


Please cite the Published Version

Edmonds, B  (2020) How Social Simulation Could Help Social Science Deal with Context. In: Advances in Social Simulation 2018, Stockholm, Sweeden.

DOI: https://doi.org/10.1007/978-3-030-34127-5_1

Publisher: Springer

Version: Accepted Version

Downloaded from: <https://e-space.mmu.ac.uk/628031/>

Usage rights:  In Copyright

Additional Information: Copyright © 2020, Springer Nature Switzerland AG

Enquiries:

If you have questions about this document, contact openresearch@mmu.ac.uk. Please include the URL of the record in e-space. If you believe that your, or a third party's rights have been compromised through this document please see our Take Down policy (available from <https://www.mmu.ac.uk/library/using-the-library/policies-and-guidelines>)

Chapter 1

How Social Simulation Could Help Social Science Deal with Context



Bruce Edmonds

Abstract Much human cognition and behaviour is context-sensitive, but context (especially social context) has largely not been explicitly represented or included in theories of explanations of behaviour. Some of this is due to the fact that the word “context” is over-used and so has a variety of subtly different meanings but more to do with the perceived difficulties of dealing with context. Quantitative social science has tended to ignore context, treating contextual variation as “noise”. Qualitative social science has often almost deified context, resisting any attempts to generalise from specific contexts. This paper suggests that agent-based modelling could play a key role in dealing with context, representing it, understanding it and thus allowing the well-founded integration of qualitative and quantitative evidence.

Keywords Context · Social science · Social simulation · Quantitative · Qualitative · Agent-based modelling

1.1 Introduction

That context is important for understanding social phenomena should be uncontroversial; yet dealing with it has been largely avoided. On one side, quantitative social scientists tend to fit data that originates from a variety of contexts with a single model (e.g. variants of linear regression) on the grounds that they are only interested in generic patterns. At the other extreme, qualitative researchers interested in rich (“thick”) descriptions of observations and experience take context seriously, ensuring that they take care to describe relevant aspects of the context in what they record and discuss but tend to resist any generalisations that cross contexts. The point is that a crucial issue is not being explicitly addressed: that of context-dependency itself.

B. Edmonds (✉)

Centre for Policy Modelling, Manchester Metropolitan University, Manchester, UK
e-mail: bruce@edmonds.name

© Springer Nature Switzerland AG 2020

H. Verhagen et al. (eds.), *Advances in Social Simulation*, Springer Proceedings in Complexity, https://doi.org/10.1007/978-3-030-34127-5_1

Let us start by making clear what everybody knows: people behave differently in different *kinds* of situation that we can effectively recognise these kinds of situation and use them to understand, and even predict, what people will do in these situations. For example, we all recognise a lecture and know the social norms, habits, conventions, roles, etc. that pertain there. If the lecture is declared finished and coffee or wine served to celebrate something, the context has changed and everybody will behave differently. To take another example, traders in a stock market behave very differently during a bull and bear market (e.g. [11]). In a bull market, it is relatively easy to make some money and traders might seek to maximise their profits and endure quite high risk. In a bear market, traders are in danger of losing their job, so it might be more important to not be the worst performer in their group above all. In both cases, understanding behaviour is much easier and more effective if you divide the case into the different contexts. To produce models of behaviour that pertain to both lecture and celebration or to both bear and bull markets involves a much more complex, ineffective and abstract model. So why don't quantitative social scientists pay any attention to this common sense knowledge? Similarly, we all are able to recognise the difference between a lecture and a celebration or traders between a bull and a bear market and, without thinking about it much, apply the appropriate knowledge to each. So why don't some qualitative social scientists accept the usefulness of cautious generalisation over particular situations?

This paper seeks to examine the difficulty of talking about and studying social phenomena in a way that includes context but also suggests some ways forward to do this. *Firstly*, there is some discussion of the different ways that the word “context” is used, distinguishing between some different uses and trying to make clear what I intend by the word. *Then*, the paper looks at and critiques some of the ways that social science deals with (or avoids) the issue of context, looking at the responses from quantitative and then qualitative approaches. *Thirdly*, it looks at two ways in which agent-based social simulation could represent or be informed by social context – via the use of context-sensitive cognitive models for agents and by using the analysis of narrative data (including information about the relevant social contexts) to help inform the specification of simulations. It ends with a plea to change how we do social science.

1.2 Talking About “Context”

Before we can look at the issues, there are many potential confusions that can arise from the term “context”. Thus I start by discussing this confusion and distinguishing between some different uses – making connections between the different meanings and trying to make my intent clear.

“Context” is a tricky word to use and a tricky phenomenon to pin down. Like other notorious c-words (“complexity” and “creativity”), it is often used as a “dustbin” concept – what one evokes when one’s normal explanation fails. It can

also be used as a “flag” to indicate that the research is qualitative since context is emphasised (almost deified) in the qualitative social sciences yet downplayed (usually ignored completely) in the quantitative social sciences. Context, as a word, is often used informally and hence has lots of subtly different usages, e.g. as documented in Hayes [9]. Finally it is not clear that a particular context can be reliably reified and talked about as *thing* at all. These difficulties may explain the reluctance of researchers to engage with context, knowing it is a notoriously slippery and difficult subject – I guess that many think it is better to avoid the swamp and only play on firmer ground. However, with a little care, I argue that the idea and its manifestations can be sensibly and usefully dealt with and the potential pay-off for the social sciences could be considerable.

1.2.1 *Situational Context*

Firstly “the context” can refer to the situation one is in [3] – for them the exterior situation *is* the context. This could be indicated by the exact coordinates and time; however this is not a very helpful notion. The details that could be potentially relevant to any series of events in any situation are indefinitely extensive. Rather it is usual to abstract from specific situations to *kinds* of situation, for example, going home on the train or shopping in a supermarket. The question “What was the context?” implies that the speaker does not have enough information about the situation some utterance or text comes from to understand it. The answer to such a question would not be to specify the precise situation but to give enough information about it to characterise the *kind* of situation one was in (e.g. “I was talking on the phone to my mother”).

1.2.2 *Cognitive Context*

The fact that we can give enough information in a few words for the recipient to be able to infer the right kind of situation indicates that such recognition is not only feasible but also normal. It is well established that many aspects of human cognition are highly context-dependent, including memory, preferences, language use, language comprehension, decision-making, perception and reasoning [12, 16]. This implies that the brain has learned to reliably recognise these kinds of situation and effectively the same kinds as others do. The cognitive correlate of the kind of situation is called the “cognitive context” [9]¹. Though most of us, as individuals, do this unconsciously and with great facility (at least after childhood), we do not

¹However “internal” factors such as emotion and current goals may also be inputs to determining this.

know how the brain does this, and it may be that it is very hard to replicate this recognition explicitly². However this ability allows for the following heuristics: to learn knowledge with respect to the cognitive context currently being recognised and give preferential access to that knowledge when the same cognitive context occurs again. Thus when we enter a lecture, we do not have to “sift” through all the social norms we have learned for the relevant one, but those relevant automatically come to mind in that situation.

1.2.3 Social Context

Although cognitive context may be infeasible to determine in many cases, there is one case where this may be much easier – that where the context has been co-determined by many individuals in a society so that everybody recognises the same kinds of situation. Examples include the lecture, a celebration, commuting within a shared vehicle, religious ceremonies and an interview. Over time, specific norms, habits, language, spaces, technologies and even clothing might have been developed for that kind of situation, allowing the particular context to be easily distinguished. Of course, the reverse also happens: the more easily a particular context is distinguished, the more easily we will recognise it and develop specific practices, technologies and methods of coordination for it. Thus, over time, some contexts can become socially entrenched, acquire their own labels and be explicitly talked about. For this reason, such “social contexts” are much easier to identify and study than context in general.

Such social contexts can be very important since they allow for very different systems for social coordination to be developed for different kinds of situation. For example, how one coordinates behaviour on a fishing boat during periods of calm might be very different from that when a storm is approaching. It is not just the parameters of the coordination that change but the whole process.

Due to the difficulties involved in studying context, and a simple wish to avoid the extra complexity that they imply, researchers have tended to, in effect, avoid dealing with context head-on. A number of common research strategies have this effect. These will now be discussed in turn.

1.3 How Social Science Deals with Context

Here I briefly review some of the ways in which social science currently deals with (or avoids) context to motivate the need for the suggestions that follow in the next section. This is by no means a comprehensive survey, since that would far too

²Thus it may not make sense to assume that “the context” can always be reified as a distinct object that can be referred to, though as we argue it sometimes can be.

lengthy, but I hope this is sufficient to make the huge lacuna real for the reader. It includes two major approaches/assumptions made in quantitative social science and a critique of how some qualitative social science avoids the problem of context-dependency by leaving any cross-context generalisation implicit.

1.3.1 *Quantitative Social Science*

1.3.1.1 Context-Dependency and Randomness


For those researchers who claim to be “only interested in what behaviour is generic”, they may choose a model and “fit” it to some data, to see how good a fit it is (or whether it has a better fit than an alternative model). The variation not captured by the models is then attributed to “noise”, which is usually represented as some kind of randomness (sometimes this is indicated by the presence of an “error” term in equations). Typically the same, relatively simple model is fitted to the whole available data set and the extent that it “explains” the data – the likelihood that this fit is not by chance assessed (the so-called significance tests). In the social sciences, these are often variants of linear correlation models, though others variants also exist, such as the use of the “POMDP” class of models [10] in natural language processing.

The problem with this approach is that the generalisation could be occurring over different kinds of situation where different kinds of strategies might be being exhibited. The generic model “averages” over these different kinds of behaviour – producing a composite behaviour that might have elements of all of them, but misses some of the essential structural information about the observed.

A simple abstract example can illustrate this problem. Say there are two kinds of situation that occur within a sample of data: type **A** and type **B**. Within type **A**, variable a is strongly correlated with outcome x , and variable b is weakly anti-correlated with outcome y . Within type **B**, the opposite occurs: variable a is weakly anti-correlated with outcome x , and variable b is strongly correlated with outcome y . If types **A** and **B** occurred with roughly equal frequency, a generic correlation model relating a and b to x and y fitted to the complete data set might come to the conclusion that a is weakly correlated with x and b is weakly correlated with y for the whole domain at a significant level. This is illustrated below in Fig. 1.1. In this way, a lot of valuable information has been lost compared to a composite model that fitted separate models for each type. Comparing the generic model to the composite model, one would find that the generic model is not as strong and it misses the fact that there are parts of the data (from a specific kind of situation) with anti-correlations. Even an approach which included a variable to say whether a point belonged to type A or B would not help unless it was able to “switch” off and on the appropriate parts of the generic model.

If one imagines fitting a generic model over a great *many* kinds of situation, the expected result would be that many variables would be correlated with many

Fig. 1.1 An illustration of averaging out context-specific trends into a generic model



others at a significant level but only explaining a relatively small level of the total variation. This is, indeed, the result of many exercises in social science that apply generic models to data that may cover many different kinds of situation.

Not only has a lot of information been lost, but any policy based on such an analysis might be ineffective or even counter-productive for sub-groups of the population. Consider the case where there were twice as many of type **A** than of type **B**; then a generic correlation model fitted to the data might be that variable a is weakly correlated with x but there is no overall correlation of b with y . If the objective of policy is to increase x and y , then the inferred strategy would be to increase a only – despite the fact that this would have the contrary impact in a third of the cases. It seems obvious that if there were a technique to detect that there were essentially two different groups and to determine what model of behaviour fitted each, this might allow for a finer-grained understanding of data that might allow the more effective targeting of policy. Of course if it turned out that there was a substantial commonality between the separate models inferred, it might make sense to combine them into a generic model.

1.3.1.2 Over-generic Cognitive Models

Many of those that produce models based upon some kind of micro-specification simply *assume* that there is some generic model of behaviour that is valid across different contexts. The idea seems to be that there *must* be some generic cognitive model, albeit complex, that changes when the input to that agent or unit changes. I call the assumption that there must be a generic underlying model “behavioural foundationalism”.

For example, neo-classical decision theory reduces all decisions between choices to a single model: that of a utility comparison on the consequences of decisions. More complex or social decisions are implemented by a more complex utility function. However, such an approach excludes any examination of the *process* by which decisions might be made³, which might well be different in different circumstances – processes that might have very different collective outcomes from each other. For example, even if a process of individual consideration of the options and a social one (maybe looking what others are doing and imitating those who are most successful) might have a similar individual outcome for that individual, they might have very different collective outcomes. If many are following the imitative strategy, then there will be spreading waves of innovation. Another example of where the only behaviour modelled is that of social norms (e.g. [7])⁴. On a more mundane level, most of the cognitive algorithms within social simulations tend to be fairly simple and include no context-dependency at all, so in a way these all build in the assumption that a single simple algorithm will be sufficient to cover the different circumstances the agents face.

Whilst it is true that in some ultimate, biological sense, humans do have roughly the same equipment for making decisions – the nervous system – it is also true that this equipment takes years of external input, training, before it is very useful for making decisions. This suggests that a generic model of decision-making would have to be similarly complex and able to learn different strategies for different kinds of situation. There is ample evidence that many aspects of human cognition are context-dependent, including memory, decision-making, language, preferences and perception [12, 16]. One suspects that neo-classical economists simply hoped that they could produce physics-like models and thus bypass the complex and messy ways people actually make decisions, enabling them to find a shortcut that dealt with analytically modellable processes. However it has not had good empirical success. There are now so many exceptions to the received pattern of economic rationality that we should start to question whether this is the right starting point or whether this dogma might be better ditched.

It needs to be pointed out that many doing agent-based simulations are just as guilty of assuming a simple generic model of behaviour as neo-classical economists. Again this seems to be justified by an assumption that a more abstract model will be more general⁵. This is no more than a convenient hope. Whilst it might well be true that adding empirically based detail into a model might make it less general, the reverse does not work – simplifying away detail does not mean one achieves greater generality. The reason for this is clear – when simplifying, one does not know a priori what detail can be safely abstracted away. If one removes something essential to the phenomena being studied, then the result is a model that does not

³Simon's [15] distinction between procedural and substantive rationality

⁴Even social norms work in a complex and context-sensitive way [20].

⁵Although one suspects that often the reasons are more based on the mundane constraints of time and complication.

work for *any* observed cases, i.e. with no generality at all. Imagine abstracting away the variables from a linear model and just leaving the constant; this has not resulted in a more general model but one that is true almost nowhere.

A problem here is that an abstract model may often *seem* to be potentially applicable to a wide range of cases, but not in a precise manner. Here the model is used as a kind of analogy – that is, what the model parts refer to is not precisely defined but is left to each interpreter to construct “on the fly” – each person will interpret it in a different way. This is in contrast to a model where its relationship to what we might observe (represented by data of some kind) is well defined. Analogies provide very useful ways of thinking about a situation but do not give reliable or testable knowledge. Their success as an analogy does not give any guarantees that a more concrete version will be able to establish a more direct relationship with anything observable – developing an analogical model does not necessarily lead to an empirically validated one. In particular, a more abstract model of behaviour that appears to have general applicability (because it is used as an analogy) may well turn out to have less scope than one that is specific to a particular kind of situation in an empirically precise manner.

It may turn out that some elements of our behaviour can be understood in a generic manner, independent of the context, but this is something that needs to be demonstrated rather than assumed because it makes our job (as researchers) easier.

1.3.2 *Qualitative Social Science*

In contrast to the above approaches, many qualitative approaches pay a lot of attention to context – this is often described and included in their accounts and is by no means an afterthought or avoided. Indeed context is often deemed so important in these studies, that *any* possibility of generalisation to a different context is avoided – each context is unique. Thus one might have some high-quality observational or ethnographic work describing individual behaviour and strategies within a specific context but without any indication as to what could be learnt from this that might be useful elsewhere⁶ – generalisation is often left to the reader here.

This is a highly defensible stance since generalisations are risky and open to criticism by others. By keeping to discussion of phenomena only *within* specific contexts, one can counter any objection with regard to the unique circumstances within the observed contexts – contexts that the presenting researcher has unique

⁶The exception is negative knowledge – counter examples to established assumptions – but this leads to the conclusion that we know nothing except specifics.

access to⁷. Here we have to opposite problem to overgeneralisation, to a situation where almost nothing is generalised at all⁸.

In order for any knowledge to be useful, one needs to have some idea as to when it is applicable. Thus although detailed qualitative observations can expand our ideas of what people do in different situations – the possibilities – to be useful, we also need to know something about to what kinds of situation we can apply this knowledge.

1.4 Some Ways Agent-Based Social Simulation Could Deal with Context

Despite the difficulty of the subject and the corresponding circumvention in much social science, there are a number of approaches whereby we might at least start to touch upon context within social phenomena. These include using machine learning algorithms to attempt to infer context from suitably rich and extensive data [4, 17] and context-sensitive visualisation approaches, staging modelling to make the scope of sub-models clearer and extending qualitative elicitation techniques to better clarify implicit indications of context. However, here I only concentrate on two approaches which involve social simulation.

1.4.1 *Implementing Context-Sensitive Agents in Social Simulations*

Whilst it is very hard to include context-dependency within analytically solvable models, there is no reason why this needs to be the case with agent-based simulation models. However this does require a bit more “cognitive” machinery. Instead of each agent having a fixed resource of knowledge or behavioural rules, it needs to have different pools of such resources that can be selected depending on the context. In other words, the memory of the agent needs to be context-sensitive, so that context-relevant knowledge and behaviours can be preferentially applied in decision-making. Although this requires some technical changes, this is quite possible to do, ending up with an architecture as illustrated in Fig. 1.2. This sort

⁷Even if others have observed the same general kind of situation, it can always be claimed that this was at a different time or involved different actors with different goals – the defensive strategy that says every context is unique.

⁸To be precise, specific observations might be accompanied by imprecise and analogical discussion, but this is also immune to being wrong (except in maybe missing out a favorite dimension of a reader) due to its informality. Also it does not help in the identification of context or any other indications of when knowledge can be reliably used elsewhere.




Fig. 1.2 The basic context-sensitive architecture for agents

of architecture also has major advantages in terms of the feasibility of learning or reasoning, since each of these is restricted to the relevant set of knowledge for that context. It also allows a well-structured integration between context recognition (which may leverage “fuzzy” machine learning techniques) with reasoning and belief update algorithms (which tend to be crisp and derived more from the field of artificial intelligence).

However such an architecture does impose an extra burden in terms of specifying a lot more knowledge and/or behaviours for the agents for all the relevant contexts. In the simplest case, where one knows what the relevant contexts are and how to recognise them, then the different behaviours can be simply programmed into the agent, along with how to recognise each context, e.g. the approach of the CYC project [13]. Of course, one does not need a specialised architecture to do this – one could just program in more complex rules – but a specialised context-dependent memory might facilitate the process and its checking. In the more complex case, one may not know all the relevant contexts, in which case the agents might need to induce these themselves. This is more complex but possible [6, 8, 19].

Since the behavioural rules at the micro-level of agent-based simulation can be quite specific, I suspect there are quite a few existing agent-based models that have taken some aspects of context-dependency into account without necessarily calling it context-sensitive (e.g. [1, 2, 14]). Each agent in such simulations does detect the kind of situation it is in, and so it will behave differently in different situations. However, these do not distinguish context from other inputs that might influence behaviour and hence do not distinguish what can and cannot be shared between what kinds of situation.

1.4.2 Approaching Context from Qualitative Narratives

One source of information about context that could be exploited to inform the design of social simulations is that of qualitative evidence – that is, the text of observational or ethnographic work or those deriving from the relevant subjects (e.g. in interviews or their online posts). However, such evidence is often treated with suspicion by those who want to be seen to be doing “science”. Thus I start this section with a little discussion as to why qualitative research could inform simulation as well as sketching how this might work to a limited degree.


A fundamental value of science is that evidence “trumps” theory, in the sense that if evidence and theory clash then, generally, it is the theory that should be either discarded or modified. One corollary of this is that evidence should not be ignored without a very, *very* good reason. Thus neither qualitative nor quantitative evidence should be ignored. Of course, one should judge the significance of data with respect to its nature and how it was derived, for example, in terms of its relevance, reliability, subject dependence, precision, biases due to observation procedure, distortions in the process of derivation and communication and context-dependency. The quality of data is thus judged in a multitude of ways – different sets of data having different characteristics. Thus qualitative evidence might have a high degree of relevance and precision concerning events that occurred but be subject to different interpretations. Quantitative data is not necessarily more reliable just because it is expressed in a formal, precise form, but it may be if the process by which it is derived is carefully controlled and well founded.

Neo-classical economics has been notorious for ignoring a lot of evidence as to how people make economic decisions. Often this is done via an “as if” argument, which can be roughly expressed as follows: “we know people do not act in this way, but *en masse* we can treat them *as if* they do”. In the last couple of decades, experiments have shown that people do not act as if the theory of neo-classical economic decision-making would suggest, e.g. due to different frames of reference [18] or in-group bias [21].

One problem about using qualitative evidence is that it has been difficult to use qualitative data in conjunction with formal modelling methods. However agent-based modelling is well placed to use qualitative evidence to inform the menu of behavioural strategies that people might use in different situations. There is now a growing stream of work on methods to improve the process of the analysis of textual narrative data into behavioural rules (suitable for an agent in an agent-based simulation) – that is, make this more transparent and systematised [5]. Once these behaviours have been incorporated into a simulation at the micro-level, the simulation can be run and then measured to produce numbers that can be compared to macro-level quantitative data [22]. The agent-based simulation can be inspected and experimented upon to understand the process by which this occurs, and the coherence of the qualitative assumptions and the quantitative evidence investigated. Furthermore, a careful analysis of narrative data can suggest some of the context-dependency of behaviour, and if the agents in the model have a context-dependent architecture (as discussed above), this can be incorporated in a *systematic* manner into the model (Fig. 1.3).

Thus explicitly recognising and including context-dependency in formal simulation models can facilitate the integration of qualitative, quantitative and formal modelling approaches in the social sciences. In this way, some of the wealth of qualitative ethnographic, observational and interviewing work that is done in the social sciences can be used to enrich formal simulation models directly, and it does so in a way that allows the quantitative and the qualitative to be assessed together and against each other. The complexity of social phenomena will require all our resources to unpick and understand – facing context-dependency can aid the use of a wider range of evidence without abandoning rigour.

Fig. 1.3 An example (from [5]) of a process of narrative analysis separately identifying context, scope and narrative elements



1.5 Concluding Discussion

Before the advent of cheap computing power, analytic mathematical models were the only formal models available. Solving them or computing answers from them was onerous, so that only simple models were feasible. Their simplicity effectively ruled out context-dependency, leading to a focus on what generic models might tell us. Some of those who appreciated the complexity and context-dependency of social phenomena understandably reacted to this oversimplification and went to the opposite extreme, almost deifying context.

Now that we have cheap computing power, none of this is necessary. We no longer have to distort the phenomena we study in order to achieve useful formal models – we are now free to choose the *most appropriate kind* of formal model – which may well be a computational model such as an agent-based simulation. Cheap computational devices have also resulted in there being a lot more data around about social phenomena – official and informal. We are starting to have enough data to distinguish the different contexts and their associated behaviours – we no longer have to fit generic models to it due to data paucity and limits to the complexity of what we can store/manipulate. Now it is relatively easy to capture, retain, process and compare such data. Finally, we can start to use qualitative and formal methods together – enriching each other. There is no longer any need to ignore context or oversimplify what we observe to obtain and use formal models.

This has been a long time coming, since the old habits derived from a pre-computational age die slowly. However, the age of context-dependent modelling and manipulation is now within our reach. We now longer *have* to avoid it and hope for the best but can start to grapple with its complexity and so make better use of

our data (throwing less of it away as noise) and knowledge (bringing more of it to bear down on problems in an more integrated manner). It has the potential for more meaningful, more accurate and more useful models of social phenomena. It will seem odd to future generations that we have been so slow to do this.

Acknowledgements The author acknowledges funding from the EPSRC, grant number EP/H02171X/1, as well as discussion with Emma Norling and a great number of people at the Using and Modelling Context conference series.

References

1. Alam, S.J., Geller, A., Meyer, R., Werth, B.: Modelling contextualized reasoning in complex societies with “endorsements”. *J. Artif. Soc. Soc. Simul.* **13**(4), 6 (2010). <http://jasss.soc.surrey.ac.uk/13/4/6.html>
2. Antunes, L., Nunes, D., Coelho, H., Balsa, J., Urbano, P.: Context switching versus context permeability in multiple social networks. *EPIA*. **2009**, 547–559 (2000)
3. Barwise, J., Perry, J.: *Situations and Attitudes*. MIT Press, Cambridge (1983)
4. Edmonds, B.: 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 (2002)
5. Edmonds, B.: Using qualitative evidence to inform the specification of agent-based models. *J. Artif. Soc. Soc. Simul.* **18**(1), 18 (2015). <http://jasss.soc.surrey.ac.uk/18/1/18.html>
6. Edmonds, B., Norling, E.: Integrating learning and inference in multi-agent systems using cognitive context. In: Antunes, L., Takadama, K. (eds.) *Multi-Agent-Based Simulation VII*, 4442, pp. 142–155 (2007)
7. Fehr, E., Fischbacher, U.: Social norms and human cooperation. *Trends Cogn. Sci.* **8**(4), 185–190 (2004)
8. Harries, M.B., Sammut, C., Horn, K.: Extracting hidden contexts. *Mach. Learn.* **32**, 101–112 (1998)
9. Hayes, P.: Contexts in context. *Context in Knowledge Representation and Natural Language*, AAAI Fall Symposium, November 1997, MIT, Cambridge (1995)
10. Kaelbling, L.P., Littman, M.L., Cassandra, A.R.: Planning and acting in partially observable stochastic domains. *Artif. Intell.* **101**, 99 (1998)
11. Kim, K.A., Nofsinger, J.R.: The behavior of Japanese individual investors during bull and bear markets. *J. Behav. Financ.* **8**(3), 138–153 (2007)
12. Kokinov, B., Grinberg, M.: Simulating context effects in problem solving with AMBR. In: Akman, V., Bouquet, P., Thomason, R., Young, R.A. (eds.) *Modelling and Using Context*, *Lecture Notes in Artificial Intelligence* 2116, pp. 221–234 (2001)
13. Lenat, D.B.: CYC – a large-scale investment in knowledge infrastructure. *Commun. ACM.* **38**(11), 33–38 (1995)
14. Nunes, D., Antunes, L., Amblard, F.: Dynamics of relative agreement in multiple social contexts. *EPIA*. **2013**, 456–467 (2013)
15. Simon, H.A.: *Administrative Behaviour, a Study of Decision-Making Processes in Administrative Organization*. Macmillan, New York (1947)
16. Tomasello, M.: *The Cultural Origins of Human Cognition*. Harvard University Press, Cambridge (1999)
17. Turney, P.D.: The identification of context-sensitive features: a formal definition of context for concept learning. In: *ICML-96 Workshop on Learning in Context-Sensitive Domains*, (Bari, Italy, 1996), pp. 53–59 (1996)

18. Tversky, A., Kahneman, D.: The framing of decisions and the psychology of choice. *Science*. **211**, 435–458 (1981)
19. Widmer, G.: Tracking context changes through meta-learning. *Mach. Learn.* **27**, 259–286 (1997)
20. Xenitidou, M., Edmonds, B.: *The Complexity of Social Norms*. Springer (2014)
21. Yamagishi, T., Mifune, N., Liu, J.H., Pauling, J.: Exchanges of group-based favours: Ingroup bias in the prisoner’s dilemma game with minimal groups in Japan and New Zealand. *Asian J. Soc. Psychol.* **11**, 196–207 (2008)
22. Moss, S., Edmonds, B.: Sociology and simulation: statistical and qualitative cross-validation. *Am. J. Sociol.* **110**(4), 1095–1131 (2005)