3 Methodology and tools

3.1 Introduction

The previous chapter presented the context of the research. Providing a simple definition of management and presenting demand side management, the chapter also detailed the approaches traditionally used, exclusively qualitative or quantitative. The description of the actors involved in water management in England and Wales, and the scenarios for the environment present the sketch of the target system, what we will attempt to represent.

This chapter presents the various methods to address the issue of scenario evaluation and the associated tools. Following a reflection on the way to represent the issue, it is argued that modelling is an appropriate method. Describing the nature, advantages and limits of the different modelling techniques, it is argued that the properties of the target system lead to a unique choice of modelling tool. It provides the description of the different stages for the research, which starts with two stages. First devising the model / adjusting the parameters, and then evaluating the scenarios and their assumptions.

3.2 Problem representation

The aim here is to assess the various scenarios of water management for households. It is then necessary to define more precisely the possibilities that can be used to represent the object of the study, as well as justifying the selected tool, and describing the resulting model.

In order to communicate the various questions we want to investigate, a representation, or model, is needed. According to the Oxford English Dictionary, a model is "*the representation of a structure*", while modelling is "*the devising or use of abstract or mathematical models*". Therefore, as soon as one tries to communicate about the object of an analysis, or a structure, the representation used is a model. Sometimes fairly accurate, as is the current representation of the orbits of planets, it can also be more vague, as could be an oral description of a landscape for example. Both are descriptions and representations, but the precision and objectivity of the language used are different.

Modelling can be descriptive or analytic. The latter intends to infer some general truth from the representation, while the former generally does not allow that. It could be for example because of the lack of information about the structure that is represented, or the lack of objectivity of the language itself.

Often modelling has been understood as mathematical modelling. The reason is most likely that historically mathematics is the most commonly used form of formal language. Formal languages give precision. Although they are not the only ones in the analytic class of models, they are the most used, for the reasons stated above

3.3 Modelling techniques and model properties

The purpose of the enquiry must drive the selection of the tools involved, as some can be more suitable than others. A model is the representation of a system. As such, it can be devised using common language, for example in surveys, or using more formal languages, such as mathematics, or programming languages. Selecting formal modelling for a research means using formal tools for analysis of the system involved rather than surveys to describe it. They are not exclusive (rather interactive even), as will be shown later.

The reasons for modelling are in general (Kwasnicki (1999)) an attempt to:

- Understand and explain a given phenomenon
- Forecasting (more widely prediction and retrodiction)
- Supporting decisions to achieve well-defined goals
- Design a system with optimal performances

As seen in the water management literature, the most common modelling techniques are qualitative or quantitative. Examples were given in the literature review, mainly with surveys and econometrics. I will now describe the use of these approaches in the case of a formal model and in which measure they are appropriate.

The quantitative class of models is composed of analytical models. Typically using algebra, they offer the possibility to achieve a generality of results. It is also used to find optimal solutions to systems, should they exist. Disadvantages of quantitative models are that they might need to use extreme simplifications or very high level formalism, which is consequently difficult to follow for anybody else than the modeller. In our case, it is also the fact that they refer to ideal systems. Ideal systems are not common in social sciences. Our knowledge of the components and interactions accounting for a phenomenon is often imperfect. Including in a model assumptions that are thought to hold, or assuming that a model is complete when the knowledge it is built upon is not, can only lead to results with a limited validity (if any).

If it is possible there is no better representation of a system than the system itself. The advantage is an unbeatable realism and a straightforward verification. Still, its strength is also its problem. Many systems cannot be studied individually and easily. They might not be replicable, or the timescales involved might be too different from the modeller's, and the parameters for the system could also prove very difficult to tune / constrain.

The simulation approach is the third way. Computer simulations can use expert knowledge from various sources. The experiments are repeatable and controllable. Timescale is easy to set up, although it can generate issues with some computational capabilities. The major inconveniences of simulations are the necessity of running many simulations, the lack of generality, and the derived time constraint.

More importantly simulations are an iterative process. Unlike analytical methods that start from the properties of a state to express a solution, when simulating there is a starting point and rules of evolution that result in a situation with unplanned characteristics. Both approaches can be used, but the former is only suitable if the purpose is to find a (potentially optimal) solution (here a state of the world) with known characteristics. Simulations on the other hand are appropriate when investigating a process and undertaking a dynamical analysis. In cases where cognition and subjectivity play a significant role, the generality of results expected from analytical models can not be achieved. Among other reasons, the presence of cognition makes path-dependency appear in the model.

There is an open question about the validity of modelling and simulating itself. Due to the capacities of today's resources (be they human brains or microchips and memory), a trade-off is necessary. If one wants to represent a society, the agents composing it could be in a continuum from numerous and simple to few and complex.

With simple, numerous agents, absence of sophisticated behaviour could turn the whole set of results into useless numbers, the system obtained being too different from the one observed.

Another option is to create a system with fewer agents, but more elaborate behaviour. Thus the modeller has less data to analyse, which makes the treatment easier. The challenge is to define 'fewer' agents, as there is obviously a limit on how much fewer. When attempting to solve problems with 1, 2 or 3 dimensions, a scientific approach is likely to find the/a solution(s). Increasing the number of aspects or dimensions of the problem could make it too complicated (or not practical) to find these solutions. In some fields of research, objective solutions or equilibrium/a could become difficult to find as the space of possible expands with the dimensions / aspects of the problem. There is also the possibility that a problem becomes too constrained to be solved, and a solution might not exist. In game theory for example, any successful negotiation model is with two players only. The type of model developed in this research is not as basic, and does not have a small, limited set of possibilities.

From the range of methods that could be used, the analytical one is too idealistic for social purposes, and the more qualitative one would involve building a system upon which the modeller would not have enough control.

The conclusion is that to investigate modern socio economic processes, computer simulations seem to be the most promising tool.

Modelling is about representing a system. It is hence reasonable to expect that the outputs of the model are comparable to the observed data. Consequently, if the latter present some statistical properties, one would expect a faithful model to generate data in accordance. Of course, it is not a sufficient condition for validation, although it seems a necessary one.

Social phenomena can present the characteristics of complexity.

Pavard and Dugdale (2000) define a complex system as "*a system for which it is difficult, if not impossible to restrict its description to a limited number of parameters or characterising variables without losing its essential global functional properties*." The properties of such complex systems are: non-determinism, limited functional decomposability, distributed nature of information and representation, emergence, and self-organisation.

It is commonly assumed that the behaviour of a system can be foreseen. But as complexity generates non-straightforward behaviour (Edmonds (2000)), if a system is recognised as complex, this assumption cannot hold. In the present case, the resulting data will be analysed further below and some statistical tests undertaken, to demonstrate that the data does not seem to have the commonly assumed property of normality.

Natural and social systems generate data characterised by the presence of occasional and unpredictable events.

The characterisation of an event in the case of social systems can be done through the changes that occur. The time series or cross-sectional analysis of a specific indicator will reveal changes between two situations, referenced by the stage, or object studied, or location (commonly shown for the number of references in a paper, the values of stocks in finance, the number of links pointing to or originating from a web page web links, or town size distribution).

For example, in the case of modelling water demand, one can consider water consumption as a random variable, and treat these data via some statistical method in order to know more about the eventual underlying distribution. Software like SPSS can ease this task. It can apply a test like the Kolmogorov-Smirnov to a set of data, and assess whether the sample observed could have come from a set with a particular (known) distribution. This includes a test for normality. When putting the water consumption data from the Fairlight region to this test, the results are non-ambiguous. The probability that it is part of a normal sample is very close to zero.

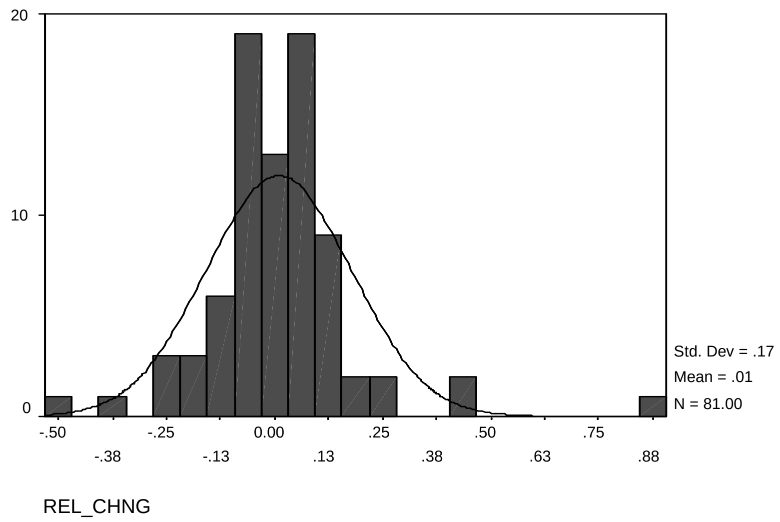

Figure 3: Comparison of distributions for a given sample

Figure 3 displays two distributions. The bar chart is the actual distribution of observed relative changes in a simulation run, while the curve represents the theoretical distribution for such values should the sample be normally distributed. The standard deviation and mean used to generate the theoretical distribution are taken from the sample of observed data. Although this is a standard limitation of the KS test, it results in overestimating the chances that the underlying distribution is found to be normal. This only tends to make the opposite assumption more difficult to prove with this method, which therefore retains its validity in this particular case.

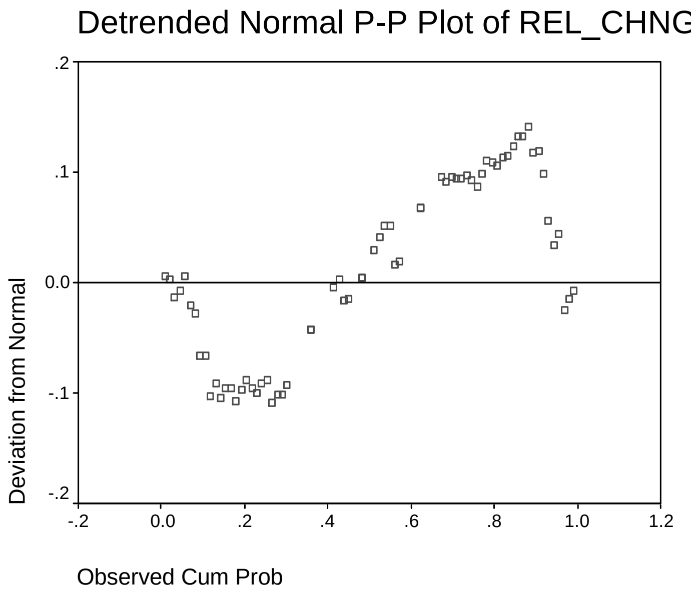

Figure 4: Comparison of de-trended cumulative probability with normal

Figure 4 is a probability-probability (P-P) plot, used to see if a given set of data follows some specified distribution. It is constructed using the theoretical cumulative distribution function of the specified model, and it should be approximately linear if this specified distribution (here normal) is the correct model. As this particular graph is corrected for trends, if the model is correct, the dots would be aligned with the horizontal line.

Not surprisingly the tests tend to confirm the assumption that this sample does not fit a normal pattern. The property shown, a relatively fat tail and thin peak, is known as leptokurtosis. This means that there is in the sample an excess of data values near the mean, and far from it.

Instead of normal distributions, even with time-dependent mean and variance considered by economists, physicists like Per Bak (1997) used the power law devised by Pareto in 1893. If the data observed fit this kind of distribution, the consequences are important. The probability density function of the Pareto distribution has two parameters, α as the “peakedness” parameter, in the interval (0,2], and β as the “skewness” parameter, in the interval [-1,1]. The issue is that these parameters have critical values. When α is equal to 2, the characteristic function of the paretian distribution reduces to that of the normal distribution. But for $\alpha < 2$, there is no finite variance for the distribution, and for α less or equal to 1, there is no finite mean.

Moss (2001) has investigated some different means of generating such a distribution. Three explanations are: a normal distribution with predictable time varying parameters, a stable Pareto distribution with infinite variance generated by self organised critical social process, or a non stable distribution generated by a self organised critical social process.

Because the model self-organises around the critical state and remains around that state thereby to produce power law distributed data of extreme events, this phenomenon was called self-organised criticality[10]. Some necessary conditions in which self-organised criticality (SOC) emerges were summarised by Jensen (1998) as those where:

- Model components (cells, agents, etc.) are metastable in the sense that they do not change their behaviour until some level of stimulus has been reached.
- Interaction among the model components is a dominant feature of the model dynamics.
- The model is a dissipative system.
- The system is slowly driven so that most components are below their threshold (or critical) states most of the time.

The social embeddedness is coined in Granovetter (1985) and defined by Edmonds (1999), as “the extent to which modelling the behaviour of an agent requires the inclusion of other agents as individuals rather than an undifferentiated whole”. It means that formally, it is more relevant to model an agent as a part of the total system of agents and their interactions as opposed to modelling it as a single agent that is interacting with an essentially unitary environment.

[10] Because the model self-organises into the critical state and remain in that state thereby to produce power law distributed data of extreme events.

The model devised must then be able to capture qualitative behaviour, allow quantitative results to be tested against the Environment Agency's own, and describe households' characteristics and behaviour while generating aggregated data that can be tested against commonly accepted theories or observation. All this without being based on paradigms / theories that are not validated.

System dynamics, interviews, surveys, cellular automata amongst other techniques fail on one or more of these aspects. Two methodologies might be able to match these requirements: microsimulations (or microanalytical simulation models), and Multi Agent Systems.

Microsimulations are frequently used to analyse the effects of financial and social policy interventions. In microsimulation, a sample is generated with characteristics similar to those observed, generally using probability distribution, or statistical relationships. The microsimulation model then uses probabilities of transition for individuals from one state to another and then generates aggregate changes in the artificial society.

Individual behaviours are not based upon rules, but upon probabilities. Individuals change because they are "told to", not because of a specific reason. There is no explicit decision-making process. As put by Boman and Holm (2004), "[t]he whole purpose of such models is to represent observables and facilitate policy experiments, sometimes with the help of theory and theoretical concepts; and if the model fails in prediction (as they normally do), there is no other excuse for its construction."

In essence, one could argue that microsimulation is the methodology used by the Environment Agency: a model built upon statistical relationships with careful characterisation of initial population, and transition states.

On the other hand, Multi Agent Systems are based around interactions. Agents represent entities that interact with other agents and with their environment. Decisions they make can be justified from their knowledge and beliefs. Major benefits of this approach are the fact that they have no underlying assumptions and it is possible to validate parts of the model (provided an appropriate indicator can be used), which can make a model useful, even if its “fails” its predictions.

That is why simulation experiments are so far the only way to understand self organised criticality. The SOC literature acknowledges that there is not yet an analytical method to create models that generate the appropriate processes.

3.4 Presentation of Multi Agent Systems

The multi-agent approach originates from Distributed Artificial Intelligence research, where multiple heterogeneous components of a system should interact in order to reach a global goal. The use of entities at the micro level then starts in the ‘70s, with the work of Hewitt, who defined actors, which are interactive, as self contained object that can execute tasks simultaneously with others.

Nowadays, there are different approaches to Multi Agent Systems. One is led by Wooldridge and Jennings, and considers a Multi Agent System a closed one, where the components / agents must strictly behave in a controllable and predictable way (Wooldridge and Jennings (1998)).

Another one consists in representing a system and observing the emerging behaviour at macro level.

While both approaches are bottom-up, *i.e.* start with representing the micro entities of the system in order to represent / generate a bigger one, the kind of systems they can deal with are not equivalent. As expressed in Moss (2000), the former is focused on “tidy” systems, while the latter is used for “messy” systems. The author characterises a system as tidy if its boundaries and the relationships represented are clear and well understood, as well as their carefully monitored development. They are systems that software engineer devise and use, as they allow total control of the process, by limiting for example the interactions, communications or information for the agents involved. The so-called messy systems on the contrary have fuzzy boundaries, and components whose relationships are difficult to represent.

The latter is a more appropriate idea for representing a society, while the former is perfectly suitable for devising Multi Agent Systems that must achieve a specific task.

Having shown that the system that must be represented can not be considered as tidy, due to the presence of interactions, social embeddedness and complexity, one could argue that the development of Multi Agent System would be able to tackle these issues.

The basic element of a Multi Agent System is the agent. It is defined as "a physical or virtual entity

1. which is capable of acting in an environment,
2. which can communicate directly with other agents,
3. which is driven by a set of tendencies (in the form of individual objectives or of a satisfaction/survival function which it tries to optimise),
4. which possesses resources of its own,
5. which is capable of perceiving its environment (but to a limited extent),
6. which has only a partial representation of this environment (and perhaps none at all),
7. which possesses skills and can offer services,
8. which may be able to reproduce itself,

whose behaviour tends towards satisfying its objectives, taking account of the resources and skills available to it and depending on its perception, its representations and the communications it receives" (Ferber (1998)).

A classical opposition was drawn between reactive and cognitive agents: cognitive agents can form plans for their behaviours, whereas the reactive are classically those that just have reflexes. Through the use of endorsement, the agent we are dealing can be considered as cognitive, although Ferber (1998) tries to show how both approaches can converge, while emphasising different aspects: one focuses on the building of individual intelligence whose communication is organised, whereas the other imagines very simple entities whose co-ordination emerges in time without them being conscious of it.

From that broad definition, a Multi Agent System (MAS) is a system comprising the following elements:

1. an environment, that is a space which generally has a volume
2. a set of passive objects (*i.e.* can be perceived, created, destroyed and modified by the agents)
3. a set of agents, which are specific objects and represent the active entities of the system
4. a set of relations linking agents to each other
5. a set of operations making it possible for active agents to perceive, produce, consume, transform and manipulate passive objects
6. operators with the task of representing the application of these operations and the reactions of the world to this attempt at modification.

According to Ferber (1998), the main application of Multi Agent Systems of the moment can be seen as problem solving (as an alternative to centralised problem solving), multi-agent simulation (widely used to enhance knowledge in biology or in social sciences), construction of synthetic world (used to describe some specific interaction mechanisms, and analyse their impact at a global level in the system), collective robotics (defining the robots as a Multi Agent System where each subsystem has a specific goal and deals with this goal only), and kenetic program design (a very efficient modular way to program)

In Artificial Intelligence, the agents are in general characterised as reactive, cognitive, proactive, situated, or communicating for example. In social simulations, the choices are determined by the system that must be represented, together with the choices of the modeller and the experts that have analysed it. Therefore, the agents in this project are only qualified within the standard Multi Agent categorisation when they need to be.

The MAS framework is considered as in between a strict Object oriented approach and a strong agent approach like an AI agents based framework, and a good balance between generality and ease of agent-based application development (Silva, Romao et al. (2001)).

The recent literature has seen a proliferation of Multi Agent-based research. Yet, the current trend among modellers is not a reason to choose a method. It is nevertheless the case that a Multi Agent System has properties that are appropriate to the issues this research is dealing with. First, it allows representing a system in a way that is compatible with the complexity conditions. Second, it is possible to observe emerging patterns of behaviour, via for example an indicator on a macro level. Finally, generation of the model, as a representation of a system, can be done using a participative approach ensuring that smaller elements of the model, as well as interactions amongst them, are devised properly (Barthélémy, Moss et al. (2001), Edmonds, Barthélémy et al. (2002)).

Consequently, using MAS simulations allows a focus on processes, which is a necessity in the study that is presently developed, as information diffusion and behaviour emergence are of high importance in the assumptions involved. That is why a Multi Agent System is used in this work.

The appropriateness of a method does not mean that the modelling process cannot fail, or be biased. The modeller has to ensure that the modelling method itself, *i.e.* the various modelling stages, is as sound as possible.

3.5 Representation of subjective values

In the model developed, agents will make decisions according to their perception of the current situation. For increased realism, this perception must be subjective. Because the underlying logic for decision can be seen as a maximisation of satisfaction, or utility, it has often been represented via a utility function with specific properties, or preferences. But there are other ways to implement such an evaluation process. One of them is endorsements.

One of the strength of endorsements is that the information an agent obtains is not only evaluated according to its nature, but also according to its origin.

While it might be possible to create a utility function that would have all these attributes, it will certainly be complicated. Endorsements are simple.

Information, as well as its origin, is stored in the agent’s memory. When the agent needs to make a decision, it weighs up the information using its personal endorsement weights. This set can be personal or common with other agents, or possibly follow a given distribution.

By combining the origins of the information with the weights representing how important or reliable that source is, an agent can then compute a subjective value for this information. If the information is itself (or triggers) an action, then the agent decides to select the action associated with the highest value (maybe from a combined set of endorsements).

3.6 Modelling stages

Verification and validation of that model will be crucial issues when building a model for social simulations. They must be present at every stage of the modelling process. It is then necessary to express that process and the relevant steps with respect to the targeted purposes of that representation.

3.6.1 Modelling theory

Because there can be some implicit references to the modelling process in the model structure, this is a description of the different stages and the explicit and implicit parts in them.

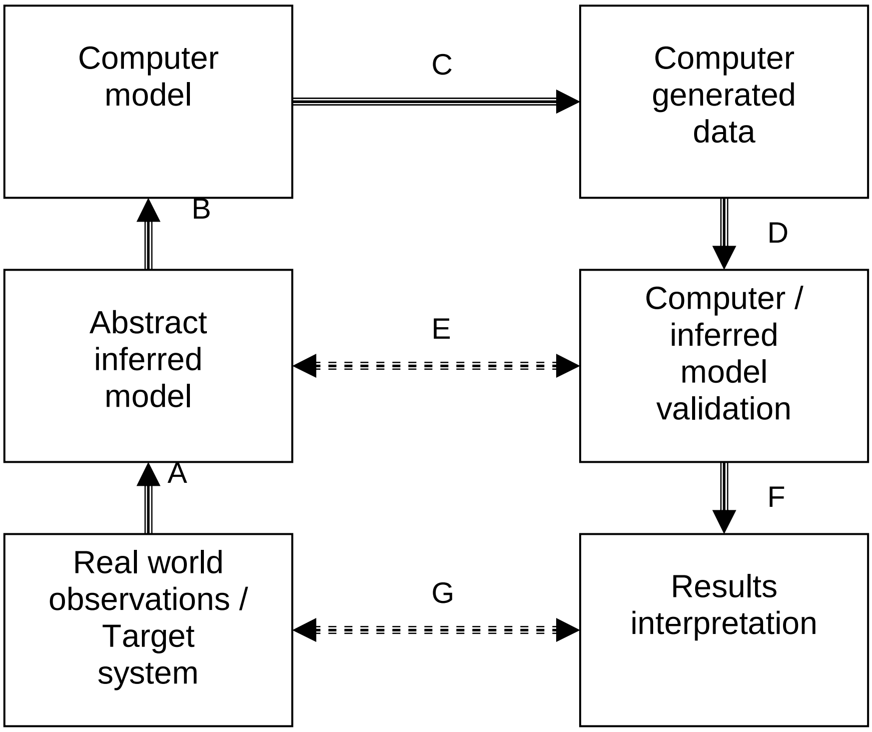

Figure 5: The modelling stages

The different links are named A - G for easier visualisation. Each relates to a modelling step.

Link A is the creation of an abstract model that will represent the parts of interest. Link B is the translation of that abstraction into a computer model (*i.e.* lines of code). Link C is the running of the computer model in order to generate data. Link D is then the analysis of these data eventually with respect to some indicator. Link E represents the validation of these data against the abstract model, and link F is the interpretation for the real world of the conclusion. Link G is the comparison of the generated assumptions, including the potential effects of the results on the “Idea” of the target system.

Some guidance was given by Edmonds (2000) on how to ensure that the complete modelling cycle is as strong as possible. He proposes several criteria in order to help evaluate the rigour of the modelling.

Abstraction, referring to the link A, has to be correctly specified .The parts or aspects of the target system that are supposed to be represented must be explicitly defined, and of course the abstract model must remain relevant.

Design, referring to link B, is the process of writing up the formal model. It has to clearly remain linked to the abstract model, since it is its translation into a more formal representation. Specially, the parts derived from the abstraction and those necessitated by some logical or computational constraints ought to be explicit.

Inference, referring to link D, is in our case the transformation of computer-generated data into a more generic rule, or phenomenon. Caution is here necessary. One must check that the expressed outcomes are effectively a necessary result of the model specification and design. Dependence upon particular parameters or settings has to be looked at.

Analysis, referring to link E, has to be clear. The eventual limitations of the technique should be expressed here, as well as the set up for further testing and replications.

Interpretation, referring to link F, must be justified and relevant.

Application, referring to link G, must verify whether the conclusions applied to the target system are justified, and whether the rest of the steps are sound enough to justify the conclusion itself.

One of the key points to bear in mind is that a model must be thought of as a representation of reality. It is necessary to bring the attention to the fact that it is not *the* representation, but more a way to focus on some parts of it.

This should help the modeller to devise a representation with a method allowing keeping a close (scientific) eye on the different stages.

Simulations are used on many occasions, from hard sciences like physics to software applications in finance, or the game industry. The aims can be very different though, and not only because of the various application fields. In the present case, the model represents an artificial society. Amongst the different uses of simulation in the case of artificial societies described in Hales (2001), two of them are involved in this work: theory building and reverse engineering.

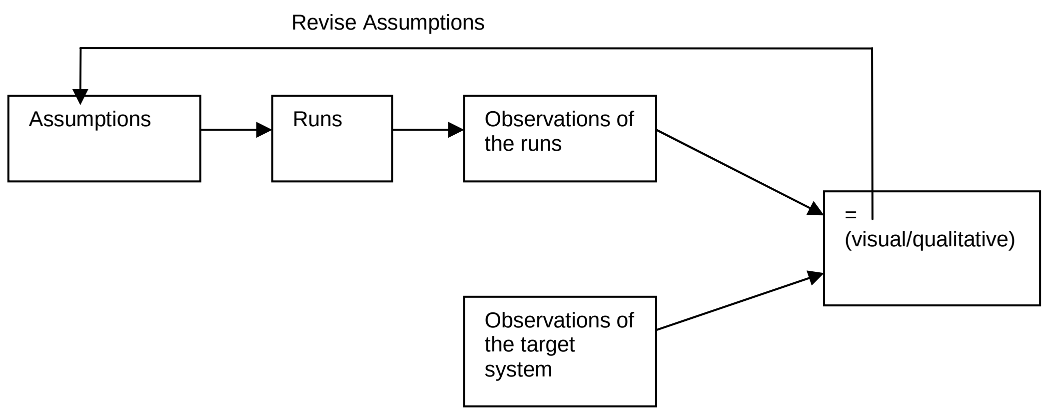

Figure 6: Reverse engineering[11]

[11] Courtesy of Dr D. Hales, original in Hales (2001)

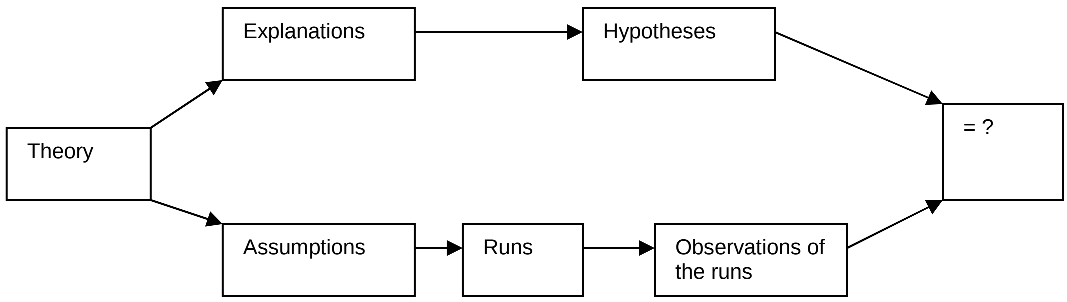

Figure 7: Theory building / testing[12]

The above figures refer to artificial society experimentation, defined as "a set of assumptions used to construct the society; a set of runs produced by execution of a computer program which embodies it; a set of measurements of observations of the runs; a set of explanations (a theory) which attempts to link the assumptions and the observations and a set of hypotheses linked to the set of explanations based on the assumptions and the observations". In our case, the two above figures have to be put together to represent our purpose, *i.e.* building a model and validating our assumptions.

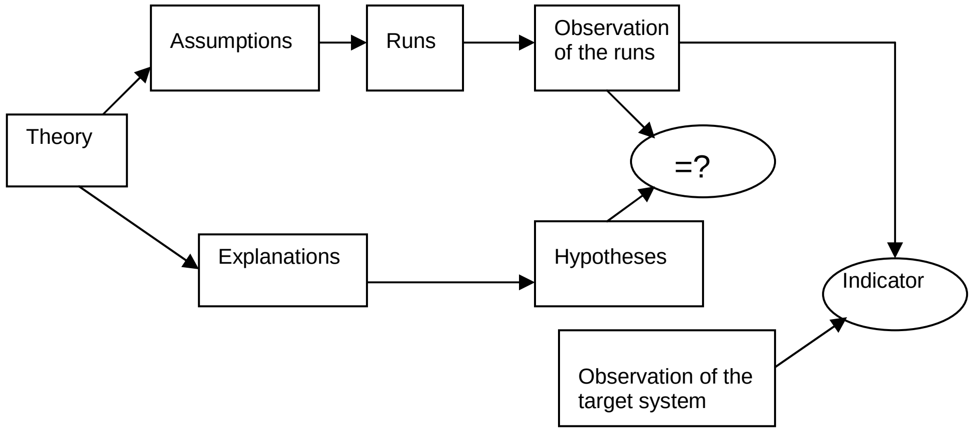

Figure 8: Model building and validation[13]

The indicator refers to the validation itself. This is actually going to be achieved via the comparison of statistical signatures of the distributions. As a preliminary assessment of the closeness of the simulated data and the observed data, some standard parametric statistics can be used to capture

[12] Courtesy of Dr D. Hales, original in ibid
[13] Courtesy of Dr D. Hales, original in ibid

main characteristics, but the nature itself of the distribution will be tested using non-parametric statistics.

These diagrams are trying to visually express the aims of this study. They show that there are two different steps. First, the reverse engineering part is the creation of a model that can generate specific scenarios. Second, the theory testing stage involves running the simulations in order to put these alternatives to the test.

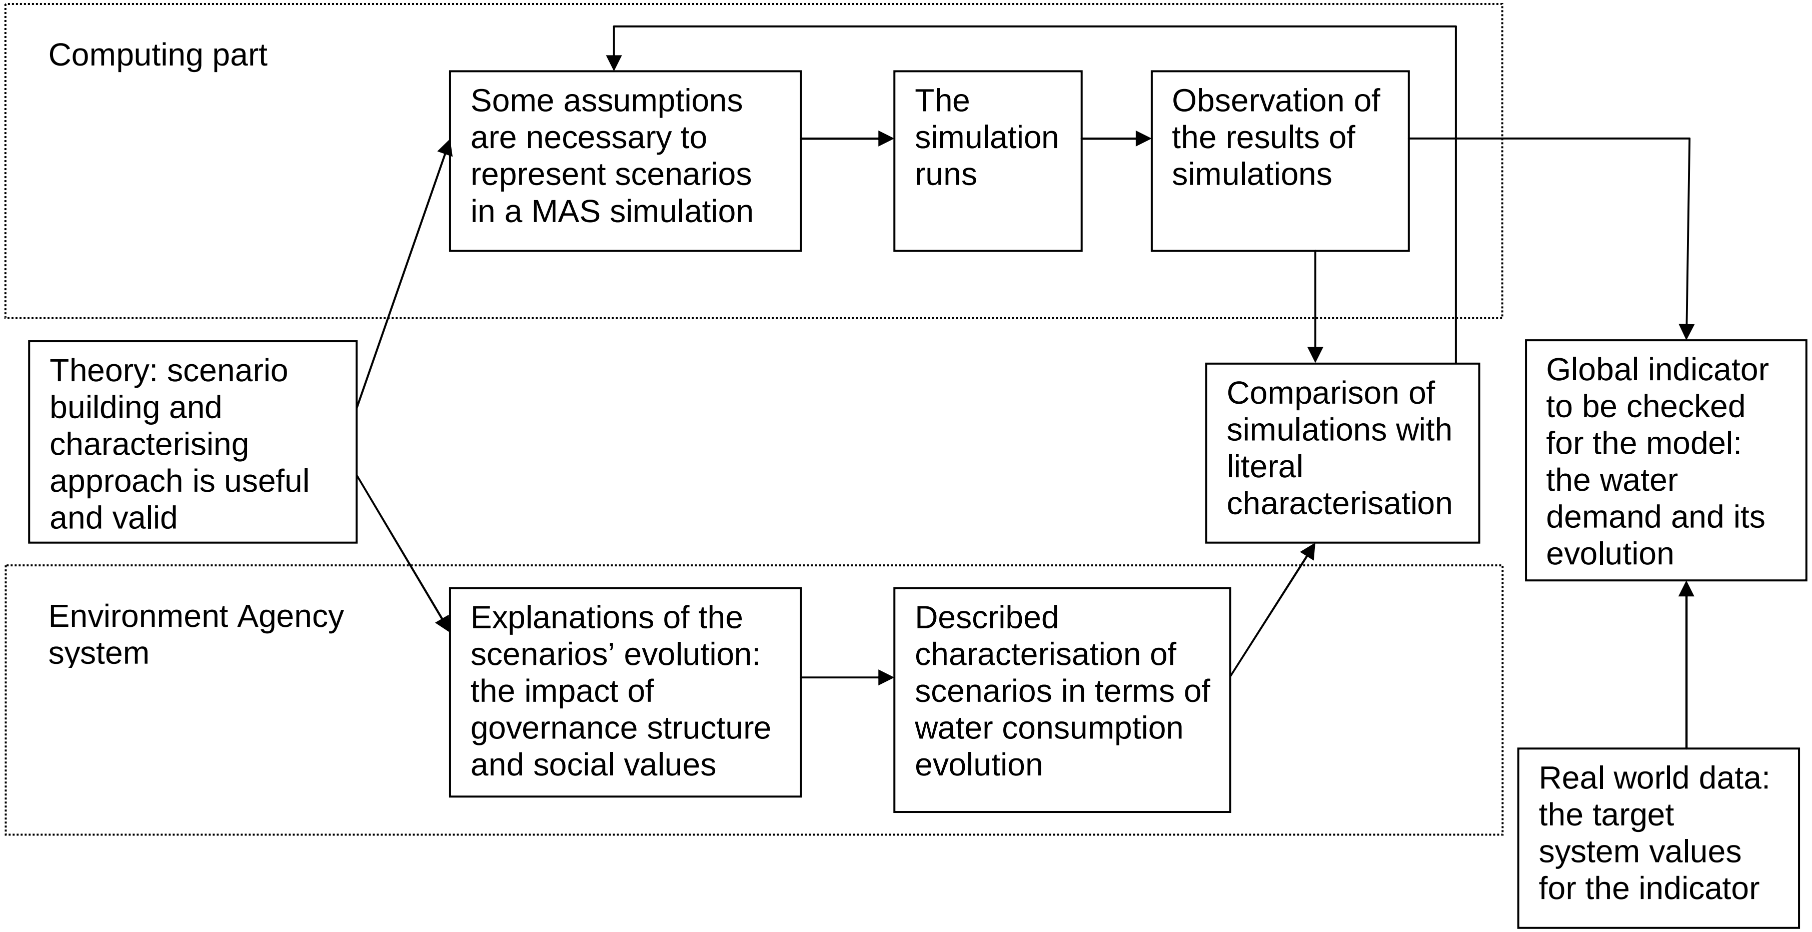

Figure 9: Overall links amongst model parts

It is important to keep consistency when building and using the model. In this research, consistency has two aspects. First, the verification of the model, expressing how appropriate the rules written are to the theoretical model described. Second, consistency is associated with validation, as it aims towards matching results obtained from the simulation with both foresight and objective observation.

Consistency is aimed at within the framework of this research and is twofold: via verification for the internal consistency of the model, and via observation and comparison of data for the consistency of the results with actual observation.

The internal consistency of the model (including for example identified and / or unique choices in decision processes, assumed knowledge every time step, and agent-specific knowledge) is partly ensured by using the right tools. In the present case, SDML and its underlying Strongly Grounded Autoepistemic Logic avoids logical issues with the formal representation. This is a step further than many languages, which mostly ensure the code is consistent.

The following section presents and discusses the participatory method and the tools allowing consistency to be maintained.

3.6.2 The limits of integrated assessment

3.6.2.1 Necessity of participation

The FIRMA (Freshwater Integrated Resource Management with Agents) and CCDEW (Climate Change: Demand for Water) projects tried to improve the understanding of consumption behaviour of potable water for household in the UK. They both used external expert knowledge to represent the best they could the various influences upon consumption behaviour (Moss, Pahl Wostl et al. (2001)). They were driven by two slightly different goals. CCDEW was investigating the impact of climate change upon households. In order to build appropriate plans, the water companies and regulators were concerned with the effects of climate change, including (or not) global warming, different rain patterns, and extreme climatic events, such as droughts and floods. FIRMA was more oriented towards the way to address the various problems faced by different European countries. Modelling came as a part of the integrated assessment, a plan to involve all stakeholders, making of the participatory component a major and necessary step to help understanding, while devising a typical scheme to implement this methodology in other similarly complex issues.

These projects have already provided useful conclusions for this research. First of all, the CCDEW project tends to demonstrate that the influence of climate change remains the same whatever the social values and governance structures. This backs up the relevance of the study of the scenarios. If the impact of climate change on the consumption is the same for every scenario possible, then the most important issue is certainly to investigate more about these various possibilities. The scenarios devised by the EA correspond to some “most likely” outputs, based on typical behaviours. Hence, checking the validity of such output scenarios with respect to the typical assumptions is worthwhile, and has not been addressed in the CCDEW research.

The FIRMA project has shown that a wide range of issues could be addressed using participation. Countries and research centres involved have addressed issues as different as for example drought management, negotiation regarding water use, flood control and river course remediation. It showed that the combination of multi agent modelling and of participation could result in useful debates, education of parties, and improved understanding.

3.6.2.2 Difficulties in having the stakeholders understand

Stakeholders involved in the FIRMA project have different goals, different schedules, and different points of view. One of the successes of the project is to have provoked meetings, and to have generated a tool that could make these various parties (or agents) make explicit their goals and constraints, confront each other, and ask some questions that were not asked before, or realise that they could be asked in a different way. Many of them were used to statistical models, and could not see at first the use of multi agent models. They expected researchers to come up with a model that they could make theirs or forget. As managers facing the uncertainty of the future, their prime concern was to find a model that would provide figures and if possible accurate ones. Most imagined models as predictors, in the narrowest sense, as a certain future. While some stakeholders involved reacted positively by realising the kind of meaning a multi agent model brought, some remained tied by more down to earth, pragmatic, economical, commercial expectations, and (understandably) statistical figures. This comes because of a trade off between generality and accuracy. While with hard sciences these two properties can sometimes cohabit, the presence of complexity prevents this happening. And in social science complexity, meta-stable behaviours and self-organised criticality are system properties that seem to be more common than the presence of a general equilibrium as described by Debreu (1959). In the case of statistics, the basic assumptions are very strong, and so is the output: as accurate as the assumptions it is based upon. The point of MAS in this case being to map the relevant system as closely as possible, there are few unlikely assumptions, and little certainty. Consequently, the results of such models are unlikely to give results as accurate as statistical models.

3.6.2.3 Integrated assessment requires interest

Therefore, due to the lack of "usability" of such a Multi Agent System for them, many were not interested in continuing the experience, unless forced to. That is one limit of the integrated assessment exercise. Participation requires some necessity for the results, or some curiosity. Also many stakeholders use an approach that has been clearly criticised, or that holds only thanks to unrealistic assumptions, which tends to hold them back. This is why integrated assessment, although successful in other cases, is not used in the present work.

As scenarios will be analysed here and their assumptions tested, the main interest of some stakeholders still lies in the attribution of a probability to each of the scenarios. They could hence build their plans / forecasts upon it. As each stakeholder (specially the water companies) has a unique region with some specificity, there is an obvious interest in developing techniques to reduce that uncertainty, such as Monte Carlo experiments, to avoid the need of specified probability density functions for model variables.

Based upon a model devised via an integrated assessment, the particular study of scenarios and their evaluation will be undertaken.

Hence the following modelling is based upon a model devised through integrated assessment. Unfortunately integrated assessment could not be used further due notably to the difficulty to have experts assess something they created or devised themselves. A lack of available resources meant the stakeholders involved initially were unable to provide further input. Therefore, all algorithms later presented will be accompanied by the reasons for their selection, but they will not have benefited from this iterative participatory process. It is worth noting that in the end, they have been presented and sometimes commented by stakeholders that did not comment negatively on them.

There is one exception to the fact that stakeholders need to show an interest in the study. In this particular research, the Environment Agency was interested, and discussions with the staff in general and Rob Westcott in particular have been extremely useful[14]. But because of its nature as a (non departmental) public body, when a member of the public submits a question, it becomes the organisation's duty to provide a reply.

In such cases, when a response is ensured, the critical issue remains to identify which section, department, group, or individual would be the most suitable to establish communication. The Water Demand Management team is not only a part of the Environment Agency it is also the successor to the National Water Demand Management Centre which created the scenarios in the first place.

As a stakeholder, the Environment Agency showed openness by having a critical discussion of the scenarios, and provided help in order to develop this model to assess them. Their interest did not lie in answers outside the scope of the tool. This is the reason why they are the most relevant stakeholder, whose views regarding the purpose of this study are presented in section 6.5

3.7 MAS as a framework needs appropriate tools

Benefits of MAS are important but require an appropriate tool. The language used in this work plays an important part in the modelling process. First as a declarative language it allows the representation of agents, and of rules of behaviour, and second, its internal consistency ensures an easier verification.

The MAS framework considers objects and the way they interact with each other and with their environment. In such a case, being able to analyse the specific focus, *i.e.* the process of interactions, in an easy-to-follow and accurate representation is not unimportant.

This presentation of SDML demonstrates how this tool is formally suitable, methodologically convenient, and logically sound for such an analysis

3.7.1 SDML: presentation of an appropriate tool

Programming a simulation model sets some requirements on the language that ought to be used. It should keep the model valid, usable, and extendable (Axelrod (1997)). Validity refers to the internal structure and consistency of the model, also called verification in our case. Usability refers to ease of following the various runs and interpreting the output. Extendibility refers to the possibilities for a future user to adapt the model by implementing or changing some of its features.

Modelling in this study is done with a specific declarative language that presents all of the necessary features. SDML stands for Strictly Declarative Modelling Language. It is an object-oriented language written in Smalltalk, and using a visualWorks engine, it is specifically developed in the Centre for Policy Modelling (Moss, Gaylard et al. (1996)). I will now present some basics, helping to understand the structure, and hence the programming references made in this research.

3.7.1.1 Under SDML: Smalltalk and VisualWorks

Smalltalk is a pure object-oriented language. While C++ makes some practical compromises to ensure fast execution and small code size, Smalltalk makes none. It uses *run-time binding,* which means that nothing about the type of an object need be known before a Smalltalk program is run.

Compared to the widely used C++, Smalltalk programming has several advantages. The use of libraries and inheritances allow a fast development. This object-oriented dimension makes the development process more fluid, so that "what if" scenarios can be easily tried out, and classes' definitions easily refined.

Unlike C++, which has become standardised, The Smalltalk language differs somewhat from one implementation to another. The most popular commercial

[14] It is from a discussion with Rob Westcott that the principles for an algorithm allowing for a longer memory have been sketched, and later implemented. See section 5.3.4 for the relevant analysis.

"dialects" of Smalltalk are VisualWorks from *ParcPlace-Digitalk, Inc.*, Smalltalk/V and Visual Smalltalk from *ParcPlace-Digitalk Inc.*, VisualAge from *IBM* and VisualWorks.

Only the latter is of interest for us for now.

VisualWorks was developed by ParcPlace, which grew out of the original Xerox PARC project that invented the Smalltalk language. VisualWorks is *platform-independent,* so that an application written under one operating system, say, Microsoft Windows, can work without any modification on any of a wide range of platform supported by ParcPlace, from Sun Solaris to Macintosh. VisualWorks also features a *GUI* (Graphic User Interface) builder that is integrated into the product and used in SDML.

3.7.1.2 SDML components

A brief overview of the components of SDML is given, although with simplicity in mind it might not be technically sufficient, or rather ambiguous to the aware reader.

A SDML program is made of modules, objects, definitions, rules, and rule bases.

A module can be saved as a separate file from the program itself. A single file can contain a module hierarchy if necessary. A module contains all the definitions, objects and rules. Several modules can be loaded into SDML, and as for standard objects, one module will inherit the contents of its parent, and allow access to its own to its eventual “child”.

The object part of the language means that we can define properties for some type of agent (they can be thought of as a “mould”), and later generate many objects (or more precisely instances) of this type. Some types already exist, and it is straightforward to create new ones with the properties that are relevant to the actual modelling.

Types (or classes, their equivalent in other object oriented languages) can be defined by the modeller. They can represent some characteristics or properties of an entity. The LoopingAgent will allow them to go through multiple time levels, while an equally important one is the ParallelAgent, from which many agents will inherit. It provides the capacity for multiple agents to act simultaneously, but due to the everpresent logic core of the language, it also prevents them from accessing information that would not be available before they start acting. This ensures the strict containment of information and a very rigorous behaviour with respect to the timeframe of the simulation.

There is necessarily a unique “root” agent, called the UniversalAgent, that every subagent will inherit from.

An agent’s actions can be of various natures. They are presented as rules, and use clauses.

One can define new appropriate clauses and their associated syntax. Some basic ones are already present, but this flexibility comes in very helpful. The syntax can be multiple. In this eventuality, the engine will fetch the corresponding components. The whole point in creating definitions that have a particular syntax is to be able to keep track of them. They are kept in a database.

This database is what allows us to keep track of the different values of particular arguments. This enables the possibility of backtracking decisions and assertions. When observing a simulation result, the user is able to analyse every value of every object, and devise queries that would return the exact set of values or parameters that were at the origin of that decision or event.

An object can be active, or inactive in a simulation. In our case, an inactive object will only have properties, while an active object will also have rules. Rules are composed of an antecedent and a consequent. They can be thought of as a “if (antecedent) then (consequent)” sequence. Antecedents define the conditions for firing a rule, while the consequent generates the set of arguments for the definition. Every entry in the database refers to initial conditions or has been asserted by a rule firing.

The antecedent to a rule is composed of assertions. The SDML engine will then retrieve the values of these assertions, and evaluate the eventual modifications. If the logical result of that computation is true, then the consequent will be asserted in the rulebase of the agent / type for which the clause was created.

As a declarative language, SDML facilitates exploration and analysis of the dynamics of the simulation. Internal consistency of a model is provided by the fact that SDML is based on the Strongly Grounded Autoepistemic Logic devised by Konolige.

In rule-based programming languages there are occasions when the ordering of rule-firing is underdetermined but may alter the results. In declarative programming, where rules represent relations this might result in inconsistent results (e.g. rule A fires, then rule B but the results of rule B invalidate the firing of rule A). SDML, like many other declarative languages uses a careful inference engine that ensures that rule-firing is consistent relative to a well defined sound and consistent logic. In SDML's case it is a fragment of Konolige's "Strongly Grounded Autoepistemic Logic".

A practical decision rather than one based on high principles, this logic was chosen because it allows the sort of inferences that support the production of social simulations. The inference engine of SDML is not complete in that it is possible to write rules which SDML will not be able to solve. However, the engine of SDML is optimised for the sort of rules that social simulators use, so this almost never happens in practice. The overwhelmingly important property that the logic confers on SDML is this: IF the simulations runs and finishes without SDML reporting inconsistencies THEN we know that the rule-firing was logically consistent.

A distinctive feature that is used in the model developed is the presence of a meta-agent. It is not an agent *per se*, but can be pictured more as the thinking part, or brain of an agent. It can devise rules that its agent will use during the course of a simulation, authorising for example a changing structure of preferences throughout the simulation period.

Because it is declarative, at every stage the program will look for values that conform to the rules, if possible. It can be forced to make assumptions. These assumptions are explicit and consequently easily traceable, therefore not hidden in any way and eventually subject to debate and discussion.

3.7.2 SDML nature and use

The representation of stochastic processes is of course possible. For that reason, there are random number generators. But in order to backtrack the various choices made, these random numbers can be uniquified by SDML, and can hence be retrieved by providing appropriate arguments for any posterior query.

Simulations generated by SDML show a trajectory within the space of possibilities. With a finite set of possibilities, mapping of these trajectories can still be immense. This is why the use of SDML is restricted. In the current case, it can be used to find trajectories, or to represent a sample of trajectories that can be obtained with a set of initial conditions. It is difficult to conceive that the results provided by SDML could be used in order to assert with certainty a set of properties to the result of a process with specified initial conditions.

Instead, what SDML can provide is a way to put assumptions to the test, and a logic-based example of what some process, associated with its representation via conditions and rules is likely to generate.

Many formal tools such as statistics or game theory attempt to provide strong assertions, representing imperfectly a phenomenon and using limited and constrained techniques. These techniques sometimes rely on assumptions that can be unlikely or unrealistic (Moss (2001)).

As this is not the case for Multi Agent Based Social Simulations, the nature of the tool invites a change of use. Modelling, as a representation of a system, can be considered as part of the answers that are sought for. The model itself as well as the modelling is then not used as input to a decision process, but as part of a decision process.

That is the purpose for which integrated assessment has been created, as an iterative and reflective process to link knowledge (science) and action (policy). Such a modelling method allows its use within a framework such as the integrated assessment.

The very nature of computer simulation can be seen as more than just an input. As expressed by Varenne (2003), the status of simulation can be an experience, a theory, or something intermediate. It is that latter stance that this research is attempting to emphasise. The particular status of the tool depends on the field and the spirit it is used in. For example, in artificial societies, a simulation is by definition an experiment, while in biology, the observed growth of a virtual plant will be strictly theory.

Varenne therefore defines computer simulation as treatment step by step by a computer of either a mathematical model without analytical solution, or an inference engine based upon rules. The latter is obviously relevant to this study.

3.8 Conclusion

This chapter intended to demonstrate that there is an improved alternative to modelling techniques commonly used. The improvements is twofold. Not only is it possible to find a technique (and the associated tools) that will capture both quantitative and qualitative aspects of the phenomenon observed, this technique will not depend on any underlying assumption, and will therefore be usable where statistics are not.

When observed with a particular tool or method, a society is most likely to be averaged via statistics, or detailed via surveys, that will themselves be treated as a representative sample. Most representations with the former require assumptions, as described earlier. In the later, the data gathering process itself must be very rigorous, using well-designed methods and experienced enquirers.

As seen in the previous chapter, the analysis of water demand has mainly been done within a single view, either qualitative or quantitative. Despite some reservations in some of these studies, they represent two options in the knowledge of household behaviour for water demand. The SDML implementation allows ignoring these dichotomies and including both qualitative and quantitative components in the model.

Section 3.2 argues that in order to communicate the various questions we want to investigate, a representation of the phenomenon involved, or model, is needed. A model is then defined as "*the representation of a structure*", while modelling is "*the devising or use of abstract or mathematical models*".

Section 3.3 presents facts regarding the data in the chosen case of water demand. It argues that social phenomena can present the characteristics of complexity, and demonstrates that sets of observed data do not comply with an often assumed normality. The statistical analysis of the observed data showed a property called leptokurtosis. Presenting the power law distribution as an alternative to take into account leptokurtosis lead to defining the property of self organised criticality. One of the potential causes for SOC is social embeddedness, which is a major characteristic of the target system in the present case. A critical consequence is that, as acknowledged by the SOC literature, there is not yet an analytical method to create models that generate processes with this property.

Section 3.4 introduces the concept of agent, and then of Multi Agent Systems. This section argues that Multi Agent Systems allow representing a system in a way that is compatible with the complexity conditions, observing emerging patterns of behaviour, and generating a model using a participative approach.

Section 3.6 presented the six modelling stages, their sequence and how they are linked and can be compared. It also discussed the necessity of integrated assessment, a method that was used to generate the initial versions of this model. Mostly, it emphasizes the importance of stakeholder participation, which can only be secured if there is sufficient interest from the stakeholders.

Finally section 3.7 presents the Strictly Declarative Modelling Language (SDML), arguing that its characteristics make it an ideal choice to implement the scenarios. The added value provided by using Multi Agent Systems as a framework and SDML as a tool is that this approach can be considered in two separate aspects. The first one is the possibility to investigate and assess both quantitatively and qualitatively, while the second one is the multiplicity of scales that can be used.

Now the problem is presented, and its properties analysed in order to select the most appropriate methods and tools to tackle the issues they raise, it is necessary to provide a more detailed view of the assumptions, and how these tools are used. This is the object of the following chapter.