Demonstrating the Role of Stakeholder Participation: An Agent Based Social Simulation Model of Water Demand Policy and Response

Scott Moss*, Thomas E. Downing and Juliette Rouchier***
**Centre for Policy Modelling, Manchester Metropolitan University*
***Environmental Change Institute, University of Oxford*

1 Introduction

This paper reports the successful implementation of a prototype model integrating representations of both natural and social systems and developed in collaboration with stakeholders. The model integrates a hydrological model parameterised to represent the effects of precipitation and temperature on water availability in the Thames region of southern England (including Oxford, London and the Southeast) with a model of demand for water by households. It is the first of a series of models to be developed as part of a project on Freshwater Integrated Resource Management with Agents (FIRMA). Models are being developed for five regions of Europe where the consequences of climate change for water use and availability could well be significant.

While the model reported here is itself of some interest to both the integrated assessment[1] and social simulation communities, it was developed as a key step in the demonstration of a new methodological approach. This approach rests on stakeholder participation in the model design and validation stages together with a compositional validation procedure. The participation of stakeholders is essential because the role of models in this methodological approach is to explicate and articulate decision-makers' presumptions in their formulation of expectations of the outcomes of their decisions. This role of models includes conventional "what-if" analysis but it has a deeper role in the identification by stakeholders of relevant social processes. The point is not just to change some parameters or exogenous variables in order to observe modelled consequences. It is more important to engage the stakeholders in a process of evaluating their (frequently implicit) judgements or (perhaps ill-considered) assumptions about how other stakeholders and social processes will generate desired outcomes.

---

[1] Integrated assessment is "a structured process of dealing with complex issues, using knowledge from various scientific disciplines and/or stakeholders, such that integrated insights are made available to decision-makers" (Rotmans, J. (1998)).

The purpose of the model reported here is to provide a tool of analysis for the development of water demand policy. The particular policy question addressed with the model is the ability of public authorities to induce households to reduce water consumption by exhortation rather than legislation and enforcement. This question is of some interest in Britain where, until water supply was privatised in the early 1990s, exhortation was indeed sufficient to induce households to reduce water consumption during periods of drought. Privatisation, subsequent large salary increases and bonuses for senior water company managers and news reports that up to 40 per cent of mains water was lost by leakage are widely thought to be responsible for a change in public attitude so that exhortation has not been effective in the most recent periods of drought.

The model supports the representation of the determination of household water consumption as a result of individuals' preferences, how those preferences change as a result of interaction among neighbouring households and the influence of exhortation by policy authorities. Individuals are represented by means of intelligent software agents that are programmed to reason about their preferences on the basis of their observations of and communications with other software agents. The policy authority agent exhorts agents representing domestic water consumers to reduce their consumption in response to drought conditions.

2 Methodological issues

Concern with the social sources and consequences of climate change have brought to the fore problems of analysing situations that arise from complex interactions among already complex systems. Though there have long been extensive attempts to model meteorological and social systems separately, noone would suggest that any of these models support even modestly accurate forecasting over any but the shortest term. Integrated models of both systems can only compound the failures of accuracy. Consequently, it is hard to see any likely advantages in attempts to develop forecasting models for (natural and/or anthropogenic) climate change and any social sources or consequences.

At the same time, it is clearly prudent to think about adverse effects on human society of possible changes in the environment and how human activity can ameliorate the changes or limit their effects. This is not "thinking the unthinkable" but instead amounts to "thinking about the unknowable". The FIRMA project is intended to bring to bear on such ill-understood issues any clear and firm understandings held by scientists and stakeholders in particular (and particularly sensitive) water-related areas. It is not to be *presumed* that these understandings are in some sense correct understandings – only that they are clear and firmly believed. One important purpose of models in this regard is that they can be used to clarify such understandings and to ascertain whether understandings and beliefs of individuals are coherent as well as whether (and why) the beliefs of different stakeholders in the same processes cohere.

Inevitably where water is concerned, stakeholders concerns reflect or engender social conflict. Examples are not hard to imagine or to document. Two examples being investigated by members of the Firma project givbe a flavour of such potential conflict:

- Extracting river water upstream for agricultural irrigation will reduce the flow for downstream leisure activities such as fishing and boating.
- Deepening a river to expand its navigational potential may exacerbate the effects of dry weather on riverside nature reserves.

In point of fact, such conflicts can and frequently are resolved by the creation of new (or modification of old) social institutions. But such resolutions are by no means inevitable. Clarifying the assumptions and interests of conflicting interests could either sharpen the conflict or help the various stakeholders to understand each others' concerns and thus to reach sustainable compromise. If such clarification is wanted, then formal modelling is a natural approach. *Whether* such clarification is wanted is a matter which is beyond the purview of this paper or, indeed, the competence of its authors.

The fact is that the integrated assessment community does develop models. Some of these (*e.g.*, the DICE model (Nordhaus. (1994)) are offered as forecasting models while others (*e.g.* the Targets model (Rotmans and de Vries (1997)) are offered as models of the consequences of different social norms. The involvement of stakeholders in the modelling process supports the objective of the FIRMA project to clarify issues and identify inconsistencies among the assumptions of the stakeholders as well as to help the stakeholders to identify any lack of coherence between their assumptions and the expected outcomes of their own or others' actions. The prototype model reported here has been developed with some stakeholder participation and is certainly intended to clarify how social relations and processes will either support or limit the effectiveness of policy measures.

2.1 Modelling the uncertain

An essential issue in this venture is whether and how ill-understood and highly uncertain systems and their properties can usefully be modelled at all. We define useful modelling as a process of specifying and implementing models that are seen by participating stakeholders as plausible and that enable them to explore the consequences of their own positions and the actions they might undertake. Evidently, the stakeholders must believe that the models accurately capture important social and natural relationships and, in order to form such a belief, they must understand the structure and specifications of the models themselves. To replace systems that are too complex to understand with models that are too complex to understand is hardly an advance.

The procedure for such modelling with stakeholder participation must therefore be one in which the stakeholders identify their relevant assumptions and can see them implemented in a model. The consequences of these assumptions are then identified by simulation experiments. If they are surprising to the stakeholdersbut not known to be empirically wrong, it is essential to identify the reasons for those outcomes and to engage the stakeholders in that identification and the interpretation of the reasons. This part of the procedure might indicate the need for more detailed modelling in order to unwrap the unexpected results and evaluate whether the results are plausible or some implausibility has crept into the implementation or whether the underlying assumptions warrant revision.

2.2 Compositional Validation

Validation is the process by means of which software systems are demonstrated to satisfy the requirements of the users. Within the framework of the methodology of the Firma project, validation is the process by means of which model users develop confidence that the simulation models accurately capture their own assumptions about their environment and anthropogenic effects on the environment as well as key determinants of their own behaviour and that of other stakeholders. A presumption of the Firma project methodology is that such user confidence will be enhanced by the ability of the user to investigate the micro level behaviour that, through social interaction, leads to macro level observation. This has two advantages. One is that knowledgeable users can evaluate the accuracy of agent behaviour as a representation of the behaviour of individuals or organisational units known to them. The second is that a description of the behaviour of component entities of the system as well as a description of system behaviour supports the identification of system properties that are not simply the aggregates of individual behaviour but result from the interactions among individual behaviours.

The process of validation, whereby models at each grain of analysis are validated by means of the validation of their component entities is *compositional validation*. It is closely related to the compositional verification of software engineers (*e.g.* Brazier, *et al.*, 2000) whereby a top-level requirements analysis is undertaken and then decomposed into incrreasingly fine grained software agents until the properties of the finest grained agents can be formally proved. Compositional validation is less formal though it must be the case that models at each grain of analysis are formally consistent with models at both coarser and finer grain. The extent of the decomposition of the models is determined by the needs of the users. Each model is decomposed until the users feel able to validate the representations of the relevant individuals or organisational units and understand how the interactions among these entities leads to representations of phenomena they observe or otherwise believe to be plausible in the real world.

Compositional validation and compositional verification have a common purpose in rendering complex systems comprehensible and open to informed assessment. The difference is that compositional verification is a formal tool of the knowledge engineer intended to give confidence that the logical properties of the software design are consistent, sound and decidable while compositional verification is a descriptive tool of the social scientist intended to give confidence to modellers and users alike that representations of individuals, social entities and social processes are accurate (if not precise) descriptions of their targets.

3 Policy issue: exhortation in demand management

The particular policy issue addressed by the model reported here is the effectiveness of exhortation in managing domestic water use. In the UK, exhortation was apparently effective during dought periods until the water providers were privatised, the water company managers were then awarded large remuneration packages and, in the midst of a drought and restrictions on water use with campaign exhorting households to conserve water, it became public knowledge that nearly half of mains water was lost through leakage *and* the water companies claimed it was uneconomic to repair the leaks. In effect, the public view of the water companies was based on a perception of greed: the companies required households to conserve water so that the companies did not have to repair leaky mains so that, in turn, the companies' managers could receive large increases in their remuneration. So it is clear that one issue is the reputation and public confidence in the water suppliers. In addition, there are likely to be cultural issues. Whereas, at least formerly, exhortation was a useful and successful element in water demand management strateby in the United Kingdom, our Dutch partners assure us that exhortation has been ineffective in the Netherlands in similar conditions.

Evidently, the conditions in which exhortation will be an effective tool of demand management is not a modelling issue, though modelling can be used to describe appropriate social processes that will support demand management by exhortation. In the model reported here, we investigate social relations that support the effectiveness of exhortation. These relations are neither necessary (there are doubtless other relationships that could be described and have the same effect) not sufficient (since we know that cultural and political factors can and do predominate). The point is to describe a plausible mechanism that has some independent support; that is, if you like, validated by some other research community.

The key concept on which the model is built derives from the social psychological consistency principle. This principle, for which there is substantial descriptive evidence, has it that individuals tend to agree the most with those whom they like the best and tend to like best those with whom they agree the most. There is a strong correlation between shared attitudes and attractiveness. (Brown, 1977; Byrne *et al.*, 1986).

This principle was given expression in the Thames prototype model by specifying household-agents who could communicate directly with, and observe some behaviour of, a limited set of "neighbours". The household-agents were located on a grid and their visible neighbours were all other household agents within a specified number of cells in each of the four cardinal directions (up, down, left and right). The extent to which each household might be influenced by policy authorities' exhortations, by other households or simply by their own preferences were set at random. Among their visible neighbours, each household agent came to value most the examples set by neighbours whose water use pattern was closest to their own.

In general terms, the Thames prototype model was developed to assess the creation and strengthening of sources of information and thevalues placed on each information source by the agents representing domestic water users. The modelling of the influence of information sources closely followed the consistency principle. As reported below, this approach took us a considerable way towards a model of domestic water use that was found plausible by key stakeholders from the water supply industry and the relevant government authorities. But it did not take us all of the way. As a result of discussion with the stakeholders, the model was revised to replicate better their observations of water use patterns over the course and in the aftermath of recent droughts in the United Kingdom. The modifications to the model not only improved the plausibility of the output but also provides a natural means of extending the model to take account of technological changes in domestic water-using activities.

4 The model: common elements

The first and the revised models shared overall structure, the hydrological model, though essentially the same, was extended in the revised model and the social model was implemented in an identical grid structure with identical parameters. The substantive difference between the two models was in the representation of cognition. The reasons for this difference are themselves instructive and may be of more general interest.

4.1 Model structure

The model was implemented in SDML which supports agents and agent containers. The containers can order the sequence in which agents fire their rules or the agents they contain can fire their rules effectively in parallel. The agent thamesWorld is the container directly or indirectly of all of the components of the model. One of these components is the simple agent thamesGround while the other is itself a container agent called firmaModel. The agent thamesGround implements the hydrological model, determining the amount of water in the soil from temperature and precipitation data. The soil water is determined first at each time step of the simulation and asserted to the database of thamesGround. The agents contained by firmaModel, in particular policyAgent representing the policy authorities, can read the clause specifying the current level of soil water from thamesGround's database. policyAgent has rules to determine from the soil water level whether drought conditions prevail and, if so, whether domestic consumers should be exhorted to restrict their use of water and, if so, to what amounts. In all of the reported simulation runs, there were 100 agents representing domestic water users. Appropriately for a European project and reflecting the input of our French co-author, these agents were labelled as *citoyens* – citoyen-1, citoyen-2, *etc.* Each of these *citoyens* agents itself contained an special type of agent called in SDML a meta-agent.

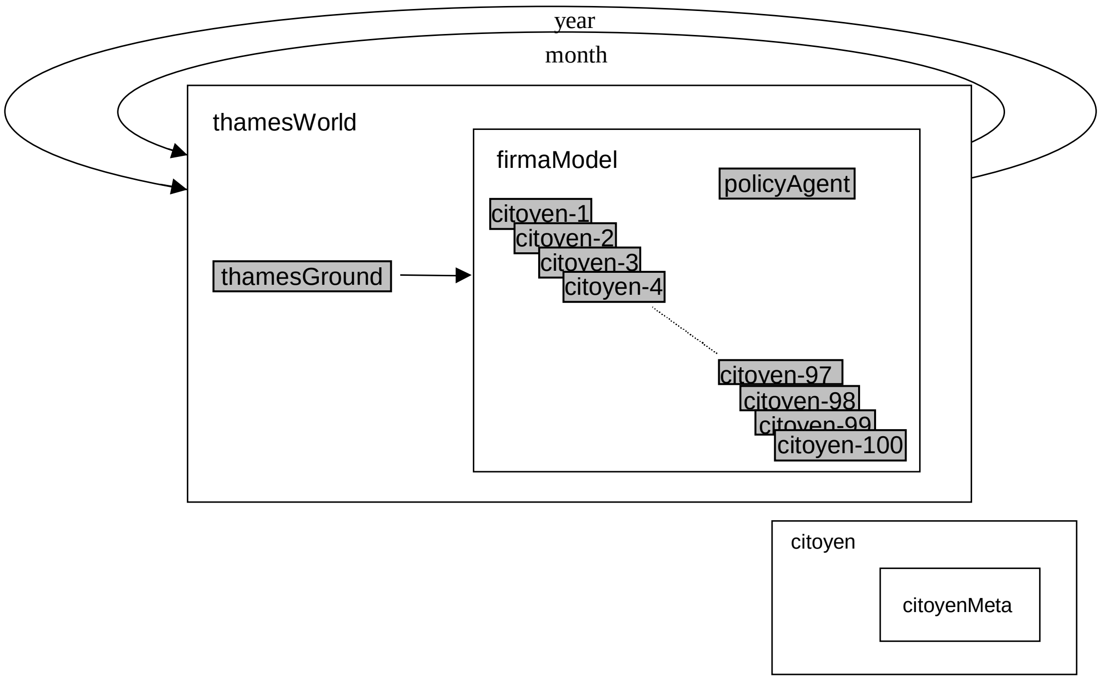

Figure 1: Thames water demand model structure

To understand the characteristics of the meta-agent it is necessary first to understand something of the way in which agents in SDML function.

Every agent has a set of databases and a set of rulebases. There is a database and a rulebase for every time step and for every time level. In the Thames model, there are two time levels for thamesWorld – the year and the month so that monthly data can be used and seasonal weather characteristics represented. Rules have antecedents and consequents. If the antecedents of a rule can be instantiated, the consequents are asserted to a database. If the rule is in a rulebase corresponding to (say) the time level *year*, then the consequent clauses of the rule are asserted to the level *year* database. As a result, the asserted clauses remain current for the remainder of the year in which they were asserted. If the rule is in the lowest level rulebase, then unless specifically stated to be of longer duration, the consequents are asserted to a database for the current time step at the lowest time level.

So much determines the duration of the currency of an asserted clause. There is also the issue of which agent's database will store the clauses. This is determined by which agent type defines the clause. Every agent is an instance of an agent type. The relevant segment of the *type hierarchy* of the Thames prototype model is reproduced in Figure 2.

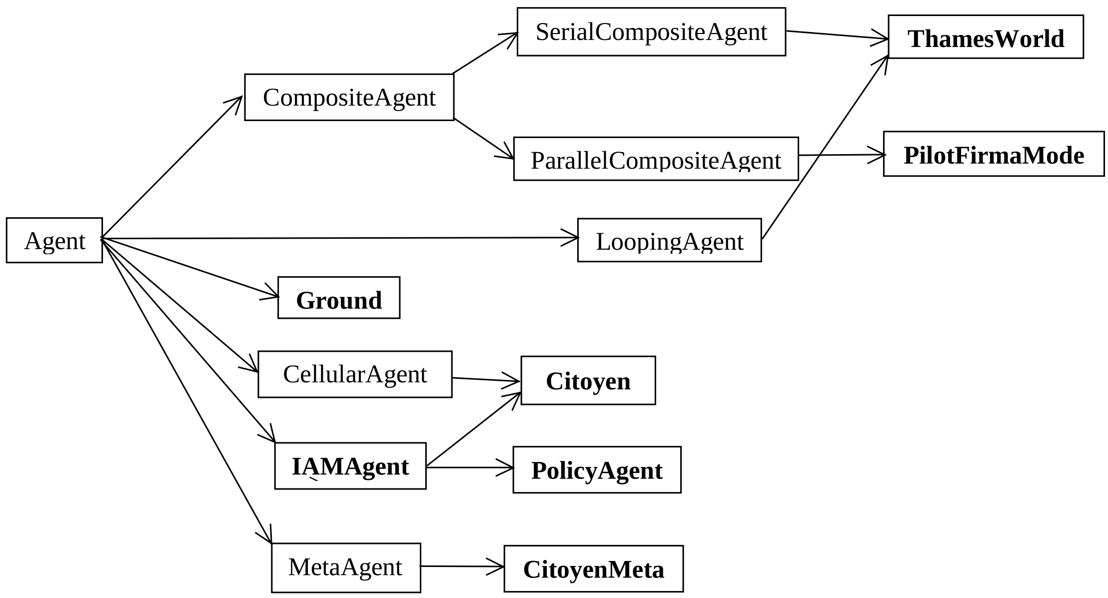

Figure 2: Type hierarchy segment for Thames prototype model

Evidently, the various citoyen-*n* are instances of type *Citoyen*, policyAgent is an instance of type *PolicyAgent* and thamesWorld is an instance of type *ThamesWorld*. Also, the agent thamesGround is an instance of type *Ground* while firmaModel is an instance of type *PilotFirmaModel* Each instance of type *Citoyen* contains an instance of type *CitoyenMeta*.

The syntax of every clause that can be asserted to a database must be defined in advance. The clause definitions are attached to a type of agent. If the agent contains other agents, they can be allowed to read and/or write to the database of their containers. If a clause is defined in the type of the container, then the assertion of an instance of the clause will write it to the database of the container rather than to the database of the asserting agent.

A meta agent differs from other agents in that the meta agent treats the rulebase of its container as a database. That is, a meta agent can read the rules held on the rulebase of its containing agent and it can write rules to the containing agent's rulebae as well. In the Thames prototype model, the instances of CitoyenMeta write rules for their containing instances of Citoyen to determine the water use by the Citoyen instances each month. Consequently, all of the reasoning about water use and the decisions about how much water to use in which conditions is encapsulated in the specification of the CitoyenMeta.

4.2 The hydrological model

The purpose of the hydrological model is to simulate the occurrence of drought conditions on the basis of real precipitation and temperature data. The policy agent determines there to be a drought when soil water is less than 85 per cent of the soil water retention capacity for two or more consecutive months.

The current soil water content is determined by the following rule antecedent:

```
and
    PandLastMonthsSoilwater ?waterIn\     ; precipitation and inherited soil water
    PET ?pet\                             ; potential evapotranspiration
    time month ?month\
    kc ?kc ?month\                        ; monthly correction factor to PET
    is ?pet_kc ?pet * ?kc\                ; corrected PET
    greater ?waterIn ?pet_kc\
    AET ?aet\                             ; actual evapotranspiration
    is ?waterInMinusAET ?waterIn - ?aet\  ; water left after AET
    max ?calcSoilWater ?waterInMinusAET 0\
    AWC ?awc\                             ; maximum soil water
    (if
        greater ? calcSoilWater ?awc\     ; if ?calcSoilWater > AWC
        = ?soilWater ?awc\                ; soil water = AWC
        = ?soilWater ?calcSoilWater)\     ; else, = calculated value
```

The table of correction factors for the calculated value of potential evapotranspiration (PET) is:

Jan	Feb	March	April	May	June	July	Aug	Sept	Oct	Nov	Dec
0.8	0.8	0.8	0.9	1	1.1	1.1	1.1	1.05	0.85	0.8	0.8

The value of PET is calculated by a backward chaining rule[2] from the modified Thornthwaite algorithm from temperature and hours of daylight per day. The rule is:

```
and
    time month ?month\
    latitude ?lat\
    unadjustedPET ?upet\                      ; backward chaining rule
    daysInMonth ?days ?month\                 ; from table
    meanDayLength ?mdl ?month ?lat\           ; backward chaining rule
    is ?pet ?upet * (?days / 30) * ?mdl / 12\; correct for length of month and day
```

The value of the unadjusted PET at temperatures above freezing is calculated as:

PET	Temperature (T) range
$- 415,8547 + 32.2441T - 0.4325T^2$	$26.5 \leq T$
$16.5\ (9\ T / H)^{a}$	$0 \leq T < 26.5$
0	$T < 0$

where H is heat defined as

$$H \equiv \left(\frac{T}{0.7} \right)^{1.514}$$

and the exponent a is

$$a \equiv 6.75\text{e-}7\ H^{3} - 7.71\text{e-}5\ H^{2} + 0.01792\ H + 0.49239$$

The day lengths are calculated from the day relative to the winter solstice and the latitude.

The model also yields the runoff of water but this value was not used in the Thames water demand model.

4.3 *The social model*

Because nobody knows everybody but everybody knows somebody, the model was constructed to capture the fact that every individual has a limited by usually non-empty set of acquaintances. This was achieved by situating agents of type *Citoyen* on a toroidal grid and enabling each agent to know every agent within a set number of cells in each of the four cardinal directions. In all of the simulation runs reported here, the grid size was 16^2 (= 256) cells, there were either 80 or 100 instances of

---

[2] A backward chaining rule is given the consequent and then finds instantiations of the antecedent. A forward chaining rule is one where the antecedent is instantiated and then the consequent is asserted.

Citoyen distributed at random on the grid with no more than one agent to a cell and each agent had a horizon of 4 cells in the sense that it knew all agents in the four cells to the right, to the left, up and down.

Two abstract types of water consumption event are defined: private consumption events that are not visible to any but the consuming agent and public consumption events that are publicly visible. Each agent is able to observe the public consumption events of all agents known to it. Each agent decides how much water to use during each event and the frequency of use events measured by the number of events per month. This is also the information viewed by the agent's neighbours.

In addition, each agent of type *Citoyen* can read the exhortatory policy pronouncements of the policy agent. These pronouncements take the form of recommended frequencies of each activity (public and private) as well as consumption per use event. As indicated, such pronouncements are made only during periods of drought defined as months in which soil water is below 85 per cent of capacity and has been below that level for at least two consecutive months. The frequency and consumption per event recommended by the policy agent are both determined by the previous month's use and the current shortfall of soil water below 85 per cent of capacity. So if, for all domestic households, the previous average frequency of a water using activity were f and the use per event were u and the current soil water magnitude is 75 (capacity = 100), then the policy agent would recommend that each household engage in the consumption activity on $0.75 \times f$ occasions (rounded down to the nearest integer) and that the use per event be $0.75 \times u$. In every drought month when the soil water dimishes, the frequence and user per event demanded by the policy authority is reduced according to the above formulae. The policy agent does *not* change the demanded frequence or use per event when drought conditions persist but the soil water increases. When the drought ends – because soil water exceeds 85 percent of capacity – the policy agent simply ceases to offer any exhortation.

This behaviour by the policy agent captures at coarse grain the behaviour actually observed with regard to the authorities' public pronouncements and statutory restrictions in the UK during periods of drought.

4.4 Agent cognition

The cognitive processes of the agents of type *Citoyen* were specified to capture observed patterns of demand for goods and services by representing social interactions that have been well documented by observation and experiment by social psychologists in a manner that also reflects the fact that individuals change their behaviour in response to a significant weight of evidence and social pressure but not in response to small changes in price, incomes or any other variable. The representation of the determinants of behaviour are described first and then the specific representations developed to capture the role of exhortation and social interaction.

The representation of the decision-making process is the *endorsements* mechanism devised by Paul Cohen (1985) and used to good effect in several social simulation studies (e.g. Moss and Sent, 1999; Moss, 1998) to capture observed behaviour. The endorsements mechanism entails the use of rules to attach tokens to objects. Two types of object were endorsed by household agents in the Thames prototype model: actions and other household agents. The actions available to the household agents are water-use actions consisting of use frequencies and consumption per use-event. The other household agents are, of course, the other household agents (instances of Citoyen) that they can see.

There is an endorsement scheme for each type of object to be endorsed. So in the Thames model there is one endorsement scheme for water-use actions and one for visible household agents. For all households, the tokens with which actions and agents, respectively, can be endorsed are identical.

The endorsements on actions are *globallySourced*, *neighbourhoodSourced*, *selfSourced* or *bestEndorsedNeighbourSourced*. The first three are straightforward. Globally sourced actions are actions that the policy agent is recommending publicly. Neighbourhood sourced actions are actions that are observed to be taken by visible neighbours while self sourced actions are actions preferred by the individual independently of any external influences. The fourth endorsement, *bestEndorsedNeighbourSourced*, is accorded to the action observed by the neighbour that is most highly valued among all of the visible neighbours of the agent. That agent is the neighbour that is most like the endorsing agent and is determined by the agents endorsements scheme.

The endorsements on other agents are *closestActivityConsumption*, *closestActivityFrequency* or *closestActivityFrequencyAndConsumption*. The idea here is that other agents can be valued because their consumption per use-event is closer to one’s own than the consumption per use-event of any other agent or because the frequency with which an agent engages in an activity is closest to one's own frequency of engaging in that activity. There is an added appreciation of any agent whose consumption is most like one's own in both frequency of events and consumption per event. The motivation for these endorsements is the common finding by social psychologists that similarity of attitudes and mutual (not necessarily romantic) attractiveness are highly correlated. That is, individuals tend to share the attitudes of those the like and to like those whose attitudes they share. The attitudes that lead to similar consumption patterns are taken here to induce personal attraction which in turn reinforce the similarity of attitudes. In this model, attitudes towards authority and towards social norms determine patterns of water consumption.

It is, of course, essential to be able to evaluate which actions and which neighbours are the most attractive. This requires some means of comparing different collections of endorsement tokens. Cohen's original endorsements approach allocated endorsement tokens to classes of importance. The action chosen would be that which had the most endorsements of the highest class or, if several had the same number in the highest class, the action that was tied in the highest class but had the most endorsements in the second highest class. If there was another tie at the second highest class of endorsements, the third or if necessary the fourth or lower class would be used to break the tie.

A more general approach, and that used here, is to define a number base b and evaluate each endorsed object according to the formula

$$V = \sum_{e_i \geq 0} b^{e_i} - \sum_{e_i < 0} b^{|e_i|}$$

where e_i is a (usually integer) value associated with the ith endorsement token. Negative values of endorsement tokens indicate naturally enough that they are undesirable. The higher the value associated with an endorsement token, the higher the class of tokens containing that particular token. The value of b is the importance of an endorsement token relative to the value of a token in the class below. If the base is 2, then an endorsement of class three contributes 8 to the endorsement value of an object while an endorsement of class two contributes only 4. For values of b larger than the number of tokens in any class used to endorse any object, the results from this evaluation scheme are the same as from Cohen's evaluatin scheme. For smaller values of b it is possible for a large number of lesser endorsements to outweigh a small number of endorsements of greater value.

There is no limit to the number of endorsement valuation schemes that could be concocted. Indeed, there are some existing schemes whereby numerical values are calculated to determine choices of individual agents that could readily be transformed into endorsement schemes. Several good examples are to be found in the models of Rouchier (2000). While the scheme used here has yielded well validated simulation outputs in a wide range of applications, the relative virtues of other schemes is evidently a necessary and appropriate topic of further research.

For each instance of *Citoyen*, the value of each endorsement token was set at random as was the numerical base (the value of *b*) of the endorsement scheme. The desired consequence of this approach was to allow for heterogeneity among agents as well as different agent populations in each simulation run. Some agents would be most influenced by external authority, some by their neighbours and some would be mainly self-directed. On average, a third of the agents would fall into each category.

A useful set of experiments would specify the proportions of agents in each category with results that could be validated by surveys or focus groups to determine the actual proportions in different cultural contexts and the critical proportions that make exhortation effective. Clearly, there must be some minimum proportion of agents who respond to public exhortation if the exhortation is to be noticed, much less effective. Equally clearly, some agents must be influenced by their neighbours if the behaviour of the agents influenced by exhortation is to proliferate through the rest of the community. Moreover, the greater the proportion of agents who are self-directed, the less likely are there to be agents who respond to exhortation and the less will be density of agents influenced by their neighbours.

The point of further experimentation with this model will be in part to verify that there are critical proportions of agents influenced by authority, by neighbours or not at all. While it is conceivable that the critical proportions found in the simulation experiments are reflected in reality, a more robust finding would be statistical patterns corresponding to the presence or absence of the proportions of each type of agent and their concentrations in the grid. Such statistical indicators have been found in other simultion studies (e.g. Moss, Critical incidents, canonical environments).

4.5 Simulation results

The results obtained from the first version of the Thames prototype confirmed that, for randomly generated endorsement schemes, the combination of agents who are influenced by authority and word of mouth communication among socially influenced agents can be effective in managing demand. A typical pattern of drought and demand is depicted in Figure 3 where the lower saw-toothed line is the number of consecutive months of low so8il water and the upper line is the simulated domestic water use. By inspection, the volume of water consumption declined during the dry periods which was when the policy agent was exhorting households to restrict their consumption.

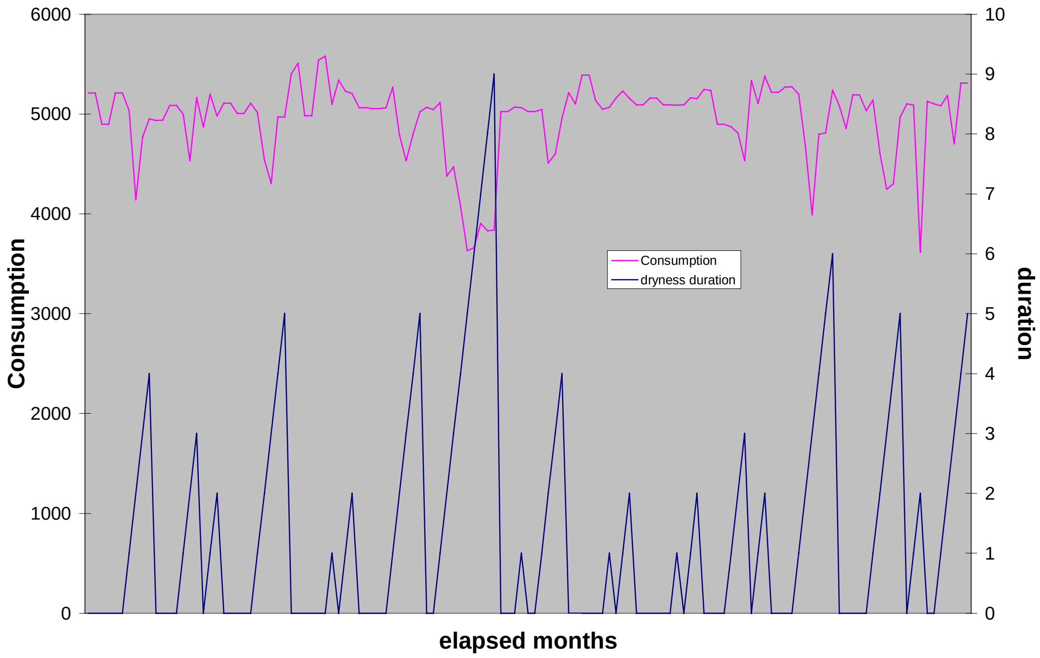

Figure 3:Drought and simulated domestic water use

Figure 4 reports a representative segment of the time series reported in Figure 3, using the soil waterdata simulated for the period from March, 1990 to April, 1992. The dip in consumption from April, 1990 until January, 1991 corresponds exactly to the period of exhortation by the policy agent. There are no lags and no residual consumption dip after the end of the dry period and policy agent's exhortation.

This model and simulation results were presented to stakeholders from the relevant ministry of government, regulatory authorities and water supply companies at a meeting in Oxford late in 1999. Some of the stakeholders pointed out that the simulated consumption patterns differ from what they have observed in that normal levels of water use recover more slowly than indicated by the model. The particulars of the representation of cognition were therefore changed in the revised model to accommodate that observation.

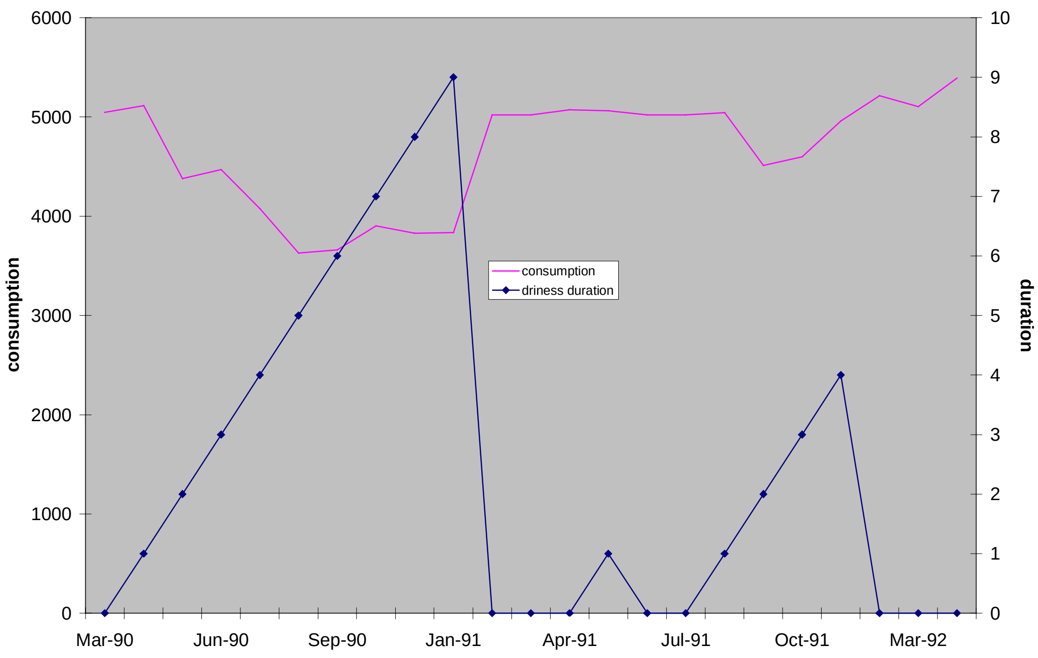

Figure 4: Drought and simulated domestic water use, March '90 – February '92

5 The revised model

Capturing the statistical signature of domestic water use in relation to stakeholder observation was achieved in two steps.

The first was to change the objects endorsed as actions. In the first version of the model, the "action" was identified with the whole rule specifying the conditions in which a particular frequency and use-per-event would be adopted. The conditions were two aspects of the state of the environment: temperature and driness. The temperature aspect could take on of the values in {*warm*, *cold*} and the driness aspect could take one of the values in {*wet*, *dry*}. The environment was *dry* when soil water was less than 85 per cent of capacity. This was calculated from the hydrological model. The {*warm*, *cold*} aspect had no effect and, so, was set randomly.

In the first version, the policy agent exhorted household agents to adopt the whole of an *if*-*then* rule of the sort

If (environmentalState [warm dry])

then (activityFrequencyConsumption ?activity ?freq ?cons)

Household agents in the first model version would remember both the *if* part of the rule and the *then* part of the rule and any endorsements would attach to the rule. This was the main reason why domestic consumption reverted immediately to normal when exhortation ended. The reason for the end of exhortation was that soil water increased to at least 85 per cent of capacity and so the driness aspect of the environment changed from *dry* to *wet*. The rules that were adopted during the drought simply ceased to be relevant until the next drought. The rules themselves remained in memory, however, so that at the next onset of driness, water deman declined immediately.

Consequently, the first change made to the representation of cognition in the rivised model was to change the form of the exhortation from an *if-then* rule to the *then* part of the rule – the consequent – alone. As part of the same change, the household agents were respecified to note, remember and endorse the consumption frequency and use-per-event of their visible neighbours rather than the consumption *and* prevailing environmental conditions. This was an essential revision to enable the drought-incuded consumption patterns to persist in the aftermath of the drought.

As a result of this respecification of agent cognition and communication, the effect of drought and exhortation was a secular decline in water use. That is, each exhortation resulted in a reduction in water use that was not subsequently reversed. This was clearly too effective a change since water consumption does revert to normal after a drought and period of voluntary or statutory restriction even if it does so more gradually than in the first model.

The reversion to normal consumption was obtained by defining a normal level of consumption for each agent in each water using activity and also by allowing for most of the endorsements attaching to each object to be forgotten over time. The consumption norms were generated at random for each instance of *Citoyen* at the time of its creation. The norms for frequency and use-per-event were endorsed as such by each agent and were remembered throughout the simulations. All other endorsements would be remembered with a probability that is directly related to the importance of the endorsement token (the magnitude of the token's value) and inversely related to the lapse of time since the endorsement was attached to its object. This mechanism, suggested by Anderson's (1993) work on memory and recall, sets the probability of remembering an endorsement as vt^{-d} where v is the value of the endorsement token, t is the number of months since the token was attached to the action or agent and $d>0$ is the factor determining the rate of decay of the probability of remembering the endorsement. The value of *d* was determined at random with a maximum set by the model operator for each experiment.

These changes in model specification had the intended effect as confirmed by the consumption and duration of soil driness reported in Figure 5 where we see that consumption recovered over a period of months after the drought spikes in the chart.

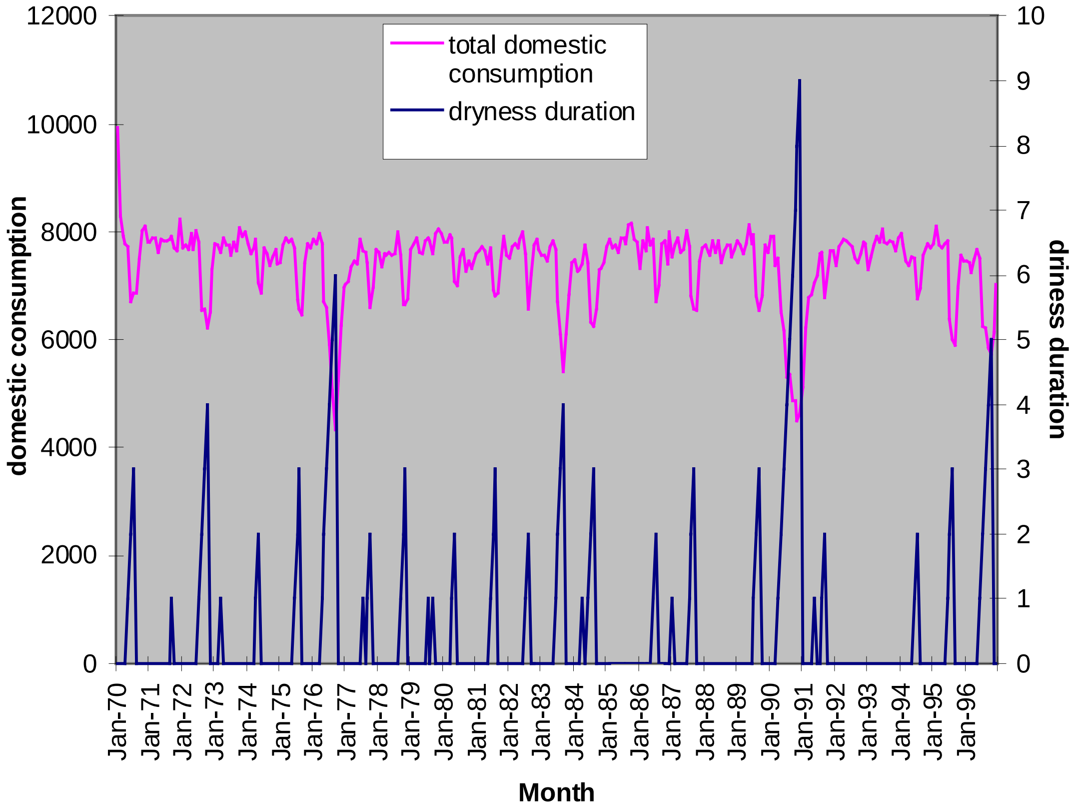

Figure 5: Simulated consumption and calculated drought, 1970-1996

The relationship of consumption to exhortation in the simulation experiments is reported in figure 6 covering two periods of drought: a relative short period in the summer of 1989 and the longer drought during much of 1990 and into 1991. In both cases, simulated domestic users reduced consumption more or less in line with exortation. An interesting result to be confirmed by data and domain experts is the reduction in consumption was longer lasting and the recovery more gradual as a result of the longer period of exhortation and deeper cuts in water user in 1990-91 in comparison with the relatively brief drought and period of exhortation in 1989.

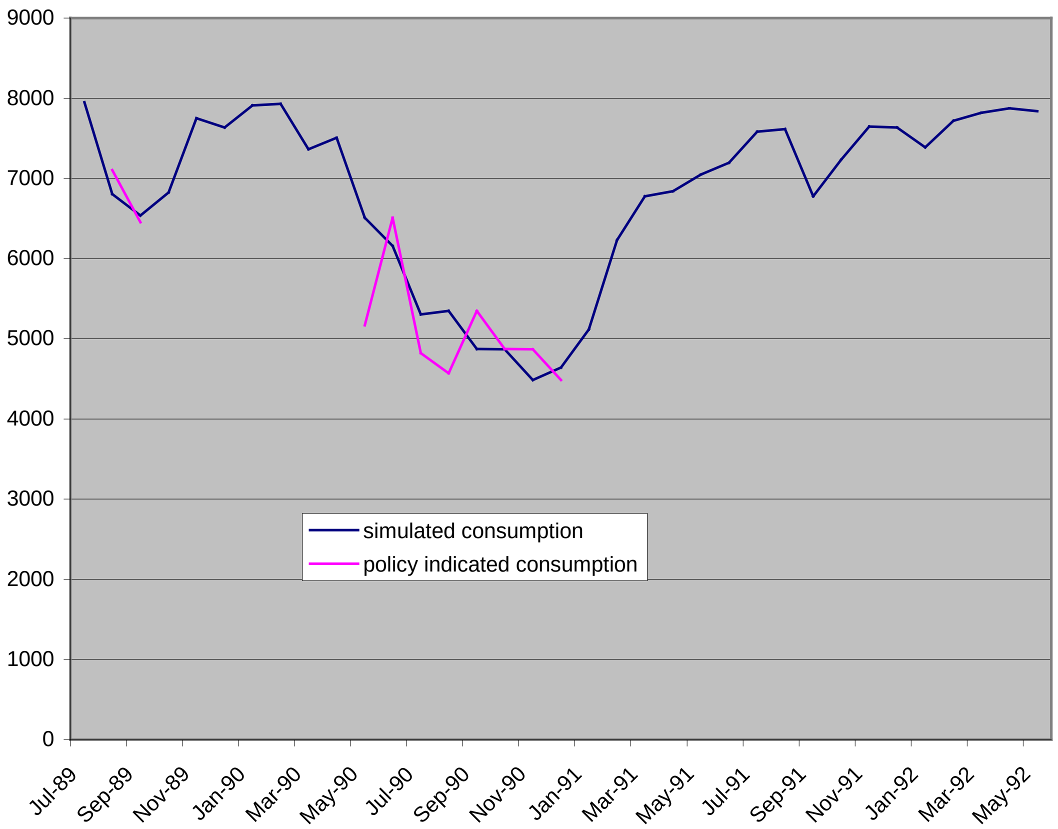

Figure 6: Exhortation and domestic use (simulated)

6 Future development path

Future work based on the prototype model reported here will be both abstract and applied. The more abstract issues turn on the results with respect to exhortation and persistence. Do longer droughts and periods of exhortation result in longer periods of recovery to normal levels of water use? The data is just becoming available and will naturally be used to validate (or invalidate) the representation of cognition in the revised model. If the representation is validated, it will be useful and perhaps important to apply that representation of socially determined consumption to the analysis of the demand for other goods and services.

More concretely, provided the abstract relations are validated, the next step in the development of the model is to identify actual water uses, to categorise them as public or private and to substitute empirically driven specifications of those actual uses for the abstract water consumption activities generated in the prototype models. It will also be useful to identify consumption pattern with the demongraphic characteristics found in different neighbourhoods to model water use at a fine grain and relate the simulated and validated consumption patterns modelled at fine grain to more coarse grained models of domestic water use.

A further development will be to replace the randomly generated norms of water use in abstract activities to technologically appropriate ranges of frequency and use-per-event in actual activities. This will also provide a straightforward means of representing technological change affecting water use.

References

[1] Anderson, J.R. (1993), *Rules of the Mind* (Hillsdale NJ: Lawrence Erlbaum Associates).

[2] Brazier, F.M.T., Jonker, C.M., and Treur, J. (2000), “Compositional Design and Reuse of a Generic Agent Model”, *Applied Artificial Intelligence Journal*, vol. 14, pp. 491-538.

[3] Byrne D, Clore G.L, Smeaton G (1986), “The Attraction Hypothesis - Do Similar Attitudes Affect Anything”, *Journal of Personality and Social Psychology* 51(6), pp. 1167-1170

[4] Cohen, P.R. (1985), Heuristic Reasoning: An Artificial Intelligence Approach (Boston: Pitman Advanced Publishing Program).

[5] Moss, Scott (2000a), “Applications-Centred Multi Agent System Design (with special reference to markets and rational agency)”, *Proceedings of the Fourth International Conference on Multi Agent Systems*, (IEEE Computer Society), pp. 199-206.

[6] Moss, Scott (2000b), “Canonical Tasks, Environments and Models for Social Simulation”, *Computational and Mathematical Organization Theory*, vol. 6, pp. 249-275.

[7] Nordhaus, W. (1994) *Managing the Global Commons: The Economics of Climate Change* (Cambridge MA: MIT Press).

[8] Rotmans, J. (1998) Methods for IA: the challenges and opportunities ahead. Environmental Modeling and Assessment 3(3): 155-179

[9] Rouchier, J. (2000), *La confiance à travers l’échange*, PhD thesis, lUniversité d’Orléans.

[10] Rotmans, J and de Vries, H.J.M. (1997) Different Perspectives on Global Futures: The TARGETS Approach. Cambridge University Press, Cambridge.