The Room Around the Elephant: Tackling Context-Dependency in the Social Sciences

Bruce Edmonds

Abstract Context is crucial for understanding social phenomena, but is not being addressed. Contexts can become socially entrenched and acquire their own labels, allowing different social coordination systems to be developed for different kinds of situation. Three ways to avoid context are discussed. Fitting data to mathematical models which 'explain' the data using significance tests avoids the problems of context, but may average over different contexts inappropriately. 'Behavioural foundationalism', which assumes a generic model of behaviour that is valid across different contexts, avoids the context problem by producing models based on a micro-specification to see if the macro-consequences match the available data, e.g. neo-classical decision theory and some agent-based simulations. A third strategy to avoid the context problem is to retreat into specificity, providing so much detail that the context is unique with no attempt at generalisation. Three ways forward are proposed (1) using data mining techniques to look for models whose output 'fits' various target kinds of behaviour, (2) context-dependent simulation modelling, with the memory of the agent being context-sensitive, and context-relevant knowledge and behaviours being applied in decision-making, and (3) combining qualitative and formal approaches, with neither qualitative nor quantitative evidence being ignored. Agent-based modelling can use qualitative evidence to inform the behavioural strategies that people use in given situations. Simulations based on micro-level behaviours can produce numbers for comparison with macro-level quantitative data. This supports experimentation to understand emerging processes, and investigate the coherence of the qualitative assumptions and the quantitative evidence. Explicitly recognising and including context-dependency in formal simulation models allows for a well-founded method for integrating qualitative, quantitative and formal modelling approaches in the social sciences. Then some of the wealth of qualitative ethnographic, observational and interviewing work of the social sciences can enrich formal simulation models directly, and allow the quantitative and the qualitative to be assessed together and against each other. Before the advent of cheap computing power, analytic mathematical models were the only formal models available, but their simplicity ruled out context dependency, leading to a focus on what generic

B. Edmonds (✉)
Centre for Policy Modelling, Manchester Metropolitan University, Manchester, UK
e-mail: bruce@edmonds.name

J. Johnson et al. (eds.), *Non-Equilibrium Social Science and Policy*,
Understanding Complex Systems, DOI 10.1007/978-3-319-42424-8_13

models might tell us. New information and communication technologies have resulted in a lot more data on social phenomena to distinguish different contexts and behaviours. We no longer have to fit generic models due to data paucity and limits to storage and processing, or ignore context or over-simplify what we observe to obtain and use formal models. Addressing context has huge potential for the social sciences, including: better models and understanding of human behaviour; more effective ways of collecting, integrating and analysing data; and the prospect for a well-founded means of integrating the insights from quantitative and qualitative evidence and models.

1 Introduction

That context is important for understanding social phenomena is uncontroversial, yet dealing with it has been largely avoided. 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 descriptions of observations and experience take context seriously, ensuring that they include a lot 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 addressed: that of context-dependency itself.

Let us start by making clear what everybody knows: people behave differently in different kinds of situation, but can effectively recognise situations 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 then something about 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. Kim and Nofsinger [15]). In a bull market it is relatively easy to make 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 not to be the worst performer in their group. In both cases understanding behaviour is much easier and more effective in the different contexts. 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 reality of generalising over particular kinds of situation?

This chapter highlights the importance of this issue, critiques approaches that seek to circumvent it, tries to make some useful distinctions and show some practical ways in which we might seek to understand it. In the process it indicates the huge potential for the social sciences, including: better models and understanding of

human behaviour, better recognition of the potential for self-delusion in terms of progress, more effective ways of collecting, integrating and analysing data, and even the prospect for a well-founded means of integrating the insights from quantitative and qualitative evidence and models.

2 About Context

Context is a tricky word to use and a tricky phenomenon to pin down. Like other notorious C-words ('complexity', 'creativity') it is often used as a 'dustbin' concept—what one evokes when one's normal explanation fails. Its occurrence in a paper can be effectively a 'flag' to indicate that the research is qualitative since context is emphasised in the qualitative social sciences yet downplayed (usually ignored completely) in the quantitative social sciences. Context is used largely informally and hence has lots of subtly different usages, as documented in [13]. Finally it is not clear that a particular context can be reliably reified and talked about as an object at all. These difficulties may explain the reluctance of researchers to engage with context, knowing it is a notoriously slippery and difficult subject—better to avoid the swamp and only play on firmer ground. However, with a little care, the idea and its manifestations can be sensibly and usefully dealt with, and the potential pay-off for the social sciences is immense.

'The context' can refer to the exact situation one is in Barwise and Perry [3]. This could be indicated by the exact coordinates and time, but 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, e.g. 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 to understand an utterance or text. 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').

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. This implies that the brain has learned to reliably recognise these kinds of situation as effectively the same kinds as others do. The cognitive correlate of the kind of situation is called the *cognitive context* [13].[1] Though we, as individuals, do this unconsciously and with great facility (at least after childhood) we do not know how the brain does this and it may be that it is very hard to replicate this

[1] 'Internal' factors such as emotion and current goals may also be inputs to determining this.

recognition explicitly.[2] However this ability allows for the following heuristic: to learn knowledge with respect to its cognitive context, give preferential access to that knowledge when the same cognitive context occurs. Thus when we enter a lecture we do not have to 'sift' through all the social norms we have learned since the relevant ones automatically comes to mind in that situation.

Although cognitive context maybe 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 across 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 more 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. 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 when a storm is approaching.

As with many social phenomena (e.g. social norms [22]), social contexts are manifested both at the level of observable practice as well as having cognitive correlates. In this case there may be both the unconscious correlate, the kinds of situation automatically recognised by individuals, as well as an explicit and conscious recognition of well-entrenched social contexts. The former might be recognisable through the fact that many aspects of behaviour change together with that context, since different habits, norms, expectations, and knowledge may be associated with the same context. Some of these will be specific to an individual, but others will be shared by more than one person. However even if they are shared only some of them will be entrenched to the degree that there is explicit conscious recognition of them.

Due to the difficulties involved in studying context, and a simple wish to avoid the extra complexity implied, researchers have tended to avoid dealing with context head on. A number of common research strategies have this effect. These will be discussed in turn before considering some positive ways in which context could be tackled.

[2]Thus it may not make sense to assume that 'the context' can always be reified as a distinct object that can be referred to.

3 Avoidance Strategy 1: 'Signal Plus Noise' Data-Fitting

Researchers who claim to be 'only interested in generic behaviour' may choose a model and 'fit' some data to it[3] to see how good the fit is (or whether it has a better fit than an alternative model). The variation from the models is then attributed to 'noise', which is usually represented as some kind of randomness. Typically the same relatively simple model is fitted to the whole available data set and an assessment made of the extent that it 'explains' the data and the likelihood that this fit is by chance (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 [14].

The problem with this approach is when generalisation occurs over different kinds of situation where different kinds of strategies are being exhibited. Then the generic model is averaged over different kinds of behaviour—producing a composite behaviour that might have elements of all of them, but misses essential structural information about what is being observed. Some of the variation due to the different contexts being 'averaged' over will be represented by the noise from this generic model.

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* be strongly correlated with outcome *y*. If types *A* and *B* occurred with roughly equal frequency, a generic correlation model relating *a* & *b* to *x* & *y* fitted to the complete data set might come to the conclusion that *a* is weakly correlated with *x* and *b* weakly correlated with *y* for the whole domain at a significant level. This is illustrated below in Fig. 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 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' on and off 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 others at a significant level, but only explain a relatively small level of the total variation. This is, indeed, the result of many exercises 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 the analysis might be much more 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

[3]There are different measures of fit of varying sophistication, e.g. likelihood measures, but they all quantify the difference between the data and a model in different ways.

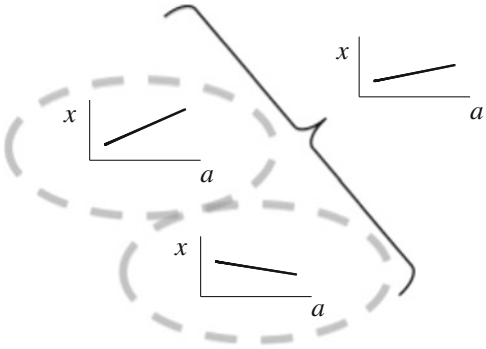

Fig. 1 An illustration of averaging out context-specific trends into a generic model

B. Then a generic correlation model fitted to the data might be such 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. If there were a technique to detect that there were two different groups and the determine what model of behaviour fitted each this might give a finer grained understanding of data allowing a 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.

4 Avoidance Strategy 2: Behavioural Foundationalism

Producing models based upon a micro-specification to see if the macro-consequences match the available data assumes 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. The assumption that there must be a generic underlying model will be called *behavioural foundationalism*.

For example, neo-classical decision theory reduces all decisions between choices to a single model: that of a utility comparison of the consequences of decisions. More complex or social decisions are implemented by more complex utility functions. However, this approach excludes any examination of the process by which decisions might be made.[4] These may be different in different circumstances—processes that might have very different collective outcomes from each other. For example, a process of individual consideration of the options and a social one of looking what others are doing and imitating those who are most successful might have a similar individual outcome but a very different collective outcome.

[4]This is Simon's distinction between procedural and substantive rationality [19].

If many are following the imitative strategy then there may be waves of innovations spreading across the population [9–11].

Whilst in a biological sense humans have roughly the same equipment for making decisions—the nervous system—this equipment takes years of external input and training before it is 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 (e.g. [16, 20]). One suspects that neo-classical economists hoped that they could produce physics-like models, by-passing the complex and messy ways people actually make decisions, to find a short-cut that deals with analytically modellable processes.[5] 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 abandoned.[6]

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 general.[7] This is no more than a convenient hope. Whilst it might 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 work for any observed cases, i.e. with no generality at all.[8]

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 computational model is used as a kind of analogy, i.e. what the model refers to is not precisely defined but is left to each interpreter to construct 'on-the-fly', with each person interpreting it in a different way. This is in contrast to a model where its relationship to what might be observed is well defined. Analogies provide very useful ways of thinking about a situation but do not give reliable or testable knowledge. Success as an analogy does not give any guarantees that a more concrete

[5] I have sympathy for those wishing for formal models, since this allows for a developmental process of comparison, critique and improvement between researchers, and to make such models analytically tractable did require strong assumptions. However, the advent of computer simulations removes this difficulty.

[6] One might argue that wheels should be square, its just that they need the corners rounding off, just as one might have argued that planetary orbits should be circular, they just need adjustment using epicycles of other circles However at some point the evidence rather than tradition needs to predominate if we are to be a science.

[7] Although one suspects that the real reasons are more constraints of time and complication.

[8] 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 with a greatly restricted scope.

version will be able to establish a more direct relationship with anything observable. In particular, a more abstract model of behaviour that appears to have general applicability when used as an analogy may have less scope than one that is specific to a particular kind of situation when used as a model with a defined relationship to real-world data. 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 justified rather than imposed.

5 Avoidance Strategy 3: Retreat into Specificity

In contrast to the previous strategies, 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 that any possibility of generalisation to a different context is avoided, since 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 that might be useful elsewhere.[9] Generalisation is often left to the reader with only the mildest of generalisations made by the researchers themselves.

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 researcher has unique access to.[10] Here we have the opposite problem to over-generalisation, the situation where almost nothing is generalised at all.[11]

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, we then need to know something about the other kinds of situation where they might also exhibit this behaviour for this knowledge to be useful.

[9]The exception is negative knowledge, counter examples to assumptions made by others, but this just leads to the conclusion that we know nothing except specifics.

[10]Even if others have observed the same general kind of situation it can always be claimed that this was at a different time or involving different actors with different goals.

[11]To be precise, specific observations might be accompanied by imprecise and analogical discussion without specific referents but this is also immune to being wrong (except in many missing out a favourite dimension of a reader) and does not help in the identification of context and when knowledge can be reliably used elsewhere.

6 Way Forward 1: Data-Mining Clues to Social Context

As indicated above, one of the difficulties in dealing with context is identifying what cognitive contexts are relevant to which observed situations. Sometimes, if the context is socially entrenched then this may be obvious but at other times it will not be, even to the participants themselves, since knowledge of context is implicit rather than explicit knowledge. Context is sometimes not neatly defined, and may be recognised in a complex but imprecise fashion.

However, the recent accessibility of fine grained data on social phenomena and advances in the field of data-mining and knowledge discovery might provide an approach that would give clues to the existence of scopes of social contexts. The idea is as follows:

1. Use data-mining techniques to identify local patterns/models that fit the data, where 'local' means that the patterns/models have an explicitly identified scope (the conditions in terms of when they hold with respect to the data) [12].
2. Do step 1, for various different target variables—that is look for models whose output 'fits' various target kinds of behaviour.
3. The result of steps 1 and 2 should be a set of overlapping models 'predicting' many different aspects of the behaviour encoded in the data.
4. Within this set look for models which have roughly the same scope—that is conditions (a) within which models seems to fit the data well and (b) which marks a change of behavioural pattern in many different dimensions.
5. The conditions indicated by this 'meta-cluster' of scopes, suggests a context (Fig. 2).

The result might not be clear with a confusion of overlapping scopes of models with no obvious cluster. In this case a commonly recognised context would not be suggested since a context suggests a whole bundle of specific knowledge and habits. There may be 'areas' of the data in which no local models were found, in which case one would have to conclude that there are no accessible patterns there.[12] Also the data needs to be sufficiently fine grain and multi-dimensional for such a technique to gain purchase, but increasingly large and detailed data sets are becoming available.

The suggestions of context that may result from such a process can be combined with qualitative evidence about context in the situations being studied. The use of qualitative evidence is discussed below.

[12]This may be for many reasons: the behaviour might just be too complex, not regular, or different for different purposes. A negative result does not mean that there are no meaningful patterns, just that we cannot detect any.

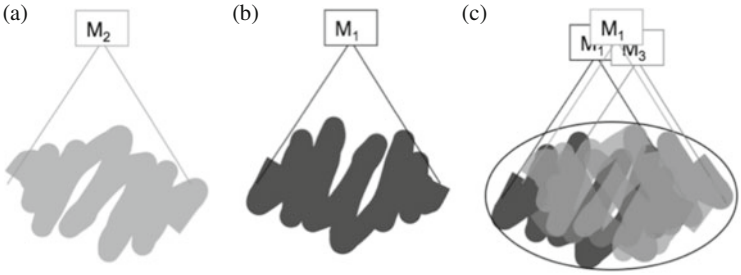

Fig. 2 Three models with different behavioural targets but coinciding scopes suggesting a context (after Edmonds [4]). (**a**) Scope of Model 1. (**b**) Scope of Model 2. (**c**) Plus Scope of Model 3 (abstract to a context)

7 Way Forward 2: Context-Dependent Simulation Modelling

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 requires 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 are applied in decision-making. Although this requires some technical changes, it is quite possible to do, ending up with an architecture as illustrated in Fig. 3. This sort of architecture 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 recognition algorithms (which tend to be crisp, 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 into the agent 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.[13] Of course, one does not need a specialised architecture to do this—one could just add 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, and the agents might need to induce these themselves. This is more complex but possible [8, 21].

Since the behavioural rules at the micro-level of agent-based simulation can be quite specific, some existing agent-based models have implicitly taken some aspects of context-dependency into account [1, 2, 18]. Each agent in such simulations

[13]This explicit approach is that of CYC [17].

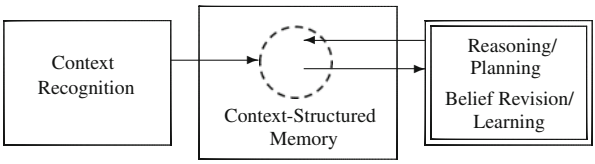

Fig. 3 The basic context-sensitive architecture for agents

does behave with respect to its local environment, and so will behave differently in different situations. However these do not distinguish context from any other perception that might influence behaviour, and hence do not specifically distinguish what can and cannot be shared between what kinds of situation.

8 Way Forward 3: Combining Qualitative and Formal Approaches

A fundamental value of science is that evidence 'trumps' theory, in the sense that if evidence and theory clash then it is the theory that should be 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 judged in a multitude of ways, with 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 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 then as if they do'.[14]

[14]Examples of economic assumptions that contradict common-sense have been extensively discussed, but this attitude is neatly encapsulated by a review received by Scott Moss and the author from the *Journal of Economic Dynamics and Control* which read in its entirety: "We do not see the point of not assuming that agents know the perfect model of the economy. This is not economics".

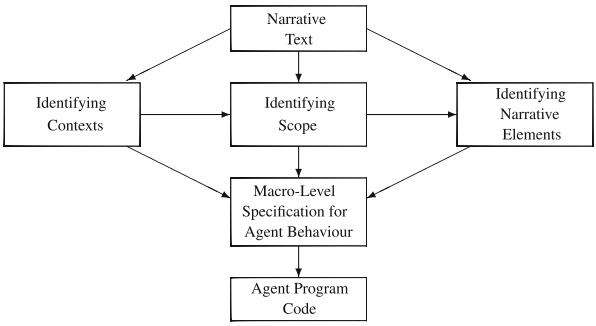

Fig. 4 An example (from Edmonds [5, 6]) of a process of narrative analysis separately identifying context, scope and narrative elements

In the last couple of decades experiments have shown that people do not act as the theory of neo-classical economic decision making would suggest, even in tightly controlled experiments with limited information and simple economic gain.

One problem about using qualitative and quantitative data together is that it has been difficult to use qualitative data in conjunction with formal modelling methods. A second problem is that qualitative evidence is often context-dependent and hence hard to incorporate into a generic model. However agent-based modelling is well placed to use qualitative evidence to inform the menu of behavioural strategies that might be used in given situations. There is now a growing stream of work on methods to make the process of the analysis of textual narrative data into behavioural rules suitable for an agent in an agent-based simulation transparent and systematised.[15]

Once these behaviours have been incorporated into a simulation at the micro-level, the simulation can be run to produce numbers that can be compared to macro-level quantitative data. 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. 4).

Thus explicitly recognising and including context-dependency in formal simulation models allows for a well-founded method for integrating qualitative, quantitative and formal modelling approaches in the social sciences. This enables some of the wealth of qualitative ethnographic, observational and interviewing work that is done in the social sciences to be used to enrich formal simulation models directly, 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, but facing context-dependency can aid the use of a wider range of evidence without abandoning rigour.

[15] See special issue of the *Journal of Artificial Societies & Social Simulation* [7].

9 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 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 over-simplification 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 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 about social phenomena, both official and informal. We are starting to have enough data to distinguish the different contexts and their 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, but can retain and use 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 over-simplify what we observe to obtain and use formal models. It has been a long time coming, since the old habits derived from a pre-computational age change slowly. However, the age of context-dependent modelling and manipulation is now within our reach. We no 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 on problems in an more integrated manner). It holds out 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 con-textualized 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. In: EPIA 2009, pp. 547–559 (2000)
3. Barwise, J., Perry, J.: Situations and Attitudes. MIT Press, Cambridge (1983)
4. Edmonds, B.: Learning and exploiting context in agents. In: Proceedings of the first International Joint Conference on Autonomous Agents and Multiagent Systems, (AAMAS), Bologna, July 2002, pp. 1231–1238. ACM Press, New York (2002)
5. Edmonds, B.: Complexity and context-dependency. Found. Sci. **18**(4), 745–755 (2013)
6. Edmonds, B.: Towards a context- and scope-sensitive analysis for specifying agent behaviour.In: Kamiński, B., Loloch, G. (eds.) Advances in Social Simulation, Advances in Intelligent Systems and Computing, vol. 229, pp. 319–332. Springer, Berlin (2013)
7. Edmonds, B.: Using qualitative evidence to inform the specification of agent-based models. J. Artif. Soc. Soc. Simul. **18**(1) (2015). http://jasss.soc.surrey.ac.uk/18/1/18.html
8. 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, vol. 4442, pp. 142–155. Springer, Berlin (2007)
9. Fagiolo, G., Dosi, G.: Exploitation, exploration and innovation in a model of endogenous growth with locally interacting agents? Struct. Chang. Econ. Dyn. **14**, 237–273 (2003)
10. Fowler, J.H., Smirnov, O.: Dynamic parties and social turnout: an agent-based model. Am. J. Sociol. **110**(4), 1070–1094 (2005)
11. Gilbert, N., Pyka, A., Ahrweiler, P.: Innovation networks - a simulation approach? J. Artif. Soc. Soc. Simul. **4**(3) (2001). http://jasss.soc.surrey.ac.uk/4/3/8.html
12. Harries, M.B., Sammut, C., Horn, K.: Extracting hidden contexts. Mach. Learn. **32**, 101–112 (1998)
13. Hayes, P.: Contexts in context. In: Context in Knowledge Representation and Natural Language, AAAI Fall Symposium, Nov 1997. MIT, Cambridge (1995)
14. Kaelbling, L.P., Littman, M.L., Cassandra, A.R.: Planning and acting in partially observable stochastic domains. Artif. Intell. **101**, 99–134 (1998)
15. 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)
16. 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, vol. 2116, pp. 221–234 (2001)
17. Lenat, D.B.: CYC—a large-scale investment in knowledge infrastructure. Commun. ACM **38**(11), 33–38 (1995)
18. Nunes, D., Antunes, L., Amblard, F.: Dynamics of relative agreement in multiple social contexts. In: EPIA 2013, pp. 456–467 (2013)
19. Simon, H.A.: Administrative Behaviour, a Study of Decision-Making Processes in Administrative Organization. Macmillan, New York (1947)
20. Tomasello, M.: The Cultural Origins of Human Cognition. Harvard University Press, Cambridge, MA (1999)
21. Widmer, G.: Tracking context changes through meta-learning. Mach. Learn. **27**, 259–286 (1997)
22. Xenitidou, M., Edmonds, B.: The Complexity of Social Norms. Springer International Publishing, Switzerland (2014)