Context and Social Simulation

Bruce Edmonds
Centre for Policy Modelling
Manchester Metropolitan University

1 Introduction

Context is everywhere in the human social and cognitive spheres but it is often implicit and unnoticed. Some possible reasons for this will be sketched below. However, when one is involved in trying to understand and model the social and cognitive realms it becomes an important factor. This paper is an analysis of the role and effects of context on social simulation and a call for it to be squarely faced by the social simulation community.

This paper starts with brief review of some the different kinds of context and what they are. It then considers two, *essentially different* issues: firstly the issue of the context of a model and secondly the issue of including aspects of context-recognition and dependency in agents within simulations. I do realise that putting both of these in one paper is inviting confusion, but I am keen to air both issues. Thus the rest of this paper is in three sections:

- Section 2 is about context in general, including some of the different conceptions of it and its difficulties;
- Section 3 discusses the context *of* a simulation;
- and Section 4 talks about the representation of context *within* a simulation.

2 About Context

"Context" is used in many different senses and has many different analyses. It is somewhat of a "dustbin" concept, in that if a theory or idea does not work the reason may be assigned to "the context". Thus to many (e.g. linguists) context is a subject that is to be avoided due to its difficulty. I cannot touch on all the approaches to and models of context in the literature, but will give a brief introduction to context in general, including four conceptions of it, and a few of the issues surrounding it before proceeding to the two main arguments.

2.1 Situational Context

The situation context is the actual situation where some events or other described phenomena takes place. This could include the time and location, but could include all that is the case about that situation, including: who was there, the knowledge of those people, the history of the place and all the objects present. In this sense the context is indefinitely extensive, it is notionally includes all the circumstances in which an event or utterance occurs.

Such a context may be able to be specified adequately (if rather uninformatively) by giving the time and place of the events[1], but the relevant details might not be effectively retrievable from this. For example, the fact "I was reminiscing about our summer holiday" might well not be detectable from the time and place except by the person doing the reminiscing. Thus when talking about the situational context it is

[1] As is essentially the approach in (Barwise and Perry 1983).

almost universal to abstract from this to what is relevant about that context, or what might be commonly understood (and hence safely not described but left implicit). Thus the phrase "the context" (as in the question "what was the context?") may mean "those factors that are relevant to understand this particular occurrence" even though it may refer to the situational context in general. Thus to understand what someone is saying to you, you might ask "what was the context?" and get a description of the circumstances, e.g. "I was on the train".

2.2 Linguistic Context

Whilst the situational context could include anything, at least in theory, the linguistic context is composed of the words that surround an utterance or phrase. This typically indicates the words that precede or frame the target of understanding, but could also include common knowledge that could be reasonably be expected to be known by the listener/reader, e.g. elements of the relevant culture. Sometimes this is taken to be the same as all that which is necessary to understand some natural language.

Historically this has been the last resort of the linguist in attempting to pin down the meaning of an utterance – what one appeals to if there seems to be no detectable foreground features to explain the meaning. However more recently more positive attention has be focused on context in linguistics. For example, Peter Gardenförs (1997) has said (pragmatics being close to contextual considerations in linguistics):

> *Action is primary, pragmatics consists of the rules for linguistic actions, semantics is conventionalised pragmatics and syntax adds markers to help disambiguation (when context does not suffice).*

Clearly the linguistic context could refer to almost any of the language or culture that surrounds an utterance, and hence is not something that can be captured in its entirety. Often context is thought of as linguistic context because the interactions that are being considered are linguistic communication.

2.3 Cognitive Context

Clearly many aspects of human cognition are context-dependent, including: visual perception, choice making, memory, reasoning and emotion (Kokinov and Grinberg 2001). What seems to occur is that the human brain categorises kinds of situation which it is able to later recognise, largely without conscious effort. A lot of recall, learning and inference is with respect to a recognised kind of situation. This abstraction of a situation in the brain – the recognised kind of a situation – is the cognitive context. It is the cognitive correlate of the situational or linguistic context. Such cognitive contexts could be identified using a description of the kind of situational context that invokes them or else by the set of all the knowledge, norms, expectations, habits etc. that are immediately accessible once recognised.

It is essential that different contexts can be effectively and reliably recognised but this does not mean that they have to be consciously recognisable as distinct contexts and labelled, they may be unconsciously recognised by all the members of a community but never named; maybe they their features are distinctive and consciously recognisable but too complex and fuzzy to be completely specified.

Dividing thought about the world into these cognitive contexts, which are learnt and recognised in a rich, automatic and largely unconscious manner, and the more formal and conscious learning, recall and reasoning that is done *within* such contexts seems to be an effective heuristic for thinking about the world we live in. It is far from

obvious that such a heuristic will always be possible, or even helpful. The assumptions and advantages of this heuristic are discussed below.

2.4 Social Context

Many of the cognitive contexts we have learnt seem to correspond to recognisable kinds of social situation. Examples include: greeting, lecturing, and a political discussion. Once established these seem to be self-perpetuating, in that habits, conventions, norms, terms etc. can be developed by people who recognise the context, but in turn this might mean that the context is more recognisable as an important kind of situation which has its own characteristics. Thus social contexts can be co-constructed over time and passed-on (mostly by experience) to others.

When people are asked to describe the context, they will often do it in social terms. Thus it is that the social context, although it is a special case of situational context is closely linked to the synchronised cognitive context that participants have learnt to associate with situation, because it is often the social aspects that are important in terms of communication and understanding. It is because of the context-dependency of human cognition that when the social context is recognised, experienced inhabitants of that context will know what set of norms, habits, terms, etc. are associated with it and automatically bring them to bear in their social organisation. Thus one of the consequences of the context-dependency of our cognitive capabilities is the prevalence and importance of social context in our understanding of the world[2].

2.5 The "Context Heuristic"

The "context heuristic" is a way of dealing with the world by a cognitively limited being. It is the way in which types of situational, social or linguistic context are associated with a cognitive context. The outline of this heuristic is described now.

In all the above uses, the "context" is associated (directly or indirectly) with the set of implicit 'background' assumptions, constancies, features, knowledge, terms etc. within which the explicit 'foreground' reasoning, events, processes etc. are conceptualised to *occur*. The 'background' factors are those that are either so constant that they can safely be ignored (Zadrozny 1997), or aspects that hold for a kind of situation that can be effectively recognised. The foreground features are those that vary, whose interactions and relationships are investigated, talked about, or simply noted.

For this to be a useful distinction, it is necessary that "packages" of foreground aspects are generally associated with a sets of situations that can be reliably indentified. This identification may not be a neat or formalisable inference, but seems often to be a complex, rich and vague recognition process that is not immediately obvious or conscious. *If* the reality the actor is dealing with is structured in this way then the following heuristic will be applicable.

- The different "contexts" are recognisable in a sufficiently reliable way
- Knowledge, habits etc. that applies to the situations recognised as these contexts are learnt and associated with the context

[2] Of course, as with language (Deacon 1998), it may be that some of the survival value of our brains is that it allows the co-construction of social context with which to associate sets of applicable norms, behaviours, knowledge, terms etc.

- When a situation that corresponds to the context is encountered it is recognised and the "package" of knowledge that is associated with it is available to be reasoned about etc.
- Both the learning and the later application of the content are done taking the background aspects that are associated with the context for granted, allowing a focus on a much more constrained set of foreground aspects

The principle advantages of this heuristic are that is makes learning, recall and reasoning limited to what is relevant *within* a context[3]. Thus it makes these processes feasible since a more limited set of 'foreground' knowledge has to be considered (Edmonds 2002). In other words this heuristic solves the "Frame Problem" (McCarthy and Hayes 1969). (Greiner et al. 2001) points out that trying to apply generic reasoning methods to context-dependent propositions and models, will be either inefficient or inadequate. This heuristic also allows the same situation to be considered from the point of view of different cognitive contexts. So that if a package of knowledge from one cognitive context does not seem to allow a decision to be made then another cognitive context may be sought which does allow this (Edmonds 2002). In other words, what the best cognitive context is for any problem or task can be flexibly determined – there may be more than one cognitive context relevant for any particular situation. For example a situation might be recognised as both "an interview" as well as "an evacuation" context, if one is being interviewed for a job when the fire alarm goes off. Flipping between contexts, so that one can bring different sets of knowledge, assumptions, terms, behaviours etc. to bear, can be a powerful tool for understanding or decision making.

Human cognition seems to work using a combination of rich, unconscious and fuzzy recognition of contexts along with a more precise, conscious and limited reasoning within the currently associated context[4]. Thus it allows these two, very different kinds of cognitive system to be integrated (Edmonds and Norling 2007).

Clearly this heuristic is a contingent one – it *may* be applicable in cases where reality is usefully dealt with in this way, for example in many human social occasions, but in other cases may not be a useful way to proceed. It also relies on the fact that the relevant context can be recognised with reasonable reliability by different people which, presumably in turn relies on there being some underlying commonality between similarly recognised situational contexts (Edmonds 1999a).

2.6 Identifying and Talking about Context

One of the difficulties in discussing context is that they may well not (a) be accessible to us (b) identifiable even if they are accessible or even (c) definable in precise terms even if we can identify them. Thus although, in some way, the brain abstracts its stream of information to some properties of its state that it can use, at a later time, to recognise and retrieve knowledge that is relevant to the same kind of situational context, there is no reason to suppose that we can safely reify these properties that

[3] Talking about knowledge *within* a context draws upon a common analogy which is useful in talking about context. In this analogy the context is thought of as a (possibly overlapping, fuzzy, hierarchically structured) set of containers – the contexts – which once accessed allow their contents to be made accessible.

[4] Although the outline of the heuristic seems clear, how the brain does this is not. Thus the mechanism(s) by which context is recognised is unknown as well as how contexts might be related to each other, e.g. hierarchically or by association.

would correspond to the cognitive context. Rather we often have to try and deduce what the cognitive contexts are by introspection and other observation.

Despite this, we often talk about contexts as if they were discrete "things", however it needs to be understood that *for our conscious selves* they may not be the case. Thus "the" context is an abstraction of the aspects of those background features that define it, whether or not this is a meaningful or reifiable entity for us. To simplify the discussion I will generally talk about contexts in the sections below as if they are well defined identifiable entities, but the caveats just mentioned need to be always taken into account. This difficulty means that *the* context for any situation is often not made explicit or represented – those involved may well not be aware of the cognitive context they are assuming.

The fact that the relevant cognitive context may not be directly accessible to our consciousness does not mean that it is *totally* immune to being partially identified or uncovered, just that this might be unnoticed, non-obvious, complex, fuzzy and only partially inferable. For example, although we may not be aware of what brought to mind a particular person in a situation, on introspection we *might* be able to work out that some music brought to mind a past event in which that person figured. Thus we may be able to work out something about what sort of cognitive context is relevant but still not be able to characterise it completely[5].

2.7 Context and Causality

An important claim here is that *causation* is essentially a context-dependent abstraction. This argument is dealt with more in (Edmonds 2007), but will be summarised here. That is, in order to be able to *effectively* learn and reason about the world using fairly definite (i.e. sufficiently well-defined as to be reasoned about) models one has to separate out the foreground causes from the background ones (which can be abstracted to or relegated to a context).

If this were not the case there is nothing to stop the number of causes is unlimited of almost any event being unbounded. So, for example, if a man breaks a leg while walking down some steps, the cause could be indentified as: the distance down (hence the energy imparted to the bone), the makeup of his bones, the event that distracted him and caused him to trip, the friction on the stair, the strength of gravity on earth, the fact that he had to go to work with a cold, the design of the building, that the meeting he was in ran late so he had to hurry, the fact he is uncoordinated due to his genes, his lack of fitness due to laziness etc. etc. In almost any situation there will be an indefinite number of factors that could be included in any set of relevant causes since the world is densely connected and usually complex (especially the social world). This is what is called "causal spread" in (Wheeler and Clark 1999) or "social embeddedness in (Edmonds 1999b).

In this example, what would be identified as the cause would depend on what could be assumed as constant or irrelevant to the situation. Normally we assume that the situational context is an earthly one, so the strength of gravity would not be counted as a cause, however if the man was a spaceman born and brought up on the moon it might be. This decision as to what is the important, foreground factor that is worth

[5] It is this fuzzy and indefinable aspects of cognitive context, that it is something that is recognised in a rich, complex and defeasible manner that makes context impossible to simply express within a statement, e.g. if B is true in context A this does not mean there is any statement $A' \rightarrow B$ where A' defines the conditions that correspond to context A.

calling the cause, depends on what we can assume away as being normal to the kind of context we are considering. If the assumed cognitive context implies that we are on earth, then the strength of gravity is not considered as a cause. The assumed commonality between considered situations allows the context heuristic to work by recognising the situations which have such constancy in common as a single kind of cognitive context.

However if the kind of context is somehow given (explicitly or otherwise) then attributing causality can make sense. For example an inquiry into a road accident *can* discover the cause because what is normal and given is well understood by all. Pearl's formal structural model of causality (Pearl 2000), has the assumption that every *possible* cause has been identified and then gives a principled way to determine what is and is not a cause of any particular event. This can happen *only if* the set of possible causes is delimited, which is a consequence of setting the context.

3 The Context *of* a Simulation

A simulation is a representation of the relevant causation in some target system – the causation deemed relevant by the modeller. This is only possible due to the fact we can safely ignore many facts about what we are modelling, including many potential causes, and hence not have to represent them all in our simulations. Although some of what is included or not is a deliberate decision by the modeller, most of the possible causes are eliminated as a consequence of the assumed cognitive context of the modeller. All simulations are conceived of within a particular cognitive context, which may be related to the situational contexts of what is being modelled.

However if a simulation was not applicable to *any* situation or within *any* cognitive context other than the very specific one it was conceived in, there would not be much use for it by anyone other than its author. Thus, for any particular simulation, there is presumably a range of situations in which it *is* applicable. The intended range of situations that the modeller is considering will include the specific cognitive context in which the simulation is conceived. As a short-hand this intended range of situations will be identified as the context *of* a simulation, since the modeller will have some context in mind that corresponds to this.

If the assumed context of a simulation and the model were perfect then all the aspects that are assumed to be constant or irrelevant for this context will, in fact, be so. However the constancy or irrelevance might well not be perfect, allowing some effect from what has been assumed away into the context of the simulation into outcomes. This "leakage" of factors from outside the context of a simulation can be identified with "noise", of which random noise is only a special case (Edmonds 2009). The inclusion of randomness, noise terms and the like, is often a recognition of the imperfect nature of the assumed context of the simulation – the randomness is a representation of the effect or a proxy for its unknown nature.

Identifying what is the best cognitive context for a modelling project is a crucial step. That is choosing a naturally recognisable *set* of situations in which the background features *can* safely be ignored and the foreground situations be sufficiently defined and few that they are amenable to formal modelling. Another way of putting this is that the intended scope of a model is important – a scope that presumably corresponds to a cognitive context in the mind of the modeller. Even if the model is developed with a single situation in mind, there will be aspects of that situation that will be deemed safe to leave out of the model. In such a singular case the scope is the set of

situations similar to the one focused upon but where these unimportant factors are varied (otherwise all aspects would have to be included in the model which is almost never feasible). Completely global models, where there is no restriction *at all* on their applicability are extremely rare, and almost certainly non-existent in the social sciences.

However it is frequently the case in papers describing simulations, that the intended context of a simulation is not made clear, so one has to guess when and if a given simulation is thought (by the author) to apply to any given case a reader might be interested in. Of course it may be that this cognitive context *is* readily recognisable by other modellers. However the exact limits of a model's intended applicability are often not clear and mistakes can be made. Thus indicating (as far as it is feasible and known) the intended context of a model is extremely useful, allowing others to check they understand the intended scope of a presented simulation.

3.1 Shifting Between Contexts

Clearly knowledge is only usefully recorded or remembered if it can be applied in a different situation to that in which it was acquired. In simulation terms a simulation model is developed with one target of representation in mind but that simulation might be applied to another target. Where both original and subsequent target lie within similar, or identical, situational contexts this is likely to be an appropriate application. This is illustrated in Figure 1.

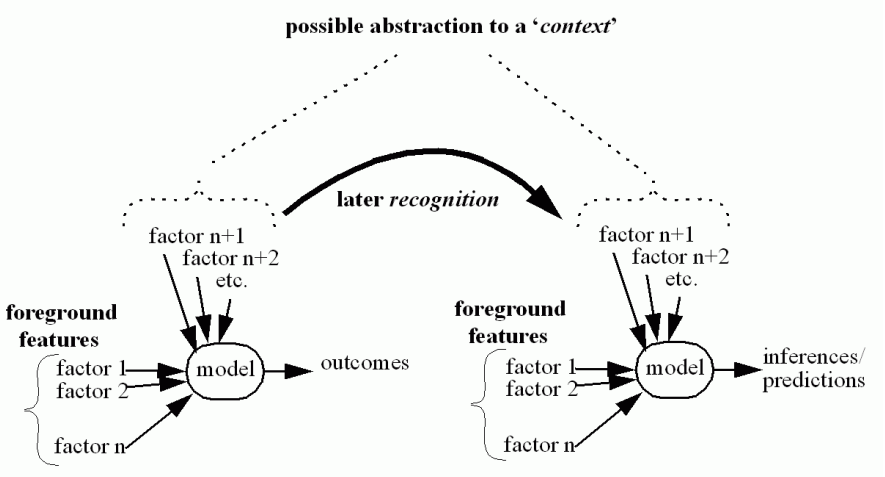

Figure 1. Knowledge transfer using a model

However, if one is applying a simulation model that is developed for one situational context (or within a particular cognitive context) to a different context, then this can subtly invalidate the model. The reason for this is that some of the aspects of the context that the model relied upon in its formulation to be irrelevant or constant, might not be so in the new context. An example of this is in (Edmonds and Hales 2005) where even the interpretation of the 2D Schelling simulation of segregation is shown to be context-dependent.

A clear and (to me) astounding case of careless context shift is when an algorithm from Artificial Intelligence or Machine Learning is transplanted as the cognitive process of an agent in a social simulation without regard for its plausibility in the new context or the sensitivity of the simulation results to the cognitive model chosen. A very clear case of this is the use of Genetic Algorithms as discussed in (Chattoe 1998).

3.2 Transcending Context

Generally knowledge is deemed less useful the more context-specific it is. There is a long tradition in seeking general truths, going back at least to ancient Greece[6]. Thus there is a long-established desire to transcend specific context. However taking a simulation model that was designed (and hopefully validated) for a particular context and seeking to generalise it to a more general scope is not at all easy.

One of the sources of this difficulty is that the scope of a model relies (at least in part) upon the context of the model and this can rarely be satisfactorily reified, or made completely explicit. If one could formalise the context that a model applied to one could add that into the foreground assumptions. Thus if one knew model M held in context A and model M' (a variant of M) held in context B then one could form the composite model: *if A then M and if B then M'*. This is the idea behind (McCarthy 1971). However it should be clear from the above analysis why this is rarely feasible – the contexts A and B might be inaccessible, unknown, imprecise or merely too complex to be able to reify in this way.

If one is simply trying to adapt a model so it is applicable to a more general context, then the model now has to cope with a new source of variation – that is, those background factors that were constant in the more specific context but not in the more general one. Thus every generalisation step one takes involves adding in more complication to the model structure. The more factors that are taken into account in a model the more evidence is needed to adequately validate it – and it is often evidence that is the limiting factor. This is the opposite of what many seem to assume – that simplifying a model will lead to greater generality. Whilst it may be true that making a model represent more features of a specific context does involve more complexity for less generality, the opposite is far from the case. The reason for this is one simply can not tell if a more general and adequate model is possible, and even then which elements one should simplify away and how.

The most systematic way of trying to generalise a model is by investigating possible additional factors, one at a time and seeing if they can be safely ignored (if this is possible). This involves adding the factor explicitly into the foreground features of a model, then testing it to see if it makes a significant difference to the results. If it does not then one could postulate that the appropriate context could be that which results from not having to assume that factor is constant. If it does make a difference then one has to include the new feature into the model, but then one knows how the model has to differ to be applicable over and above within the original context[7].

(Terán 2004) suggests a hierarchy of frameworks or contexts in which to compare simulation models, pointing out that very important difficulties arising because of the lack of an appropriate contextual framework to compare different models.

More fundamentally though, is the question of whether adequate and *general* simulation models are possible at all. It is often assumed that they are, but on no evidence other than some simply models work as a model of the ideas we have about social phenomena (Edmonds 2001a), in other words as an analogy expressed in computational form. These only seem more general because we are not aware of the

[6] There may well be many different examples of seeking general truths in other, and possibly earlier cultures, but I am not sufficiently familiar with these to judge.

[7] This paragraph glosses over the difficulty of "adding in a new feature" into a model. This might be very difficult to work out and, indeed, can involve changing the whole structure of the original model.

subtle adjustments to context that people use when applying analogies. Since we are largely unconscious of how we adapt an analogy to 'fit' a new situation, essentially making a new model each time we do, it may seem that the original model has high generality, but that is because the model has different mappings each time it is applied, which are not so apparent.

There is simply no evidence that simple and adequate models of social phenomena are possible, and some reasons to think they are not. Since it does seem to be the case that we do co-construct social contexts via mutual interaction (e.g. a lecture) and human cognition *is* context-sensitive for many of its processes, then we should expect that different rules, behaviours, norms, language games etc. will hold in each of these. In other words, that social science will be more like zoology, with a plethora of different species, mechanisms, and kinds of interaction each viable within its own biological context (which is close to the idea of a 'niche'), rather than physics, with its relatively simple and widely applicable models.

3.3 Ignoring Context

Given the difficulties and added complexity of dealing with context in social simulation a natural question is: *when and how can we safely ignore context?* A general answer to this is not known, and in fact may itself be context-dependent. However the above analysis suggests some factors.

In general, it is dangerous to assume that a simulation or model that is designed with one set of situations in mind (or within one cognitive context) will be applicable in other situations. The more similar the new situations are to the original context of simulation then the more plausible this is. However, any difference might make the simulation inapplicable, so this is something that would need checking.

If the scope of a cognitive context is sufficiently wide and recognisable so that one knows that in a certain domain all that happens is safely within that context, then that level of context can be forgotten. Thus we all easily and reliably recognise what is living and what is inanimate, despite the fact that what precisely characterises Life is hard to pinpoint. Thus within the context of biology we can take a lot of things for granted. However that does not mean that ignoring sub-contexts within biology is warranted – as in the *in vitro* vs. *in vivo* distinction.

There are some special cases where the effects relevant to different contexts cancel out or are negligible. Thus it is true that, for many practical pusposes, the microscopic movements of gas particles are sufficiently numerous and random so that they effectively cancel each other out at the aggregate level. In this case one *can* simply use the simple gas laws to predict macroscopic properties of a gas, to a high degree of accuracy and precision. However this case is a special one, where one has both empirical and theoretical reasons for its legitimacy. *Assuming* this is the case for social phenomena in the absence of evidence or other good reason is dubious – people are not like the particles in an ideal gas. The ideal gas is effectively random, but the individuals that compose many social systems are subtly coupled and/or coordinated, so that even if they appear to be acting as *if* randomly this may not be the case – e.g. in stock markets. Also just because they act *as if* randomly in one sense, does not

mean other aspects will be random. The law of large numbers simply cannot be *assumed* for social systems or their simulations[8].

3.4 Practical Implications of the Context *of* a Simulation

Context is inextricably bound up with the way we think, the way society is structured and the way we model – it can not be safely ignored. Taking on board the contextuality of our subject matter will not make us less scientific, but rather more scientific (Edmonds 2007). Ignoring context will mean that our simulation models are either (a) deeply misleading or (b) are no more than analogies expressed in a computational form.

The broad implications of the above considerations for modelling practice are as follows:

- It is useful to describe, as far as possible, the set of situational contexts the simulation is designed for (or the cognitive context it is conceived in). This will rarely be possible in total but the more information that the modeller is able to give the more likely that mistakes as to applicability or mistakes in the identification of the intended context will be more likely to be identified.
- That applying a model developed with one cognitive context in mind and applying it to another is not easy. It is easy for us to *use* the model as an analogy, adapting the mappings and meanings as with an analogy, but such a shift can invalidate a model (for any particular purpose) in subtle but critical ways. This includes a shift to a more general simulation context.
- There is no easy way to transcend or safely ignore the context of simulation except in some special cases – none of which are in social simulation. Context can not be easily included *in* the model nor its effects discounted.

Talking about and describing context is standard in many of the social sciences – it is time this became as common in social simulation.

4 Context *in* a Simulation

Given that context-dependency seems to be fundamental to human cognition and human social behaviour, it is a notable fact that very few social or cognitive simulations represent any of the processes for dealing with such context-dependency. That is to say, the agents in social simulations tend to be endowed with cognitive processes which are not sensitive to, recognise or use context. In other words, agents in social simulations tend not to have anything that might act as a cognitive contexts. If the situation in which the agents are being represented can be considered as a single and fairly simple set of situational contexts, then this is reasonable since one only has to capture the behaviour and interactions within that.

However many simulations are intended not as a representation of something more general than those corresponding to a single cognitive context but aspire to be a more general theory of social interaction. In this case, one has to assume that either the simulation is to be taken only as an analogy or that the simulator does not think

[8] A test for the effective randomness is to take systems of increasing size and see if the variance as a proportion of the size drops towards zero as in (Kaneko 1990).

people's behaviour, norms etc. will be sufficiently similar between situational contexts so the context-free representation is adequate[9].

In the former case where the simulation is used only as an analogy, then this is valid because humans are experts at applying analogy in a context-dependent manner, adjusting its assumptions and form to be appropriate to its domain of application.

In the later case, where an essentially context-independent algorithm is used to represent a highly context-dependent process must, at least, be the legitimate target for doubt. Whilst the psychological realism that is necessary in a social simulation does depend upon the purpose of the simulation and the level of aggregation (Gilbert 2006), it is certainly not the case that the results of a simulation are robust against changes in the cognitive model being used (e.g. Edmonds & Moss 2001).

There are not many simulations which represent some aspects of context-dependency in their agents, but there are a few: (Edmonds 1998) used a cognitive learning model specifically because it included some aspects of context-dependency; (Schlosser et al 2005) argue that reputation is context-dependent, (Edmonds and Norling 2007) looks at the difference that context-dependent learning and reasoning can make in an artificial stock market, (Andrighetto et al. 2008) shows that learning context-dependent norms is different from a generic adaption mechanism, and (Tykhonov et al. 2008) argue that the definitions of trust mean that trust is also context-dependent. These show that, at least in some cases, that context-sensitive cognition can make a difference. The fact that it can make a difference is not very surprising given the apparently important role it plays in human cognition, means that there is a burden of justification on those who claim it is unnecessary – explaining why it *can* be safely ignored in their simulations.

There are approaches to including cognitive context within the learning and decision-making of agents, for example; (Edmonds 2001b) which suggests a particular algorithm and approach to learning appropriate cognitive context (but did not achieve the co-development of cognitive context due to the anti-cooperative environment they were embedded within; and (Andrighetto et al. 2008) use an approach based on social norms, whereby some of the habits and knowledge of agents are dependent upon the social context, in the sense of which group they are part of. However it must be said that cognitive contexts that implement cognitive dependency are thin on the ground.

4.1 Practical Implications of Context *in* a Simulation

The lack of agents endowed with the cognitive ability to recognise social context must limit or change the social complexity that results when they interact. In particular, the co-development of social contexts will be lacking, where the recognisability of a distinct social context will allow new and specific habits, norms etc. to be developed for that situation, enabling that social context to become more recognisable etc. This will limit the ability of such simulations to capture some classes of social phenomena where the co-development of social context is a key part. Thus it may be, for example, that such things as a "jittery market" might correspond to a co-developed cognitive context, recognised and reinforced by the market traders in that market[10].

[9] These are the charitable assumptions, of course. More often one suspects that the simulator has simply not thought about the difficulties involved in such an enterprise.
[10] As well as many other factors, of course.

Thus this suggests that:

- That a simulation composed of agents with essentially non-context cognitive models might be giving deceptive results, especially in cases where the agents are learning and/or making decisions in a wide variety of situations.
- Sometimes less "smooth" learning and inference algorithms in the agents in a simulation, that mimic some aspects of context-dependency, as observed in the humans that are being modelled, might well produce a simulation that matches the observed outcomes better.

In other words, the cognitive model encoded in the agent can matter. One can not hope that an "off-the-shelf" model based on something from another context, like AI or machine learning, will be good enough.

Acknowledgements

Many thanks to all those at the Centre for Policy Modelling and attendees of the International Conference on Modelling and Using Context for discussions about modelling and context.

References

Andrighetto, G., Campennì, M., Conte, R., and Cecconi, F. (2008). Conformity in Multiple Contexts: Imitation Vs Norm Recognition,. In *World Congress on Social Simulation 2008* (WCSS-08) George Mason University, Fairfax, USA.

Barwise, J. and Perry, J. (1983). *Situations and Attitudes*. Cambridge: MIT Press.

Chattoe, E. (1998) Just How (Un)realistic are Evolutionary Algorithms as Representations of Social Processes? *Journal of Artificial Societies and Social Simulation* 1(3):2 http://jasss.surrey.ac.uk/1/3/2.html

Deacon, T.W. (1998) *Symbolic Species: The Co-Evolution of Language and the Brain*. W. W. Norton & Co.

Edmonds, B. (1998). Modelling Socially Intelligent Agents. Applied Artificial Intelligence, 12, 677-699.

Edmonds, B. (1999a) The Pragmatic Roots of Context. CONTEXT'99, Trento, Italy, September 1999. Lecture Notes in Artificial Intelligence, 1688:119-132.

Edmonds, B. (1999b). Capturing Social Embeddedness: a Constructivist Approach. Adaptive Behavior, 7:323-348.

Edmonds, B. (2001a) The Use of Models - making MABS actually work. In. Moss, S. and Davidsson, P. (eds.), Multi Agent Based Simulation, Lecture Notes in Artificial Intelligence, 1979:15-32.

Edmonds, B. (2001b) Learning Appropriate Contexts. In: Akman, V. et. al (eds.) Modelling and Using Context - CONTEXT 2001, Dundee, July, 2001. Lecture Notes in Artificial Intelligence, 2116:143-155.

Edmonds, B. (2002) Learning and Exploiting Context in Agents. Proceedings of the 1st International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), Bologna, Italy, July 2002. ACM Press, 1231-1238.

Edmonds, B. (2007) The Practical Modelling of Context-Dependent Causal Processes – A Recasting of Robert Rosen's Thought. Chemistry and Biodiversity, 4(1):2386-2395.

Edmonds, B. (2009) The Nature of Noise. In Squazzoni, F. (Ed.) Epistemological Aspects of Computer Simulation in the Social Sciences. LNAI 5466:169-182.

Edmonds, B. and Hales, D. (2005) Computational Simulation as Theoretical Experiment, Journal of Mathematical Sociology 29(3):209-232.

Edmonds, B. and Moss, S. (2001) The Importance of Representing Cognitive Processes in Multi-Agent Models, Invited paper at Artificial Neural Networks - ICANN'2001, Aug 21-25 2001, Vienna, Austria.

Published in: Dorffner, G., Bischof, H. and Hornik, K. (eds.), Lecture Notes in Computer Science, 2130:759-766.

Edmonds, B. and Norling, E. (2007) Integrating Learning and Inference in Multi-Agent Systems Using Cognitive Context. In Antunes, L. and Takadama, K. (Eds.) *Multi-Agent-Based Simulation VII*, 4442:142-155.

Gärdenfors, P., The pragmatic role of modality in natural language. in 20th Wittgenstein Symposium, (Kirchberg am Weshel, Lower Austria, 1997), Wittgenstein Society.

Gilbert, N. (2006) When Does Social Simulation Need Cognitive Models? In Sun, R. (ed.) *Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation*. Cambridge University Press, 428-432.

Greiner, R., Darken, C. and Santoso, N.I. (2001) Efficient reasoning. *ACM Computing Surveys*, 33 (1). 1-30.

Hayes, P. (1995). Contexts in Context. Context in Knowledge Representation and Natural Language, AAAI Fall Symposium, November 1997, MIT, Cambridge.

Kaneko, K. (1990) Globally coupled chaos violates the law of large numbers but not the central-limit theorem. *Physics Review Letters* 65:1391–1394

Kokinov, B. and Grinberg, M. (2001) Simulating Context Effects in Problem Solving with AMBR. in Akman, V., Bouquet, P., Thomason, R. and Young, R.A. eds. *Modelling and Using Context*, Springer-Verlag, 2116:221-234.

McCarthy, J. and Hayes, P.J. (1969). Some philosophical problems from the standpoint of artificial intelligence. *Machine Intelligence*, 4:463-502.

McCarthy, J. (1971) Generality in Artificial-Intelligence - Turing Award Lecture. Communications of the ACM, 30 (12). 1030-1035.

Pearl, J. (2000) Causality. Cambridge University Press.

Schlosser, A., Voss, M. and Brückner, L. (2005). On the Simulation of Global Reputation Systems. Journal of Artificial Societies and Social Simulation 9(1)4 <http://jasss.soc.surrey.ac.uk/9/1/4.html>.

Terán, O. (2004) Understanding MABS and Social Simulation: Switching Between Languages in a Hierarchy of Levels. Journal of Artificial Societies and Social Simulation 7(4):5 http://jasss.soc.surrey.ac.uk/7/4/5.html

Tykhonov, D, Jonker, C, Meijer, S and Verwaart, T (2008), Agent-Based Simulation of the Trust and Tracing Game for Supply Chains and Networks. Journal of Artificial Societies and Social Simulation 11(3)1 <1 http://jasss.soc.surrey.ac.uk/11/3/1.html>.

Wheeler, M. and Clark, A. (1999) Genic representation: reconciling content and causal complexity. British Journal for the Philosophy of Science, 50, 1, pp. 103-135.

Zadrozny, W. (1997). A Pragmatic Approach to Context. Context in Knowledge Representation and Natural Language, AAAI Fall Symposium, November 1997, MIT, Cambridge.