8.2.2 The effect of adding a TA to Lesson #1, #3 and #4

In Lessons #1 and #3 a TA was added and in Lesson #4 the inactive TA (who had been out of the room for most of the lesson) was activated by increasing the value of the TASOL parameter. The results for the three lessons are shown in Figure 8-12, Figure 8-13 [19] and Figure 8-14. For each lesson, three plots are shown, summarizing the key aspects involved. In all cases:

(a) as TASOL was increased (the TA offers more help) TAH time increased – confirming that the parameter does control TAH appropriately;

(b) Independent %Productivity and TH time were fairly constant with small fluctuations.

[19] The additional TASOL=0.01 value was included because that was the value used during validation.

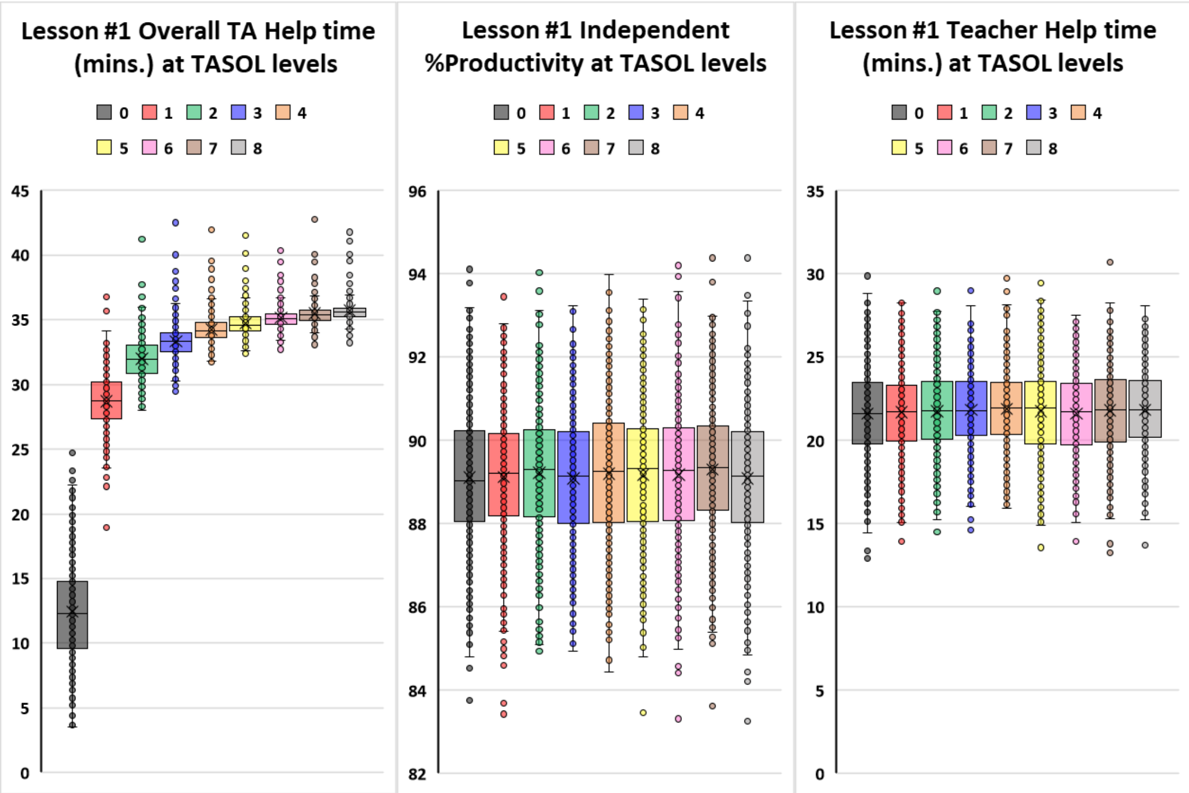

(a) (b) (c)

Figure 8-12 Lesson #1 metric distributions for several TASOL levels
(a) total student-TAH (b) Independent %Productivity (c) TH

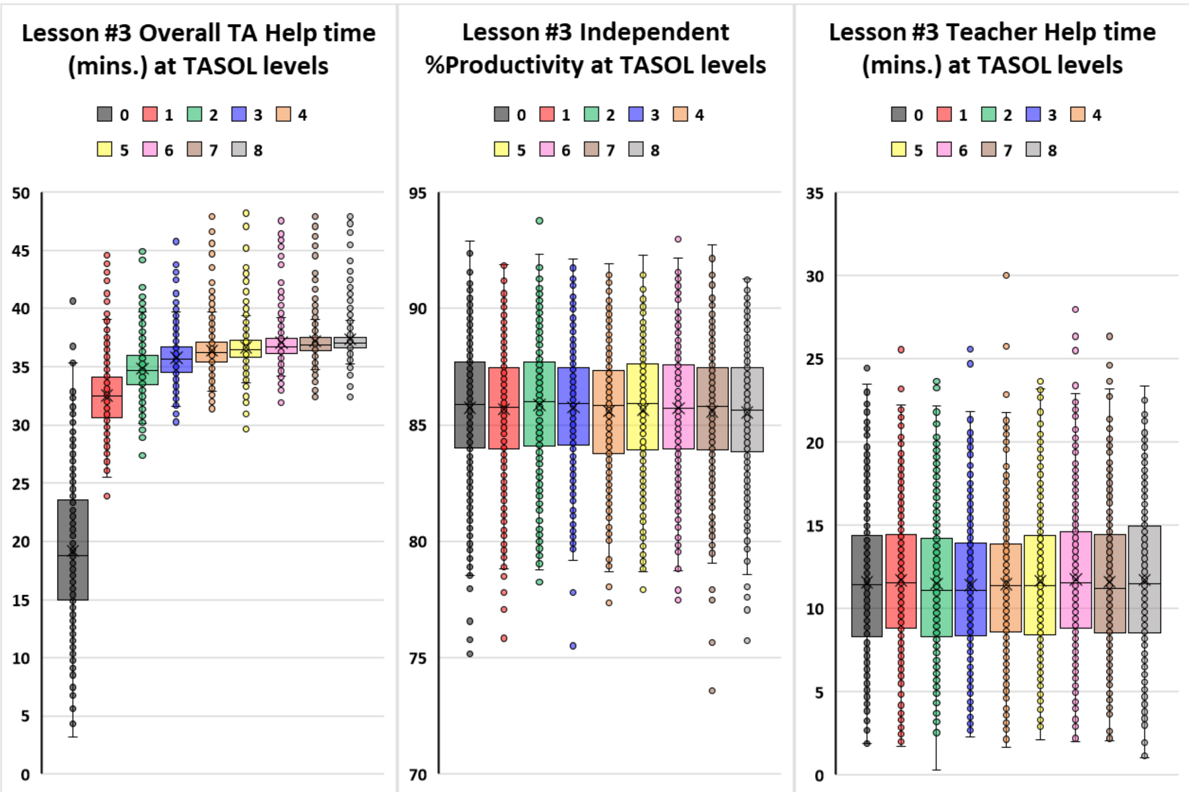

(a) (b) (c)

Figure 8-13 Lesson #3 metric distributions for several TASOL levels
(a) total student-TA Help (b) Independent %Productivity (c) Teacher Help

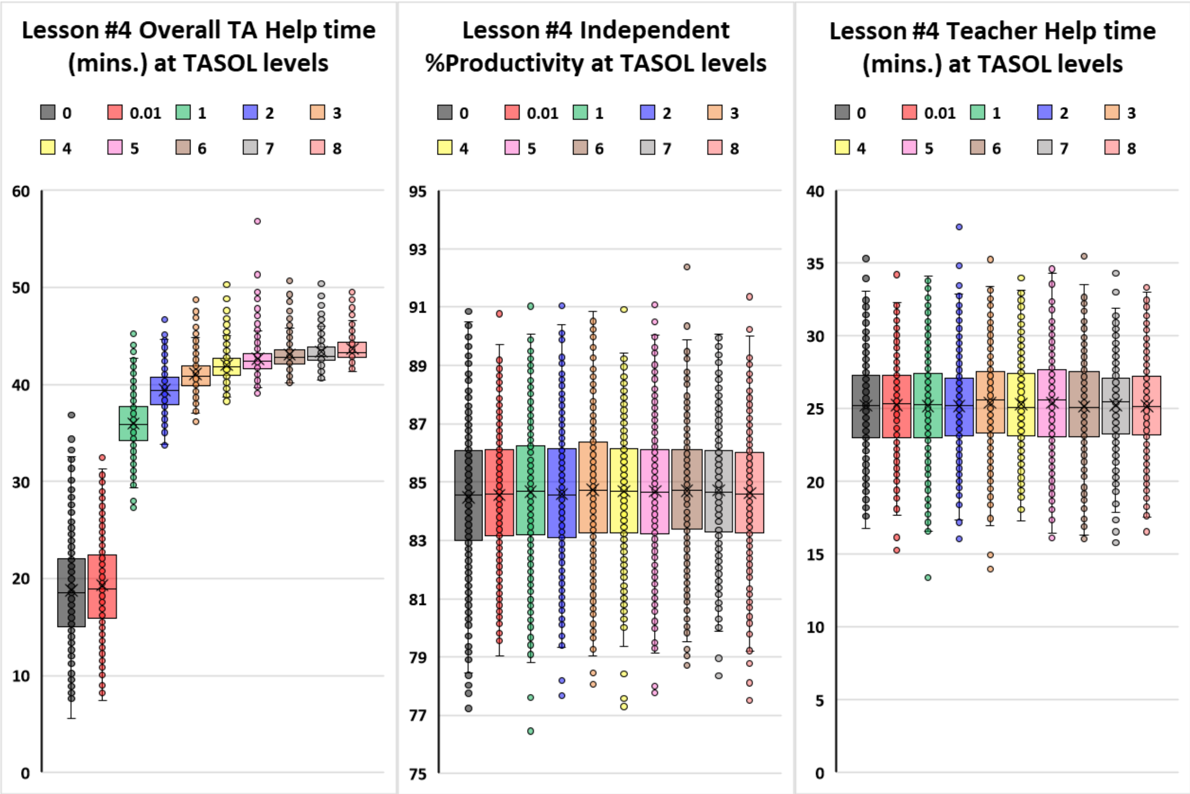

(a) (b) (c)

Figure 8-14 Lesson #4 metric distributions for several TASOL levels
(a) total student-TA Help (b) Independent %Productivity (c) Teacher Help

These results indicated that neither Independent %Productivity nor the amount of TH were directly related to the inclination of the TA to offer assistance. But this does not mean that the TA has no effect on the dynamics of a lesson. The state trajectories during different replications vary enormously. This is due to the presence of the TA plus the changes that this triggers in the random number sequences that all agents are using. Different random numbers lead to different state choices. As an example, Figure 8-15 shows three versions of Lesson #1. In each panel, time flows up and the teacher trace is the thin plot on the left-hand side. The first panel (a) contains the empirical lesson traces, the second panel (b) shows one simulation replication (particular pseudo-random number generator (PRNG) seed) of the lesson, and the third panel (c) shows the same replication (same PRNG seed) but with a TA added to the lesson. In this replication it happened that the TA helped the student in seat 6 (arrowed) extensively.

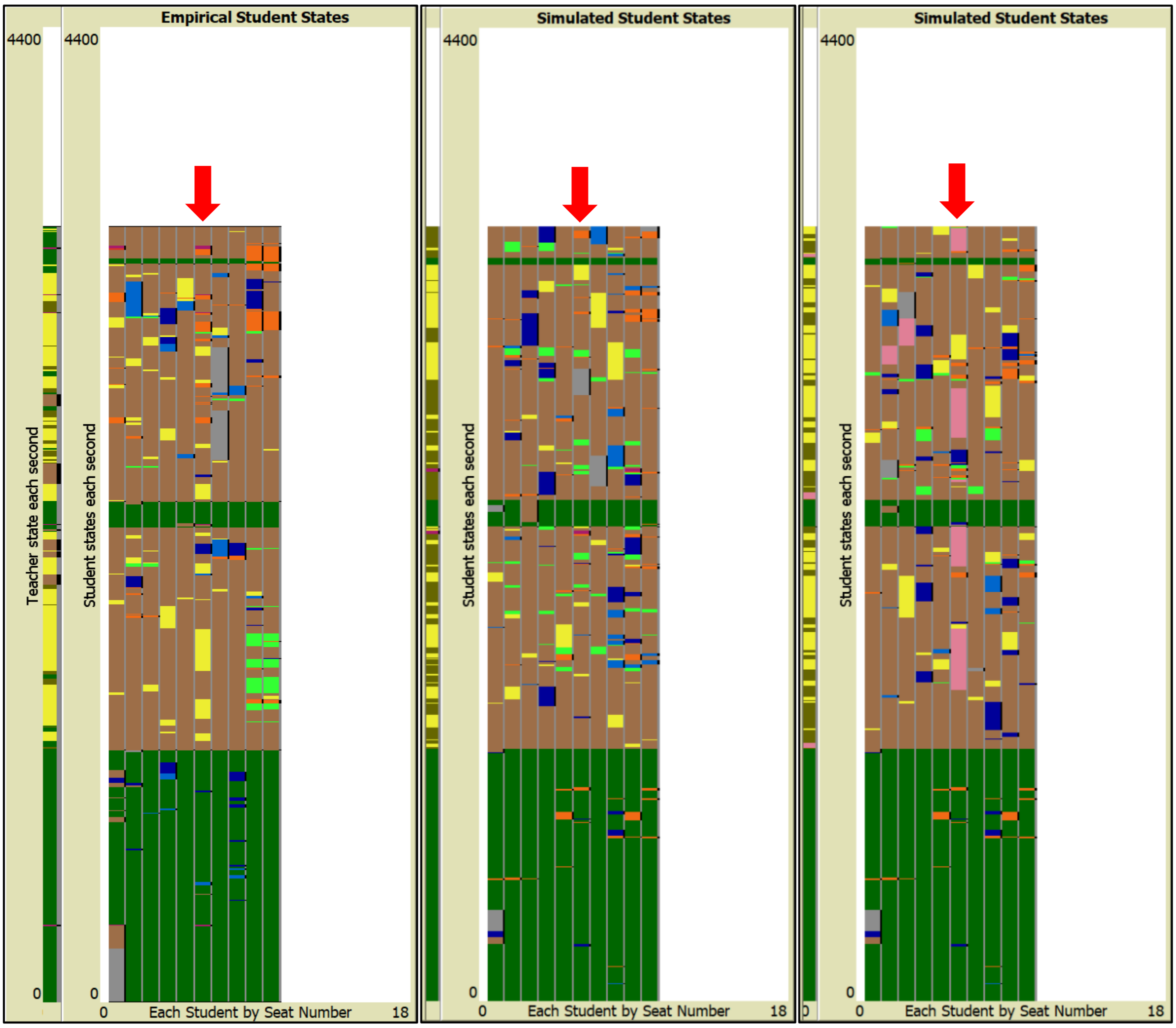

(a) (b) (c)

Figure 8-15 Example Lesson #1 replication
(a) empirical lesson replay (b) simulation with specific PRNG seed and no TA
(c) simulation with same PRNG seed but with a TA

It cannot be said that the presence of a TA alone caused the multitude of differences between (b) and (c). The addition of a TA (with TASOL = 0, so the TA only responds to students and never proactively offers help) altered the random number sequence so that everyone was choosing states differently. But notice that during the first whole-class teaching (green section) there is no visible difference in the simulation traces: only when independent working starts does the presence of the TA matter. In this case, the teacher talks to the TA (pink section in teacher's trace on the left of panel (c)) at the start of this lesson section. In this replication the output metrics (see Table 8-2) were improved with the TA (productivity increased, disengagement decreased), even though TH increased, and TD decreased (to zero).

Table 8-3 compares the change in productivity from no TA to having a TA with two different levels of TASOL over 500 replications. This shows that individual replications can experience increases or decreases in Independent %Productivity, sometimes with more than 1% change. In other words, it is possible that changes in productivity are a consequence of the randomness built into the model, rather than the presence or absence of a TA.

Table 8-2 Lesson #1 comparison of metrics without and with a TA (TASOL = 0)

Metric	Without TA	With TA
Independent %Productivity	90.4	92.2
Independent %Disruption	2.4	2.0
Independent %Other	7.2	5.8
Overall %Productivity	92.9	94.0
Overall %Disruption	2.0	1.8
Overall %Other	5.1	4.2
Teacher Help (mins.)	15	20
Teacher Disciplining (s)	27	0
TA Help	0	17

Table 8-3 Lesson #1 changes in productivity after TA added, at two different TASOL levels

Measure	TSOL=2 changes	TSOL=8 changes
No. of increases	338	247
No. of decreases	162	252
Average %increase	2.32	1.88
Average %decrease	-1.56	-1.84
Average %change	1.07	0.00
Maximum %increase	7.11	6.19
Maximum %decrease	-6.89	-7.10
Median %change	1.22	-0.02

8.2.3 Conclusions to research question 2

The research question being investigated was: Does the simulation model provide realistic explanations of the effects on overall student productivity of providing or withdrawing a TA who gives individual support to any student? It seems that the mere presence or absence of the TA and the amount of time the TA spends helping are not directly related to the amount of overall Independent %Productivity. What matters is the change in behaviour of the teacher and students as a consequence – as discussed in Experiment 1. However, as stated in the previous subsection, there is a possibility that the results were a consequence of the randomness built into the model, rather than the presence or absence of a TA. As discussed in section 8.1.3 in the case of the teacher helping students, to increase productivity, non-productive (i.e., disengaged) time has to be replaced by productive time. If the TA helps mainly students who would have been productive anyway, then there will be little effect on Independent %Productivity. The greatest impact would be when the TA always helped a student who would otherwise have disengaged. %Productivity values depends on which individual students were provided TA support, and that is random.

The results are naturally due to the rules in the model. Whereas the model includes mechanisms for the teacher to influence student misbehaviour choices, there are no rules that lead students to adjust their inclinations to misbehave because of the TA: it was a modelling assumption that the TA did not discipline the students. Students are more inclined to misbehave when the teacher is helping another student, but there is no such mechanism linked to the TA. Hence the simulation results are consistent with the model design. A more precise conclusion to this experiment is: the presence or absence of a TA who is not involved in student disciplining does not appear to directly influence the amount of overall Independent %Productivity.

In conclusion, the CLSM did provide plausible explanations for the influence or lack of influence a non-disciplining TA might have on student productivity. It also raised questions about the consequences of which students are helped (by either the TA or the teacher). As in the discussion about teacher help in section 8.1.3, if the TA targets those students who tend to misbehave then this may well increase productivity, but that would be unfair to well-behaved students.

It is important to reiterate that the CLSM is a model of behaviour, not learning. The activities of the TA may significantly boost individual learning without significantly altering the amount of overall productive time. In actual lessons, a teacher might change his/her lesson structure, planned activities and behaviour to maximize the benefits of having a TA (Hattie, 2012).

8.3 Experiment 3: Seating arrangements and student behaviour

This experiment addressed research question 3:

Does the simulation model provide realistic explanations of the effects on lesson behaviours and lesson outcomes of different student seating arrangements?

This experiment investigated the consequences of different student seating arrangements for student %Productivity, %Disruption, %Other and the time the teacher spent disciplining (TD). In addition, the results were compared to what an experienced teacher (the author) expected to happen. The general thinking amongst teachers is that, by separating students who frequently chat, the amount of overall disruptive behaviour would decrease. Another strategy is to try to put students together who would likely work productively together. During the case study both strategies had been described by the teachers and observed in practice. The teachers all acknowledged that they change the seating in an attempt to reduce distractions and increase focus, but also realized that some students just seemed to distract whoever they were next to. However, reducing disruptive behaviour does not automatically increase productive behaviour: students may instead increase their passive disengagement because now there is no-one to chat to. Another strategy teachers sometimes use is to deliberately seat chatty students next to each other. This sometimes happens when teachers have observed that these students distract whoever they sit next to so they might as well distract each other – and it is often easier to manage them this way. The teachers in the case study were also acutely aware of how sensitive the issue is with students and their parents, with some strong feelings about who should or should not be placed besides or near whom. This was one of the motivations for looking at lesson simulation.

For each of the four validated lessons, two seating rearrangements were selected, one that was expected to reduce student disruptive behaviour compared to the original seating (Reseat 1) and one that was expected to increase disruptive behaviour compared to the original seating (Reseat 2). For these two scenarios and the original lesson, 500 replications were run. For the two lessons that had empirical data for all the students (Lessons #1 and #3), the expectations for each seating rearrangement were compared to the empirical data; for the other two lessons (Lessons #4 and #5) it would have been inappropriate to compare the empirical data – which includes only monitored students – to whole-class results, hence the two reseatings were compared to each other. Afterwards, the simulation outcomes were compared to the teacher's expectations.

8.3.1 The chosen reseatings for each lesson

This subsection explains the reseatings chosen for each lesson. They were chosen by the author, being highly familiar with reseatings in his 14 years as a teacher. In the screen captures of the empirical lesson replays below, the students are labelled by seat number and this is their order in the state trajectories on the left-hand side of the figures (the thin plot on the far left being the teacher's state trace). Reseat 1 (expected to reduce student disruptive behaviour) is indicated by the blue arrow ➡; Reseat 2 (expected to increase student disruptive behaviour) is indicated by the red arrow ➡.

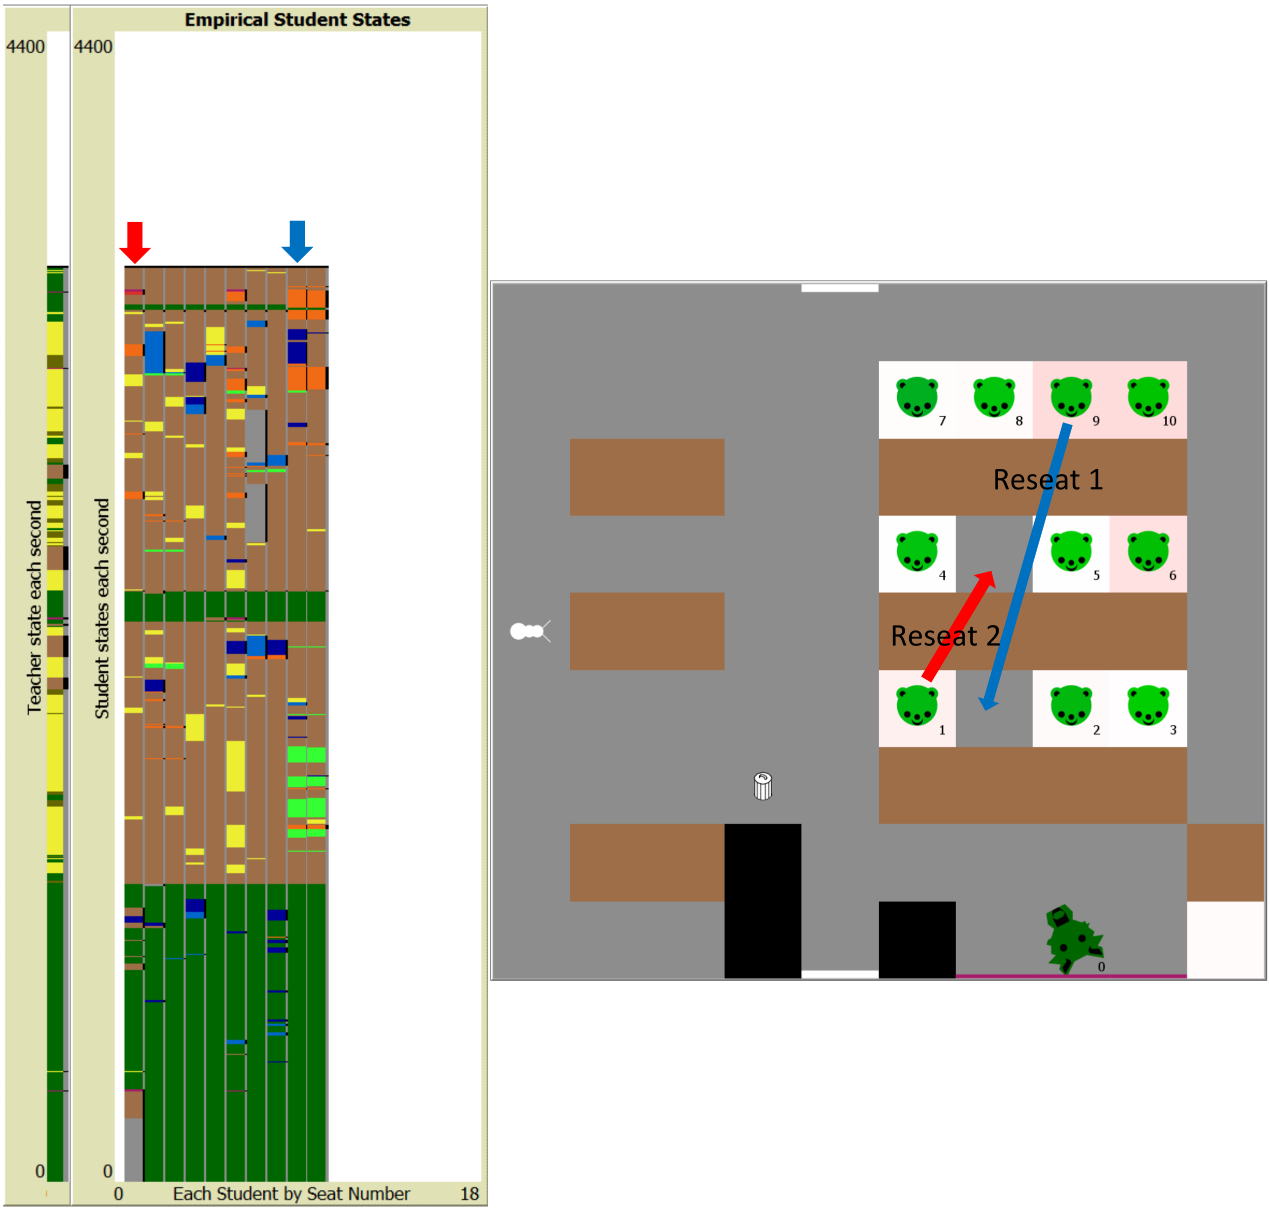

Figure 8-16 Lesson #1 replay and seating arrangement showing Reseat 1 (blue arrow) and Reseat 2 (red arrow)

From the empirical Lesson #1 replay in Figure 8-16, it can be seen that student 9 chatted quite a bit with their neighbour (pinkish background at end of the lesson and orange in their state trace). Assuming this was a common occurrence so the teacher decides to separate students 9 and 10 by bringing 9 to the front (closer to the teacher), leaving it more difficult for student 10 to chat. This is Reseat 1. As a contrast, Reseat 2 was expected to increase the amount of chatting by moving student 1 closer to other students and further to the back of the class, away from the teacher.

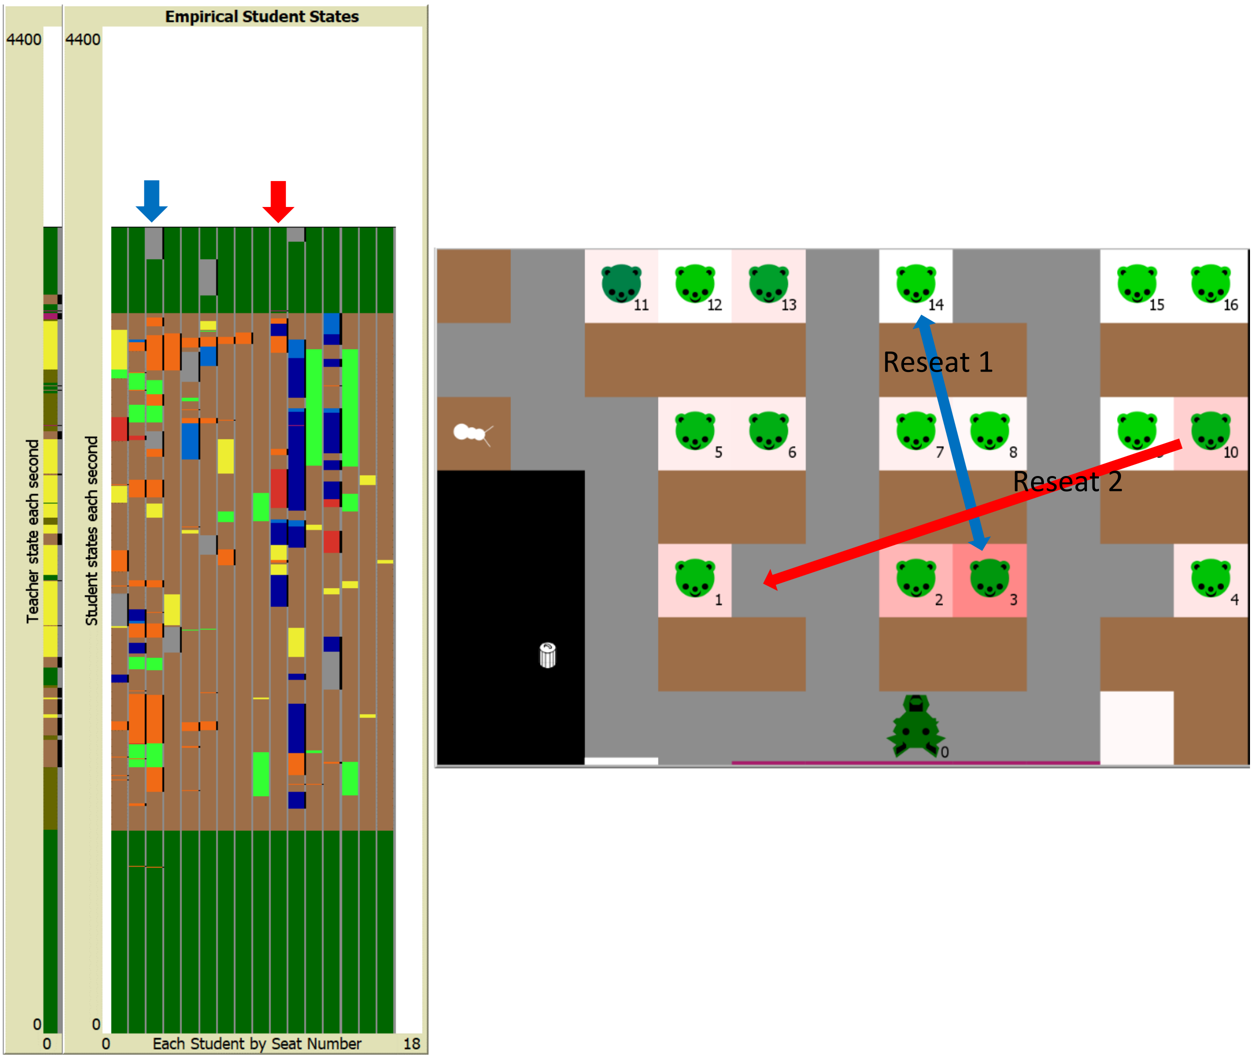

Figure 8-17 Lesson #3 replay and seating arrangement showing Reseat 1 (blue arrow) and Reseat 2 (red arrow)

From the empirical Lesson #3 replay in Figure 8-17, it can be seen that student 3 chatted quite a bit with their neighbour – some orange and red in their state trace and their seat background is very pink at the end of the lesson. Assuming this was a common occurrence so the teacher decides to isolate student 3 by <u>swapping</u> them with student 14. This is Reseat 1. As a contrast, Reseat 2 was expected to increase the amount of chatting by moving the chatty student 10 next to the slightly chatty student 1 (although student 10 would now be closer to the teacher and the front of the class).

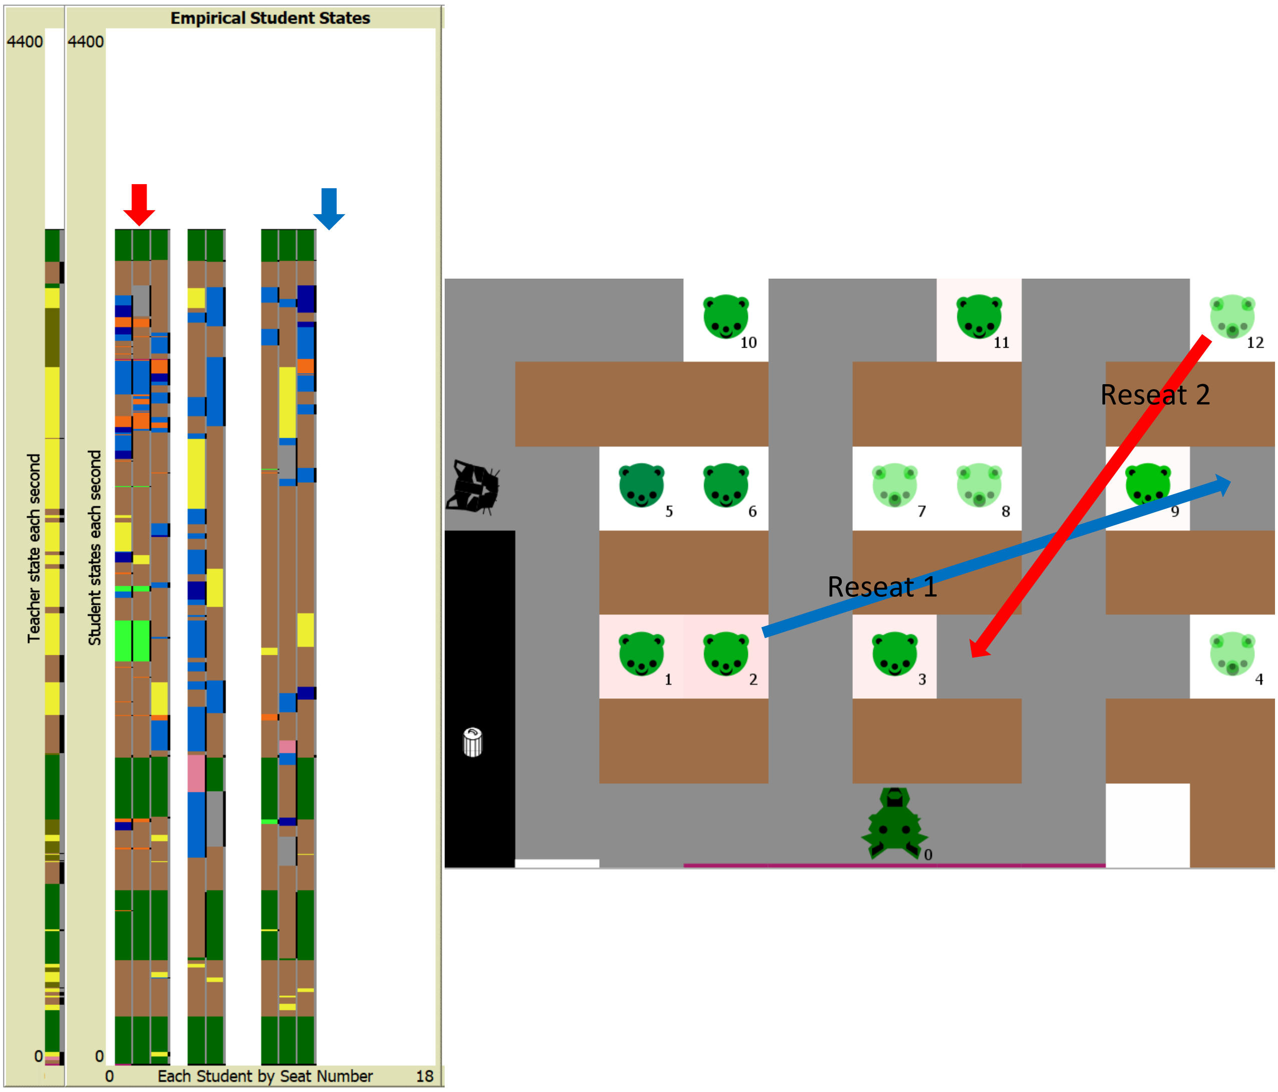

Figure 8-18 Lesson #4 replay and seating arrangement showing Reseat 1 (blue arrow) and Reseat 2 (red arrow)

From the empirical Lesson #4 replay (Figure 8-18), it can be seen that there was not much disruptive behaviour but student 2 chatted a little with their neighbours. Assume this was a common occurrence and so the teacher decides to move student 2. This is Reseat 1. As a contrast, Reseat 2 was expected to increase the amount of chatting by moving student 12 closer to others who chat, although this was also closer to the teacher and right at the front of the classroom. Note that student 12 was not monitored in this particular lesson (so there is no state trajectory in the left-hand plot and the student is 'ghosted' in the animation).

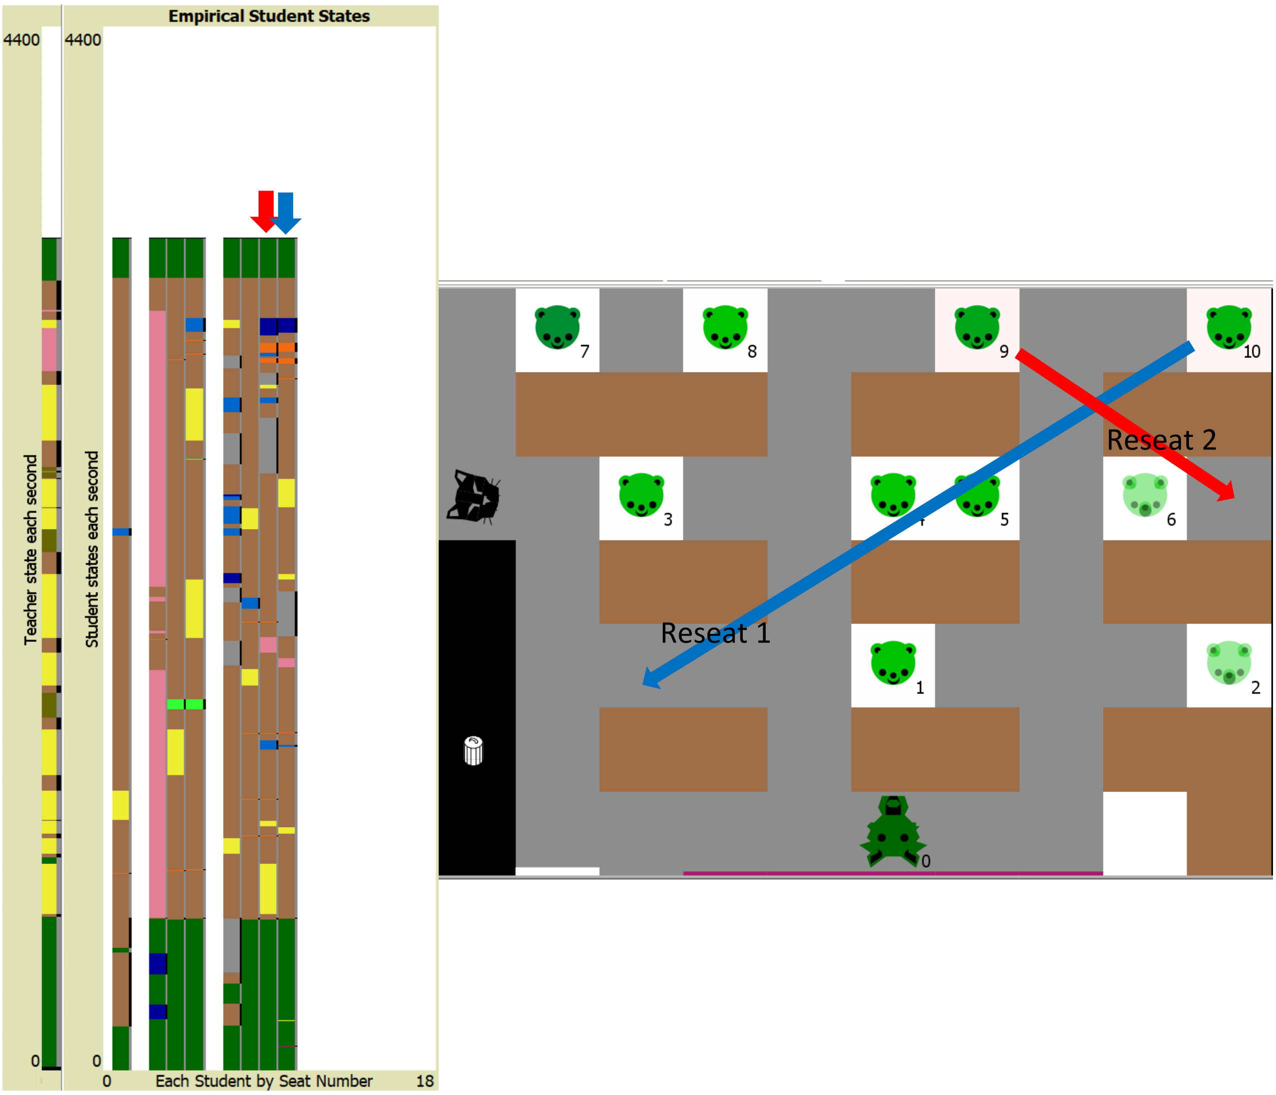

Figure 8-19 Lesson #5 replay and seating arrangement showing Reseat 1 (blue arrow) and Reseat 2 (red arrow)

Lesson #5 (Figure 8-19) presents a rather unrealistic scenario for reseating considerations in that there was almost no chatting in the empirical lesson. Nevertheless, suppose that the two students 9 and 10 did often chat in lessons and the teacher wanted to try two alternative new arrangements. As mentioned earlier, there are many other reasons for rearranging seating. For example, it may help the teacher provide more individualised support if a student is brought to the front, or the teacher might want to see how the students in seats 6 and 9 work together.

8.3.2 The effects of the seating rearrangements

First the results of Lessons #1 and #3 are discussed. In these experiments, the expectations associated with the seating rearrangements are compared to the empirical lesson data. Lesson #1 results are shown in Figure 8-20 . %Other and TD seem barely affected while %Productivity and %Disruption - called %Distracting in these plots and referred to as %Dist - show small variations. But, as was pointed out at the start of section 8.1, even 0.02% disruption alters the atmosphere in a lesson. Lesson #3 results are shown in Figure 8-21. These appear to show negligible changes due to the reseatings.

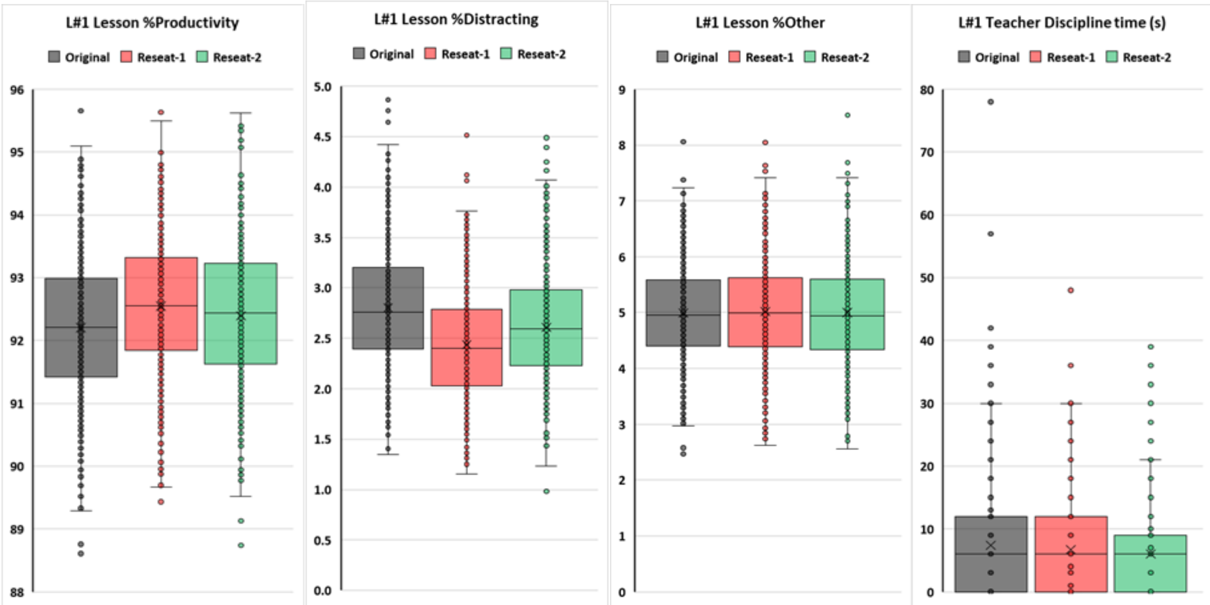

Figure 8-20 Lesson #1 reseating results for %Productivity, %Distracting, %Other and TD

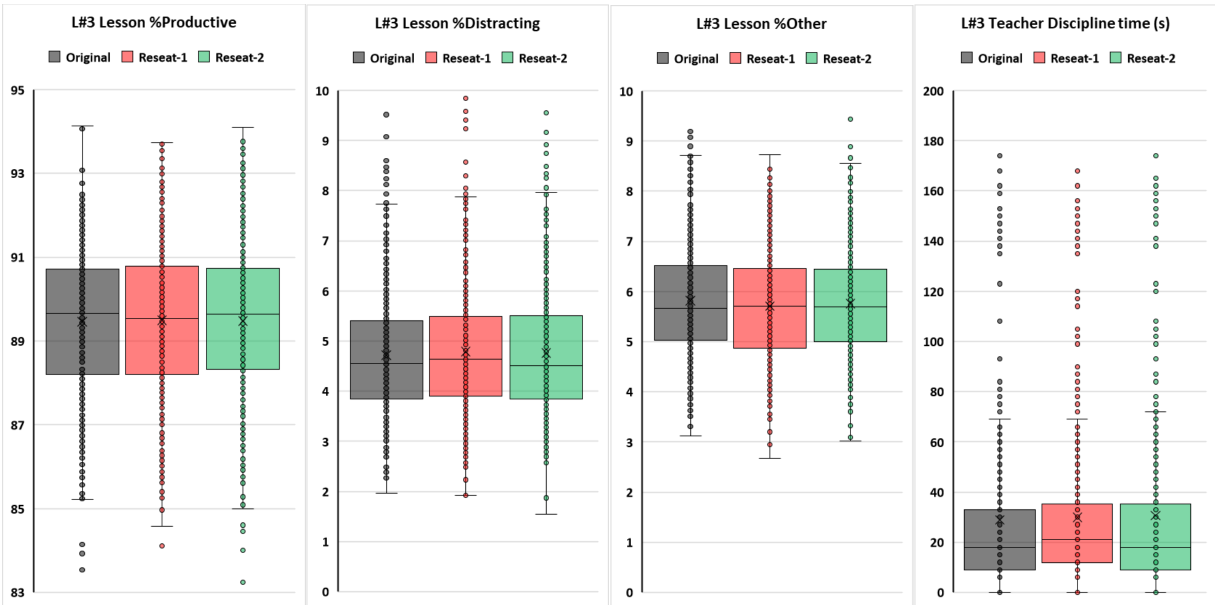

Figure 8-21 Lesson #3 reseating results for %Productivity, %Distracting, %Other and TD

However, the goal of this investigation was to evaluate whether the reseatings had the expected effects on productivity, disruption and other behaviour. Table 8-4 summarizes the results for Lesson #1. Considering just the direction of the expected change, of the 6 expectations (predictions) made, 3 were met and 3 were not. Table 8-5 summarizes the results for Lesson #3. Of the 6 expectations (predictions) made again 3 were met and 3 were not.

Table 8-4 Lesson #1 expectations for the reseatings

Reseat 1	Expectation	Empirical	Simulation	Difference	Expectation met?
%Prod	up	92.201	92.545	-0.343	yes
%Dist	down	2.803	2.437	0.366	yes
%Other	-	4.996	5.018	-0.023	-
TD time (s)	down	7.378	6.676	0.702	yes
Reseat 2	**Expectation**	**Empirical**	**Simulation**	**Difference**	**Expectation met?**
%Prod	down	92.201	92.392	-0.190	no
%Dist	up	2.803	2.610	0.193	no
%Other	-	4.996	4.998	-0.002	-
TD time (s)	up	7.378	6.040	1.338	no

Table 8-5 Lesson #3 expectations for the reseatings

Reseat 1	Expectation	Empirical	Simulation	Difference	Expectation met?
%Prod	up	89.463	89.497	-0.034	yes
%Dist	down	4.719	4.792	-0.073	no
%Other	-	5.818	5.711	0.107	-
TD time (s)	down	28.908	29.988	-1.080	no
Reseat 2	**Expectation**	**Empirical**	**Simulation**	**Difference**	**Expectation met?**
%Prod	down	89.463	89.477	-0.014	no
%Dist	up	4.719	4.761	-0.042	yes
%Other	-	5.818	5.762	0.056	-
TD time (s)	up	28.908	30.668	-1.760	yes

To summarize, for the two lessons in which the simulation outputs could be compared to the empirical lessons, 6 out of the 12 expectations were met and 6 were not.

For the other two lessons, where comparison with the empirical lessons was inappropriate due to the presence of unmonitored students, the comparisons were to establish whether Reseat 1 caused less disruption and more productivity than Reseat 2. The results are shown in Figure 8-22 and Figure 8-23.

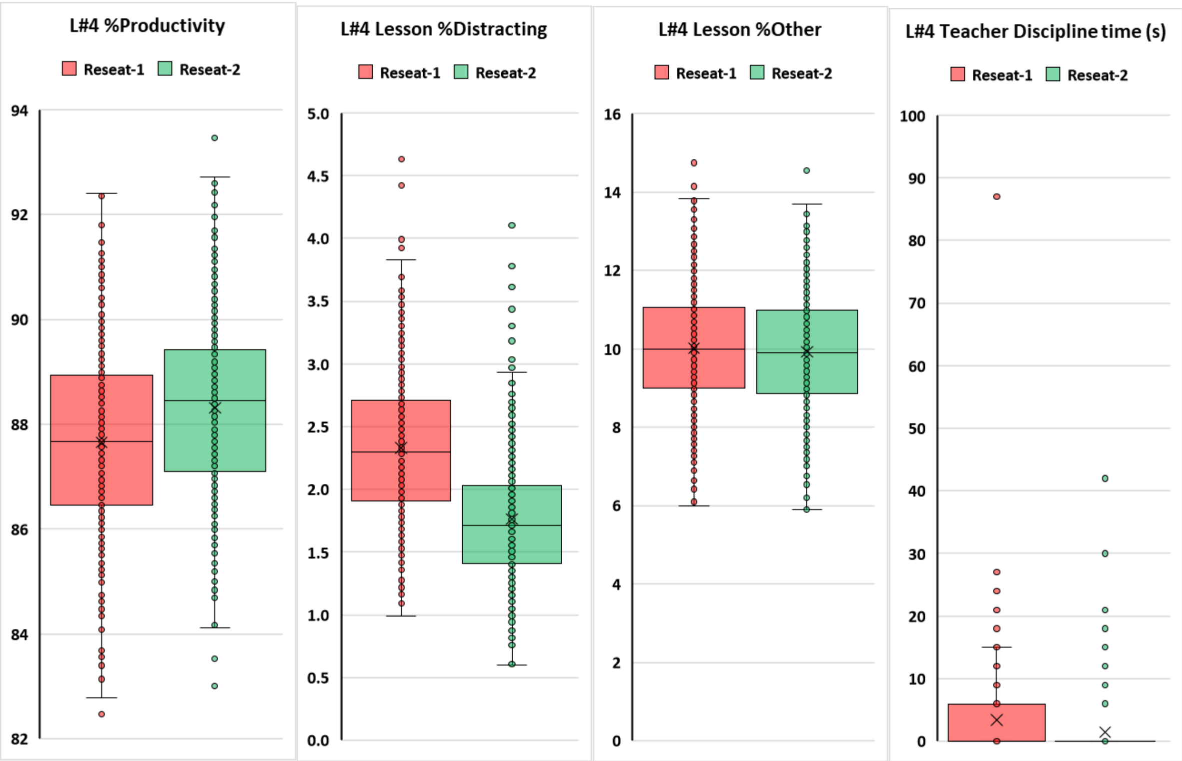

Figure 8-22 Lesson #4 reseating results for %Productivity, %Distracting, %Other and TD

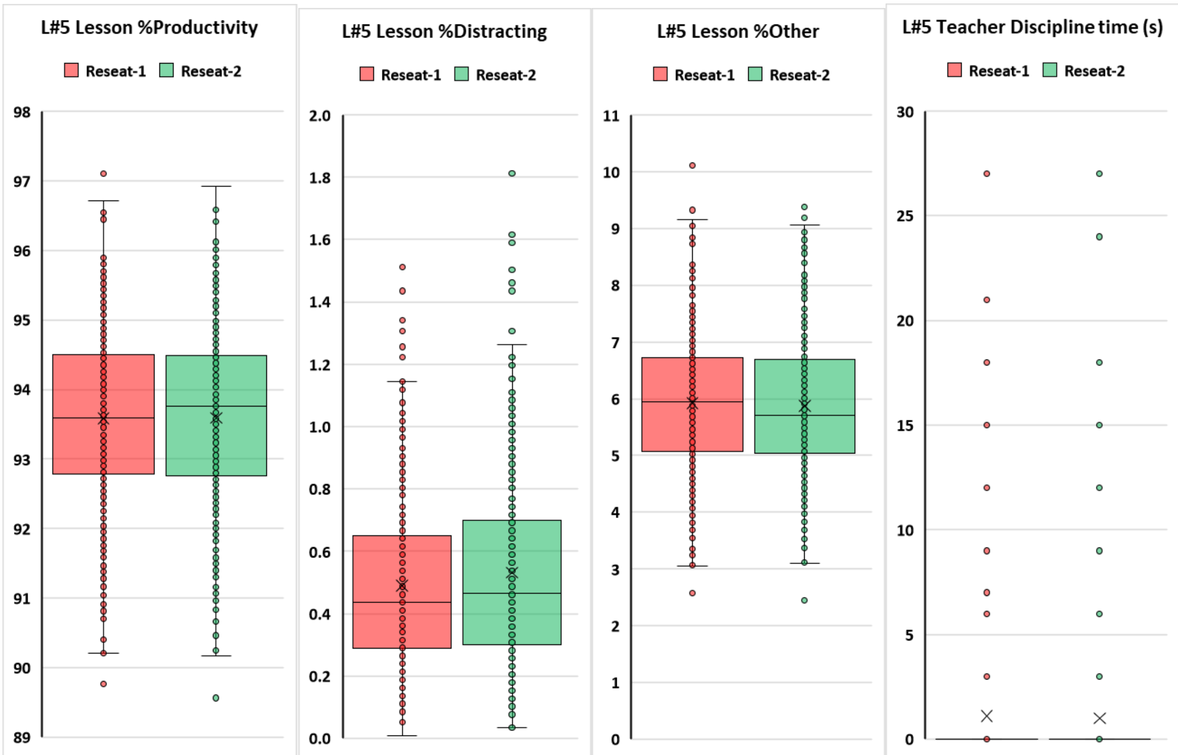

Figure 8-23 Lesson #5 reseating results for %Productivity, %Distracting, %Other and TD

In Lesson #4, %Other and TD seem barely affected while %Productivity and %Distracting show small but tangible variations (as mentioned above). The results for Lesson #5 show negligible changes due to the reseatings. Checking again whether expectations were met, Table 8-6 and Table 8-7 show that only 2 out of the 6 expectations were met.

Table 8-6 Lesson #4 expectations for the reseatings

Lesson #4	Expectation	Reseat 1	Reseat 2	Expectation met?
%Prod	Reseat 1 greater than Reseat 2	87.655	88.313	yes
%Dist	Reseat 1 lower than Reseat 2	2.332	1.764	no
%Other	-	10.014	9.923	-
TD time (s)	Reseat 1 lower than Reseat 2	3.426	1.416	no

Table 8-7 Lesson #5 expectations for the reseatings

Lesson #5	Expectation	Reseat 1	Reseat 2	Expectation met?
%Prod	Reseat 1 greater than Reseat 2	93.586	93.592	no
%Dist	Reseat 1 lower than Reseat 2	0.491	0.532	yes
%Other	-	5.923	5.875	-
TD time (s)	Reseat 1 lower than Reseat 2	1.106	1.008	no

To sum up, considering just the directions of the expected results, expectations were met on only 8 out of 18 occasions, i.e., 44% of the time. Note that TD time is in seconds and the differences between the means were 1 or 2 seconds, not a significant amount in a 1-hour lesson. %Other values had been included in case a decline in disruption was accompanied by an increase in passive disengagement, resulting in no increase in productivity. However, there was no consistent evidence for such a relationship.

8.3.3 Conclusions to research question 3

The four lesson models had been validated and were therefore assumed to be operating within their 'domain of applicability' (Schlesinger *et al.*, 1979, p. 104). However, the results were that 56% of the time the seating changes did not have the anticipated outcomes – which is worse than chance. The research question being investigated was: Does the simulation model provide realistic explanations of the effects on lesson behaviours and lesson outcomes of different student seating arrangements? But the results raised the question as to whether the simulations were producing unrealistic outcomes and the teacher's predictions were better, or whether the simulations were realistic and it was the teacher who was not a reliable predictor of the consequences in real lessons? To answer this question definitively would require empirical testing of the rearrangements and the teacher's predictions. If the simulation model happens to be realistic, still the effects appeared small. But, as has been emphasised, even a very small percentage change in disruption can have a huge effect on lesson quality. This is why teachers go to the bother of making seating changes – the benefits of a small reduction in disruption can be worth it. The results of this experiment would imply that a teacher's predictions about the consequences of reseatings may be unreliable, and therefore that it is entirely possible that in many cases a reseating merely shifts the location of the disruptive behaviour. Teachers in the case study reported that often their seating reorganizations failed for that reason. Perhaps the consequences of a seating rearrangement are just too difficult to pin down, for teachers and simulation models. This leads one to conclude that the simulation model did provide realistic explanations, probably as reliable as that of a teacher, but, as stated, this would need further research to be confirmed.

This experiment demonstrated a potential benefit of lesson simulations as outlined in section 1.1. With a simulation tool one can try out scenarios without disadvantaging anyone and without wasting resources. In fact, one can test an intervention hundreds of times, whereas a teacher may be able to test an intervention only a few times. In principle, simulations could be run on all seating permutations and the best ones selected. If a seating rearrangement produced promising outcomes the teacher could actually try it for several lessons to build up a picture of the results (i.e., not just one lesson, just as one would not run only one simulation replication).

There are probably several factors that could affect the consequences of a reseating. For example, the consequences may well depend on the particular student's strength of influence in the class and how influenceable the students around are. However, these social factors are not included in the CLSM. The specific location of a disruptor may also be a factor and this can be investigated using the CLSM. Experiment 4 follows up on this experiment to explore if the location of disruptors has much impact on overall disruption.

8.4 Experiment 4: Experiments with artificial classes

This experiment addressed research question 4:

Does the simulation model provide realistic results in experiments with a class of completely artificial students?

To create artificial students, one would ideally like to derive profiles from empirical data that quantified several characteristics (as in *simSchool* (Deale and Pastore, 2014)). However, for even a very simple binary categorization on just three characteristics, high or low disruptivity, high or low passive disengagement and high or low one-to-one support (from the teacher or the TA), there was insufficient empirical data to populate each of the resulting eight categories. Instead, only two profiles were constructed: an ideal student profile and a disruptive student profile. A profile is a set of the three PMFs (explained section 3.3.2). The disruptive student profile was the profile of a real student who was disruptive in classes, student 802. The ideal profile was a set of idealized PMFs (based on the empirical data), assuming a student sometimes chatted, sometimes disengaged, sought help a typical amount of time and generally behaved as required. Thus both profiles were realistic. Figure 8-24 summarizes for each student over all their lesson time the percentage of productive versus disruptive behaviour. Student 802 (circled in red) was disruptive approximately 12% of lesson time and productive 81% of the time, with the 8% balance being passive disengagement time (which was not extreme, unlike the four students to the bottom left). The ideal student profile (positioned at the green star) had 5% disruptive and 90% productive behaviour with 5% for passive disengagement.

Two experiments were conducted. The first experiment (Experiment 4a) followed up on what was observed in Experiment 3 and explored why the locations of disruptors might or might not have much consequence. It investigated the consequences of replacing one or more ideal students with a disruptive student. The results of this experiment raised a question which motivated a follow-up experiment (Experiment 4b) to investigate the proportion of disruptive vs other (passive disengagement) behaviour.

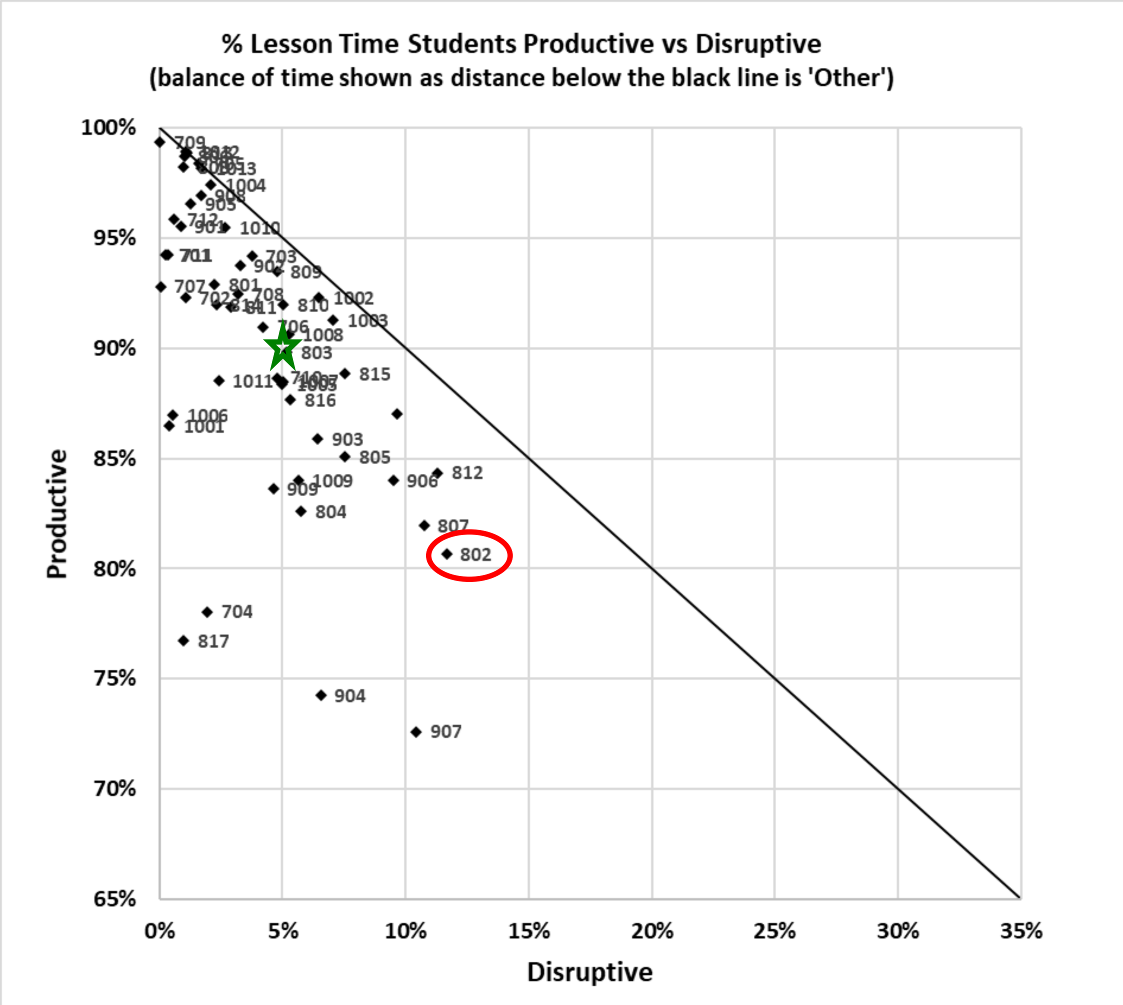

Figure 8-24 Percentage of lesson time students spent productively vs disruptively with student 802 (circled in red) and the 'ideal' student (green star)

8.4.1 Experiment 4a: Investigating the effect of disruptive students

To investigate the consequences of replacing one or more ideal students with a disruptive student, an artificial lesson was constructed by taking Lesson #5, its lesson plan, seating arrangement, teacher, TA and validated parameter values, but replacing various original students with artificial students. Using the parameter set for Lesson #5 on the artificial lesson produced results that satisfied all the acceptability criteria. The lesson plan had only three sections: 11 minutes whole-class teaching at the start, 45 minutes independent working and 3 minutes whole-class teaching at the end. The three students to be replaced initially are shown in Figure 8-25. In one of the scenarios, student 1 is moved to the back of the class, indicated by the red arrow.

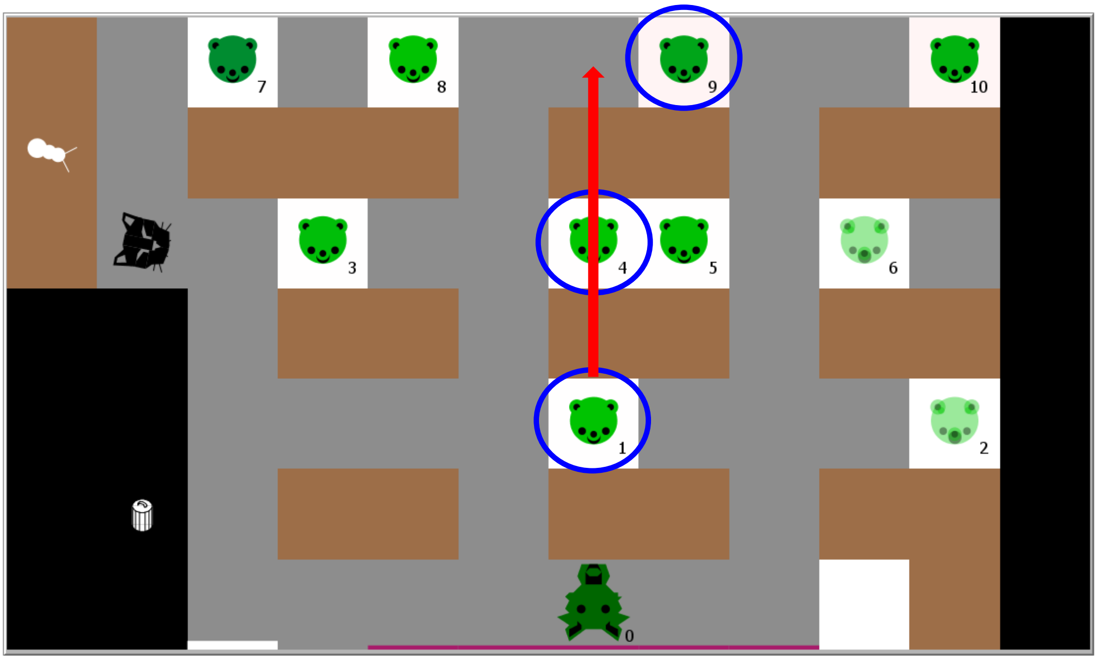

Figure 8-25 The students who became disruptive (labelled by seat number)

This experiment compared the following seven scenarios:

	Label	Scenario	Use or expectation
1.	ideal	the ideal class – no disruptive students	baseline for comparison
2.	original	the actual Lesson #5 with the real students	for comparison
3.	seat 1	replace the student in seat 1 with a disruptive student	increased disruption
4.	seat 4	replace the student in seat 4 with a disruptive student	increased disruption
5.	seat 9	replace the student in seat 9 with a disruptive student	increased disruption
6.	all 3	replace all three students in their seats	greatly increased disruption
7.	3 together	replace all three students and put them in close proximity (move student in seat 1 to back – red arrow in Figure 8-25)	even more disruption

The results of 500 replications of each scenario are summarised in Figure 8-26.

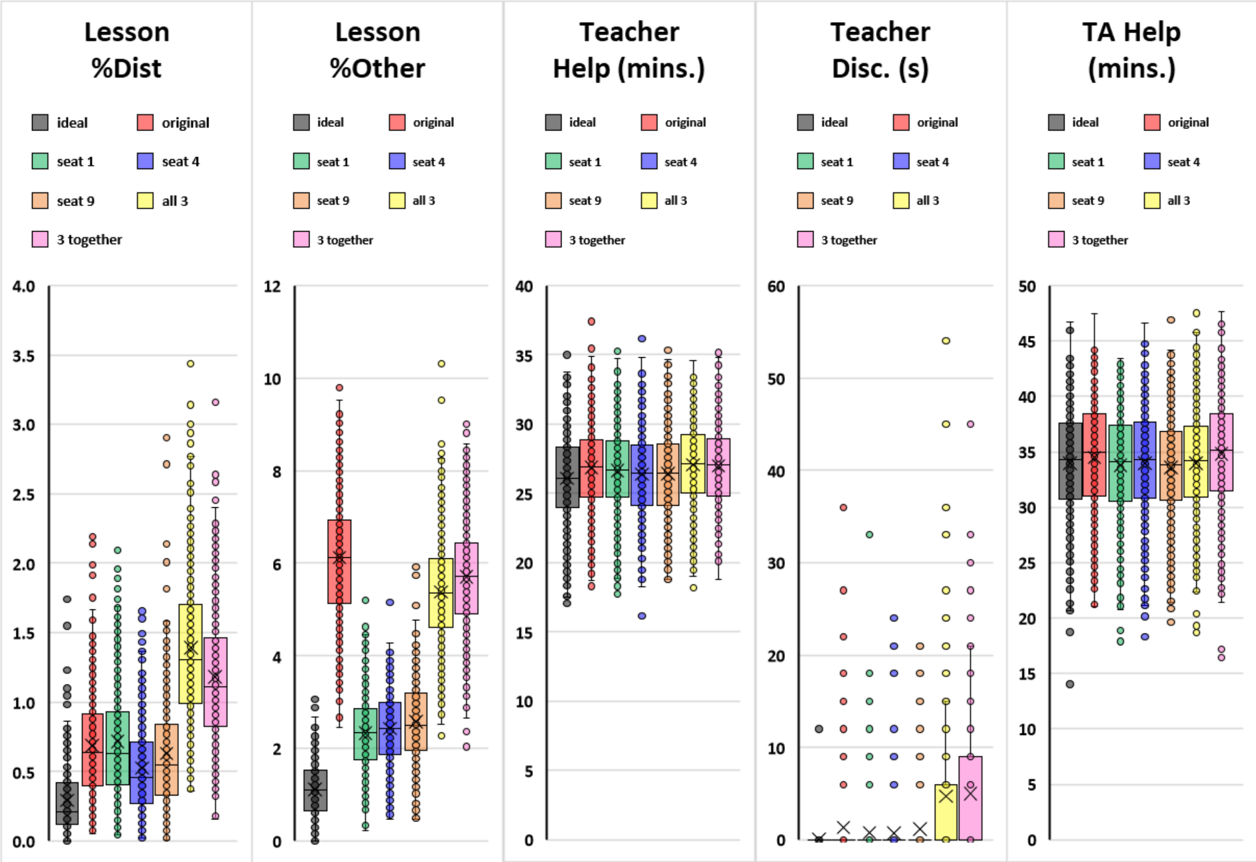

Figure 8-26 The effect of the disruptive students on metrics %Dist, %Other, TH, TD and TAH

Firstly, note that the amount of teacher and TA help were barely affected by the changes, not surprising since the teacher parameters did not change, therefore confirming that the model performed as required. Secondly, scenarios 6 and 7 (yellow and pink) resulted in increased discipling by the teacher – as one would expect if students misbehave more. The main focus though was %Disruption (labelled %Dist for %Distraction in the plots). Relative to the ideal class, adding disruptive students always increased %Dist, particularly when there were three disruptive students. %Other (passive disengagement) also increased relative to the ideal class - possibly because the disruptive student profile also had higher scores for passive disengagement than the ideal student profile (8% vs. 5%).

What was <u>not</u> anticipated was that disruption (%Dist) would decrease from scenario 6 (the three disruptors separated) to scenario 7 (the three disruptors seated together). It had been expected that putting all the disruptors together would increase overall disruption, not lead to a decrease. To investigate this result, the details of each student were inspected. Table 8-8 summarizes the relative changes in the metrics per student (labelled by seat number, with disruptive students highlighted). An empty cell means there was no significant change, a single arrow means there was a slight change and two arrows means there was a small but distinct change. Note that %Dist includes student states 1 (being disciplined by the teacher), 2 (unproductive, away from own desk and distracting others), 3 (in own seat chatting, distracting, socialising, turning around etc.) and 5 (listening to teacher in response to the class being disciplined). And %Other includes states 4 (intentionally unproductive, not participating, but not distracting others, at or away from own desk), 7 (not sure if productive: just sitting, not disturbing others - maybe thinking or waiting for instructions or help) and 17 (left the classroom). (The student states are explained in section 3.1.)

Table 8-8 The relative changes between scenarios 6 and 7 for each student

Student Seat	%Dist	%Other	Teacher Help	Teacher Disc.	TA Help
1	↓	↑↑			
2					
3					
4					↓
5					
6					↓
7					↑
8					
9	↓	↑		↑	↑
10					

Compared to scenario 6, in scenario 7:

- the student in seat 4 behaved overall as before, with a little less TA help (and very slightly less chatting and passive disengagement).
- the student in seat 9 and the student who moved from seat 1 changed their behaviour; both chatted less and disengaged passively more, especially the student from seat 1.
- the student in seat 9 was disciplined more.
- the other students all had negligible changes in overall %Dist and %Other, and all students had negligible changes in Teacher Help time.

From these data, it appears that seating disruptive students 1 and 9 closer together led to their disruptive behaviour being converted to other disengaged behaviour. There are several mechanisms that could explain this. Figure 8-27 shows the factors incorporated into modelling student misbehaviour.

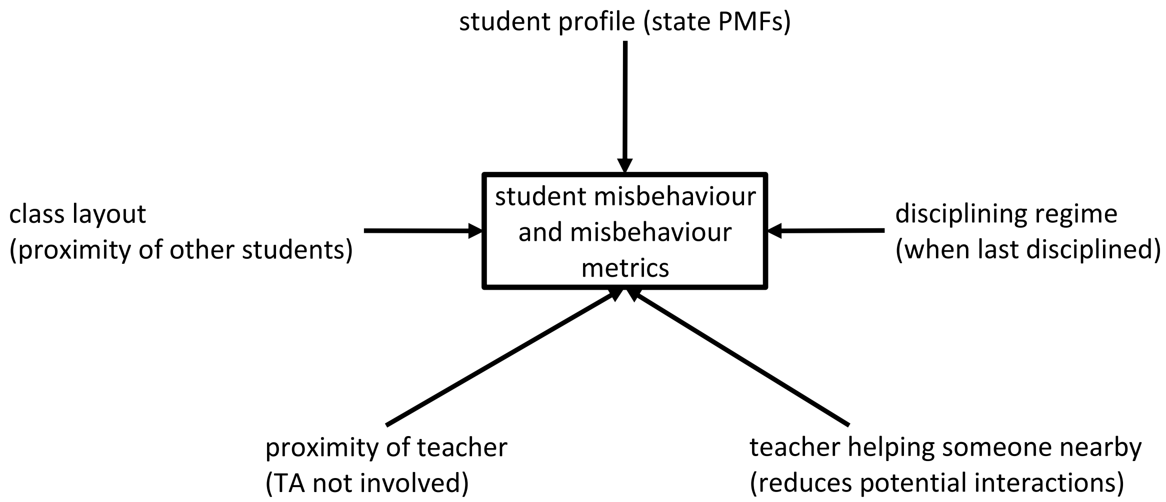

Figure 8-27 Factors affecting student misbehaviour and lesson misbehaviour metrics

The model has rules that explicitly make students less inclined to misbehave, especially disruptively, after they have been disciplined. The rules also make the teacher stricter with disruptive behaviour than with passive disengagement. In scenario 7 (with all three disruptive students in close proximity), the student in seat 9 was disciplined more so this probably accounted for their decrease in chatting and increase in passive disengagement. Another factor was the proximity of the teacher. The model has rules that decrease the likelihood of misbehaviour the closer the teacher is. When the teacher was helping any student near the middle rear of the class, that affected all the students nearby, including the disruptive ones, reducing the likelihood of their chatting (or passively disengaging). Having the teacher helping the student next to you is not only a disincentive for you to misbehave, it removes one student from the list of those you could chat with. The model rules explicitly prevent a student interacting with a student receiving help (from the teacher or TA). Class layout also plays a role in the amount of disruption. If everyone is quite spread-out then the spread of chatting is impeded. But students can always make a comment briefly out loud to no-one in particular or disengage passively.

This analysis concerns the model rules and is not irrefutably inevitable real-life behaviour. But all the explanations given are plausible, and the results do match the observations and the teacher and student accounts in the case study: a student who needs a break (perhaps because they are struggling or have worked for a long time) might try to interact with someone, and, if that fails they may take a break by resting.

However, the plots in Figure 8-26 show that in all scenarios, %Other was more prevalent than %Dist. This was unexpected because in the student profiles the probabilities of other (passive disengagement) versus disruptive behaviour are 8% versus 12% for the disruptive students and 5% versus 5% for ideal students. These results may have been due to the factors discussed above but warranted further investigation in a second experiment.

8.4.2 Experiment 4b: Active and passive disengagement in a full classroom

The purpose of this experiment was to determine whether the amount of disruptive behaviour compared to passive disengagement depended on the 'density' of the student seating. The previous experiment had 10 students who were quite spread out in the classroom. Experiment 4b used the same room layout as 4a but with the maximum number of students (19) (see Figure 8-28), so most students had a neighbour to chat with most of the time. In this experiment the lesson plan was copied from a Year 7 Maths lesson with three sections: 15 minutes whole-class teaching, 38 minutes independent working and 6 minutes whole-class teaching. No TA was used. A parameter set was contrived by setting all the parameter weightings to have no effect on the students' base state probabilities (which were again of only two types, ideal or disruptive). The teacher was given a low inclination to offer help (TSOL=0.5) and a moderate tolerance of misbehaviour (TMT=50).

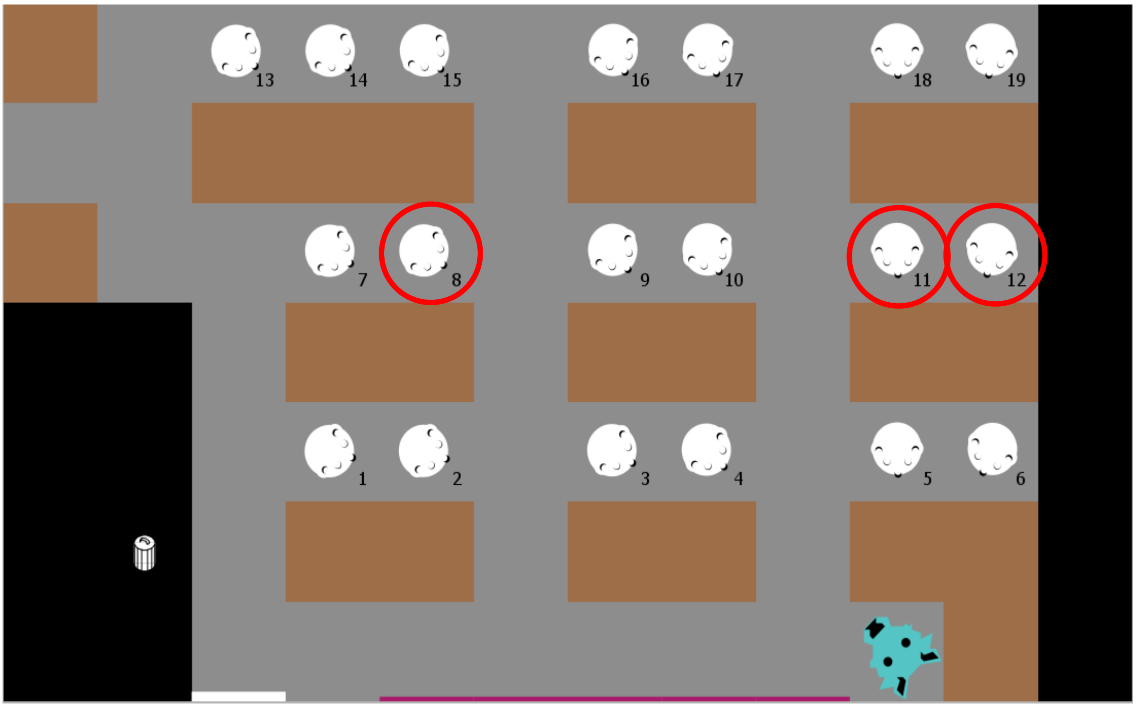

Figure 8-28 Full class experiment with students selected for disruptive behaviour highlighted

The experiment compared the five scenarios below. The expectation was that disruption (%Dist) would increase from scenario to scenario until scenario 5, when the teacher would discipline more and the students would be inclined to disengage passively – so %Other was expected to increase while %Dist would decrease.

	Label	Scenario	Use or expectation
1.	ideal	the ideal class – no disruptive students	baseline for comparison; extremely high ASD
2.	apart	two students apart (seats 8 and 11) disruptive	slight increase in disruption (small compared to whole class productivity)
3.	together	two students together (seats 11 and 12) disruptive	larger increase in disruption (still small compared to whole class productivity)
4.	misbehave	all students disruptive	maximum disruption; %Dist increases more than %Other; reduced ASD
5.	stricter	all students disruptive but teacher stricter (TMT=20)	more disciplining, perhaps whole class; reduced disruption but increased passive disengagement; less reduced ASD

The results of 500 replications of each scenario are shown in Figure 8-29.

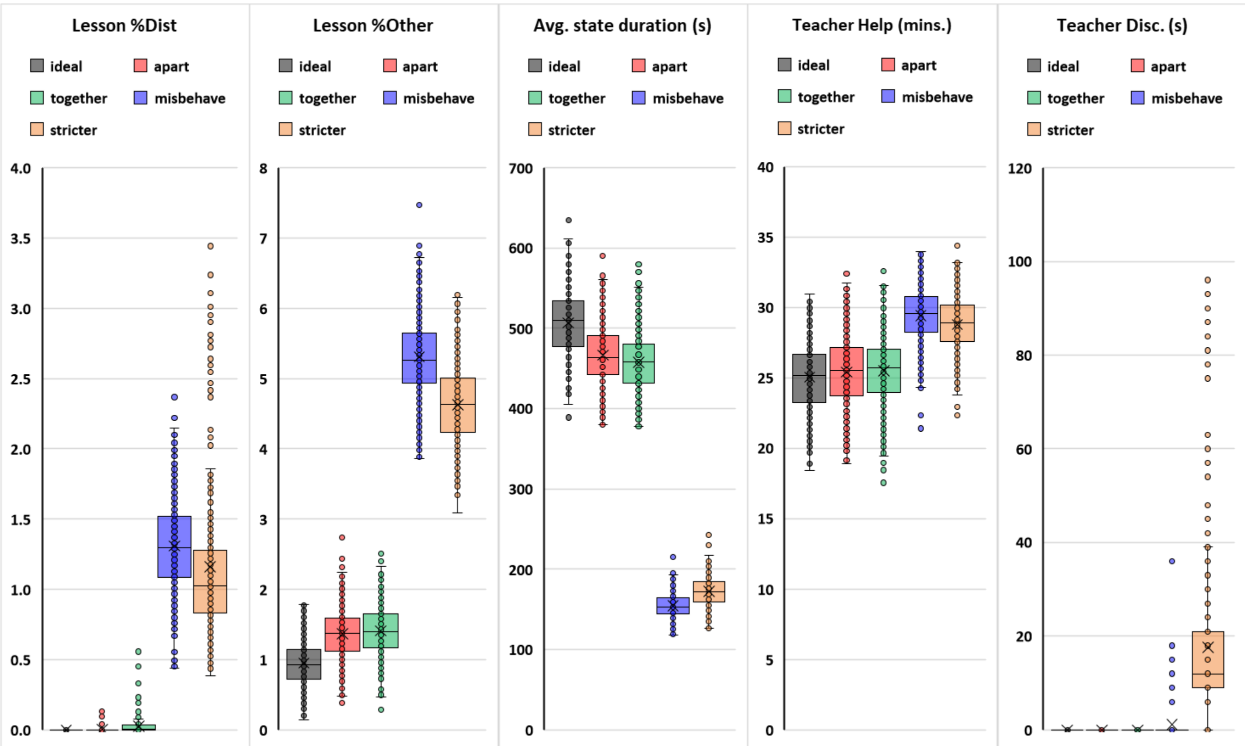

Figure 8-29 The metrics %Dist, %Other, ASD, TH and TD in the five scenarios

Below are key observations on the results and the expectations:

- Two students being disruptive had little overall effect on %Dist, with two students together (scenario 3) having more impact than two apart (scenario 2).
- In all scenarios, the amount of passive disengagement (%Other) was still more than the amount of disruptive behaviour (%Dist).
- There was an enormous difference in all metrics between two students disrupting and a whole class of disruptive students (scenario 4).
- When the teacher was stricter, having a lower tolerance of misbehaviour TMT=20 (scenario 5), TD increased markedly and all misbehaviour decreased.
- ASD was extraordinarily high with ideal students, dropping to empirically realistic levels with an entire class of disruptors (scenarios 4 and 5). This was considered realistic because ASD is higher when students behave as required and lower when students misbehave (where there is a lot of state changing).
- When all the students were disruptive students (scenarios 4 and 5), the teacher spent slightly more time helping individuals.

This all seemed reasonable except %Other was always greater than %Dist even though in the student profiles the ratios of the probabilities of other (passive disengagement) behaviour to disruptive behaviour are 2:3 for the disruptive students and 1:1 for ideal students. Experiment 4b had produced similar results to Experiment 4a. This indicates that student seating density was not a significant factor in these outcomes.

8.4.3 Conclusions to research question 4

The first of these two experiments, Experiment 4a, followed up on what was observed in Experiment 3 and explored why the locations of disruptors might or might not have much impact on the output metrics. It investigated the consequences of replacing one or more ideal students with a disruptive student. Seven scenarios were compared. Relative to the ideal class, adding disruptive students always increased %Dist, particularly when there were three disruptive students. %Other (passive disengagement) also increased relative to the ideal class. However, an unexpected result was encountered: disruption decreased from scenario 6 (three separated disruptors) to scenario 7 (three disruptors seated together). It had been expected that putting all the disruptors together would increase overall disruption because a positive response to a chatting request would be more likely. After inspecting the details of each student, it appeared that seating the disruptive students closer together led to their disruptive behaviour being converted to other disengaged behaviour. Several mechanisms were proposed to explain this. For example, the model has rules that make students less inclined to misbehave, especially disruptively, after they have been disciplined. The rules also make the teacher stricter with disruptive behaviour than with passive disengagement. In scenario 7 (with all three disruptive students in close proximity), the student in seat 9 was disciplined more so this probably accounted for their decrease in chatting and increase in passive disengagement. Also, the proximity of the teacher is very significant: the model has rules that decrease the likelihood of misbehaviour the closer the teacher is. When the teacher was helping any student near the middle rear of the class, that affected all the students nearby, including the disruptive ones, reducing the likelihood of their chatting (or, to a lesser degree, passively disengaging). Having the teacher helping the student next to you is not only a disincentive for you to misbehave, it also removes one student from the list of those you could chat with. The model rules explicitly prevent a student interacting with a student receiving help (from the teacher or TA). Class layout plays a role in the amount of disruption: if everyone is quite spread-out then the spread of social chatting is impeded, although students can always make a comment briefly out loud to no-one in particular. These all seemed plausible and realistic explanations.

However, in all the simulation scenarios in Experiment 4a, %Other was more prevalent than %Dist. This was unexpected because in the student profiles the probabilities of other (passive disengagement) and disruptive behaviour were respectively 8% and 12% for the disruptive students and 5% and 5% for ideal students. This result may have been due to the factors discussed above but warranted further investigation in a second experiment.

In Experiment 4b, the objective was to determine whether the amount of disruptive behaviour compared to other, passive disengagement depended on the 'density' of the student seating. Experiment 4a had only 10 students and they were quite spread out in the classroom. Now the classroom was filled with 19 students. Five scenarios were explored:

- the ideal class – no disruptive students
- two disruptive students separated
- two disruptive students together
- all students disruptive
- all students disruptive but the teacher stricter

In all scenarios, the amount of passive disengagement (%Other) was still more than the amount of disruptive behaviour (%Dist). When all the students were disruptive students, the teacher spent slightly more time helping individuals. When the teacher was stricter, TD increased markedly and all misbehaviour decreased.

The fact that, in both Experiment 4a (students spread out) and Experiment 4b (students densely packed), %Other was greater than %Dist indicates that student seating density was not a significant factor in these outcomes. The experiments suggested several mechanisms that could contribute to %Other being greater than %Dist, but did not yield a comprehensive explanation linking the student base state probabilities with the outcomes - further experimentation is needed.

Nevertheless, the conclusion to research question 4 is that the simulation model can provide plausible results in experiments with a class of completely artificial students.

8.5 Chapter summary

The results of the four experiments are summarised below.

<u>Experiment 1</u> (Section 8.1) The CLSM provided comprehensive, plausible explanations of the effects on lesson behaviours and lesson outcomes of alterations to the teacher's inclination to offer one-to-one support and to take disciplinary action. The results highlighted that if teacher wants to increase the time spent providing individual help, then, to maintain the same level of whole-class time on-task, the teacher needs to become less tolerant of misbehaviour. Alternatively, a teacher could manage student behaviour without additional disciplining by providing less individual support and being seen to observe more. The analysis also highlighted that if the teacher were to focus their assistance on those students who tend to misbehave or looked as though they were about to disengage then this may well increase productivity – by converting non-productive time to productive time. Even just going over to a student and offering assistance would reduce misbehaviour by and around that student. But, as noted, that strategy would be unfair on well-behaved students who would also benefit from teacher help.

<u>Experiment 2</u> (Section 8.2) This experiment indicated that the presence or absence of a TA who is not involved in student disciplining does not directly influence productivity when students are working independently, and provided plausible explanations for this conclusion. What mattered

most was the change in behaviour of the teacher and students – whether the teacher does more or less helping and whether students ask the teacher for help more (as was discussed in Experiment 1). Similar to the results of Experiment 1 concerning teacher help (section 8.1.3), the CLSM highlighted that if the TA helps mainly students who would have been productive anyway, then there will be little effect on Independent %Productivity. But if the TA targets those students who tended to misbehave or looked as though they were about to disengage then this may well increase productivity, but that would be unfair to well-behaved students who would also benefit from TA support.

Experiment 3 (Section 8.3) In this experiment different student seating arrangements – intended to have an effect on lesson disruption – were simulated and the results were compared to what an experienced teacher (the author) expected to happen. In more than half the cases the teacher's expectations and simulation outcomes did not match, raising questions about the reliability of both. Teachers in the case study reported that often their seating reorganizations failed as they merely shifted the location of the disruptive behaviour. Perhaps the consequences of a seating rearrangement are just too difficult to pin down, for teachers and simulation models. The CLSM did however provide realistic explanations of the dynamics, probably as plausible as those of a teacher. This experiment demonstrated a potential benefit of lesson simulations as outlined in section 1.1: with a simulation tool one can try out scenarios without disadvantaging anyone and without wasting resources. If a seating rearrangement simulation looked like it would work the teacher could actually try it, for several lessons to build up a picture of the results (i.e., not just one lesson, just as one would not run only one simulation replication). One cannot judge the reliability of either the CLSM or a teacher without empirical testing.

Experiment 4 (Section 8.4) Two experiments on artificial classes were conducted. Experiment 4a investigated whether the locations of disruptors might or might not have much impact on overall lesson metrics. Several scenarios were compared and the results were considered realistic and explanations plausible. However, in one scenario disruption decreased instead of increasing as anticipated. It was postulated that seating the disruptive students closer together had led to their disruptive behaviour being converted to other disengaged behaviour. Several factors and mechanisms were proposed to explain this outcome. The purpose of Experiment 4b was to investigate one of these proposals, that the amount of disruptive behaviour compared to other (passive disengagement) might depend on the 'density' of student seating. However, in all the experimental scenarios, the amount of passive disengagement (%Other) was still more than the amount of disruptive behaviour (%Dist). The fact that, in both Experiment 4a (students spread out) and Experiment 4b (students densely packed), %Other was greater than %Dist indicates that student seating density was not a significant factor in these outcomes. Both experiments suggested several mechanisms that could contribute to %Other being greater than %Dist, but did not yield a comprehensive explanation linking the student base state probabilities with the outcomes - further experimentation is needed. Nevertheless, the conclusion to research question 4 is that the simulation model can provide plausible results in experiments with a class of completely artificial students.

These four research experiments demonstrated that the CLSM modelled some important aspects of real lessons, providing plausible causal explanations for student, teacher and TA interactions, and how these resulted in the flow and outcomes of a lesson (including in completely artificial lessons).

9. Discussion, further research and conclusions

This final chapter starts with a summary of the thesis. This is followed by a discussion of some limitations in the model and its development, with, in most cases, suggestions for addressing the issues in further research. The discussion is extended to consider issues with using the CLSM, and ABMS in general, to explain real-world behaviours and support decision-making. Despite these concerns, some topics for further classroom lesson research using the CLSM are proposed. The thesis ends with the overall conclusions of the research.

9.1 Thesis summary

The primary research objective was to answer the question:

How and to what extent can an agent-based model adequately represent the behaviours of, and interactions between, students, teacher and teaching assistant in classroom lessons at a UK secondary school?

To answer this, four experiments were designed. The answer would depend on how well the simulation model enabled the following four questions to be investigated (the first three of which are of on-going interest to educators).

1. Does the simulation model provide realistic explanations of the effects on lesson behaviours and lesson outcomes of alterations to the teacher's inclination to offer one-to-one support and to take disciplinary action?
2. Does the simulation model provide realistic explanations of the effects on overall student productivity of providing or withdrawing a TA who gives individual support to any student?
3. Does the simulation model provide realistic explanations of the effects on lesson behaviours and lesson outcomes of different student seating arrangements?
4. Does the simulation model provide realistic results in experiments with a class of artificial students?

The motivation behind the research (explained in section 1.1) was the belief that a classroom lessons simulation tool could have multiple potential benefits for teachers, head-teachers and other educators, such as:

- providing teachers with a suitable tool to explore alternative teaching strategies;
- testing a proposed or existing theory of classroom dynamics;

- providing insights into the dynamics of entire lessons;
- identifying the relative contribution of factors that are known to affect student behaviour.

Besides presenting the above points, Chapter 1 explained how the research project followed a typical empirically-based ABMS development methodology and included an investigative case study. Examples of other ABMS applications were given and existing classroom-related models and their limitations were reviewed. The first two stages of the CLSM development were described in Chapter 2. Following a comprehensive conceptualization of the classroom lesson system (Stage 1), a detailed, general conceptual model of behaviour in classroom lessons was formulated (Stage 2). Agent decision-making was modelled using 'production rules', with an emphasis on realism. The implementation of the conceptual model as a simulation model (Stage 3) was explained in Chapter 3. The agents, their activities and their decision-making were described and the central role of lesson plans in determining behaviour was explained. While the teacher follows a simple algorithm (follow the lesson plan handling any interruptions along the way), and the TA follows a few simple rules, the students follow more complicated logic. Some of the intricacies of agent decision-making and interactions were described and the model parameters that adjust behaviour were introduced. In Chapter 4 several key topics common to both the model calibration and model validation stages in ABMS development were covered. The lesson metrics used to compare simulated and empirical lessons were defined. The empirical data were analysed and the metric acceptability ranges defined. These ranges are one of the criteria used to decide whether simulation outputs are realistic. The main topic in this chapter was that, although a generic CLSM had been developed, the model needed instantiating for each empirical lesson. This was the primary use of the lesson parameters: to 'calibrate' the model to a specific lesson. The main reason for this is that each student's profile was based on their aggregated behaviour over all the empirical lessons, and similarly the teacher and TA had the 'average' teacher and TA tendencies over all lessons. The lesson parameters adjusted the class's, teacher's and TA's behaviours so that they became more closely matched to the specific lesson in question. Chapter 5 explained how the seven selected lesson models were calibrated (Stage 4). This process, also known as parameter estimation, comprised two steps. In the first step, 'categorical calibration' was applied, discarding any parameter set that produced any unrealistic result, that is a result that fell outside the acceptability range. In the second step, the remaining parameters were explored in finer detail over more replications. Not only did the results need to be realistic, they needed to match (as closely as possible) the empirical lesson's values. Of the many candidate parameter sets, the one that was least worst on all the metrics was chosen. The purpose of Stage 5 (model validation, Chapter 6) in the ABMS development was to extend confidence in the lesson models by having experts confirm the plausibility of the conceptual model and participate in 'face validity' tests. The lesson simulation results, already judged realistic at the overall lesson level (macro-level), were checked to see if they were close enough to the specific empirical lesson results to be considered a plausible simulation model of that specific empirical lesson. Plausibility was investigated at the micro-level of agent interactions and the meso-level of overall behaviour for individual students. Only four of the seven calibrated lessons were judged to be sufficiently like their empirical lessons that they could be used in the research experiments. In Stage 6 some basic sensitivity analyses were conducted (Chapter 7). These investigated how changes in the model parameters and in students' base state probabilities for their inclination to chat (as a relevant test-case) affected the simulation output metrics. The lesson simulation models were found to respond appropriately to both types of changes. Finally, in Chapter 8, the four research experiments were described and the results analysed (Stage 7). The results demonstrated that the CLSM modelled some important aspects of real lessons, providing plausible causal explanations for student, teacher and TA interactions, and how these resulted in the flow and outcomes of a lesson (including in completely artificial lessons).

It is also pertinent to summarize here one way in which the CLSM could be used as a lesson analytics tool. The CLSM is generic - it does not need alteration, just different parameter values for the specific lesson to be investigated. For example, suppose a teacher or headteacher wanted to investigate whether a particular class would benefit from having a TA (perhaps because there are one or two students who need a substantial amount of one-to-one support). They could ask for simulations to be run but, if there were no empirical lesson data for this class with that teacher for that subject in the specific classroom, empirical data would be needed. The CLSM needs data to initialize it: a classroom layout and student seating layout, a class of students each with their historical behaviour data, a lesson plan and a validated set of lesson parameters. So, data would be collected over three or four lessons and profiles for the students and teacher in that subject and classroom with a range of lesson plans (e.g., different balances of whole-class teaching, solo, paired and group work) would be constructed. These data would be used to calibrate and validate that lesson model. The data would be used to initialize the two simulation scenarios – one with a TA and one without – which would then be replicated a few hundred times to obtain an overall picture of the likely outcomes. Analyses would look for explanations of the results and consider how these might contribute to decision-making.

9.2 Further development to reduce model limitations

This section proposes further research that could remove or reduce some of the limitations of the CLSM. Before criticising the CLSM on educational grounds, it should be remembered that it is a model of behaviour, not learning (despite time on-task being promoted as a proxy for learning – see section 2.3), and does not incorporate factors such as the influence of student learning success, enjoyment or motivation, etc., focussing instead on student misbehaviour and teacher discipline.

9.2.1 Quantifying disruption

It was assumed that disruption to a lesson could be measured by the amount of time students spent in disruptive states (states 1, 2 and 3 – see section 3.1). The problem is it is not just the total duration of disruptive states that matter – it is their frequency, intensity and nature too. For example, the disturbance could be an interaction between two students or between a student and the teacher; it could be light-hearted or aggressive; it could affect one student or the whole class. As has been explained, teachers and students report that even a brief, few seconds, unpleasant teacher-student interaction can spoil the rest of a lesson. There is also a question about how much disruption is significant? Similar questions can be raised about passive disengagement. Although there is extensive research on the effects of different behaviours (Marzano, Marzano and Pickering, 2003; Petty, 2006; Hattie, 2012) more research would be needed to develop a comprehensive disruptivity index – which could then be a new output metric for the model.

9.2.2 The initiator of an interaction

A fundamental simplification that has affected the entire model was that agent interactions were considered directionless. This was also discussed in section 1.5.2 and is not a simple issue to rectify. During the case study, it was thought that it would be useful to have a state for students raising their hand and that this would indicate who initiated an interaction with the teacher. What was observed though, was that there were too many alternatives to assume initiation of the next event. For example, after a student raised their hand:

- the teacher could indicate to them to put it down (so they don't have their hand up for 5 or 10 minutes waiting) – and then the teacher could forget and go and help someone else;
- the student could just carry on working (perhaps skipping over that particular problem, or realizing what to do);
- the student could just sit, waiting or doing nothing or just fiddling with something;
- the student could attempt to chat to another student or respond to a chat request;
- the student could ask another student for help or respond to an offer to help.

Also, this state would not have solved the problem of identifying the initiator of a student-student interaction, as in the last two items in the list above. Without additional technology, determining who initiated an interaction would require a more intrusive observation regime (e.g., asking a person what happened), even with video recording and/or multiple observers. It might, however, be possible to use automated classroom data collection techniques (Raca and Dillenbourg, 2013, 2014; Raca, Tormey and Dillenbourg, 2013) such as teacher location and movement tracking plus teacher and student gaze tracking. If one could capture this information and collect enough data per student, then it might be possible to split existing states in two as follows:

- teacher offered assistance and teacher responded to request for assistance;
- TA offered assistance and TA responded to request for assistance;
- student asked for assistance and students responded to offer of assistance,

and create new states for the following:

- student initiated a request to chat and student responded to a request to chat;
- student initiated a request to work together and student responded to a request to work together.

This data would allow the generation of historical frequencies for these student states (exactly as for the current states). Also, if it were possible to determine in which direction a teacher is looking (Bidwell and Fuchs, 2011; Raca and Dillenbourg, 2013, 2014), one could use this in the student decision-making logic: Is the teacher looking at me?

Of course, this would all require a complete revision of the conceptual model, with new and updated rules and model constants.

9.2.3 The influence of student friendships

The influence of peers, friends and friendship networks on academic achievement is particularly important (Halliday and Kwak, 2012; Blansky et al., 2013). Hanushek et al. (2003) and Yeung and Nguyen-Hoang (2016) found that a student's performance is significantly influenced by the achievement of their peers. Hirschy and Wilson (2002) highlighted the significant influence that peer relationships have on classroom norms and how students interact to either encourage or hinder their fellow students. It is because of such research (and observations during the case study) that peer influence was included in the CLSM (as described in section 3.4). But friendship relations may be even more important. Blansky *et al*. (2013) investigated the relationships between students' academic results and their reported network of friendships. It was found that students tended to move up the class academic achievement ranking if their friends were scoring higher grades than themselves. Additionally, students would tend to move down the ranking if their friends scored lower grades than them. The authors suggested that this social network phenomenon might even provide a quick predictor of students' future academic performance. Other researchers developed an ABS that used sociograms (maps of the social relationships between students) as a means of predicting students' academic performance (García-Magariño and Plaza, 2015; García-Magariño et al., 2016, 2017).

During the case study, a simple, three-category 'Friendships' questionnaire was piloted. The intention was to establish each student's friends, non-friends and indifferent relationships and use this information in the behavioural rules. However, it soon became apparent that more would be required. It was observed and reported, by teachers and students, that relationships were volatile. It was not uncommon for friendships to reverse over a holiday, a school trip or a weekend, or even a lunch break. Sometimes best friends, who always sit together, end up fighting and the result is feuding parents who instruct the school that on no account should the two sit near each other. Then, a week later, the students are best friends again and want to sit together. In other words, although Blansky *et al*. (2013) were able to find a relationship between academic progress over a year and the students' friendship networks established through one survey, it is doubtful that collecting data on friendships once at the start of a study is ever going to be adequate for lesson simulations: it may be necessary to collect this data every observed lesson – which is not practical (the teachers in the case study said they would not countenance such an exercise).

To sum up: it seems desirable to incorporate student relationships into the model, but it is unclear how frequently that data should be refreshed or how that could be accomplished efficiently to the teachers' and students' satisfaction. This is a topic for future research.

9.2.4 Enhanced decision-making logic

Improvements could be made to the behavioural rules by enhancing the agents' decision-making logic. For example, as explained in point 34 of section 2.6, currently when an agent makes a decision about what to do next, only the current state of the other agents is considered, not the values of any of their attributes. But when considering who to propose chatting to, a student could consider the chattiness of their neighbours and choose a nearby student who currently has a high state 3 (chat) attribute score. This would be an example of an agent considering the current inclinations of other agents. Decision-making logic could also be extended by taking into account when another student was last told off: e.g., if X was told off less than 2 minutes ago then I won't try chatting to him. This would be an example of an agent considering the history of other agents.

9.2.5 Understanding model stochasticity

The model generates a wide range of results for each scenario, primarily because of the stochasticity introduced to make the simulated lessons realistic. The distributions of the output metrics (shown in all the boxplots) are the result of the different pseudo-random number sequences used in the replications. Some randomness is needed to avoid robotic behaviour and cycles, but the question is how much. As mentioned in section 7.1, investigations showed that one small variation in the pseudo-random number sequence can lead to completely different lesson trajectories, all of which seemed to be plausible. Sensitivity analyses (Chapter 7) could be extended to try to establish the minimal amount of stochasticity necessary to produce lesson simulations that passed validation tests. This would require taking each point in the model where stochasticity is inserted and systematically reducing the extent of the variability introduced. Points in the simulation model where randomness is injected have been listed in the relevant sections in Appendix A. Sensitivity analyses should also be extended to the threshold constants built into the CLSM (also listed in the relevant subsections of Appendix A). It is expected that additional empirical data that provided distributions for these constants would help more precisely constrain them.

9.2.6 Improving model validation

Conceptual model validation (evaluation of the model rules and assumptions) and the face validity tests for simulation model validation involved only three other teachers besides the author. This could be improved, both for this research and future research simply by involving more teachers. The procedures for both should be made more formal and thorough. For example, the procedures should be redesigned to rule out the possibilities that evaluators did not understand what they were being shown (e.g., the visual representation of lessons and the presentation of summary statics was inadequate and/or confusing) or were simply being agreeable or were not participating properly (e.g., choosing at random). Also, the tests depend crucially on how observant the teacher was, their attention to detail. It was also perhaps unreasonable to ask teachers to judge a lesson animation without knowing the lesson plan.

There is also a more fundamental issue concerning face validity in general: having a fallible human assess just a few simulation results is not adequate when hundreds or thousands of replications are needed to obtain a decent overview of the range of outcomes (Edmonds *et al.*, 2019). One of the reasons that an ABMS approach was followed was that it can potentially provide agent-level suggestions for causality in agent interactions and their consequences (mentioned in section 1.3). This is obviously attractive but it is not easy for a human to follow the multiple sequences of interactions that happen in a lesson. More face validity testing could be undertaken, or alternatively testing of replications could be automated – although this just displaces human activity onto checking the consistency and completeness of the logic in the automation.

A more radical action would be to redevelop the model participatively (Tako and Kotiadis, 2015), using focus groups of teachers and students to postulate the behavioural rules they seem to follow and explain why they did what they were observed to do (e.g., discipline or not discipline, chat or just sit). This would be expected to improve the generality and realism of the conceptual model.

Despite these concerns, sometimes people (teachers and non-teachers) have commented that the model is not telling us anything new, that 'we know that – it's obvious'. They point out that, for example, teachers know that when one is helping a student other students may misbehave behind one's back, or that one can never be sure that a seating rearrangement will improve the next lesson. Non-teachers speak from their experiences as students; teachers speak from their experiences as teachers and as students. This is a well-known phenomenon in ABS research: others assert that the results are trivial or obvious (Eberlen, Scholz and Gagliolo, 2017). To quote from these authors:

> This is similar to hindsight bias in classic experiments. There are two scenarios: First, by creating an ABM, we have realized that the results are actually trivial – in this case, the ABM has fulfilled its use as a thinking tool. Second, the possibly unexpected result seems suddenly more plausible than the initially predicted outcome. Differing from a real-world experiment, here, we have the support of the implemented ABM that the surprising result is actually the outcome of the dynamics programmed into the model.
>
> (Eberlen, Scholz and Gagliolo, 2017, p. 157)

However, although intended to challenge the value of the research, these comments are actually further confirmation that the model does reflect people's experience in lessons and does reflect the mental model that people have about what happens in classroom lessons. The comments increase confidence that the model (so far evaluated only by the author and the three teachers involved in model validation) is a generally satisfactory representation of lesson behaviours. It is also possibly an exaggeration to say that everyone knows: a newly qualified teacher has a lot to learn and could gain insights through lesson simulation experiments.

9.3 Issues with the CLSM and with ABMS in general

Many authors have expressed concerns about the use of ABMS, but papers published ten or more years ago are often out-of-date (and often the criticisms applied to all modelling and simulation approaches not just ABMS) – see Manzo and Matthews (2014) for this discussion. On the other hand, many researchers have recognized the potential benefits of ABMS. For example, in sociology there were concerns about the analysis of mechanisms that generated higher level outcomes from the actions of low-level entities (known as the *transformation problem*, or *micro–macro transition*) (Manzo and Matthews, 2014). ABMS have since demonstrated that one can reliably apply a generative approach to determine the potential macroscopic implications of multiple local micro-level interactions (Epstein, 2008; Manzo and Matthews, 2014).

Nevertheless, there are still many valid concerns that apply to this research project. The following list is based on points raised by, amongst others, Manzo and Matthews (2014), Gómez-Cruz, Loaiza Saa and Ortega Hurtado (2017), Eberlen, Scholz and Gagliolo (2017) and Nuno, Nuno and Agostinho (2017).

9.3.1 Agent attributes, empirical data and model parameters

In ABM there is often a concern that a model of a very complex system may have a large number of free [20] parameters. In section 2.1, classroom lessons were categorised as complex systems, and in the case study they were seen to be extremely diverse and that each one was unique. It was noted that:

- students reported they felt different and interacted differently according to their seat position and whether they sat with friends;
- behaviour in different subjects differed, e.g., science lab work vs art classes vs historical role play vs mathematics investigations;
- students had different relationships with different teachers;
- different teachers had different attitudes towards and tolerances for chatting, asking questions, behaviour in tests, the amount of help given, etc.;
- teachers typically planned a variety of activities, some of which a student might enjoy more than others or be more competent and successful in - for example whole class teaching, working alone, working in pairs, group work, competitive or cooperative activities, exercises or investigations, more quiet work or more discussion.

As declared at the outset, the intention was to construct a model that was as realistic as possible, capable of representing a wide range of behaviours, particularly those observed during the case study. For a classroom lessons behavioural simulation model to be widely applicable it would need to accommodate the changes in behaviour caused by:

- different classes and age groups with a wide range of student characteristics (such as academic level, academic potential, special educational needs, health issues);
- different teachers with different approaches to behaviour management and one-to-one support;
- different subjects - since subject lessons can be taught very differently and students find different subjects more or less attractive and have different proficiencies in them;
- different rooms with different layouts and different student seating arrangements;

[20] Defined by Edmonds *et al*. (2019) as 'not determinable using measurement from the target of modelling, even in principle'.

- different lesson plans (e.g., different balances of whole class teaching, working alone, working in pairs, group work, competitive or cooperative activities, exercises or investigations, more quiet work or more discussion).

In other words, the simulation algorithm would need the above inputs in order to generate reliable results. If any of these factors were omitted the simulation outputs would likely be less reliable. When a stochastic simulation model is given identical inputs, different replications will generate different outputs because of the different pseudo-random number sequences. The distribution of results is more likely to be reliable if all the factors above are considered. But in order to take a factor into account, empirical data is needed for that factor. Because teachers, TAs and students behave so differently in the different lessons, if one wants to model a particular class in lessons with a particular teacher teaching a particular subject in a particular classroom room with a particular seating arrangement using a specific lesson plan, then ideally one needs empirical data on such a scenario.

It was explained in section 1.3 that, in an ABMS, agent attribute values are typically used in the decision-making process. This means data are needed to supply those values. To satisfy the needs listed above requires a significant amount of detailed data. For example, suppose a latent construct[21] *strictness* was proposed for teachers. The teacher agent could be assigned a *strictness* attribute that would be used in deciding when and how much to discipline a student. This cannot be a fixed value for that teacher in all lessons. As explained, each teacher was observed to have varying tolerances for misbehaviour with different classes, in different subjects, in different rooms, with different seating arrangements and during different types of lesson activities. There would need to be a *strictness* function that took into account the class, the subject, the classroom, the seating arrangement, the type of lesson activity, as well as the factors currently incorporated in the teacher decision-making logic.

A similar situation applies to students. Suppose a latent construct *compliance* was proposed for students. Student agents could be given a *compliance* attribute to reflect their inclination to adopt the expected state and not disengage. Students were observed to have varying degrees of compliance with different teachers, subjects, rooms, seating arrangements and lesson sections.

[21] A latent construct or variable is used when one cannot observe a property directly. Instead it is inferred from an index combining measurable variables. For example, there is no ‘thermometer’ for measuring teacher strictness so it would need to be inferred from other variables.

There would not be a fixed value for a student's *compliance* attribute that suited all lessons. As explained, each student was observed to behave differently with different teachers, in different subjects, in different rooms, with different seating arrangements and during different types of lesson activities. There would need to be a *compliance* function that took into account the teacher, the subject, the classroom, the seating arrangement, the type of lesson activity, as well as the factors currently incorporated in the student decision-making logic.

Despite having empirical data from 21 1-hour lessons involving 6 subjects, 6 teachers and 52 students, these were insufficient to support the use of a *strictness* or *compliance* function, or other comprehensive functions that would need such detailed data to determine values for proposed attributes. Creating latent constructs as agent attributes on a sound theoretical and empirical basis and assigning empirically-derived values to them is an ideal approach. But it still requires sufficient data and the only data available were the 21 lesson observations recording the sequences of states for each agent. Whatever method (e.g., Principal Component Analysis) employed to derive latent constructs, it would be using that state data. As explained in section 2.4, it was decided that the student states would form attributes directly. But the same issue applies to student state attributes. Consider for example the state 3 (chatting) attribute. Students were observed to chat in different lessons to varying degrees, apparently varying with different teachers, different subjects, different rooms, different seating arrangements and during different types of lesson activities. There was still not enough data to cover each of the factors listed above. Hence a different approach was adopted: the empirical data for all lessons, teachers and classes were combined to derive distributions and average values for each type of lesson activity (lesson plan section – see section 3.2). But, as discussed above, using an average value would not assist reliable simulation – for greater confidence it would need adjustment to suit the specific scenario being simulated. Hence model parameters were introduced to adjust these distributions and averages to the particular class with a particular teacher teaching a particular subject in a particular classroom room with a particular seating arrangement.

In building a complex model following the 'keep it descriptive stupid' (KIDS) approach (Edmonds and Moss, 2005), one expects to have many parameters as there will be many aspects about the system that are unknown, possibly unmeasurable. With empirical data it is often possible to fix a parameter at a specific value or constrain it to an interval, a range of values. In this research, model parameters that have been fixed (based on empirical data or observations) have been termed model constants as they are constant across all the lesson models. (They are declared in the simulation code in a separate 'GlobalConstants' module.) Being derived from the empirical data they are thus specific to the case study people, i.e., certain people in one specific school over one particular period, and therefore may not be appropriate for another school.

There are 11 free parameters in the CLSM. (Technically, the random number generator seed is another parameter, as is the number of simulation replications. One can get quite different results from running more replications – as was seen during model calibration when after 500 replications a parameter set could be discarded despite being the best after 80.) With sufficiently detailed empirical data, the two teacher-lesson parameters, TSOL and TMT, and the TA-lesson parameter TASOL, could be replaced by agent attributes: the teacher's inclination to offer one-to-one support and tolerance of misbehaviour, and the TA's inclination to offer one-to-one support. CSE, currently considered a lesson attribute, could become redundant too. The seven class-lesson (or student-lesson) parameters currently adjust the behaviour of the entire class of students in that lesson. With more detailed data they too could probably be replaced with student-specific attributes (assuming the initiator of interactions could be recorded, where appropriate):

- SSRW by inclination to ask for assistance;
- SIW by inclination to initiate an interaction;
- IRW by inclination to respond to an interaction;
- Peer-Weight by the extent to which a student copies their peers' behaviour;

With sufficiently detailed data – which includes the relative state frequencies for each student in all the scenarios listed at the start of this subsection - the remaining three parameters (RLESW, RLDW and RLOW), which also adjust the entire class of students, could also be eliminated. If they were, then there would be no free parameters, only model constants derived from the empirical data. No calibration would be needed, or possible.

Of course, this would all require a complete revision of the conceptual model, with new and updated rules. However, parameters also act as scaling factors for another type of unknown: the relative importance of parameters and attributes in decision rules. In the CLSM, besides adjusting average behaviour to the specific lesson (teacher, class, subject, classroom, lesson activity) (explained in section 4.1), the parameters enable the relative importance of other parameters and agent attributes to be taken into account in the behavioural rules. Even with extremely comprehensive empirical data, there would still be a need for a mechanism to adjust the relative importance of factors. For example, when a student is considering whether to chat (state 3) or not, the following influences/factors need to be weighed up and combined: the student's historical state 3 relative frequency, the states of other students, the student's IRW and the current disciplinary environment (summarized by the *misbehaviour reduction factor*: a factor that takes into account the time since the student was last disciplined and the teacher's proximity (explained in section A.3.4). Hence the unknown relative weightings of factors necessitates the introduction of some free parameters.

9.3.2 The empirical data, model calibration and generality

The empirical data used to calibrate the model were specific to certain people in one specific school over one particular period. These data are not representative of all UK students, teachers and schools. Further, the model parameters were tuned to specific lessons. Currently the CLSM's generality is uncertain, whether it could be applied at other schools and how useful it would be. The model may turn out to be suitable in other settings, but only additional case studies could determine whether the rules and assumptions were generally applicable. Some degree of generality (Edmonds and Gershenson, 2015) is essential otherwise a simulation model is of limited practical value. Fortunately, all the experiments – particularly experiment 4 – indicated that the CLSM has plausible generality beyond the empirical lessons of the case study.

There is also a concern about the arbitrary ±25% tolerance that was added to the empirical metric ranges used to decide whether simulation outputs were realistic. If the metric acceptability criteria were too tolerant this would lead to a high proportion of parameter sets passing the calibration tests (as was experienced) even though their outputs were actually unrealistic. This would help explain why some lesson models would subsequently fail validation tests.

Furthermore, when comparing simulation outputs to empirical data, arbitrary (but conventional) criteria were applied (such as using a 95% confidence interval and requiring that it lie within a user-specified ±7.5% tolerance interval – see section 4.4). How similar do simulation and empirical results need to be for the model to be considered valid? Chattoe-Brown (2014) expressed the opinion that this could not be decided a priori and suggested that this would be 'best addressed by a sequence of ABM, each justified by improving on the last' (Chattoe-Brown, 2014, p. 15). As mentioned in section 4.2, this is how the CLSM was developed, by successive enhancements intended to increase simulation realism. But if there were a large difference in empirical and simulated lesson outcomes, would that mean the model was inaccurate or could it be that the observed empirical lesson was atypical? Only further case studies could answer these questions satisfactorily.

Also, the parameter sets found during calibration are almost certainly only local optima. Resource constraints (especially time) meant that a comprehensive search of the parameter space was infeasible (as explained in Chapter 5). This means that some lesson models could probably be improved and those that did not pass the validation procedures might actually pass with a better parameter set. It would also be worthwhile improving parameter estimation to search for one parameter set that would make all lesson models calibratable and pass the validation checks. In an initial investigation into this possibility, each of the validated parameter sets for the four validated lessons was applied to the other three validated lessons. Each of these 12 combinations was subjected to the validation procedures (apart from face validity tests) described in section 6.2. The results were that none of the parameter sets suited other lessons: all versions of the lesson models were rejected as insufficiently close to the empirical lessons. This is some additional justification for instantiating the CLSM per lesson, but it does not mean there definitely is no one parameter set capable of being validated for all the lessons.

9.3.3 Variability of simulation results

As has been explained (e.g., in section 9.2.5), stochasticity was built into the model to emulate some of the apparently random variability in activities/choices in lessons (known as aleatoric variability[22]). This means that each simulation run, using a different PRNG seed, generates different outcomes. Modellers merely combine the results to form a distribution for each output metric. Manzo and Matthews (2014) reported that some researchers expressed concerns with the uncertainty in the simulation outcomes, with some advising that the focus should be on comparing distributions not comparing a single measure of central tendency (as many were accustomed). This is exactly the approach taken in this research. By taking all the empirical lesson and student data, for all teachers, subjects and classrooms, an estimate of the aleatoric variability in lesson outputs was obtained. One of the checks during calibration and then validation was that the simulation distributions were realistic, i.e., fell withing the empirically-

[22] Aleatory uncertainty or variability is the intrinsic uncertainty due to the random variability that one assumes applies to human behaviour (assuming that people's behaviour is not ultimately deterministic, that people do express free will which is not predictable).

observed ranges (±25%) (see section 4.3) and adequately matched the empirical lesson they were representing, i.e., the metric distributions were not only realistic, but they also comfortably encompassed each empirical metric value.

The variability in simulation outcomes is not a problem - it is exactly what is required. As Thiele (2014, p. 4.5) expressed it: 'variation in model output represents variation in reality'. Teachers do not expect the 'same' lesson to turn out the same each time. Teachers during the case study explained that, in their experience, one could never be certain how a lesson would turn out. The same activities worked fine with one class, but not with another. Randomness was essential for realism. A range of possible outcomes was expected and desired. What was required was that the distribution of outcomes met teachers expectations. The reasons at the micro-level were also required to be plausible, e.g., teacher helped Y instead of X, or student X chatted to Z instead of Y – the sort of alternatives discussed in Experiment 1, in connection with Figure 8-9 in section 8.1.2.3.

9.3.4 Using sensitivity analysis during model calibration

Many researchers conduct sensitivity analyses to investigate the behaviour of a validated simulation model (Robinson, 2013), as was done in this research (see Chapter 7). Sensitivity analysis is seen as a check on the stability of a validated model in response to changes in parameter values and input data, and also as an exercise in identifying which parameters have the greatest impact on the output metrics (Robinson, 2013). Often, this type of analysis is considered part of the simulation validation process, as behavioural validation: checking that the simulation is behaving in a reasonable manner. However, there is an argument that sensitivity analysis should play a role during parameter estimation (calibration), before simulation model validation. For example, consider one model parameter and suppose a search of the parameter space yielded two ranges that resulted in optimal simulation output: (1.9,2.1) and (7.2,9.3) – see Figure 9-1. What parameter value is best: the more sensitive 2 or the less sensitive 8?

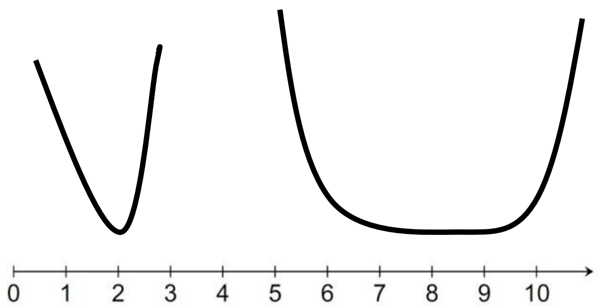

Figure 9-1 Which parameter value is better: 2 or 8?

Where different optimal values provide approximately the same simulation macro-level outputs, it may be sensible to choose a parameter value, or a smaller range of values, precisely because it is more or less sensitive. In the case of a social simulation model with many parameters, all with their own ranges, with multiple optimal solution spaces, which solution sub-space should be preferred? One could argue that the choice depends on how the real world system is expected or known to behave. For example, in a lesson where the students are constantly seeking help, the teacher's own tendency to offer help may be almost irrelevant: any value for the TSOL parameter gives the same range of output metrics. In other situations, a value may be a tipping point, a critical value (perhaps the teacher's TMT value). As has been mentioned before, in classroom lessons a small action by the teacher, disciplining too strongly or too leniently, can lead to dramatically different lesson trajectories and outcomes. The model should be appropriately less or more sensitive depending on the scenario. Ultimately, the choice of parameter values involves understanding the required sensitivity of the simulation model to the range of acceptable values, therefore it may be reasonable that sensitivity analysis be a part of parameter estimation.

9.3.5 The validity of the conceptual and simulation models

As stated in section 6.4, just because a lesson model passed the validation tests does not mean that the model is correct. Manzo and Matthews (2014, p. 446) pointed out that: 'the congruence between simulated and real macroscopic structures is not proof in itself of the realism of the microscopic and relational details designed to generate the macroscopic structures'. They also noted that a modeller should have 'accurate data on the macroscopic structure of interest' and 'sound theoretical and empirical reasons to believe in the microscopic specification of the model' (*ibid*., 448). But, as Gómez-Cruz, Loaiza Saa and Ortega Hurtado (2017, p. 323) noted for 'complex organisational systems' (of which classroom lessons are surely one): 'it is not possible to abstract these assumptions completely or univocally'. They also stated that: 'There are no standardized models that guarantee verification and validation in ABS'. There is currently no standardised procedure for validating (and verifying) an ABMS, in general or in specific contexts (e.g., social simulation or school lessons) (Nuno, Nuno and Agostinho, 2017).

To reduce the chances that the CLSM included incorrect assumptions or incorrect rules, or excluded essential rules, it was developed following an empirically-based, critical realist informed grounded theory approach (mentioned in section 1.5), using observations from a case study. It is also recognised that there could be other models, based on different assumptions, different agent attributes and different behavioural rules, that could generate equally realistic interactions and outcomes at all levels, micro, meso and macro. It is a matter for debate whether we need to be concerned or not that multiple ABMs with different behavioural rules could generate the same macro-level outcomes (Chattoe-Brown, 2014).

As mentioned in section 1.5.1 , it is generally impossible to prove that an ABMS has no errors or bugs (Galán *et al.*, 2009; Norling, Edmonds and Meyer, 2018), so standard software development practices need to be employed to reduce the chances of errors and confirm that a simulation is operating correctly. Although extensive testing took place over two years, there could still be errors in the code that have led to incorrect simulation outputs. Additional testing of the almost 10,000 lines of NetLogo code would be worthwhile.

There are many questions that this project has not addressed, but these are still to be properly addressed for ABMS in general. Even if a simulation model has passed rigorous tests, one cannot say that the model rules are real in the sense that the people actually make decisions using that logic, and there is still uncertainty about transferring simulation results and explanations to the real-world. As Eberlen, Scholz and Gagliolo (2017) wrote:

> Another pitfall of ABMs is to take the implemented procedure based on the theoretical assumptions and results model as the mirror image of the same processes and observations in the real world. While the model can serve as a "proof of concept", it cannot be conclusive evidence by itself.
>
> (Eberlen, Scholz and Gagliolo, 2017, p. 157)

9.3.6 Prediction of 'possible futures'

It might be tempting to use the CLSM for predictive purposes, even though the CLSM was developed as a theoretical explanatory model not a predictive model, for which a different methodology would have been followed (as discussed in section 1.2). Nevertheless, the experiments were subjunctive in nature, exploring ‘possible futures’ in specific ‘what-if’ scenarios (Nuno, Nuno and Agostinho, 2017). Further research could establish whether the CLM is sufficiently accurate and robust to be used for prediction, in certain restricted circumstances. The CLSM could be tested empirically, by making predictions about an intervention (e.g., a seating rearrangement) and comparing the simulation results to what happens when the intervention is actually implemented. Although one could conduct experiments with artificial classes, it would be worthwhile to conduct further case studies and collect before and after data. This type of empirical experiment has many practical challenges though, such as inconsistent student responses due to the novelty of the situation, key students being absent from some lessons, different behaviour choices (e.g., teacher helps student X instead of Y, or a student chats to the person on the left instead of the right), and deviations from the lesson plan. The proposed intervention would need repeated testing.

Whatever its purpose though, a simulation can always stimulate thought and discussion. As Epstein (2008, p. 2) wrote: ‘It is important to note that in the policy sphere … models do not obviate the need for judgment. However, by revealing tradeoffs, uncertainties and sensitivities, models can discipline the dialogue about options and make unavoidable judgments more considered.’

9.4 Further classroom research using the CLSM

There are many topics in classroom research that could benefit from simulation. For example, the CLSM could be used to investigate some of John Hattie’s arguments and research. He has argued (and demonstrated statistically) (Hattie, 2005, 2015) that class size is relatively unimportant regarding student learning. Using the CLSM, it would be possible to create classes of various sizes (with the same proportions of student profiles and the same lesson plans) and investigate the consequences of class size on overall productivity and behaviour. According to Hattie, if the teacher does not change his/her teaching approach then learning is affected. This means that one would expect simulated overall productivity to decline as class size increases. But Hattie’s main point is that a teacher should not try to teach a large class in the same way that they teach a small class: the teacher needs to adapt how they teach, what activities the students should engage in, etc. This advice could be investigated by running simulations using different lesson plans and altering the teacher-lesson parameters for one-to-one support (TSOL) and misbehaviour tolerance (TMT). Each lesson plan could break the lesson up into more or fewer sections, of different activity types.

Similarly, Hattie (2012) highlighted the relative unimportance of grouping students on ability. By creating appropriate profiles, one could easily create two groups of students that would be conventionally classified as higher and lower engagement. Simulations could be run on these separately. Then the classes could be mixed so that two mixed-ability classes were created. The same simulations could be run for these (same teacher parameter values, same lesson plans).

The CLSM could be used to perform a variety of other experiments, whether aimed at a specific class or subject or teacher, or completely hypothetical. Below are a few examples.

- How are lessons affected by the absence or presence of influential students? Which student is the most influential in the class? One could leave out each student in turn and inspect the difference in lesson outcomes.
- By adding disruptive/chatty students, one at a time, see if there is a point when a class is badly disrupted – in other words, find how many disruptors the teacher can cope with (for example, maybe 3 is manageable but 5 is disastrous). Does this vary with the teacher, the class or the subject? Or does it seem to be general?
- The model could be used as an optimization tool. For example, it could be used to find the seating arrangements that generate the greatest overall productive time. Or it could be used to find the lesson plan structure that maximizes overall class productive time and student participation, subject to constraints such as, say, having between 3 and 6 sections in the lesson.

9.5 Conclusions

The primary objective of the research project has been met. The CLSM has shown how an ABMS can adequately represent the behaviours of, and interactions between, students, teacher and teaching assistant in classroom lessons at a UK secondary school. It has also shown that it does this to the extent that it could be a useful tool in improving understanding of the dynamics of classroom lessons and investigating theories about lesson behaviours. One justification for this conclusion is, as Chattoe-Brown (2014, p. 15) expressed it: ‘we think we know how this particular aggregate pattern arose because we have been able to generate it in the ABM using only micro social processes for which we have at least some independent empirical support’. For example, the CLSM illuminated the influence on lesson disruption of factors such as the teacher’s inclination to provide individual assistance and the students' inclinations to chat. It was also able to realistically simulate completely artificial classes. It seems that the CLSM is a generic model that has faithfully represented the mental model that the author and the three validation teachers use to explain what happens in lessons. Furthermore, when the CLSM generates outcomes that appear unsurprising or obvious, this indicates that the CLSM does reflect the mental model that people in general have about what happens in classroom lessons.

This research is considered to have an impact in several ways:

(i) The models of teacher, TA and student behaviour are novel, first of their kind contributions to classroom behavioural theory; importantly, they were developed on the basis of empirical observations, not merely theory.

(ii) By constructing a complex and highly realistic model of classroom lessons, this research has extended the application of ABMSs in the area of complex social systems; others can use and modify the model.

(iii) The NetLogo-based implementation of the model (available at Ingram (2020a)) is free for anyone to use and modify and can be used to setup and simulate an enormous range of classroom lessons.

(iv) The Lesson Event Recording Tool (LERT) (available at Ingram (2018)) is a practical contribution to classroom data collection and quantitative research (having already been customized for another university research project), with several advantages over other data collection methods. Empirical lesson data collected using the LERT can be used for lesson analytics in support of educational decision-making. For example, collected data could be used to quantitatively determine to what extent what is supposed to happen in lessons is actually happening. It is possible that observing what actually happens in lessons could provide evidence to compare against what the headteacher or senior management team or governing body think happens in lessons. For example, if one student has been identified as requiring attention of some sort, data could be collected to determine if this is actually happening (regardless of what is written in the documentation, e.g., the student's individual education plan). Teachers might also be interested in an analysis of how they teach in terms of the amount of time invested, e.g., teacher-centred, learner-centred, content-centred, interactivity.

(v) The metrics and methods that were designed to quantify and compare lessons and then used to calibrate and validate the simulation model, extend the application of ABMS in the area of classroom research;

(vi) The behavioural rules took into account the relative spatial locations of others; the 2D visualizations used to represent the classroom and the interactions in lessons were novel and can be used by others.

A final note: If it were easy to predict the consequences of specific scenarios and interventions, there would be no need for lesson simulation. It is because lessons are so unpredictable that we attempt to use simulations to investigate consequences.

References

Abrahamson, D., Blikstein, P. and Wilensky, U. (2007) 'Classroom Model, Model Classroom: Computer-Supported Methodology for Investigating Collaborative-Learning Pedagogy', in *Proceedings of the 8th International Conference on Computer-supported Collaborative Learning*. New Brunswick, New Jersey, USA: International Society of the Learning Sciences (CSCL'07), pp. 49–58.

Ahoa, E. *et al.* (2010) 'Teachers' principles of decision-making and classroom management; a case study and a new observation method', *Procedia - Social and Behavioral Sciences*. Elsevier, 9, pp. 395–402. doi: 10.1016/j.sbspro.

Allday, R. A. *et al.* (2011) 'The Biasing Effects of Labels on Direct Observation by Preservice Teachers', *Teacher Education and Special Education: The Journal of the Teacher Education Division of the Council for Exceptional Children*, 34(1), pp. 52–58. doi: 10.1177/0888406410380422.

Axtell, R. (2000) *Why agents?: on the varied motivations for agent computing in the social sciences.* 17.

Balci, O. (1994) 'Validation, verification, and testing techniques throughout the life cycle of a simulation study', *Annals of Operations Research*, 53, pp. 121–173. doi: 10.1109/WSC.1994.717129.

Balci, O. (1995) 'Principles and techniques of simulation validation, verification, and testing', in *Winter Simulation Conference Proceedings, 1995.*, pp. 147–154. doi: 10.1109/WSC.1995.478717.

Balke, T. and Gilbert, N. (2014) 'How Do Agents Make Decisions? A Survey', *Journal of Artificial Societies and Social Simulation*, 17(4), p. 1. doi: 10.18564/jasss.2687.

Baumann, C. and Krskova, H. (2016) 'School discipline, school uniforms and academic performance', *International Journal of Educational Management*, 30(6), pp. 1003–1029. doi: 10.1108/IJEM-09-2015-0118.

Belhaj, M., Kebair, F. and Ben Said, L. (2014) 'Agent-Based Modeling and Simulation of the Emotional and Behavioral Dynamics of Human Civilians during Emergency Situations', *Multiagent System Technologies SE - 18*. Edited by J. P. Müller, M. Weyrich, and A. L. C. Bazzan. Springer International Publishing (Lecture Notes in Computer Science), 8732, pp. 266–281. doi: 10.1007/978-3-319-11584-9_18.

Bennett, N. G., Klimas Blanc, A. and Bloom, D. E. (2018) 'Commitment and the Modern Union : Assessing the Link Between Premarital Cohabitation and Subsequent Marital Stability', *American Sociological Review*, 53(1), pp. 127–138.

Bex, F. (2015) 'An integrated theory of causal stories and evidential arguments', *Proceedings of the 15th International Conference on Artificial Intelligence and Law - ICAIL '15*, pp. 13–22. doi: 10.1145/2746090.2746094.

Bicard, D. F. *et al.* (2012) 'Differential Effects of Seating Arrangements on Disruptive Behavior of Fifth Grade Students During Independent Seatwork', *Journal of Applied Behavior Analysis*, 45(2), pp. 407–411. doi: 10.1901/jaba.2012.45-407.

Bidwell, J. and Fuchs, H. (2011) 'Classroom analytics: measuring student engagement with automated gaze tracking', *Behavioral Research Methods*, 49, p. 113.

Blansky, D. *et al.* (2013) 'Spread of Academic Success in a High School Social Network', *PLoS ONE*, 8(2), pp. 1–4. doi: 10.1371/journal.pone.0055944.

Blatchford, P., Bassett, P. and Brown, P. (2005) 'Teachers' and pupils' behavior in large and small classes: A systematic observation study of pupils aged 10 and 11 years', *Journal of Educational Psychology*, 97(3), pp. 454–467. doi: 10.1037/0022-0663.97.3.454.

Blikstein, P., Abrahamson, D. and Wilensky, U. (2008) 'The classroom as a complex adaptive system: An agent-based framework to investigate students' emergent collective behaviors', *Proceedings of the 8th International conference for the learning sciences*. Utrecht, The Netherlands: International Society of the Learning Sciences, 3, pp. 12–13.

Bonabeau, E. (2002) 'Agent-based modeling: Methods and techniques for simulating human systems', *Proceedings of the National Academy of Sciences*, 99(Supplement 3), pp. 7280–7287. doi: 10.1073/pnas.082080899.

Borshchev, A. and Filippov, A. (2004) 'From System Dynamics and Discrete Event to Practical Agent Based Modeling: Reasons, Techniques, Tools', in *22nd International Conference of the System Dynamics Society, 25-29 July 2004*, p. 45.

Bosch, N. *et al.* (2018) 'Quantifying Classroom Instructor Dynamics with Computer Vision', in Penstein Rosé C. et al. (ed.) *The Proceedings of the 19th International Conference on Artificial Intelligence in Education*. Springer, Cham, pp. 30–42. doi: 10.1007/978-3-319-93843-1_3.

Box, G. E. P. (1979) 'Robustness in the Strategy of Scientific Model Building', in *Robustness in Statistics*. Academic Press, pp. 201–236. doi: 10.1016/B978-0-12-438150-6.50018-2.

Brouwers, A. and Tomic, W. (1999) 'Teacher Burnout, Perceived Self-Efficacy in Classroom Management, and Student Disruptive Behaviour in Secondary Education', *Curriculum and Teaching*, 14(2), pp. 7–26. doi: 10.7459/ct/14.2.02.

Brühwiler, C. and Blatchford, P. (2011) 'Effects of class size and adaptive teaching competency on classroom processes and academic outcome', *Learning and Instruction*, 21(1), pp. 95–108. doi: 10.1016/j.learninstruc.2009.11.004.

Burns, A. and Knox, J. S. (2011) 'Classrooms as complex adaptive systems: A relational model', *Teaching English as a Second or Foreign Language*. TESL-EJ, 15(1), pp. 1–25.

Buwaya, J. and Cleophas, C. (2015) 'A Game-Theoretic Calibration Approach for Agent-Based Planning Simulations', *IFAC-PapersOnLine*. Elsevier, 48(1), pp. 844–849.

Carroll, J. B. (1963) 'A model of school learning', *Teachers College Record*. Blackwell Publishing Ltd, 64(8), pp. 723–733.

Chattoe-Brown, E. (2014) 'Using agent based modelling to integrate data on attitude change', *Sociological Research Online*, 19(1). doi: 10.5153/sro.3315.

Chen, H.-J. *et al.* (2017) 'School beverage environment and children's energy expenditure associated with physical education class: an agent-based model simulation', *Pediatric Obesity*, 12(3), pp. 203–212. doi: 10.1111/ijpo.12126.

Christle, C. A. and Schuster, J. W. (2003) 'The Effects of Using Response Cards on Student Participation, Academic Achievement, and On-Task Behavior During Whole-Class, Math Instruction', *Journal of Behavioral Education*. Kluwer Academic Publishers-Plenum Publishers, 12(3), pp. 147–165. doi: 10.1023/a:1025577410113.

Clauset, K. H. and Gaynor, A. K. (1982) 'A Systems Perspective on Effective Schools', *Educational Leadership*, 40(3), p. 54.

Cotton, K. (1989) 'Educational Time Factors', *School Improvement Research Series*, (Close Up #8 November), p. 25.

Dalland, C. P., Klette, K. and Svenkerud, S. (2020) ‘Video studies and the challenge of selecting time scales’, *International Journal of Research and Method in Education*. Taylor & Francis, 43(1), pp. 53–66. doi: 10.1080/1743727X.2018.1563062.

Dillenbourg, P. (2013) ‘Design for classroom orchestration’, *Computers & Education*. Elsevier Ltd, 69(Nov), pp. 485–492. doi: 10.1016/j.compedu.2013.04.013.

Dillenbourg, P., Prieto, L. P. and Olsen, J. K. (2018) ‘Classroom Orchestration’, in Fischer, F. et al. (eds) *International Handbook of the Learning Sciences*. Taylor & Francis, pp. 180–190.

Dorin, A. and Geard, N. (2014) ‘The Practice of Agent-Based Model Visualization’, *Artificial Life*, 20(2), pp. 271–289. doi: 10.1162/artl.

Durall Gazulla, E. and Leinonen, T. (2016) ‘Why do we want data for learning? Learning analytics and the laws of media’, in *Lecture Notes in Educational Technology*. Springer, Berlin, Heidelberg, pp. 59–72. doi: 10.1007/978-3-662-47724-3_4.

Eberlen, J., Scholz, G. and Gagliolo, M. (2017) ‘Simulate this! An introduction to agent-based models and their power to improve your research practice’, *International Review of Social Psychology*, 30(1), pp. 149–160. doi: 10.5334/irsp.115.

Edmonds, B. *et al.* (2019) ‘Different Modelling Purposes’, *Journal of Artificial Societies and Social Simulation*, 22(3).

Edmonds, B. and Gershenson, C. (2015) ‘Modelling complexity for policy: Opportunities and challenges’, in *Handbook on complexity and public policy*. Edward Elgar Publishing.

Edmonds, B. and Moss, S. (2005) ‘From KISS to KIDS–an “anti-simplistic” modelling approach’, in Davidsson, P., Logan, B., and Takadama, K. (eds) *Multi-Agent and Multi-Agent-Based Simulation*. Berlin, Heidelberg: Springer, Berlin, Heidelberg, pp. 130–144. doi: 10.1007/b106991.

Epstein, J. M. (2008) ‘Why model?’, *Journal of Artificial Societies and Social Simulation*, 11(4). doi: 10.1080/01969720490426803.

Evans, A., Heppenstall, A. and Birkin, M. (2018) ‘Understanding simulation results’, in Edmonds, B. and Meyer, R. (eds) *Simulating Social Complexity: A Handbook*, pp. 205–227. doi: 10.1007/978-3-319-66948-9_10.

Ferry, B., Kervin, L. and Carrington, L. (2010) ‘ClassSim: An Approach to Educator Development Through a Simulation’, in Baek, Y. K. (ed.) *Gaming for Classroom-Based Learning: Digital Role Playing as a Motivator of Study*, pp. 238–251. doi: 10.4018/978-1-60960-195-9.ch404.

Fidalgo-Blanco, Á. *et al.* (2015) ‘Using Learning Analytics to improve teamwork assessment’, *Computers in Human Behavior*, 47, pp. 149–156. doi: 10.1016/j.chb.2014.11.050.

Fisher, C. W. (1978) *Teaching and Learning in the Elementary School: A Summary of the Beginning Teacher Evaluation Study. Report VII-1. Beginning Teacher Evaluation Study Technical Report Series*.

Galán, J. M. *et al.* (2009) ‘Errors and Artefacts in Agent-Based Modelling’, *Journal of Artificial Societies and Social Simulation*, 12(11).

Gamboa-Brooks-Gray, R. (2015) *Conceptualizing Mathematics Classrooms As Complex Adaptive Systems To Predict Student Achievement*. The Claremont Graduate University.

García-Magariño, I. *et al.* (2016) ‘A hybrid approach with agent-based simulation and clustering for sociograms’, *Information Sciences*. Elsevier, 345, pp. 81–95. doi: 10.1016/j.ins.2016.01.063.

García-Magariño, I. *et al.* (2017) 'ABS-SOCI: An agent-based simulator of student sociograms', *Applied Sciences*, 7(11), pp. 1–18. doi: 10.3390/app7111126.

García-Magariño, I. and Plaza, I. (2015) 'FTS-SOCI: An agent-based framework for simulating teaching strategies with evolutions of sociograms', *Simulation Modelling Practice and Theory*, 57, pp. 161–178. doi: 10.1016/j.simpat.2015.07.003.

Ge, Y. *et al.* (2011) 'Agent-Based Modeling for Influenza H1N1 in an Artificial Classroom', *Systems Engineering Procedia*, 2, pp. 94–104. doi: 10.1016/j.sepro.2011.10.012.

van Geert, P. (1994) *Dynamic systems of development: Change between complexity and chaos*. New York: Harvester Wheatsheaf.

Gettinger, M. and Seibert, J. K. (2002) 'Best Practices in Increasing Academic Learning Time', in Grimes, A. T. & J. (ed.) *Best practices in school psychology IV*, pp. 773–787.

Gibson, D. (2009) 'Modeling Classroom Behaviors in Software Agents', in *Digital simulations for improving education: Learning through artificial teaching environments*, pp. 119–156. doi: 10.4018/978-1-60566-322-7.ch007.

Gibson, D. and Baek, Y. (2009) *Digital Simulations for Improving Education: Learning Through Artificial Teaching Environments*. doi: 10.4018/978-1-60566-322-7.

Gilbert, N. (2007) *Computational Social Science: Agent-based social simulation*.

Godwin, K. E. *et al.* (2015) *The Variable Relationship Between On-Task Behavior and Learning*.

Goldring, E. *et al.* (2015) 'Make Room Value Added: Principals' Human Capital Decisions and the Emergence of Teacher Observation Data', *Educational Researcher*. SAGE PublicationsSage CA: Los Angeles, CA, 44(2), pp. 96–104. doi: 10.3102/0013189X15575031.

Gómez-Cruz, N. A., Loaiza Saa, I. and Ortega Hurtado, F. F. (2017) 'Agent-based simulation in management and organizational studies: a survey', *European Journal of Management and Business Economics*, 26(3), pp. 313–328. doi: 10.1108/EJMBE-10-2017-018.

Grimm, V. (2002) 'VISUAL DEBUGGING: A WAY OF ANALYZING, UNDERSTANDING AND COMMUNICATING BOTTOM-UP SIMULATION MODELS IN ECOLOGY', *NATURAL RESOURCE MODELING*, 15(1).

Halliday, T. J. and Kwak, S. (2012) 'What is a peer? The role of network definitions in estimation of endogenous peer effects', *Applied Economics*, 44(3), pp. 289–302. doi: 10.1080/00036846.2010.505557.

Hanushek, E. A. *et al.* (2003) 'Does peer ability affect student achievement?', *Journal of Applied Econometrics*, 18(5), pp. 527–544. doi: 10.1002/jae.741.

Harland, K. and Heppenstall, A. J. (2012) 'Using agent-based models for education planning: Is the UK education system agent based?', in *Agent-Based Models of Geographical Systems*. Dordrecht: Springer Netherlands, pp. 481–497. doi: 10.1007/978-90-481-8927-4_23.

Hattie, J. (2005) 'The paradox of reducing class size and improving learning outcomes', *International Journal of Educational Research*, 43(6), pp. 387–425. doi: 10.1016/j.ijer.2006.07.002.

Hattie, J. (2012) *Visible learning for teachers: maximizing impact on learning*. London: Routledge.

Hattie, J. (2015) 'Why smaller class sizes, homework and projects don't always benefit learning', (September).

Hillston, J. (2003) 'Model Validation and Verification'. Edinburgh: University of Edinburgh. doi: 10.1080/17415970008027722.

Hillston, J. (2017) ‘Performance Modelling — Lecture 16: Model Validation and Verification’. Edinburgh.

Hirschy, A. S. and Wilson, M. E. (2002) ‘The Sociology of the Classroom and Its Influence on Student Learning’, *Peabody Journal of Education*, 77(3), pp. 85–100. doi: 10.1207/S15327930PJE7703_5.

Hmelo-Silver, C. E., Liu, L. E. I. and Jordan, R. (2009) ‘About Complex Natural Systems’, *Computers & Education*, 4(3), pp. 253–280.

Hoddy, E. T. (2019) ‘Critical realism in empirical research: employing techniques from grounded theory methodology’, *International Journal of Social Research Methodology*. Routledge, 22(1), pp. 111–124. doi: 10.1080/13645579.2018.1503400.

Holstein, K., McLaren, B. M. and Aleven, V. (2017a) ‘Intelligent tutors as teachers’ aides’, *Proceedings of the Seventh International Learning Analytics & Knowledge Conference on - LAK ’17*, pp. 257–266. doi: 10.1145/3027385.3027451.

Holstein, K., McLaren, B. M. and Aleven, V. (2017b) ‘SPACLE: Investigating Learning Across Virtual and Physical Spaces Using Spatial Replays’, in *Proceedings of the Seventh International Learning Analytics & Knowledge Conference*, pp. 358–367. doi: 10.1145/3027385.3027450.

Imler, B. and Eichelberger, M. (2011) ‘Using screen capture to study user research behavior’, *Library Hi Tech*, 29(3), pp. 446–454. doi: 10.1108/07378831111174413.

Ingram, F. J. (2018) *Classroom lesson event data for teaching and learning analytics*. Lancaster University. doi: 10.17635/lancaster/researchdata/200.

Ingram, F. J. (2020a) *LesSim: An Agent-based Classroom Lesson Simulation*. Lancaster University. doi: 10.17635/lancaster/researchdata/391.

Ingram, F. J. (2020b) *LesSim Classroom Lesson Simulation Face Validity testing*. Lancaster University. doi: 10.17635/lancaster/researchdata/392.

IRIS Connect UK (2020) *IRIS Connect: CPD & Professional Development for Teachers*. Available at: https://www.irisconnect.com/uk/ (Accessed: 10 September 2020).

Jackson, T. (2005) *Motivating Sustainable Consumption: a review of evidence on consumer behaviour and behavioural change*, *Sustainable Development Research Network*. doi: 10.1260/0958305043026573.

Janssen, M. A. and Ostrom, E. (2006) ‘Empirically Based, Agent-based models’, *Ecology and Society*. The Resilience Alliance, 11(2), p. art37. doi: 10.5751/ES-01861-110237.

Kaplan, A., Gheen, M. and Midgley, C. (2002) ‘Classroom goal structure and student disruptive behaviour.’, *The British Journal of Educational Psychology*, 72(2), pp. 191–211. doi: 10.1348/000709902158847.

Karweit, N. (1984) ‘Time-on-Task Reconsidered: Synthesis of Research on Time and Learning’, *Educational Leadership*, 41(8), pp. 32–35.

Katsaliaki, K. and Mustafee, N. (2011) ‘Applications of simulation within the healthcare context’, *Journal of the Operational Research Society*, 62(8), pp. 1431–1451. doi: 10.1057/jors.2010.20.

Kim, B. and Cheong, D. (2009) ‘simClass: Simulate Your Class Before You Teach’, in Gibson, D. and Baek, Y. K. (eds) *Digital Simulations for Improving Education: Learning Through Artificial Teaching Environments*. IGI Global, pp. 289–307. doi: 10.4018/978-1-60566-322-7.ch015.

Kim, Y. and Baylor, A. L. (2016) ‘Research-Based Design of Pedagogical Agent Roles: A Review, Progress, and Recommendations’, *International Journal of Artificial Intelligence in Education*, 26(1), pp. 160–169. doi: 10.1007/s40593-015-0055-y.

Klein, G. (2008) ‘Naturalistic Decision Making’, *Human Factors*, 50(3), pp. 456–460. doi: 10.1518/001872008X288385.

Klerkx, J., Verbert, K. and Duval, E. (2017) ‘Learning Analytics Dashboards’, in *Handbook of Learning Analytics*. Society for Learning Analytics Research (SoLAR), pp. 143–150. doi: 10.18608/hla17.012.

Klügl, F. (2008) ‘A validation methodology for agent-based simulations’, in *Proceedings of the 2008 ACM symposium on applied computing*, pp. 39–43. doi: 10.1145/1363686.1363696.

Koster, A. *et al.* (2016) ‘The role of agent-based simulation in education’, *Communications in Computer and Information Science*, 677, pp. 156–167. doi: 10.1007/978-3-319-52039-1_10.

Kovalainen, M. and Kumpulainen, K. (2009) ‘Discourse-enriched instruction in the mathematics classroom: analyzing collective-problem solving’, in *Investigating Classroom Interaction Methodologies in Action*, pp. 43–71.

Larsen-Freeman, D. (2016) ‘Classroom-oriented research from a complex systems perspective’, *Studies in Second Language Learning and Teaching*, 6(3), pp. 377–393. doi: 10.14746/ssllt.2016.6.3.2.

Lee, J. *et al.* (2015) ‘The complexities of agent-based modeling output analysis’, *Journal of Artificial Societies and Social Simulation*, 18(4), pp. 1–26. doi: 10.18564/jasss.2897.

Lessonbox Ltd (2020) *Lessonbox: Using video to empower teachers and enhance learning*. Available at: https://www.lessonbox.co.uk/ (Accessed: 10 September 2020).

Li, H. and Lee, K. C. (2013) ‘An Interpersonal Relationship Framework for Virtual Community Participation Psychology: From Covert to Overt Process’, *Social Science Computer Review*, 31(6), pp. 703–724. doi: 10.1177/0894439313494665.

Liu, R., Jiang, D. and Shi, L. (2016) ‘Agent-based simulation of alternative classroom evacuation scenarios’, *Frontiers of Architectural Research*. Elsevier, 5(1), pp. 111–125. doi: 10.1016/j.foar.2015.12.002.

Lorscheid, I., Heine, B. O. and Meyer, M. (2012) ‘Opening the “Black Box” of Simulations: Increased Transparency and Effective Communication Through the Systematic Design of Experiments’, *Computational and Mathematical Organization Theory*, 18(1), pp. 22–62. doi: 10.1007/s10588-011-9097-3.

Lugrin, J.-L. *et al.* (2016) ‘Breaking Bad Behaviors: A New Tool for Learning Classroom Management Using Virtual Reality’, *Frontiers in ICT*, 3(November). doi: 10.3389/fict.2016.00026.

Macal, C. M. (2016) ‘Everything you need to know about agent-based modelling and simulation’, *Journal of Simulation*. Palgrave Macmillan UK, 10(2), pp. 144–156. doi: 10.1057/jos.2016.7.

Macal, C. M. and North, M. J. (2010) ‘Tutorial on agent-based modelling and simulation’, *Journal of Simulation*, 4(3), pp. 151–162. doi: 10.1057/jos.2010.3.

Macal, C. and North, M. (2014) ‘Introductory tutorial: Agent-based modeling and simulation’, in A. Tolk, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. M. (ed.) *Proceedings - Winter Simulation Conference*, pp. 6–20. doi: 10.1109/WSC.2014.7019874.

Manzo, G. (2013) ‘Educational Choices and Social Interactions: A Formal Model and a Computational Test’, in Elisabeth Birkelund, G. (ed.) *Class and Stratification Analysis (Comparative Social Research)*. Emerald Group Publishing Limited, pp. 47–100. doi: 10.1108/S0195-6310(2013)0000030007.

Manzo, G. and Matthews, T. (2014) ‘Potentialities and limitations of agent-based simulations’, *Revue française de sociologie*, 55(4), pp. 433–462.

Maroulis, S. *et al.* (2014) ‘Modeling the Transition to Public School Choice’, *JASSS*, 17(2).

Martin, R. and Schlüter, M. (2015) ‘Combining system dynamics and agent-based modeling to analyze social-ecological interactions—an example from modeling restoration of a shallow lake’, *Frontiers in Environmental Science*, 3(October), pp. 1–15. doi: 10.3389/fenvs.2015.00066.

Marzano, R. J., Marzano, J. S. and Pickering, D. (2003) *Classroom management that works: Research-based strategies for every teacher*. ASCD.

Matsuda, N. *et al.* (2010) ‘Learning by Teaching SimStudent: Technical Accomplishments and an Initial Use with Students’, in *Lecture Notes in Computer Science*. Springer, Berlin, Heidelberg, pp. 317–326. doi: 10.1007/978-3-642-13388-6_36.

Matsuda, N., Cohen, W. W. and Koedinger, K. R. (2015) ‘Teaching the teacher: Tutoring simstudent leads to more effective cognitive tutor authoring’, *International Journal of Artificial Intelligence in Education*, 25(1), pp. 1–34. doi: 10.1007/s40593-014-0020-1.

McDevitt, A. (2017) *Classroom Model with Learning Assistants*, *NetLogo Modeling Commons*. Available at: http://modelingcommons.org/browse/one_model/5049#model_tabs_browse_info (Accessed: 20 June 2020).

Mital, P., Moore, R. and Llewellyn, D. (2014) ‘Analyzing K-12 education as a complex system’, in *Procedia Computer Science*. (2014 Conference on Systems Engineering Research), pp. 370–379. doi: 10.1016/j.procs.2014.03.046.

Montes, G. (2012) ‘Using Artificial Societies to Understand the Impact of Teacher Student Match on Academic Performance : The Case of Same Race Effects’, *JASSS*, 15(4).

Moody, G. D. and Siponen, M. (2013) ‘Using the theory of interpersonal behavior to explain non-work-related personal use of the Internet at work’, *Information and Management*. Elsevier B.V., 50(6), pp. 322–335. doi: 10.1016/j.im.2013.04.005.

Moore, D. E. and Hayward, M. D. (2019) ‘Occupational Careers and Mortality of Elderly Men’, *Demography*, 27(1), pp. 31–53.

Norling, E. (2009) *Modelling Human Behaviour with BDI Agents*. The University of Melbourne.

Norling, E., Edmonds, B. and Meyer, R. (2018) ‘Informal approaches to developing simulation models’, in Edmonds, B. and Meyer, R. (eds) *Simulating Social Complexity: A Handbook*, pp. 61–79. doi: 10.1007/978-3-319-66948-9_5.

Norling, E., Sonenberg, L. and Rönnquist, R. (2000) ‘Enhancing Multi-Agent Based Simulation with Human-Like Decision Making Strategies’, in Moss, S. and Davidsson, P. (eds) *International Workshop on Multi-Agent Systems and Agent-Based Simulation*. Springer Berlin Heidelberg, pp. 214–228.

Nuno, D., Nuno, F. and Agostinho, C. R. (2017) ‘Verifying and validating simulations’, in Edmonds, B. and Meyer, R. (eds) *Understanding Complex Systems*, pp. 173–204. doi: 10.1007/978-3-319-66948-9_9.

Ocumpaugh, J., Baker, R. and Rodrigo, M. M. (2014) *Baker Rodrigo Ocumpaugh Monitoring Protocol (BROMP) 2.0 Technical and Training Manual*.

Onggo, B. S. S. (2010) ‘Running agent-based models on a discrete-event simulator’, in *Proceedings of the 24th European Simulation and Modelling Conference*, pp. 51–55.

ONVU Technologies Group (2020) *Lessonvu - ONVULearning*. Available at: https://www.onvulearning.com/lessonvu/ (Accessed: 11 September 2020).

Padgham, L., Singh, D. and Zambetta, F. (2015) 'SOCIAL SIMULATION FOR ANALYSIS, INTERACTION, TRAINING AND COMMUNITY AWARENESS', in *Proceedings of the 2015 Winter Simulation Conference*, pp. 3130–3131.

Parsonson, B. S. (2012) 'Evidence-Based Classroom Behaviour Management Strategies.', *Kairaranga*, 13(1), pp. 16–23. doi: 10.1100/2012/208907.

Paul, R. *et al.* (2020) 'Using Agent-Based Modelling for EER Experimental Design: Preliminary Validation Based on Student Cheating Behaviours', *Canadian Engineering Education Association (CEEA)*, pp. 1–8.

Payne, R. K. (2007) *Working with Students; Discipline Strategies for the 21st Century Classroom, Simulation Companion Edition*. aha! Process, Incorporated.

Petty, G. (2006) 'Evidence Based Classroom Management and Discipline', in *Evidence Based Teaching*. Available at: http://geoffpetty.com/geoffs-books/evidence-based-teaching-ebt/.

Piccolo, D. and Oskorus, A. (2009) 'Designing commercial simulations for teachers', in Gibson, D. (ed.) *Digital Simulations for Improving Education: Learning Through Artificial Teaching Environments*, pp. 249–272. doi: 10.4018/978-1-60566-322-7.ch013.

Poplin, M. *et al.* (2011) 'She's strict for a good reason highly effective teachers in low-performing urban schools', *Phi Delta Kappan*, 92(5), pp. 39–43. doi: 10.1177/003172171109200509.

Prieto, L. P. *et al.* (2016) 'Teaching Analytics: Towards Automatic Extraction of Orchestration Graphs Using Wearable Sensors', *Proceedings of the Sixth International Conference on Learning Analytics & Knowledge - LAK '16*, pp. 148–157. doi: 10.1145/2883851.2883927.

Qudrat-Ullah, H. (2005) 'STRUCTURAL VALIDATION OF SYSTEM DYNAMICS AND AGENT-BASED SIMULATION MODELS', in Merkuryev, Zobel, and Kerckhoffs (eds) *Proceedings 19th European Conference on Modelling and Simulation*.

Qudrat-Ullah, H. (2008) 'Structural Validation of Simulation Models : An Illustration', in *Tenth International Conference on Computer Modeling and Simulation Structural*. Tenth International Conference on Computer Modeling and Simulation Structural, pp. 537–542. doi: 10.1109/UKSIM.2008.126.

Raca, M. and Dillenbourg, P. (2013) 'System for Assessing Classroom Attention', *Proceedings of the Third International Conference on Learning Analytics and Knowledge*, pp. 265–269.

Raca, M. and Dillenbourg, P. (2014) 'Holistic Analysis of the Classroom', in *Proceedings of the 2014 ACM workshop on Multimodal Learning Analytics Workshop and Grand Challenge - MLA '14*. ACM Press, pp. 13–20. doi: 10.1145/2666633.2666636.

Raca, M., Tormey, R. and Dillenbourg, P. (2013) 'Student motion and its potential as a classroom performance metric', in *3rd International Workshop on Teaching Analytics (IWTA) (No. POST_TALK).*

Radó, M. and Takács, K. (2019) 'Relational integration in schools through seating assignments', *JASSS*, 22(4). doi: 10.18564/jasss.4115.

Ramos, A., Esslinger, K. and Pyle, E. (2015) 'The Effects of Field Experience on Delivery of Feedback.', *Physical Educator*, 72, pp. 278–287.

Read, M. and Timmis, J. (2012) ‘Techniques for grounding agent-based simulations in the real domain: a case study in experimental autoimmune encephalomyelitis’, *Mathematical and Computer Modelling of Dynamical Systems*, 18(1), pp. 67–86.

Reardon, S. *et al.* (2016) ‘Agent-based simulation models of the college sorting process’, *Journal of Artificial Societies and Social Simulation*, 19(1), p. 8. doi: 10.18564/jasss.2993.

Reynolds, C. W. (1987) ‘Flocks, herds, and schools: A distributed behavioral model’, in *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 1987*, pp. 25–34. doi: 10.1145/37401.37406.

Ringler, P., Keles, D. and Fichtner, W. (2016) ‘Agent-based modelling and simulation of smart electricity grids and markets - A literature review’, *Renewable and Sustainable Energy Reviews*. Elsevier, 57, pp. 205–215. doi: 10.1016/j.rser.2015.12.169.

Robinson, S. (1997) ‘Simulation model verification and validation: increasing the users’ confidence’, in Andradottir, S. et al. (eds) *Proceedings of the 1997 Winter Simulation Conference*, pp. 37–47.

Robinson, S. (2008a) ‘Conceptual modelling for simulation Part I: Definition and requirements’, *Journal of the Operational Research Society*, 59(3), pp. 278–290. doi: 10.1057/palgrave.jors.2602368.

Robinson, S. (2008b) ‘Conceptual modelling for simulation Part II: A framework for conceptual modelling’, *Journal of the Operational Research Society*, 59(3), pp. 291–304. doi: 10.1057/palgrave.jors.2602369.

Robinson, S. (2013) ‘Conceptual modeling for simulation’, *2013 Winter Simulations Conference (WSC)*, pp. 377–388. doi: 10.1109/WSC.2013.6721435.

Robinson, S. (2014) *Simulation: The Practice of Model Development and Use*. 2nd edn. Basingstoke, Hampshire, New York: Palgrave Macmillan.

Robinson, S., Arbez, G. and Birta, L. G. (2015) ‘CONCEPTUAL MODELING: DEFINITION, PURPOSE AND BENEFITS’, in L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. R. (ed.) *Proceedings of the 2015 Winter Simulation Conference*, pp. 2812–2826. doi: 10.1088/1751-8113/44/8/085201.

Salgado, M. and Gilbert, N. (2013) ‘Agent Based Modelling’, in *Handbook of Quantitative Methods for Educational Research*. Sense Publishers, pp. 247–265. doi: 10.1007/978-94-6209-404-8.

Salgado, M., Marchione, E. and Gilbert, N. (2014) ‘Analysing differential school effectiveness through multilevel and agent-based modelling’, *Journal of Artificial Societies and Social Simulation*, 17(4). doi: 10.18564/jasss.2534.

Sargent, R. G. (2007) ‘Verification and Validation of Simulation Models’, *Winter Simulation Conference*, pp. 124–137.

Sargent, R. G. (2011) ‘Verification and validation of simulation models’, in S. Jain, R.R. Creasey, J. Himmelspach, K.P. White, and M. F. (ed.) *Journal of Simulation*, pp. 12–24. doi: 10.1057/jos.2012.20.

Savola, L. T. (2008) *VIDEO-BASED ANALYSIS OF MATHEMATICS CLASSROOM PRACTICE: EXAMPLES FROM FINLAND AND ICELAND*. Columbia University.

Schlesinger, S. *et al.* (1979) ‘Terminology for model credibility’, *Simulation*, 32(3), pp. 103–104.

Schwieso, J. (1995) ‘Tasks and tables: The effects of seating arrangements on task engagement in primary classrooms’, *Educational Research*, 37(3), pp. 279–291. doi: 10.1080/0013188950370306.

Sergis, S. and Sampson, D. G. (2016) ‘Towards a teaching analytics tool for supporting reflective educational (re)design in Inquiry-based STEM Education’, *Proceedings - IEEE 16th International Conference on Advanced Learning Technologies, ICALT 2016*, pp. 314–318. doi: 10.1109/ICALT.2016.134.

Shvarts, A. and Abrahamson, D. (2019) ‘Dual-eye-tracking Vygotsky: A microgenetic account of a mathematics-tutorial case study as a teaching/learning collaboration’, *Learning, Culture and Social Interaction*. Elsevier, 22(June), p. 100316. doi: 10.1016/j.lcsi.2019.05.003.

Siebers, P. O. and Klügl, F. (2018) ‘What Software Engineering has to offer to Agent-Based Social Simulation’, in Edmonds, B. and Meyer, R. (eds) *Simulating Social Complexity: A Handbook*. 2nd edn, pp. 81–117. doi: 10.1007/978-3-319-66948-9_6.

Siebers, P. and Onggo, B. S. S. (2014) ‘Graphical Representation of Agent-based Models in Operational Research and Management Science using UML’, in *Proceedings of the Operational Research Society Simulation Workshop 2014 (SW14)*, pp. 143–153.

Silva, P. C. L. *et al.* (2020) ‘COVID-ABS: An agent-based model of COVID-19 epidemic to simulate health and economic effects of social distancing interventions’, *Chaos, Solitons and Fractals*, 139. doi: 10.1016/j.chaos.2020.110088.

simEd LLC (2020) *About simSchool*. Available at: https://www.simschool.org/home/simschool/#about (Accessed: 7 September 2020).

Sotirov, M., Sallnäs, O. and Eriksson, L. O. (2017) ‘Forest owner behavioral models, policy changes, and forest management. An agent-based framework for studying the provision of forest ecosystem goods and services at the landscape level’, *Forest Policy and Economics*. doi: 10.1016/j.forpol.2017.10.015.

Stallings, J. (1980) ‘Allocated academic learning time revisited, or beyond time on task’, *Educational researcher*. American Educational Research Association, 9(11), pp. 11–16. doi: 10.3102/0013189X009011011.

Strang, H. R. and Loper, A. B. (1984) ‘A Microcomputer-Based Simulation of Classroom Interaction’, *Journal of Educational Technology Systems*, 12(3), pp. 209–219. doi: 10.2190/Y2HQ-K56J-PCRB-6F2X.

Suleimenova, D., Bell, D. and Groen, D. (2017) ‘A generalized simulation development approach for predicting refugee destinations’, *Scientific Reports*. Springer US, 7(1), pp. 1–13. doi: 10.1038/s41598-017-13828-9.

Tako, A. A. and Kotiadis, K. (2015) ‘PartiSim: A multi-methodology framework to support facilitated simulation modelling in healthcare’, *European Journal of Operational Research*. Elsevier Ltd., 244(2), pp. 555–564. doi: 10.1016/j.ejor.2015.01.046.

Thiele, J. C., Kurth, W. and Grimm, V. (2014) ‘Facilitating parameter estimation and sensitivity analysis of agent-based models: A cookbook using NetLogo and R’, *Journal of Artificial Societies and Social Simulation*, 17(3), p. 11. doi: 10.18564/jasss.2503.

Thornberg, R. (2012) ‘Informed Grounded Theory’, *Scandinavian Journal of Educational Research*, 56(3), pp. 243–259. doi: 10.1080/00313831.2011.581686.

TIMSSVIDEO (1999) *The Third International Mathematics and Science Study (TIMSS) 1999 Video Study*. Available at: http://www.timssvideo.com/ (Accessed: 9 September 2020).

Troitzsch, K. G. (2018) 'Historical Introduction', in Edmonds, B. and Meyer, R. (eds) *Simulating Social Complexity: A Handbook*. Springer International Publishing, pp. 13–22.

Tsioptsias, N., Tako, A. and Robinson, S. (2016) 'Model validation and testing in simulation: A literature review', in *5th Student Conference on Operational Research*. doi: 10.4230/OASIcs.SCOR.2016.6.

Urban, C. and Schmidt, B. (2001) *PECS - Agent-Based Modelling of Human Behaviour*, *AAAI Technical Report FS-01-02*.

Vosaic (2020) *Vosaic: Video coaching, video observation, video analysis*. Available at: https://vosaic.com/ (Accessed: 10 September 2020).

Weaver, W. (1948) 'Science and Complexity', *American Scientist*. Sigma Xi, The Scientific Research Society, 51(2), pp. 180–192. doi: 10.2307/27826254.

Wei, X.-Y. and Yang, Z.-Q. (2012) 'Mining in-class social networks for large-scale pedagogical analysis', *Proceedings of the 20th ACM international conference on Multimedia - MM '12*, p. 639. doi: 10.1145/2393347.2393436.

Wheldall, K. and Lam, Y. Y. (1987) 'Rows versus Tables. II. The Effects of Two Classroom Seating Arrangements on Classroom Disruption Rate, On-task Behaviour and Teacher Behaviour in Three Special School Classes', *Educational Psychology*, 7(4), pp. 303–312. doi: 10.1080/0144341870070405.

Wilensky, U. (1999) 'NetLogo'. Evanston, IL: Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL. Available at: http://ccl.northwestern.edu/netlogo/.

Wilensky, U. and Rand, W. (2007) 'Making models match: Replicating an agent-based model', *Journal of Artificial Societies and Social Simulation*, 10(4).

Xhakaj, F., Aleven, V. and McLaren, B. M. (2016) 'How Teachers Use Data to Help Students Learn: Contextual Inquiry for the Design of a Dashboard', in Verbert, K., Sharples, M., and Klobučar, T. (eds) *Adaptive and Adaptable Learning*. Springer, Cham, pp. 340–354. doi: 10.1007/978-3-319-45153-4_26.

Xiang, X., Kennedy, R. and Madey, G. (2005) 'Verification and Validation of Agent-based Scientific Simulation Models', *Agent-Directed Simulation Conference*, pp. 47–55.

Yeung, R. and Nguyen-Hoang, P. (2016) 'Endogenous peer effects: Fact or fiction?', *The Journal of Educational Research*. Routledge, 109(1), pp. 37–49. doi: 10.1080/00220671.2014.918528.

Zohrabi, M. (2013) 'Mixed Method Research: Instruments, Validity, Reliability and Reporting Findings', *Theory and Practice in Language Studies*, 3(2). doi: 10.4304/tpls.3.2.254-262.

Appendices

Appendix A Further details of the CLSM

This chapter provides additional details to those in Chapter 3.

A.1 Teacher behaviour

The teacher agent follows the algorithm shown in Figure 3-3, which loops repeatedly through a set of production rules that determine the activity for the next time interval. An informal description of these teacher behaviour rules is shown in the box below. Written from the perspective of the teacher, they were thought to embody the selected teacher behaviour and activities. On each cycle, the algorithm moves through the list of rules until an applicable rule is found - the remaining rules are ignored. If no other rule applies, the final rule has the teacher perform the default activity specified in the lesson plan (explained in section 3.2). Briefly, the rules consider the following sequence of decisions the teacher agent needs to make:

Should I be disciplining the class?
Should I be disciplining an individual?
Should I be answering a question?
Should I be praising someone?
Should I be whole class teaching?
Should I be helping an individual?
Should I be offering help?
Should I talk to the TA?
Otherwise I should be doing what the lesson plan stipulates.

The Teacher Behaviour Rules

Do what the lesson plan says (stay in or enter the primary planned state) but handle the following exceptions, including dealing with misbehaviour and providing one-to-one support.

1. If at any time a student is seriously disruptive (student state 2 or 3 with Range 3), according to my misbehaviour tolerance, discipline them (state 1) for a suitable amount of time. Treat state 2 as more serious than state 3 so intervene earlier.

If I have been disciplining a student for a suitable amount of time, then stop and do what the lesson plan specifies; otherwise continue.

2. If there is generally too much messing about (student state 3 or 4), according to how long the messing about has persisted and according to my misbehaviour tolerance, then for a suitable amount of time remind the class to focus on the activity (state 5). Treat state 3 as more serious than state 4 so intervene earlier.
   If I have been disciplining the class for a suitable amount of time, then stop and do what the lesson plan specifies; otherwise continue.

3. If at any time a student proactively participates in the teaching activities (e.g., student state 10 or student interacts in student state 12, 13), give that student some appropriate praise (state 11) for a suitable amount of time.
   If I have been praising a student for a suitable amount of time, then stop and do what the lesson plan specifies; otherwise continue.

4. If the lesson plan says it is time to move on to a new section, then conclude what I am doing and change to the prescribed new activity state.

5. If a student asks a question during whole-class teaching, then provide an answer (state 10). If I am answering a question (state 10) check if I have answered long enough.
   If I have, stop and do what the lesson plan specifies; otherwise continue.

6. If a student asks a question while I am not whole-class teaching, then provide an answer in the context of one-to-one support (state 13). If I am providing one-to-one support (state 13) check if I have answered long enough. If I have, stop and do what the lesson plan specifies; otherwise continue.

7. If I am helping someone but I am supposed to be whole-class teaching (state 12) - so just a little help is called for - check if I have been helping long enough.
   If I have, stop and do what the lesson plan specifies; otherwise continue.

8. If the students are expected to be doing individual work (state 8) or group work (state 9) and I am helping a student (state 13) then check if I have been doing this for long enough. If I have, stop and do what the lesson plan specifies; otherwise continue.

9. If I am not helping anyone but I could, and someone has asked for help, then help that person.

10. if I am not helping anyone and I am available to help someone (which means states 4, 7, 8, 15) and it has been a while since I helped anyone (this length of time varies according to my nature to intervene and according to how long I have not been helping anyone) then:
    a. choose someone who is not being helped by the TA and appears to be doing nothing (student state 4 or 7) or is chatting (student state 2 or 3), but not someone who has recently received help, and go and offer support to him/her.
    b. otherwise (so as to allow students time to think and work)
        (i) if the TA is not busy and it is not too recent that I spoke to the TA, then confer with the TA (state 15);
        (ii) otherwise start/resume the main activity as specified in the lesson plan for that time.

11. If none of the above situations apply, the teacher starts/resumes the main activity as specified in the lesson plan for that time.

The rules require the existence of several variables and constants. Some of these apply at the agent level, some to the whole lesson. Some constants apply to all lessons - they are model constants. Figure A-1 summarizes the required data and Table A-1 defines the teacher-related model constants. Besides these variables and constants, the teacher agent was given additional attributes to enable useful animation and reporting.

During the case study, teachers were observed making decisions about when to react, how to react and for how long to react to students. The reactions may have been triggered by specific events, or by the frequency or duration of earlier events. But these triggers, thresholds and durations seemed to vary across lessons and across teachers. Even for one teacher, they would vary within a lesson. To make teacher behaviour more realistic, stochasticity was added in various places, allowing the thresholds to vary within a lesson.

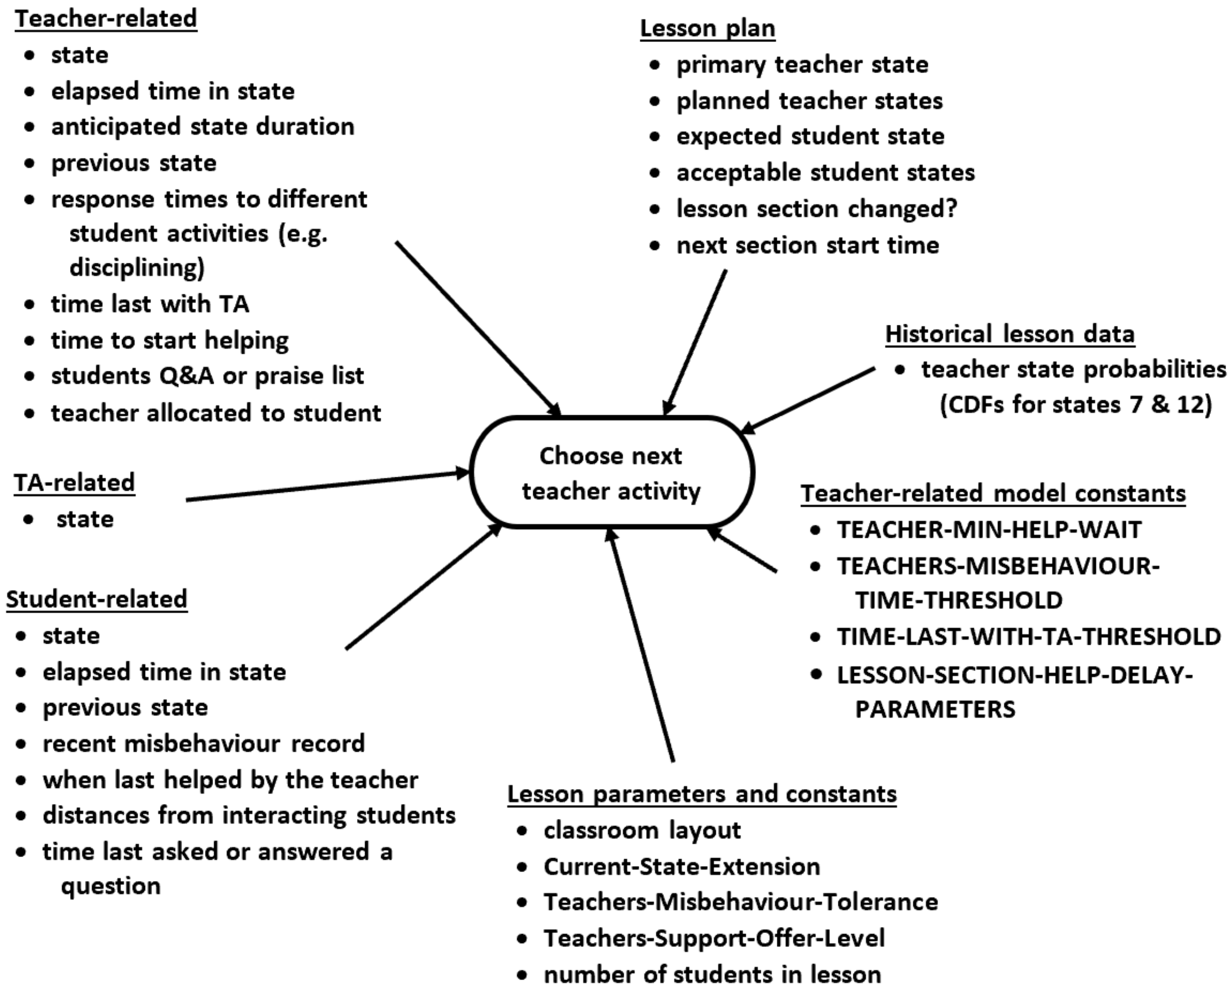

Figure A-1 The data required for teacher agent behaviour modelling

The specific places were:

- anticipated state duration – explained in Appendix 1.1.1.1.1B.1;
- when to offer help;
- which student to help;
- response time to discipline (thresholds to wait until): this is set differently for 4 types of teacher state and 3 types of student state, each combination having a different amount of randomness added, e.g., if the teacher is whole class teaching then each student will be allowed to chat (state 3) for (4 + random 3) seconds, set on an individual basis;
- additional random threshold for when to tell the class off (state 5) depending on number of students in the class;
- movement in front of class (which alters distance from individual students);

- the choice of a free location near the student being helped.

Table A-1 Teacher-related model constants

Constant name (Abbrev.)	Value	Description
TEACHER-MIN-HELP-WAIT	240 s	minimum time (seconds) between teacher helping a student again; also used to generate a random in [0,240] to add to this minimum
TEACHERS-MISBEHAVIOUR-TIME-THRESHOLD (TMTT)	600 s	the period over which student misbehaviour is remembered by the teacher – this is used to cater for persistent but intermittent misbehaviour (low-level disruption)
TIME-LAST-WITH-TA-THRESHOLD	300 s	the minimum amount of time between teacher-TA interactions
LESSON-SECTION-HELP-DELAY-PARAMETERS	λ=160, k=0.86, max=650 s	when students have been set to work, the teacher observes for a random amount of time before offering 1-to-1 help; a cumulative Weibull curve fitted to empirical data: $1 - \exp(-(^{t}/_{\lambda})^{k})$

Teachers were observed to move around the classroom. In general, teachers appeared to either:

- occupy their designated seats;
- stand in the front or to the side of the class and move randomly from there;
- pace across the front of the class;
- move to individual students and back to their preferred neutral position.

As explained in section 3.5, the simulation code therefore needed to generate realistic random movement of the teacher. The position of the teacher matters especially because student behaviour was made to vary according to the proximity of the teacher.

A.2 TA behaviour

An informal description of the TA behaviour rules is shown in the box below. Written from the perspective of the TA, they were thought to embody the selected TA behaviour and activities. If no other rule applies, the final rule has the TA assume the observing state.

The TA Behaviour Rules

1. if the teacher is whole class teaching or disciplining (teacher in state 5, 9, 10 or 12)
   then cancel any current student interaction and move into observing state 7

2. if the teacher is wanting to talk to me or is talking to me (teacher is in state 15)
   then I go into/stay in state 15 (and cancel any current student interaction if necessary)

3. if a student is wanting/having 1-to-1 support (student is in state 15)
   then go into/stay in state 15 - the student will decide when to stop

4. if I am out of class (state 16)
   then if it is time to return
   then move into observing state 7
   else stay out

5. if I am free (not helping anyone, state 7)
   AND the teacher is not talking to me
   AND the lesson plan allows TA support for students (student state 15)
   AND *I feel like offering support*
   then look for a student to help (someone in state 2, 3, 4, 6, 7, 8 or 9
   who has not been helped for *a while*
   if a student was found
   then offer help to him/her by adding them to the TA-Allocated-To list
   go into state 15
   (the student will enter state 15 when they next consider what to do)

6. if none of the above rules apply
   then move into observing state 7 (and cancel any current student interaction if necessary)

However, the TA behaviour algorithm operates quite differently from the teacher and student algorithm. Although TA behaviour can be viewed in a similar way to that of a teacher, as a job description, observations during the case study led to the viewpoint that the TA was more like a resource utilized by the students as required. TAs first respond to the teacher, then the students and only then act proactively. So, even though the TA is an autonomous agent, the teacher and students control the TA to a great extent. To reflect this, rules 1 and 2 were placed in the teacher ruleset and rule 3 in the student ruleset. The teacher or student sets an anticipated state duration for their contact with the TA and manages situations where that contact would be ended prematurely. The teacher and students also run code to organize moving the TA agent.

The TA is also affected by the lesson plan and is available for student support only when the lesson plan indicates this, normally when the students are expected to work alone (state 8) or with others (state 9 or 14). The rules required several variables and constants. Some of these apply at the agent level, some to the whole lesson. Some constants apply to all lessons - they are model constants. Figure A-2 summarizes the required data and Table A-2 defines the TA-related model constants. Besides these variables and constants, the TA was given additional attributes to enable useful animation and reporting.

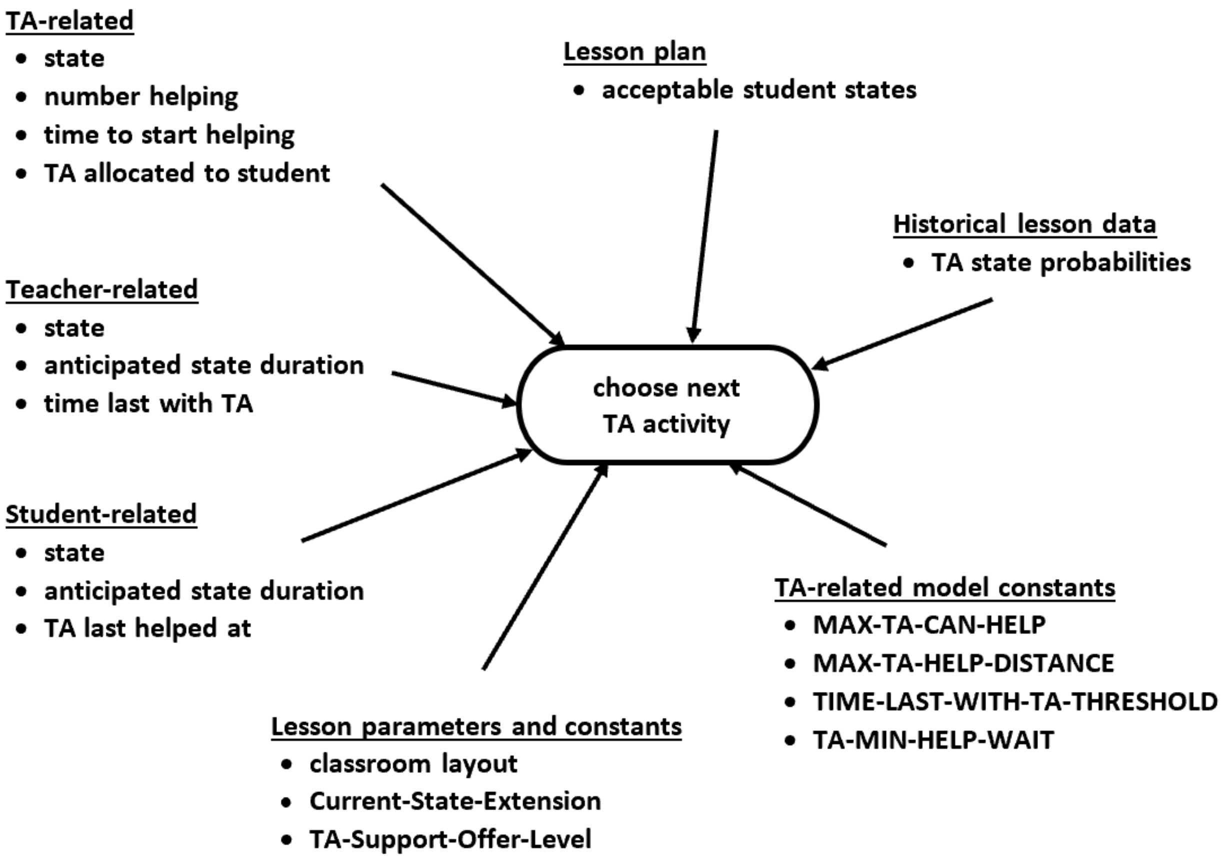

Figure A-2 The data required for TA agent behaviour modelling

Table A-2 TA-related model constants

Constant name	Value	Description
MAX-TA-CAN-HELP	3	maximum number of students the TA can help simultaneously; based on observations this has been set to 3; a student will not solicit assistance if the TA is already helping this many students
MAX-TA-HELP-DISTANCE	2 model units	a student can join another student being helped by the TA only if close enough; based on observation, this is currently set at 2 model units, approx. 2.4 m.
TIME-LAST-WITH-TA-THRESHOLD	300 s	the minimum amount of time between teacher-TA interactions
TA-MIN-HELP-WAIT	180 s	minimum time (seconds) between TA helping a student again; also used to generate a random in [0,180] to add to this minimum

To make TA behaviour more realistic, stochasticity was added in various places, allowing the thresholds to vary within a lesson. For example, in rule 5:

- *I feel like offering support* depends on the values of TA-Support-Offer-Level, the model constant TA-MIN-HELP-WAIT and a uniform random.
- the calculation of *a while* uses the current time, the time TA-last-helped-at, the values of the constant TA-MIN-HELP-WAIT and a uniform random.
- the choice of a free location near the student being helped

A.3 Student behaviour

The algorithm for student agents is more complex than those for the teacher and TA. To summarize the flowchart shown in Figure 3-4, the main stages in student decision-making are:

1. Am I forced into a state?
2. Am I waiting for someone to respond to me?
3. Am I forced to leave my current state?
4. What states are currently prohibited?
5. Recalculate all state scores and choose a state.

This section provides some details about these stages, but first summarizes the inputs required for the rules.

A.3.1 The student variables, constants, attributes and parameters

The rules require the existence of several variables and constants. Some of these apply at the agent level, some to the whole lesson. Some constants apply to all lessons - they are model constants. Figure A-3 summarizes the required data. Besides these variables and constants, the student agent was given additional attributes to enable useful animation and reporting.

The student-related model constants are explained in Table A-3. Values were chosen based on lesson observations and repeated simulation trials. To give an idea of the level of detail that is required by the model, Table A-4 explains the purpose of some of the student variables.

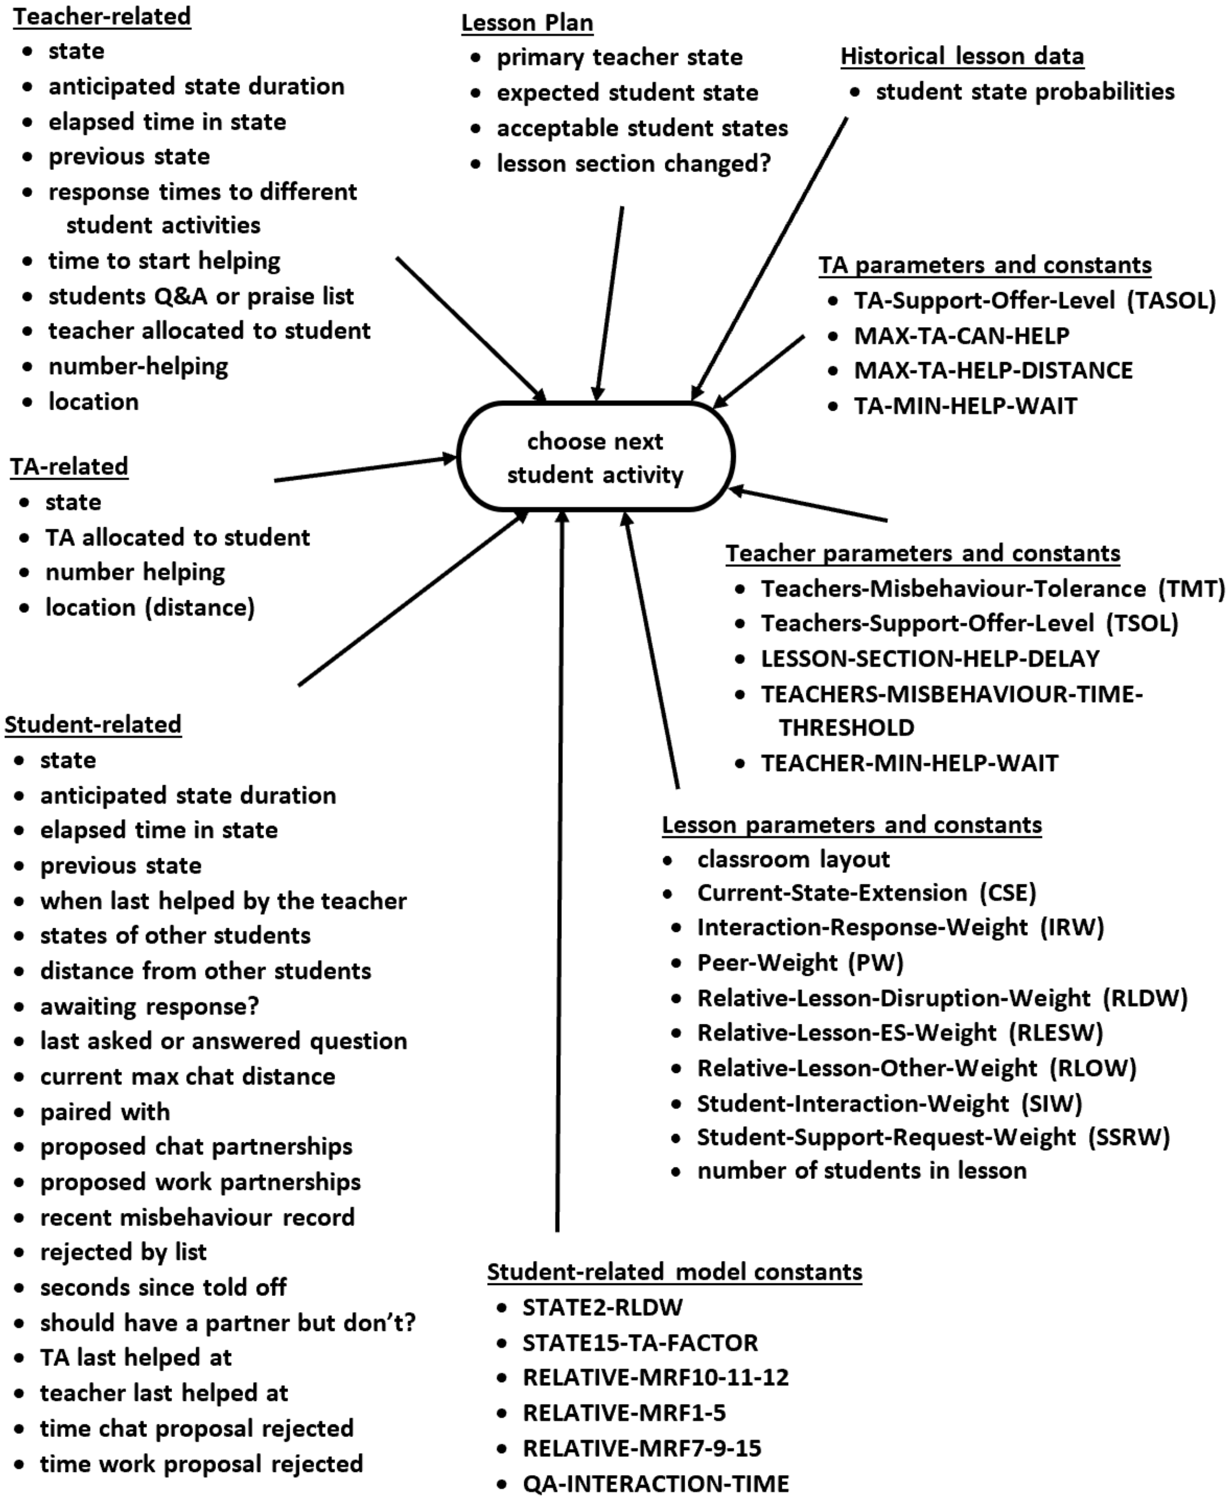

Figure A-3 The data required for student behaviour modelling

Table A-3 Student-related model constants

Constant name	Value	Description
STATE2-RLDW	0.05	disruptive states are adjusted using a single Relative-Lesson-Disruptive-Weight; however, not all the disruptive states should be equally impacted by the RLDW value; state 3 is considered to have a weighting of 1. Relative to this, state 2 (a more seriously disruptive state) is assumed to have a different response.
STATE15-TA-FACTOR	3	the increase in the likelihood of a student joining in with another student being helped by the TA (state 15) compared to just asking the TA for help
RELATIVE-MRF10-11-12	0.1	relative to the default effect of the teacher's discipline, misbehaviour is less likely when the teacher is in state 10, 11 or 12 – hence the misbehaviour reduction factor is further reduced (x 0.1)
RELATIVE-MRF1-5	0.01	relative to the default effect of the teacher's discipline, misbehaviour is highly unlikely when the teacher is in state 1 or 5 – hence the misbehaviour reduction factor is further reduced (x 0.01)
RELATIVE-MRF7-9-15	0.9	relative to the default effect of the teacher's discipline, misbehaviour is a little less likely when the teacher is in state 7, 9 or 15 (all implying some observation) – hence the misbehaviour reduction factor is slightly further reduced (x 0.9)
QA-INTERACTION-TIME	240 s	to prevent students constantly asking and answering questions a threshold is used

Table A-4 Student agent variables

Name	Description
max3distance	distance of furthest student chatting with (state 3); a distance of ≥ 2 m is considered as very disruptive
recent-misbehaviour-record	misbehaviour in past N minutes, where N is TEACHERS-MISBEHAVIOUR-TIME-THRESHOLD (TMTT)
rejected-by	list of students who refused to chat or work with me
ShouldHaveAPartnerButDont?	flag indicating that student was paired with someone but that has been unexpectedly terminated
TA-last-helped-at	the time (in seconds) when student was last helped by the TA
teacher-last-helped-at	time last helped individually by the teacher
time3rejected	time when noticed state 3 (chat) proposal was rejected by the proposed student
time9rejected	time when noticed state 9 (work) proposal was rejected by the proposed student

To make student behaviour more realistic, stochasticity was added in various places, such as:

- anticipated state duration (obtained by randomly sampling from the empirically-derived distributions);
- the order in which students processed;
- the choice of the next activity state (explained further in Appendix A.3);
- the choice of someone to interact with, chatting or working;
- when a student can ask the TA for help.

A.3.2 Rules for forced student behaviour

The first step in the student agent decision-making algorithm was to consider if a state was being forced upon the student by the teacher or the TA. These forced changes could be determined by, for example, whether the teacher or TA has targeted the student. This mechanism was implemented using two global lists, Teacher-Allocated-To and TA-Allocated-To. Other circumstances could also force a student to leave their current state, the simplest being that time is up, the anticipated state duration has passed. If a student was not already in the forced state then these circumstances will force the student out of their current state, so any pairings associated with the current state need to be cleared (procedure *ClearAllPartnerships*). In many cases, the anticipated duration of the new forced state will be determined by the teacher.

Rules for forced student behaviour

At each time instance, select and execute the first rule that applies sequentially from the following rules:

1) if I am out of the room (state = 17)
    then if it is time to return (state-duration > anticipated-state-duration)
        then set anticipated-state-duration 0
            go into state 8

2) if the teacher is telling someone off (State-Of-The-Teacher = 1)
    AND it is me (my ref is in the Teacher-Allocated-To list)
    then if I am not already in state 1 then ClearAllPartnerships
        clear my recent-misbehaviour-record
        set anticipated-state-duration = 0 (this will actually be determined by the teacher)
        go into state 1

3) if the teacher is disciplining the class (State-Of-The-Teacher = 5)
    then if I am not already in state 5 then ClearAllPartnerships
        clear my recent-misbehaviour-record
        remove my ref from the Teacher-Allocated-To list
        set anticipated-state-duration = 0 (this will actually be determined by the teacher)

go into state 5

4) if the teacher is answering a question (State-Of-The-Teacher = 10)
   AND it is my question (Student10OR11 = my ref)
   then if I am not already in state 10
   then ClearAllPartnerships (even though this clears my link with the teacher)
   reinstate my link with teacher (set Student10OR11 = my ref)
   remove my ref from the Teacher-Allocated-To list
   set anticipated-state-duration 0
   go into state 10

5) if the teacher is praising someone (State-Of-The-Teacher = 11)
   AND it is me (Student10OR11 = my ref)
   then if I am not already in state 11
   then ClearAllPartnerships (even though this clears my link with the teacher)
   reinstate my link with teacher (set Student10OR11 = my ref)
   remove my ref from the Teacher-Allocated-To list
   set anticipated-state-duration 0
   go into state 11

6) if the teacher is giving 1-to-1 support (State-Of-The-Teacher = 13)
   AND it is to me (my ref is in the Teacher-Allocated-To list)
   then if I am not already in state 13
   then ClearAllPartnerships
   set anticipated-state-duration 0
   go into state 13

7) if I am being helped by the TA (state = 15)
   AND the lesson section has changed
   OR the anticipated state duration has elapsed
   (state-duration > anticipated-state-duration)
   OR for some reason the state 15 is no longer possible
   then stop getting more help: remove my ref from TA-Allocated-To list

prevent student getting more help immediately: set TA-last-helped-at now

report no state forced

else go into state 15

8) if the TA has asked me if I want help (my ref is in TA-Allocated-To list)

AND the teacher hasn't started an interaction with the TA in the meantime

(State-Of-The-Teacher $\neq$ 15)

[I am not in state 15 - that has been checked already]

then ClearAllPartnerships

get an anticipated-state-duration

go into state 15

If none of the above rules apply then no state is forced

These rules implement many simplifying assumptions. For example, while either the student can seek 1-to-1 assistance (state 13) from the teacher or the teacher can offer it, only the teacher can conclude that assistance.

A.3.3 Rules for filtering out impossible student states

It was relatively easy to prevent student agents from choosing states that would not be realizable in a real lesson. There is a set of conditions that determine if the state is impossible. If this is the case, the score for that state is set to 0. Table A-5 summarises the assumptions that preclude states from selection (written from the perspective of a student).

Table A-5 Conditions under which the student agent will not choose a state

State	Abbrev.	Conditions under which I WILL NOT/CANNOT choose a state
1	DISC	never: the teacher forces this state on a student

State	Abbrev.	Conditions under which I WILL NOT/CANNOT choose a state
2	MESS	the teacher is whole class teaching (so the student is expected to be listening (Expected-Student-State = 12) and the teacher is assumed to be watching) OR the teacher is disciplining the class (State-Of-The-Teacher = 5) or an individual (State-Of-The-Teacher = 1) OR I am in any of these states: 1 5 10 11 13 14 15 16 17
3	CHAT	I am in any of these states: 1 5 10 11 13 16 17 OR the teacher is disciplining the class (State-Of-The-Teacher = 5) or an individual (State-Of-The-Teacher = 1)
4	NOT	the teacher is disciplining the class (State-Of-The-Teacher = 5) or an individual (State-Of-The-Teacher = 1) OR I am in any of these states: 1 5 10 11 13 16 17
5	CLOFF	never: the teacher forces this state on students
6	PREV	state not modelled
7	REST	I am in any of these states: 10 11 17
8	ALONE	I am out of the room (state 16 or 17)
9	OTHER	I am out of the room (state 16 or 17) OR the teacher is in any of these whole-class states: 5 9 10 11 12 OR I am expected to be listening to the teacher (Expected-Student-State = 12) OR working with someone is not allowed (not listed in lesson plan, so teacher hasn't allowed)

State	Abbrev.	Conditions under which I WILL NOT/CANNOT choose a state
10	EXPR	the teacher is NOT in one of these states: 7 8 10 12 OR I am out of the room (state 16 or 17) OR another student is asking/answering a question (Student10OR11 ≠ "") OR I last asked/answered a question too recently (time since last-asked-or-answered-question < QA-INTERACTION-TIME)
11	APPR	Never: as a model simplification this state was not modelled
12	ATTEN	students are not expected to be listening to the teacher (Expected-Student-State ≠ 12) AND the teacher is not teaching (state 9, 10 or 12))
13	TSUPP	the teacher is in any of these states: 1 5 9 10 11 12 14 16 17 OR the teacher is already helping someone else (Teacher-Allocated-To not empty and it doesn't contain my ref) OR I am not expected to be working alone or with others (Expected-Student-State not 8 or 9) OR I am out of the room (state 16 or 17) OR the teacher helped me too recently (time since teacher-last-helped-at) < 300) OR the teacher has said not giving help (not listed in lesson plan)

State	Abbrev.	Conditions under which I WILL NOT/CANNOT choose a state
14	GSUPP	Never: state not included in final model. However, the following rule was provided for other potential models: The teacher will not help a group I am NOT currently working with someone (state 9) or a group (state 14) OR I am not expected to be working with others (Expected-Student-State not 9 or 14) OR the teacher is already helping someone or teaching or unavailable (teacher is in any of these states: 1 5 9 10 11 12 13 16 17) OR I am out of the room (state 16 or 17) OR the teacher is insisting that students make an initial attempt (Elapsed-Time-In-Seconds < Time-To-Start-Helping) OR the teacher helped my group too recently (time since teacher-last-helped-our-group-at < 180) OR the lesson plan says no group help from teacher

State	Abbrev.	Conditions under which I WILL NOT/CANNOT choose a state
15	TASUPP	there is no TA OR the teacher is teaching or disciplining the class or talking to the TA (any of these states: 5 9 10 11 12 15) OR I am out of the room (state 16 or 17) OR I had help from the TA too recently (time since TA-last-helped-at < TA-MIN-HELP-WAIT OR the TA is helping the maximum number (MAX-TA-CAN-HELP) who can be helped concurrently OR the TA is already helping someone and I am too far away (MAX-TA-HELP-DISTANCE) to join in OR the lesson plan says no help from the TA
16	OUTTA	Never: as a model simplification this state was not modelled
17	NONE	I have already been out of the classroom 5 times OR I am in one of these states: 10 11 16 OR if another student is currently out (so assuming only 1 out at a time)

A.3.4 The students' response to the teacher's disciplining

The EffectOfDiscipline factor mentioned in section 3.4.2 depends on two functions that allow the effects of disciplining to fade over time and for the influence of the teacher to reduce as the distance from the teacher increases. The function *ToldOffEffect* calculates the effect of time:

$$\text{The ToldOffEffect is} = \begin{cases} 1 - \left(\frac{t}{2400}\right)^{1.5} & t \in [0,2400] \\ 0 \text{ otherwise} \end{cases}$$

where *t* is the number of seconds since the student was last disciplined (either individually or as part of the class being disciplined). This produces a result in [0,1], where 0 indicates a very easy-going situation (t>2400) and 1 indicates very strict discipline (t=0). The effect of being disciplined was made to reduce to 0 after 40 minutes. In reality, students can take their feelings about the events of one lesson into the next or later lessons, but this was not being taken into account.

The function *TeacherDistanceEffect* calculates the effect of the distance from the teacher:

$$\text{The TeacherDistanceEffect} = \begin{cases} 1 - \frac{d}{6} & d \in [0,6] \\ 1 & \text{otherwise} \end{cases}$$

where *d* is the floor distance in model units between the teacher and the student. The scale of the model classroom can vary. For the case study school, where desks were mostly 0.6 m x 1.2 m doubles, a model unit was approximately 0.6 m. This function produces a result in [0,1) where 1 (*d* = 0) would mean that the teacher and student were collocated and the student is definitely not going to misbehave, up towards 1 (*d* = 6 units) as distance increases, meaning misbehaviour is not dampened at all. Fixing the maximum distance at 6 model units (which in the classrooms observed equates to approximately 3.6 m) seemed to produce reasonably realistic results.

These two factors were combined into one factor that reduces the probability of misbehaviour as follows:

$$\text{MisbehaviourReductionFactor} = \sqrt{(1 - \text{ToldOffEffect}) \times (1 - \text{TeacherDistanceEffect})}$$

This results in a factor with range [0,1] where 1 causes no alteration in misbehaviour and 0 would eliminate misbehaviour.

Further adjustments are made because what the teacher is currently doing also has a significant influence on whether a student chooses to misbehave or not. For example, if students chat while the teacher is whole-class teaching that is likely to be treated more severely by the teacher than if the students were supposed to be working in pairs. To establish weightings for these modifications, the empirical lesson data were analysed.

Figure A-4 shows that empirical student disruption and disengagement varied with the state of the teacher. In fact, misbehaviour was proportionately greatest when the teacher was busy helping a student (state 13) and least when whole-class teaching (state 12). An explanation might be that when a student feels that they might be being observed then they are less likely to misbehave/disengage. When the teacher was talking to the TA (state 15) something else may have been happening: neither are helping students, but they are not necessarily observing either.

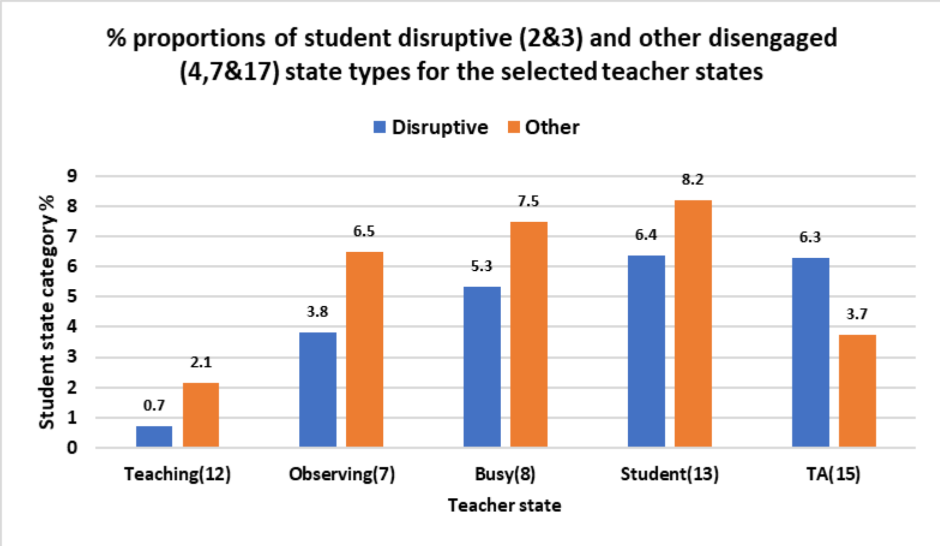

Figure A-4 The empirical % off-task student time for selected teacher states

Figure A-5 shows the empirical proportions of the various student misbehaviour states during the teacher states. These proportions were used to calculate the weightings that adjust the MisbehaviourReductionFactor.

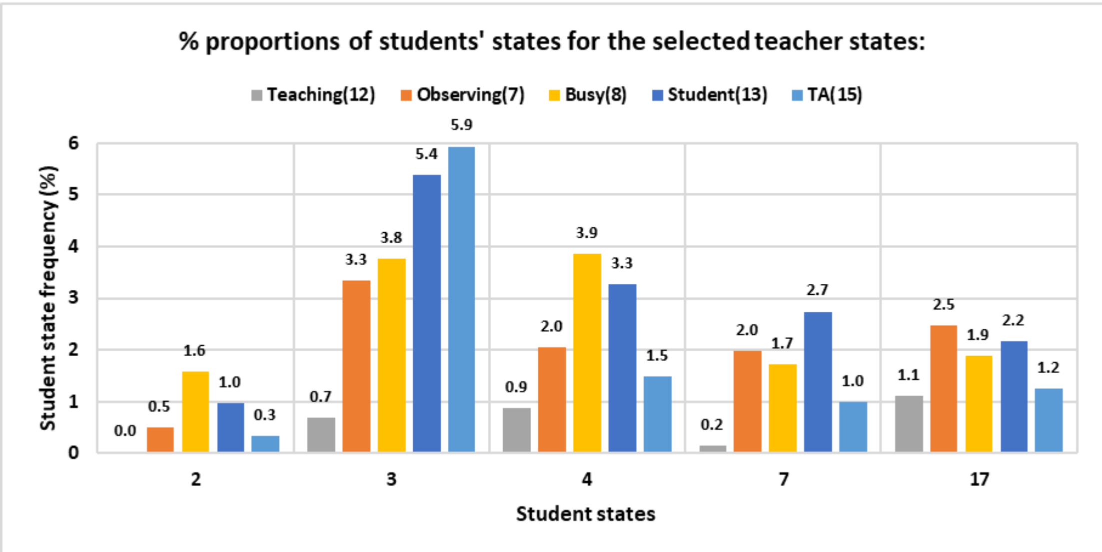

Figure A-5 The empirical % frequency of disengaged student states for selected teacher states

To summarize:

- teacher states 13 and 14 were considered to be the states in which the teacher was least likely to discipline students and were assigned a weighting of 1;
- the other teacher states were assigned multiplicative weightings (derived from the empirical data) that reduced the students' misbehaviour and passive disengagement state scores (states 2, 3, 4 and 7);
- the teacher states were grouped and each group assigned a weighting – the more likely the teacher was to respond, the less likely the students were to disengage:

Teacher state	Weighting
13 & 14	1.00
12, 8, 9, 10 & 11	0.39
7	0.25
15	0.10
1, 5, 16 & 17	0.04

- other teacher states had no additional effect included.

These adjustments result in the *EffectOfDiscipline* factor which reduces students' misbehaviour state scores. The lesson parameters also modify student state scores. The following subsection specifies the actual calculations.

A.3.5 How lesson parameters adjust student state scores

Figure A-6 indicates which student state attribute values are affected by which lesson parameters. The states that cannot be selected by students have been crossed out.

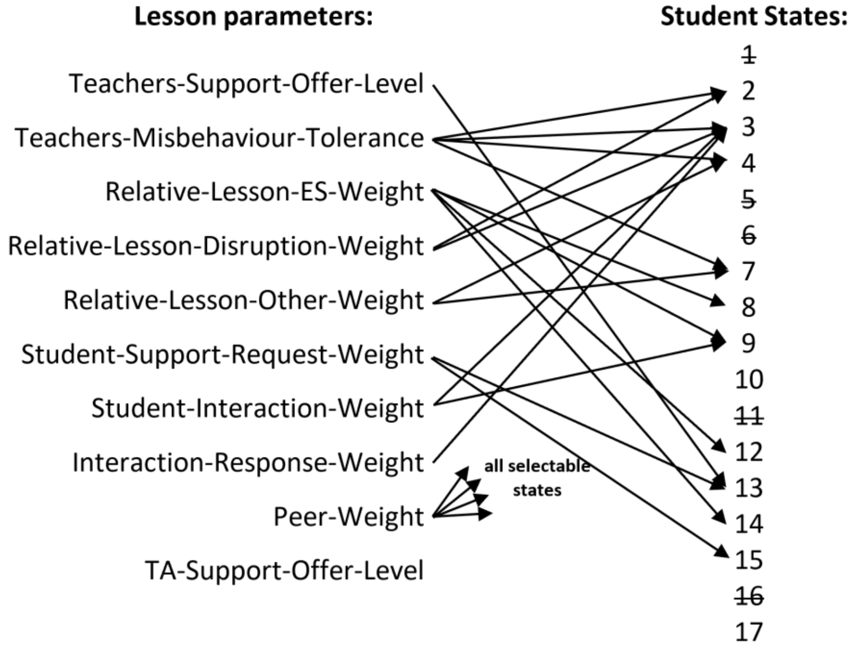

Figure A-6 Which parameters affect which student state attribute scores

The adjustment procedure (*AdjustedStudentStateProbs*) starts with a student's empirical PMF state probabilities for the current lesson section type - with the probability of impossible states set to 0 - and makes the adjustments described in Table A-6. Several abbreviations are used in this table:

tsf	teacher state factor	a factor that takes into account the teacher's current state (5 categories, explained in section A.3.4)
mrf	misbehaviour reduction factor	a factor depending on the time since the student was last disciplined and the teacher's proximity (explained in section A.3.4)
es	expected state	student's current score for the expected or acceptable state
ES	Expected State	student's empirical probability for the expected or acceptable state
sn		student's current score for state n
Sn		student's empirical state probability for state n

Table A-6 The dynamic adjustments to student state scores

State	Formula for new value
17 (out)	S17 x state17-adjustment
Current ES (one of 12, 8, 9 or 14)	ES + RLESW
2 (disturb)	S2 x tsf x (10^{RLDW}) x STATE2-RLDW x mrf
3 (chat)	s3 = (S3 x tsf x (10^{RLDW})) + SIW if s3Proposer (i.e., someone wants to chat with me) then s3 = s3 + IRW s3 = s3 x mrf
4 (disengage)	S4 x tsf x (10^{RLOW}) x mrf
7 (unsure)	S7x tsf x (10^{RLOW}) x mrf
9 (others)	s9 = S9 + SIW if s9Proposer (i.e., someone wants to work with me) then s9 = s9 + IRW
13 (help)	S13 + SSRW + (TSOL / 10)
15 (TA)	if a TA is available then if we are expected to be working with others and someone near me is receiving TA help then s15 = (S15 x STATE15-TA-FACTOR) + SSRW else s15 = S15 + SSRW
all	all non-zero state scores have peer-count fractions added (only states 2,3,4,7,8,9,10,12,13,14,15 could be involved (17 is also excluded)).
all	remove states that are currently impossible or have a negative score by setting their score to 0
all	normalize scores to [0,1]

First individual state-specific adjustments are made then general adjustments to all scores. There is also an adjustment to take into account a change in lesson plan section, because then everyone will behave differently. At the end, peer influences (explained in section 3.4) are included: every state with a non-zero score has a state count fraction times Peer-Weight added. There is also an adjustment for the probability of leaving the classroom (state 17) which is discussed in section A.3.6. Finally, the student state scores are normalised to [0,1].

A.3.6 Modelling students leaving the classroom

State 17, being out of the classroom, is different from the other states. The probability given for state 17 in the students' state PMFs is more precisely the probability that a student will be found in state 17 at any time in the lesson. It is calculated from

$$\frac{\text{number of seconds spent out of the room}}{\text{number of seconds observed}}$$

It includes the probability of leaving the room once, twice, three times, etc. However, the probability of leaving the room on another occasion is much less than the probability of leaving the room just once. This adjustment factor is stored per student in their state17-adjustment attribute.

It would be preferable to have the separate probabilities of a student going out for the first time, a second time, etc., but the empirical state probabilities do not give this directly. Table A-7 shows the number of occasions on which students were observed to leave the room, from 0 up to 7 times. With the extreme example due to one student with a temporary medical condition omitted, an exponential curve was fitted (using a least squares error method) to the four points highlighted in Table A-7 The number of occasions students were observed to leave the room.

Table A-7 The number of occasions students were observed to leave the room

Times left room in lesson	Frequency observed
0	185
1	65
2	14
3	6
4	1
5	1
6	0
7	1
Total	273

Table A-8 Adjusted frequency for leaving room

Times left room in lesson	Frequency observed	Predicted frequency
0	185	-
1	65	-
2	14	15
3	6	4
4	1	1
5	0	0
Total	271	-

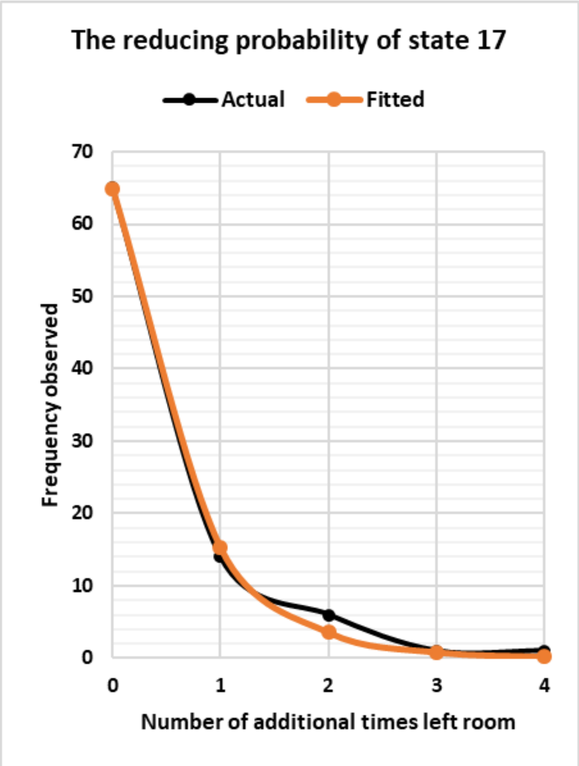

Figure A-7 Fitted curve for predicting the likelihood of a student leaving the room

A decay constant of -1.4502 was obtained, giving an exp(-1.4502) = 0.2345 rate of decline for the number of additional leaving the room events, or equivalently, the decline in the probability of additional leaving the room events.

This exponential relationship made it possible to calculate the probability of a student leaving the room again on the basis of their probability of leaving the room once (the equivalent of an initial condition) and the number of previous times the student has left the room. This was implemented as a multiplicative factor $(\exp(-1.4502))^n$, where n is the number of times the student has left the room already. For example, prob(2nd time) = $(\exp(-1.4502))^1$ x prob(1st time).

Still the probability of leaving the room just once, prob(1 time), needed to be determined. The probability (p) given in the student state 17 PMFs was the sum of the probabilities of leaving once or twice or thrice, etc., i.e.

$$p = \sum_{i=1}^{5} p(i)$$

where p(i) means the probability that the student leaves an i-th time and assuming no-one ever goes out more than 5 times.

Since prob(2nd time) is being modelled as exp(-1.4502) x prob(1st time), and prob(3rd time) is being modelled as $(\exp(-1.4502))^2$ x prob(1st time), etc., the probability given by the PMF is

$$p = p(1) + \sum_{i=1}^{4} (\exp(-1.4502))^i \times p(1)$$

$$\Rightarrow p(1) = p \div \left(1 + \sum_{i=1}^{4} (\exp(-1.4502))^i\right) = p \times 0.7660$$

Hence prob(1st time) = 0.7660p and, since p is known for each student, the other $p(i)$ can be derived for each student from this.

Appendix B The derivation of model constants and functions

Model parameters that have been fixed (based on empirical data or observations) have been termed model constants – these are declared in the simulation code in a separate 'GlobalConstants' module. The CLSM includes a variety of constants and functions that apply to all lesson models. The following subsections give an idea of how values for some of these were inferred from the empirical data, possibly with some guesswork based on experience. The first topic is the generation of realistic state durations.

B.1 State duration modelling

When an agent chooses a new state, an anticipated state duration (ASD) is required. At each simulation step, an agent checks if the elapsed time has reached the anticipated state duration. If so, the agent reconsiders what state to be in. The empirical data were used to fit functions to state duration distributions so that sensible random state durations could be generated (Janssen and Ostrom, 2006; Salgado and Gilbert, 2013). The data for teachers were aggregated to represent state durations for a stylized teacher. Similarly, student data were aggregated, thus representing average student state durations. This was done for every agent state. When the lesson plan determines the primary teacher state to be 12 or 9, then the anticipated state duration is the length of that lesson plan section – it does not need to be drawn from a distribution. Forced student states (1, 5, 11, 13 and 15) have their duration determined by the teacher or TA. For other states, state durations depended strongly on the state of the teacher and what the students were expected to be doing (as specified in the lesson plan section - described in section 3.2). It was therefore decided that anticipated state duration calculations would take into account the expected state of the students:

- state 12 listening to the teacher whole-class teaching
- state 8 working individually
- state 9 or 14 working in pairs or groups

In most cases the form of the fitted curve was a 2-parameter cumulative Weibull distribution function of time *t*: $R(t) = 1 - \exp(-(^{t}/_{\lambda})^{k})$. This cumulative function can be interpreted as giving the probability that the person would leave a state after a given time, *t*. The longer the time, the more likely the person is to have left that state.[23] The inverses of these fitted functions were used to create realistic state durations.

For example, the left-hand chart in Figure B-1 shows the distribution histogram (in 5 s intervals or bins) for students chatting while the teacher is whole-class teaching. The plot on the right shows the cumulative probability mass function points in black, with the fitted function in red:

$$1 - \exp\left(-\left(\frac{t}{8}\right)^{0.54}\right)$$

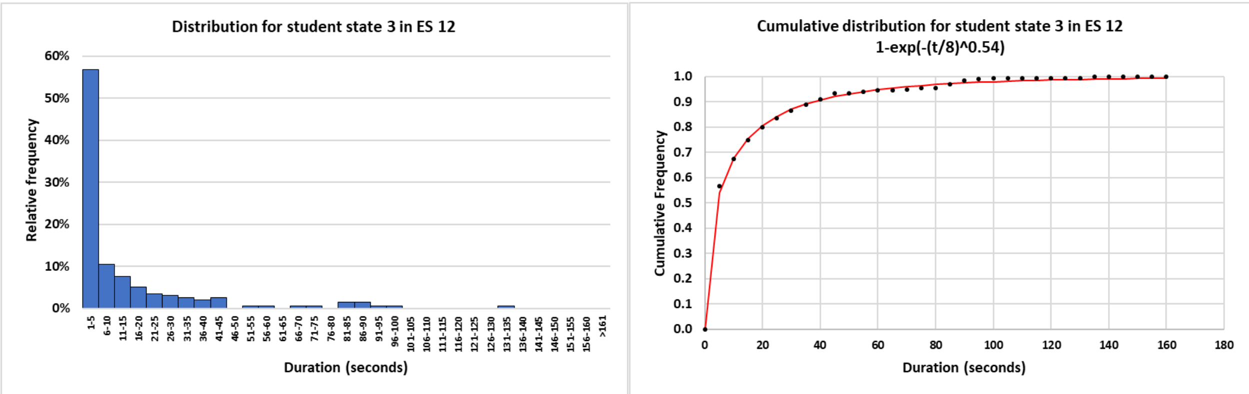

Figure B-1 An example of the type of curve fitted to the distributions in order to produce realistic random state durations

A realistic anticipated state duration, t, was calculated using the inverse function and a uniform pseudo-random $p \in [0,1)$:

$$t = \lambda\left(\ln\left(\frac{1}{1-p}\right)\right)^{\frac{1}{k}}$$

As another example, Figure B-2 shows the data used in the development of the function to determine realistic durations for the time a teacher spends giving one-to-one support to students.

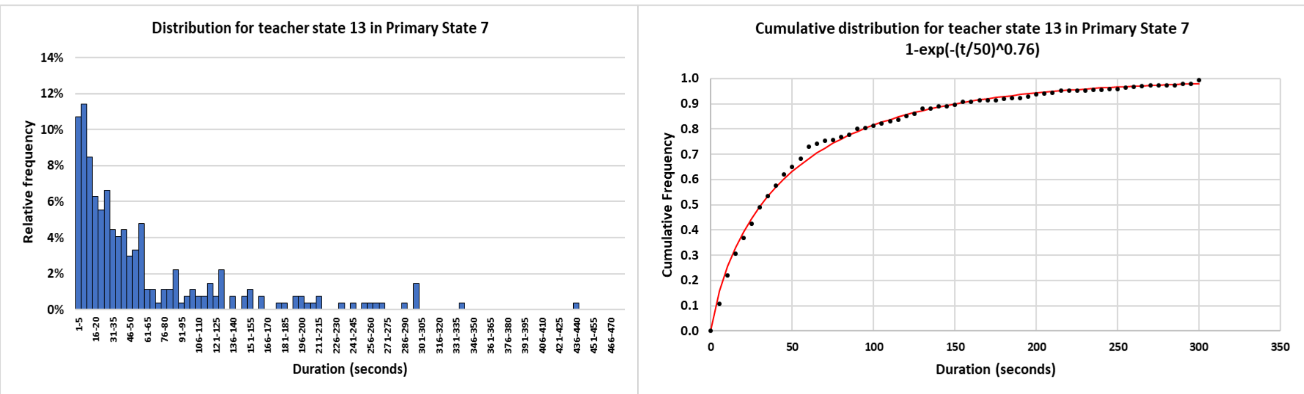

Figure B-2 Fitting a function for the time a teacher gives one-to-one support to a student

[23] This type of function is often called a reliability function, where leaving a state is considered to be some sort of failure event. Survival analysis is another name for time-to-event analysis.

B.2 The Time-To-Start-Helping lesson variable

When students have been set to work at the start of independent working lesson sections, the teacher (and the TA) often observed the students for a random amount of time before offering 1-to-1 help. The reason is to give the students time to try to figure out for themselves how they will proceed. The Time-To-Start-Helping is the clock time calculated for when the teacher and TA will start offering help. It is calculated at the start of each non-whole-class lesson section. A function for generating realistic random times was derived by fitting a cumulative Weibull curve to the cumulative empirical data (for all teachers in all lessons aggregated) – see Figure B-3. This yielded the function: $1 - \exp\left(-\left(\frac{t}{160}\right)^{0.86}\right)$. The inverse of this function is used to calculate a suitable length of time to wait, with a maximum of 650 s.

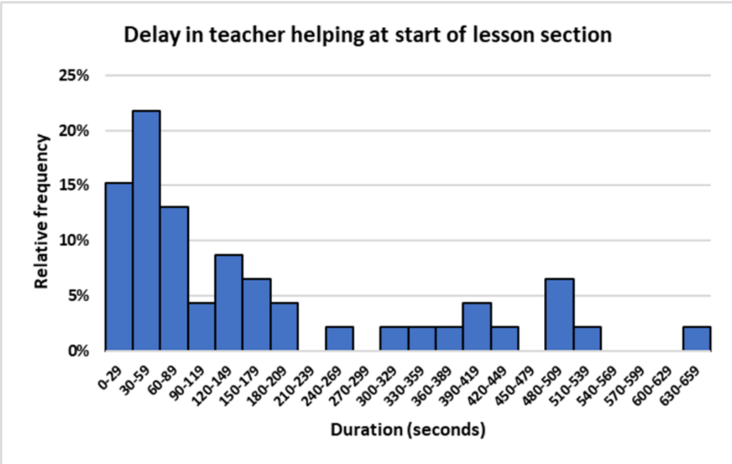

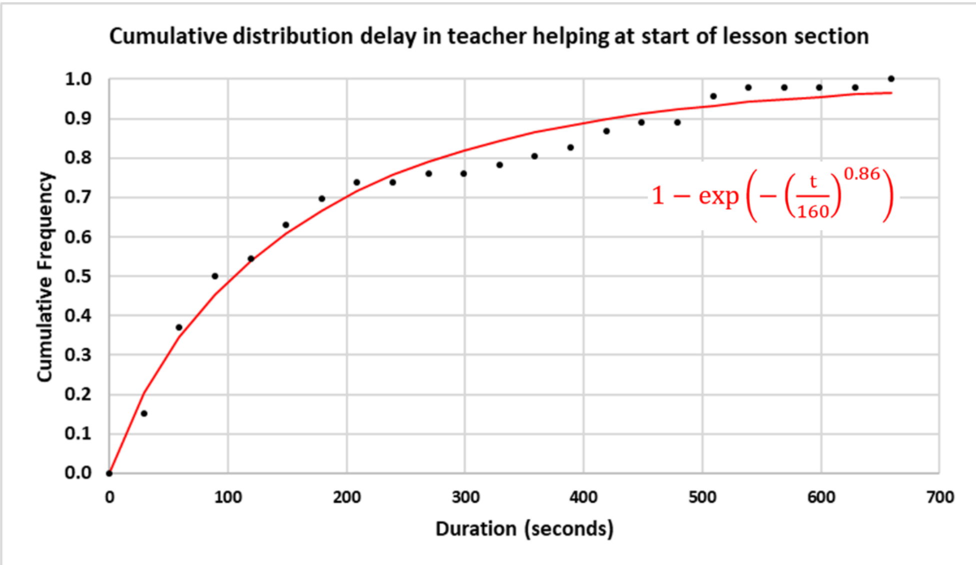

Figure B-3 Finding a suitable function for generating realistic delays in helping students (upper) distribution of empirical times (lower) cumulative distribution of times plus fitted curve

Figure B-4 shows a comparison of the empirical distribution of times and 1000 samples using the function. The results were considered sufficiently realistic.

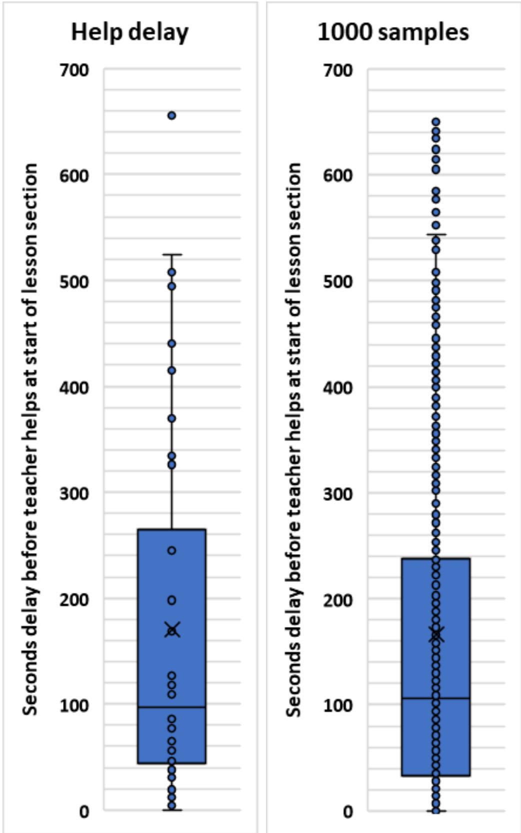

(a) (b)

Figure B-4 A comparison of the empirical delays before helping and the results of the delay function
(a) empirical data (b) 1000 samples from the function

B.3 TA-MIN-HELP-WAIT and TEACHER-MIN-HELP-WAIT

When the teacher or TA gave one-to-one support to a student there was typically a gap before that student received further help. The need was to set a realistic delay so that a student would not receive repeated, consecutive episodes of assistance. Also, to make the simulation more realistic, this minimum would need to vary randomly – otherwise fixed cycles of behaviour could easily arise. If one student had a history of extensive TA or teacher help, then it was possible that the simulation would use the higher state 13 and 15 scores to generate unrealistic levels of repeated student-TA or student-teacher assistance. Hence a delay was proposed to mimic the observed practice of the TA and the teacher not immediately offering further help.

The empirical data on the teacher and TA delays are shown in Figure B-5. The data for all the lessons were aggregated in order to have sufficient data. Each lesson contained several different lesson sections, and behaviour varied widely across lesson sections and across lessons. The data represented a stylized TA and teacher response.

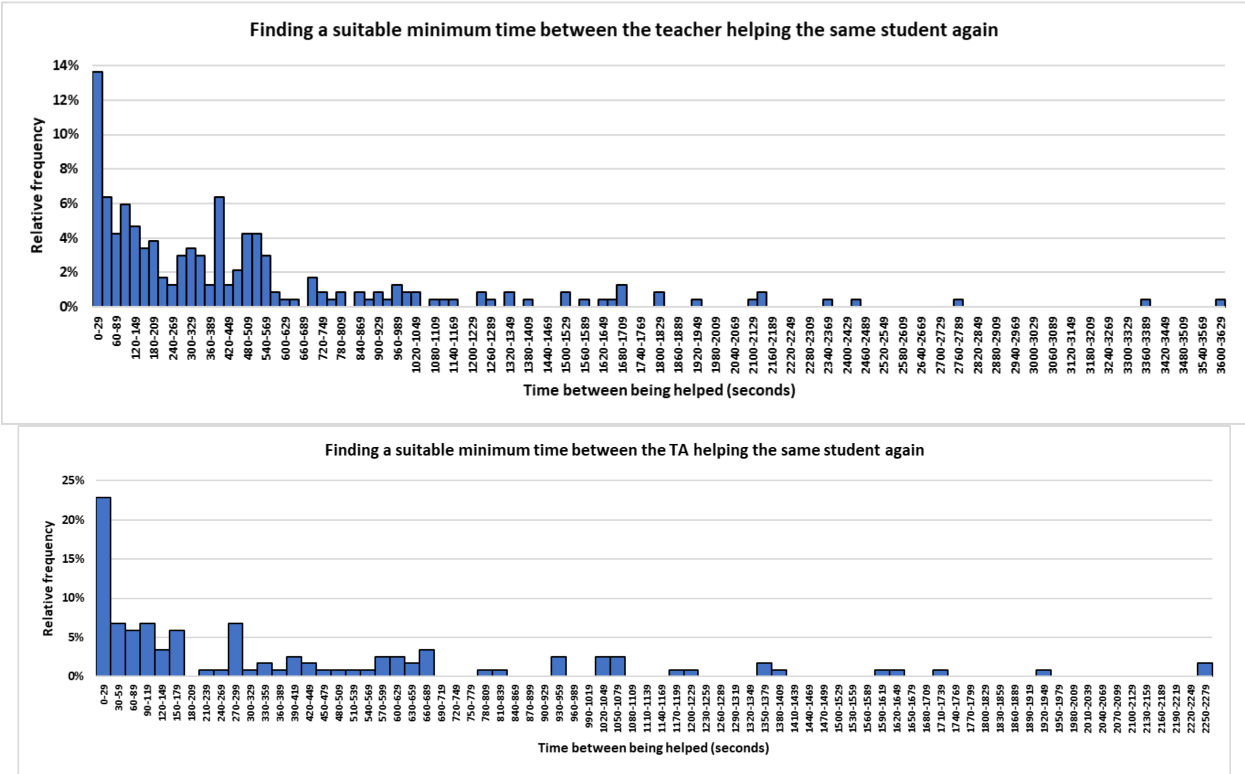

Figure B-5 The empirical delay between episodes of students being helped (upper) the teacher's delay (lower) the TA's delay

They distributions show a wide range of times, exposing several issues with the data and data collection:

- the observer may have thought that an interaction had completed but it continued so it was recorded again – leading to one long episode being split into two;
- the observer may record the same event more than once, being unsure whether it had already been recorded – again leading to shorter intervals;
- the interaction can be interrupted – again leading to shorter intervals;
- sometimes going back to a student is a behaviour-management tactic, used to prevent disengagement (perhaps just by showing that the student is being observed).

Because the data seemed quite unreliable, a pragmatic simplification was adopted: take a fixed minimum duration and add a random amount:

TEACHER-MIN-HELP-WAIT = 4 minutes (240 s) with an additional up to 50% at random.

TA-MIN-HELP-WAIT = 3 minutes (180 s) with an additional up to 50% at random.

B.4 TIME-LAST-WITH-TA-THRESHOLD

There were often interactions between the teacher and the TA, particularly when the teacher was not whole-class teaching. To prevent the simulation generating unrealistically frequent interactions a minimum amount of time between interactions was introduced. There were very little data on this (see boxplot in Figure B-6) but a pragmatic decision was made and TIME-LAST-WITH-TA-THRESHOLD was set at 300 seconds, with no randomness added as the state choice is already random.

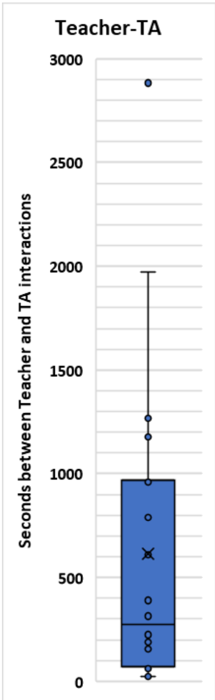

Figure B-6 Empirical data on length of time between teacher and TA interactions

B.5 The teacher's misbehaviour response delay

As outlined in subsection 3.4.1, to make the teacher's response to student disengagement vary realistically, the teacher's response is delayed by an amount that varies according to the state of the teacher and the type of student disengagement. When the teacher is very vigilant, the response time is shorter; when students are being more disruptive the response time is also shorter. Only student state 2 (serious disruption), state 3 (chatting) and state 4 (intentional passive disengagement) required disciplining. Table B-1 shows how the combinations of teacher and student states were allocated delays and that these were adjusted by the Teachers-Misbehaviour-Tolerance (TMT) parameter. The meaning of (m + random n) is create the sum of m and a pseudo-random number in [0,n). As outlined in section 3.4.2, the distance between interacting students is a major factor in how disruptive that interaction is for the rest of the class. A student attribute *max3distance* was used to note the maximum distance between a student and the others they are interacting with. A cut-off of two model units (approximately 2.4 m), was used to distinguish between major and minor disruptions. The names of the variables, s23233-response-time and 3031-response-time, are a little obscure – they refer to the data collected in the orange-coloured columns in the Lesson Event Recording Tool (see Figure 1-4 in section 1.5.2) and refer to the range code (3, 2, 1 or 0) for state 3: chatting.

Table B-1 Calculating random response times according to the teacher and student states and separations

	student state and response delay variable		
teacher state ↓	state 2 or state 3 with max3distance > 2 units **s23233-response-time**	state 3 with max3distance ≤ 2 units **s3031-response-time**	state 4 **s4-response-time**
1, 5	TMT × (1 + random 2)	TMT × (2 + random 2)	TMT × (4 + random 2)
10, 11, 12	TMT × (2 + random 2)	TMT × (4 + random 3)	TMT × (6 + random 4)
0, 7, 9, 15	TMT × (3 + random 2)	TMT × (6 + random 4)	TMT × (18 + random 8)
6, 8, 13, 14, 16	TMT × (4 + random 4)	TMT × (10 + random 10)	TMT × (20 + random 20)

Appendix C Model calibration results

It was explained in section 4.1 that each lesson needed to be considered a separate instantiation of the CLSM, which meant that each lesson needed its own set of parameter values. Table C-1 contains a summary of parameterisation results for the seven selected lessons. It is interesting to note that six of the seven parameterisations were found following an initial manual search rather than a coarse grid search. For one of these (Lesson #2), an exception was made to the criteria for selecting the winning parameter set. Although all the metrics for all 500 replications for the best coarse search parameter set stayed within the empirically-established metric acceptability ranges, the distance metrics were large, showing that the simulated lessons were not that close to the empirical lesson. On the other hand, the best manual search parameter set had 5 (1%) of the 500 replications drop below the minimum TH (teacher help time) threshold, but the distance metric scores were a third of the size of the grid search results – indicating that, apart from a few plausible outliers, the manually established parameter set caused the simulation to be much more like the empirical lesson. The decision was to choose the better match with the empirical lesson, with the occasional lesson (1%) having less teacher help than actually happened in the case study lessons.

Table C-1 Parameter estimation results

Lesson# ➔	1	2	3	4	5	6	7
Lesson ID	16062207GEO	16063010SCI	16070508MAT	16070510MAT	16070607MAT	16070608MAT	16070809SCI
TA present?	N	Y	N	(N)*	Y	Y	N
Local search: no. of parameter sets tested	59,049	2 regions 26,244 + 177,147 = 203,391	78,732	116,640	177,147	39,366	59,049
Coarse grid search: no. of parameter sets tested	177,147	177,147	177,147	177,147	177,147	177,147	177,147
Selected parameter set (method)	57942-local	11626-1st local	48541-local	91657-local	98505-local	26354-local	85126-coarse
Mean distance over 500 replications (between simulation & empirical)	397	740	889	735	558	670	417
Replications out of range?	0	5 TH (1.0%)	0	0	0	0	0
Visual stabilization of individual cumulative mean metric values?	Y	Y	Y	Y	Y	Y	Y
Stabilization of overall distance metric visually & numerically within 95% CI?	Y	Y	Y	Y	Y	Y	Y

Lesson# ➔	1	2	3	4	5	6	7
Parameter estimation success?	**Y**	**Y**	**Y**	**Y**	**Y**	**Y**	**Y**
Parameters:							
Current-State-Extension	0.8	0.6	1.5	0.9	0.9	1.1	0.7
Teachers-Support-Offer-Level	0.6	0.44	0.1	0.4	0.5	0.35	2
Teachers-Misbehaviour-Tolerance	50	17	60	50	40	75	15
Relative-Lesson-ES-Weight	0	0	0.2	-0.2	-0.1	0.1	1.3
Relative-Lesson-Disruption-Weight	1.5	0.97	1.6	1	0.33	1.49	1.5
Relative-Lesson-Other-Weight	0.8	0.89	1.1	0.82	0.8	1.31	1.5
Student-Support-Request-Weight	0	0.05	-0.02	0	0.01	0.01	0
Student-Interaction-Weight	0.2	0.1	0.05	0.07	0.1	0.1	0.5
Interaction-Response-Weight	0.2	0.2	0.35	0.15	0.1	0.2	0.5
Peer-Weight	0.2	0.0	0.1	0.2	0.25	0.2	0.2
TA-Support-Offer-Level	0	2	1	0.01	4	5	0

* Note that in Lesson #4, although there was a TA, the TA was not operating as usual. At the teacher's instruction, the TA provided no help except to one student (and that was during whole-class teaching), apart from a few minutes to one other student. Effectively there was no TA.

Appendix D Analyses of relative importance of the lesson parameters

Further to the discussion in section 7.3, the following charts show the results of using the SPSS package's linear modelling procedures to calculate normalized predictor importance scores for each of the selected lessons. The SPSS software used the *leave-one-out* method in which one predictor at a time is removed from the final full model and the result is ranked on the residual sum of squares. The value obtained is the normalized, relative importance of each parameter. This method enabled interactions and correlations to be taken into consideration. The upper plot for each lesson shows the lesson parameters most influential in the 80-replication rejection stage; the lower plot shows the most influential parameters in terms of finding values that match the empirical lesson values. Where no parameter was found to be a significant predictor there is no dark-blue bar.

<table>
<tr><th>Lesson</th><th>Predictor Importance (from SPSS Linear Modelling)</th></tr>
<tr><td>Lesson #1

16062207GEO</td><td>

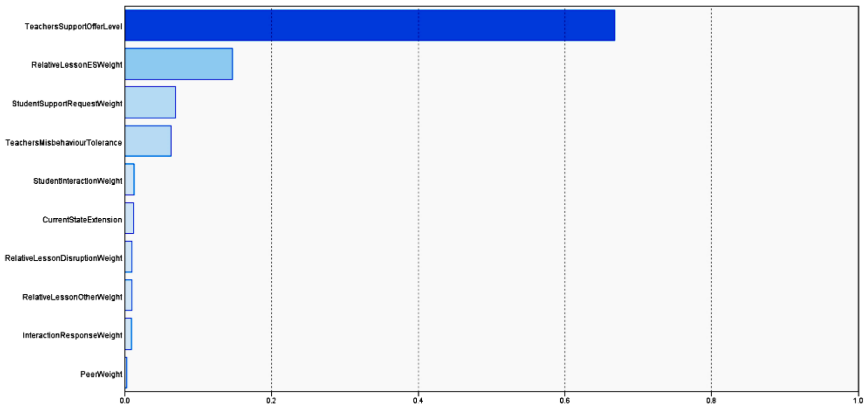

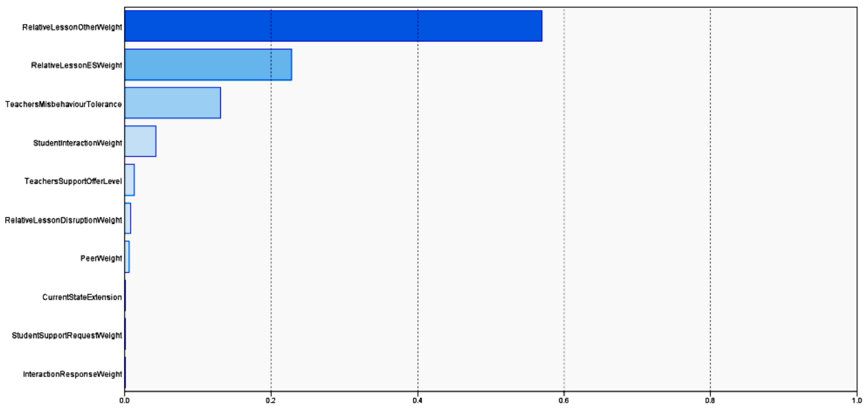

</td></tr>
</table>

<table>
<tr>
<td>Lesson #2

16063010SCI</td>
<td>
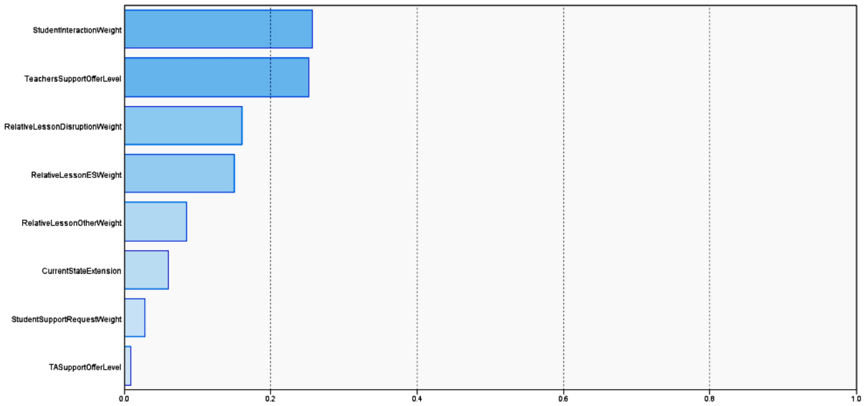

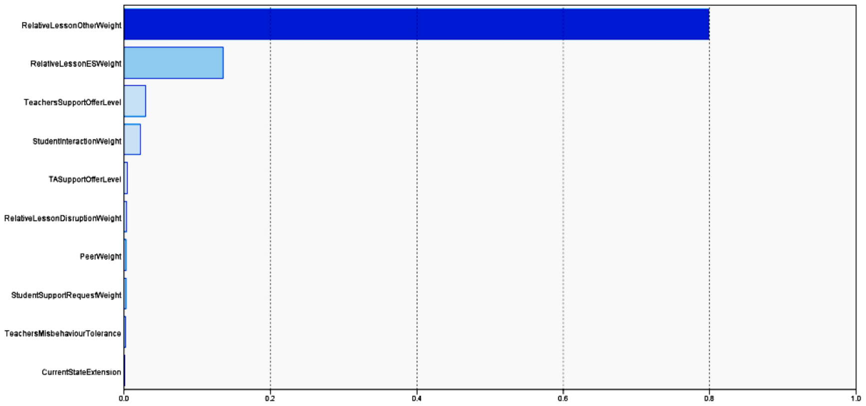

</td>
</tr>
<tr>
<td>Lesson #3

16070508MAT</td>
<td>
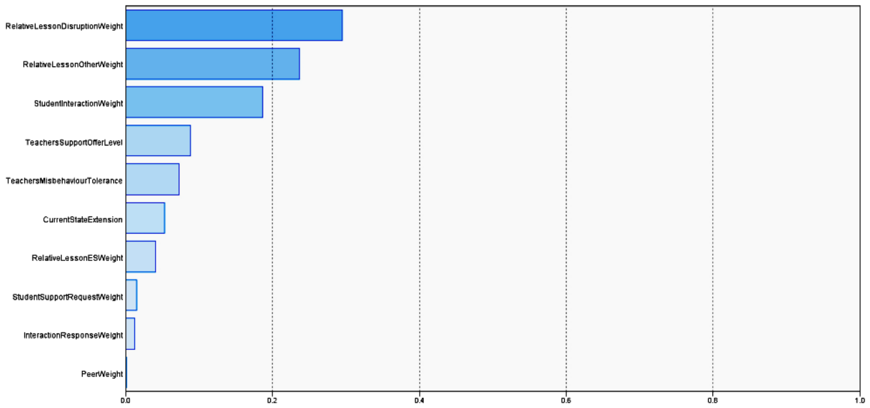

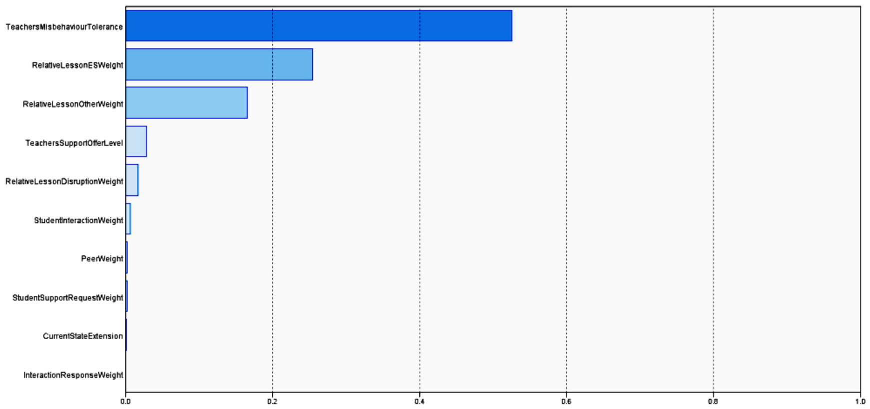

</td>
</tr>
</table>

<table>
<tr>
<td>Lesson #4

16070510MAT</td>
<td>
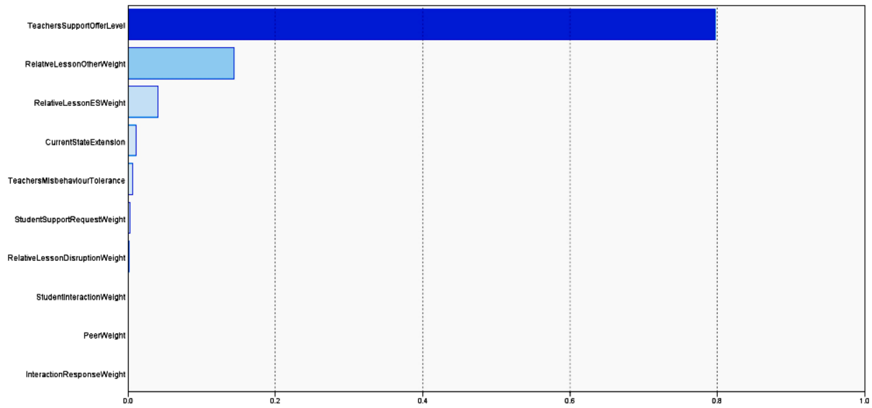

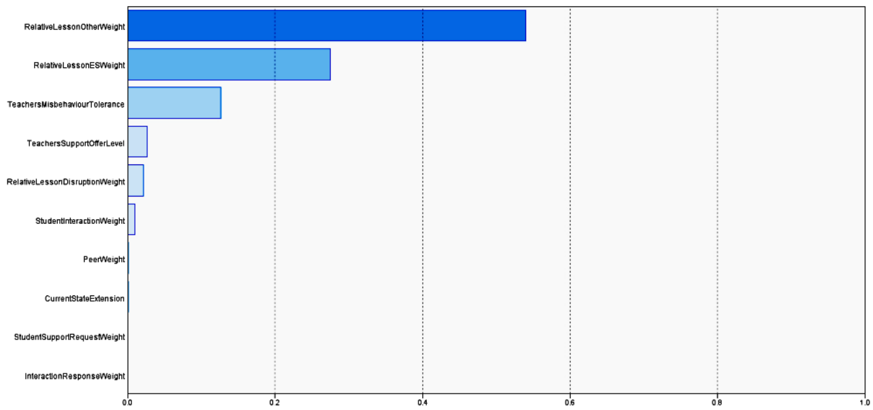

</td>
</tr>
<tr>
<td>Lesson #5

16070607MAT</td>
<td>
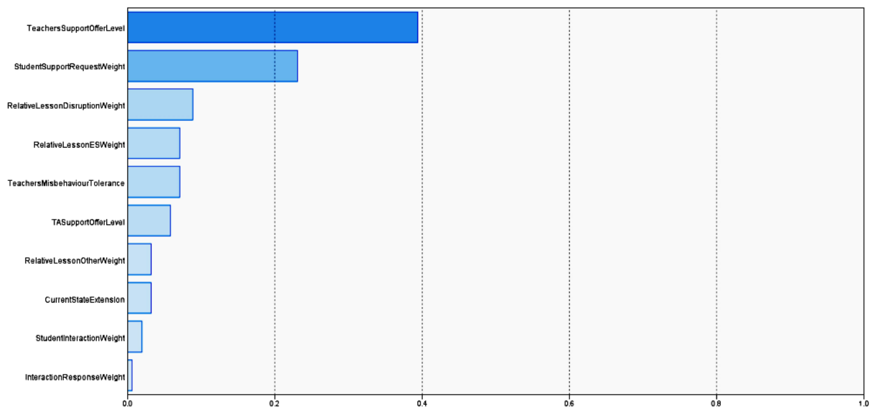

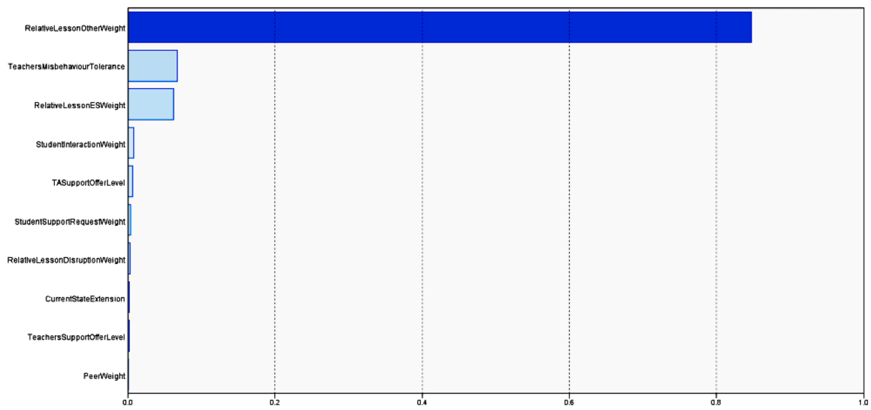

</td>
</tr>
</table>

Lesson #6 16070608MAT	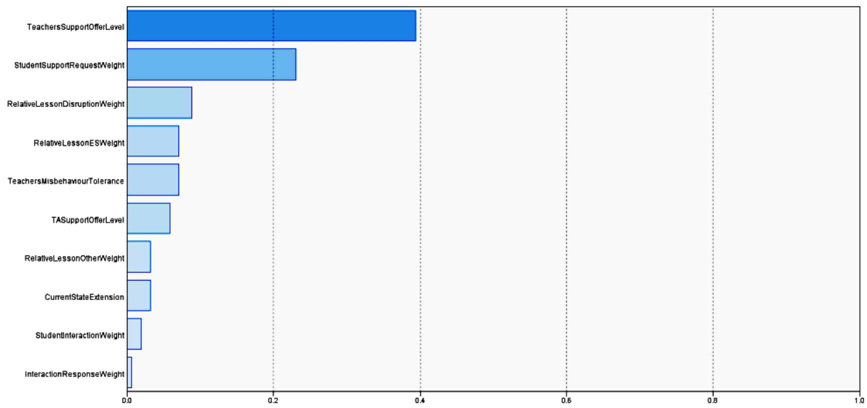 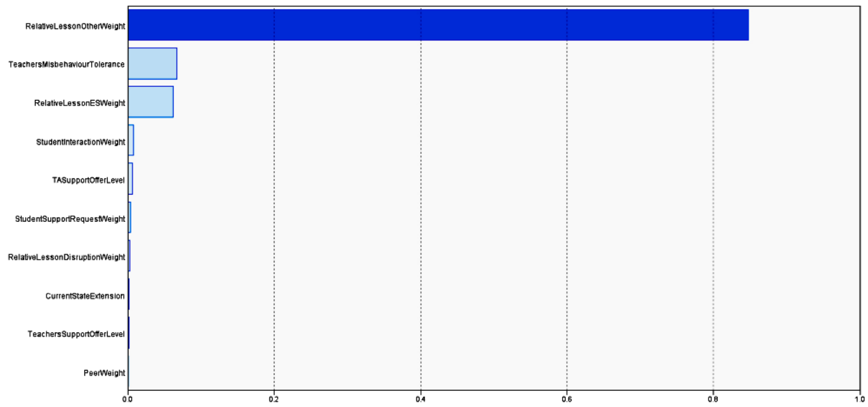
Lesson #7 16070809SCI	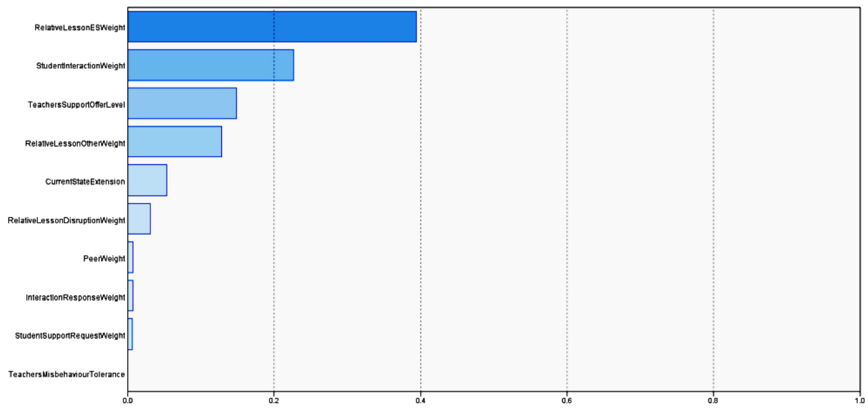 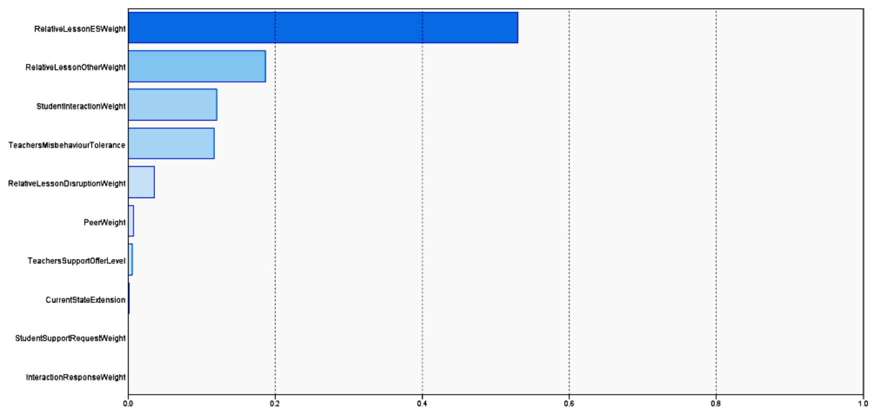

Figure D-1 SPSS analyses of the relative influence of the lesson parameters

The SPSS *syntax* code that generated the plots is shown in the box.

SPSS syntax for the automatic linear modelling:

```
LINEAR
  /FIELDS TARGET=count6_sum INPUTS=CurrentStateExtension
TeachersSupportOfferLevel
    TeachersMisbehaviourTolerance RelativeLessonESWeight
RelativeLessonDisruptionWeight
    RelativeLessonOtherWeight StudentSupportRequestWeight
StudentInteractionWeight
    InteractionResponseWeight PeerWeight TASupportOfferLevel
  /BUILD_OPTIONS OBJECTIVE=STANDARD USE_AUTO_DATA_PREPARATION=TRUE
CONFIDENCE_LEVEL=95
    MODEL_SELECTION=FORWARDSTEPWISE CRITERIA_FORWARD_STEPWISE=AICC
REPLICATE_RESULTS=TRUE SEED=54752075
  /ENSEMBLES COMBINING_RULE_CONTINUOUS=MEAN COMPONENT_MODELS_N=10.

LINEAR
  /FIELDS TARGET=distance INPUTS=CurrentStateExtension
TeachersSupportOfferLevel
    TeachersMisbehaviourTolerance RelativeLessonESWeight
RelativeLessonDisruptionWeight
    RelativeLessonOtherWeight StudentSupportRequestWeight
StudentInteractionWeight
    InteractionResponseWeight PeerWeight TASupportOfferLevel
  /BUILD_OPTIONS OBJECTIVE=STANDARD USE_AUTO_DATA_PREPARATION=TRUE
CONFIDENCE_LEVEL=95
    MODEL_SELECTION=FORWARDSTEPWISE CRITERIA_FORWARD_STEPWISE=AICC
REPLICATE_RESULTS=TRUE SEED=54752075
  /ENSEMBLES COMBINING_RULE_CONTINUOUS=MEAN COMPONENT_MODELS_N=10.
```