

The geometric modelling of social frames and contexts

K. Kitto^a, **F. Boschetti**^b, **P. Bruza**^a

^a*Information Systems, Queensland University of Technology, GPO Box 2434, Brisbane, Qld, 4001*

^b*CSIRO Marine Research, Private Bag 5, Wembley WA, 6913*

Email: kirsty.kitto@qut.edu.au

Abstract:

How do humans respond to their social context? This question is becoming increasingly urgent in a society where democracy requires that the citizens of a country help to decide upon its policy directions, and yet those citizens frequently have very little knowledge of the complex issues that these policies seek to address. Frequently, we find that humans make their decisions more with reference to their social setting, than to the arguments of scientists, academics, and policy makers.

It is broadly anticipated that the agent based modelling (ABM) of human behaviour will make it possible to treat such social effects, but we take the position here that a more sophisticated treatment of context will be required in many such models. While notions such as historical context (where the past history of an agent might affect its later actions) and situational context (where the agent will choose a different action in a different situation) abound in ABM scenarios, we will discuss a case of a potentially changing context, where social effects can have a strong influence upon the perceptions of a group of subjects. In particular, we shall discuss a recently reported case where a biased worm in an election debate led to significant distortions in the reports given by participants as to who won the debate (Davis et al., 2011). Thus, participants in a different social context drew different conclusions about the perceived winner of the same debate, with associated significant differences among the two groups as to who they would vote for in the coming election. We extend this example to the problem of modelling the likely electoral responses of agents in the context of the climate change debate, and discuss the notion of interference between related questions that might be asked of an agent in a social simulation that was intended to simulate their likely responses.

A modelling technology which could account for such strong social contextual effects would benefit regulatory bodies which need to navigate between multiple interests and concerns, and we shall present one viable avenue for constructing such a technology. A geometric approach will be presented, where the internal state of an agent is represented in a vector space, and their social context is naturally modelled as a set of basis states that are chosen with reference to the problem space.

Keywords: Context, decisions, interference, incompatibility, geometric modelling.

1 INTRODUCTION

How do humans respond to their social context? This question is becoming increasingly urgent in a society where democracy requires that the citizens of a country help to decide upon its policy directions, and yet those citizens frequently have very little knowledge of the complex issues that these policies seek to address. Frequently, we find that humans make their decisions more with reference to their social setting, than to the arguments of scientists, academics, and policy makers.

For example, Davis *et al.* (2011) have used an elegant experiment to show the manner in which a biased ‘worm’ (or line that tracks performance) in an election debate might have a significant effect in an election campaign, and could even affect the outcome of the election itself. Generally a ‘worm’ tracks the response exhibited by a sample group of 20-30 voters, who watch the debate and use a handset to register their approval or disapproval of the current speaker. This data is averaged, and then broadcast by a TV provider to the viewers of the debate, adding a sense of drama to the contest. Davis *et al.* (2011) generated an artificial worm for the first UK election debate in 2010, which was biased in one of two directions; either pro-Clegg or pro-Brown. They showed each biased worm to a different audience (sampled from a population of students) and found a consistent biasing effect in the direction of worm bias when their subjects were asked to decide who had won the debate (see Figure 1).

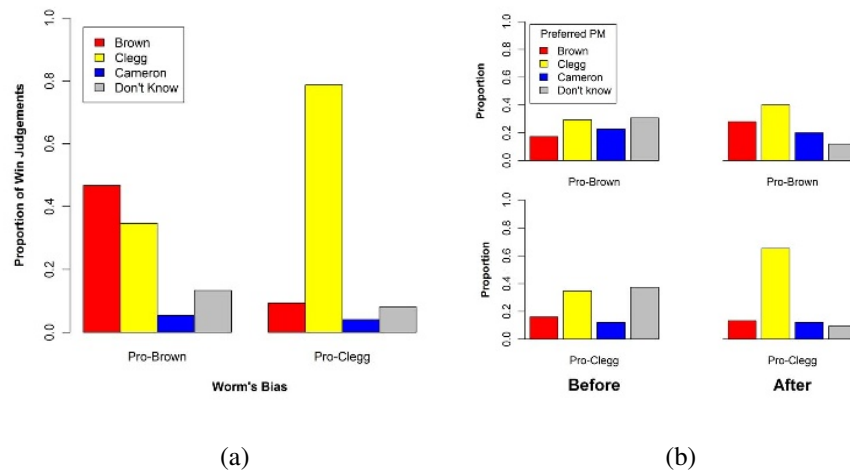


Figure 1. The results of a biased worm (a) in participants choice of who won the debate; and (b) participants predictions as to who they would vote for in the coming election. (From Davis *et al.* (2011) which is distributed under a Creative Commons Attribution License.)

This is a quick and simple quantitative example of something that is widely thought to be the case; the context in which an audience makes its decision as to the ‘winner’ of an election debate can affect that decision, at least at a statistical level. Such effects are likely to prove important when it comes to the modelling of large scale socio-environmental systems, but few quantitative models that explore such behaviour appear to exist. For example, attempts to incorporate the responses of humans to new policy settings attempting to mitigate climate change will require sophisticated models of human decisions in context, but our current modelling techniques are not likely to meet this challenge. How can the context of a decision be adaptively modelled, and different decision contexts compared in a society of decision makers? On an associated front, a wide range of results from psychology suggest that humans do not obey the standard assumptions of probability theory when making decisions (Busemeyer *et al.*, 2011) which makes statistical predictions of their likely responses problematic, especially if we utilise standard probabilistic models.

In this paper, we propose a new modelling approach, that allows us to model the decisions of a human *in a changeable context*. This technique allows for a natural treatment of contextual effects such as the one

found by Davis *et al.* (2011), but has the added advantage of naturally fitting with the psychological data (Busemeyer *et al.*, 2011). The extension of this model to a system of multiple agents will be proposed, and future avenues of research discussed.

2 A PERCEPTION IN CONTEXT

In this section we shall propose a simple model, inspired by the qubits of quantum theory (QT), which captures the manner in which a biased worm can influence the perceptions of an audience about the performance of a speaker.

Let us consider an agent, Alice, who is watching an election debate. We designate her ‘unbiased’ state as occurring when she sees a fair worm (i.e. an unadjusted one). In this state, Alice will have a particular probability of perceiving a current speaker as either performing well or not, but she will not reflect upon this perception until asked to declare who she believes is currently winning the debate.

For the sake of simplicity, we shall assume that there are only two protagonists in the debate, we shall call them Brown and Clegg. When we ask Alice which of the two protagonists she believes is performing best, Alice will respond either with Brown or with Clegg, but before this point we model her as being in a state of *superposition*. This gives a weighted sum of her two possible responses: she might believe that Brown is performing best, we denote this option with the notation $|0\rangle$, or; alternatively, she might believe that Clegg is the best speaker ($|1\rangle$).




Figure 2. An agent attempts to decide upon a course of action. (a) While watching the debate they are in a superposed state; (b) when asked to nominate the winner of the debate, they ‘collapse’ into a pure state $|1\rangle$ (or alternatively $|0\rangle$). (c) The changing context of a decision. The probability of Alice deciding that Brown is winning the debate changes between the two biased worm contexts b and c .

Within a given context, p , Alice will have a certain probability of deciding that Brown is the best speaker. However, a change in context might change her choice. In particular, a pro-Clegg biased worm might lead to her deciding that Clegg was performing better. QT provides a very natural way of incorporating such a change in context into the current state of our agent *geometrically*. We represent the unbiased Alice as

$$|A\rangle = a_0|0_p\rangle + a_1|1_p\rangle, \text{ where } |a_0|^2 + |a_1|^2 = 1, \tag{1}$$

a situation that is illustrated in Figure 2(a). Here, $\{|0_p\rangle, |1_p\rangle\}$ define an orthonormal basis, the inner product of these basis vectors returns 0 or 1: $\langle 0_p|0_p\rangle = \langle 1_p|1_p\rangle = 1$ and $\langle 1_p|0_p\rangle = \langle 0_p|1_p\rangle = 0$. These basis vectors *define the current context* of our agent, which in turn effects their decisions in a process of *measurement*. Measurement of the state (1) is defined with respect to a projection operator V , where, for the two dimensional case outlined above

$$V = |0_p\rangle\langle 0_p| + |1_p\rangle\langle 1_p| = V_0 + V_1. \tag{2}$$

According to the quantum formalism, if we ask Alice who is winning the debate, then she will decide that

it is Brown with a probability of

$$P(b) = \langle A|V_0|A \rangle \tag{3}$$

$$= \langle A|0_p \rangle \langle 0_p|A \rangle \tag{4}$$

$$= (a_0^* \langle 0_p|0_p \rangle + a_1^* \langle 1_p|0_p \rangle) \times (a_0 \langle 0_p|0_p \rangle + a_1 \langle 0_p|1_p \rangle) \tag{5}$$

$$= |a_0|^2 \tag{6}$$

and similarly, that Clegg is winning with a probability of $P(c) = |a_1|^2$. The assumption in (1) that the squared coefficients of the basis vectors sum to 1 allows for the treatment of these values as probabilities since $0 \leq P(b) \leq 1$ (and similarly for Clegg).

This approach makes use of a geometrical notion of probability (Isham, 1995). In standard probability theory, probabilistic outcomes arise from our lack of knowledge as to what has actually occurred; Alice has made her choice but we do not know what it is. In contrast, the quantum structure of probability arises through reference to the Pythagorean theorem; Alice has a current state of mind, but it may yield very different outcomes in a different context. This alternative structure can be seen with reference to Figure 2(a), where the ‘length’ of Alice’s state of mind (defined as 1 in equation (3)) is related to her probabilities of action $|a_0|^2$ and $|a_1|^2$, via a right angle triangle relationship. Probability is thus profoundly different within this framework. Rather than arising from our lack of knowledge about the agent’s intended action, probability can arise from a genuinely undecided agent. A context must be present before predictions can be extracted, and this will influence our agent’s choice.

We shall now make use of a key feature of the quantum formalism in the description of our agent’s post-decision behaviour. According to QT, the measurement of a quantum system ‘collapses’ the state of that system into an eigenstate of the measurement that was performed. Thus, in a quantum inspired model, if Alice responds to a query that she thinks “Clegg is winning the debate”, then afterwards she will be in the state $|1_p \rangle$. Indeed, if we were to *immediately* ask her again who she thought was winning then we could predict with certainty (i.e. with probability 1), that she would make the same decision. This state is represented in Figure 2(b), and ensures that any further questions we ask of Alice will be historically contingent upon the questions that we have already asked. Of course such a situation is ideal in social analysis and be unlikely to work out quite as neatly as is the case with a physical quantum system. Indeed, it is likely that once Alice’s state has collapsed to this new state her context will have changed. We leave this complication for further work, noting that the example above is intended more as an explanation of the formalism than as a realistic example.

We see that in this model, context is not just a label for missing information (as is frequently the case in more standard treatments of contextual behaviour (Brézillon, 1999)). We can illustrate this point by extending it to represent the biased worms. Consider Figure 2(c), which is an elaboration of Figure 2(a), and represents Alice’s opinion of the debate in three different contexts. The first is the unbiased state that was represented by equation (3). However, we can now see two other contexts represented using two different *basis states*: the pro-Brown one (red), and; the pro-Clegg one (yellow). While Alice’s state has not changed, we can immediately see through reference to the diagram that the probability of Alice deciding that one or the other of the protagonists is winning the debate has changed with the shift in context (i.e. Alice’s exposure to a biased worm). Thus, in the case of a pro-Brown biased worm, we see that Alice’s state is now represented as:

$$|A \rangle = b_0|0_b \rangle + b_1|1_b \rangle, \text{ where } |b_0|^2 + |b_1|^2 = 1, \tag{7}$$

and the pro-Clegg worm leads to a situation where

$$|A \rangle = c_0|0_c \rangle + c_1|1_c \rangle, \text{ where } |c_0|^2 + |c_1|^2 = 1. \tag{8}$$

Recalling that a basis state of $|0 \rangle$ refers to a situation where Brown is judged the best performer, we see that in the pro-Brown basis, Alice has shifted towards a higher probability of judging Brown to be the winner (a scenario that is represented graphically in Figure 2(c) by the greater length of b_0 , i.e. $|b_0| > |a_0|$). A similar situation occurs in the pro-Clegg basis (recall that $|1 \rangle$ denotes a judgement of Clegg

as the best performer, and so the observed value $|c_1| > |a_1|$ increases the probability of a participant in the trial judging Clegg the winner in the Clegg biased context). Thus, in shifting the angle of our basis, we can naturally incorporate the worm's bias into this model, so modelling something of the effect discovered by Davis *et al.* (2011). Note that the shift in perception depicted in Figure 1(a) shows that on average the subjects judged Clegg to be performing better than Brown. Alice (as she is depicted in Figure 2) is one of the minority of subjects who prefer Brown in her unbiased state, and thus her shift towards Clegg in the pro-Clegg case still leaves her fairly un-impressed by Clegg; she is almost equally likely to decide that Clegg or Brown are winning the debate even when biased towards Clegg by a pro-Clegg worm.

3 A MULTI-AGENT MODEL




Figure 3. A set of agents $\{|A\rangle, |B\rangle, |C\rangle \dots\}$ all participating in the election debate experiment, each with a different unbiased state. The biased worms are represented using the red and the yellow axes.

This simple model can be naturally extended across a set of multiple agents. Here we shall briefly sketch out the basic idea of the multi-agent model, leaving many of the details to future work. In Figure 3 we have drawn a collection of agents, all with a current state (which is dependent upon the individual but can change in time), and forming an impression as to who is winning the debate. For the sake of simplicity, we have chosen to represent all three bases (i.e. all three contexts of worm bias) as the same across all agents. This is likely to be an unrealistic assumption, however we shall keep it for the sake of development, noting that it could be dropped in future modelling. The extension to different frames for each agent will be straightforward computationally, if somewhat non-trivial psychologically. We shall return to this point at the end of this paper. We note that both an agent's state, and the frame representing a worm's bias would be allowed to change over time in this model. We can quickly see by an examination of Figure 3 that while each agent has their 'own' private internal state, population effects can emerge as a result of a biased worm which will affect the responses given by individuals in the population. Thus, this simple model can illustrate the manner in which the context of a biased worm might affect population level statistics as to who is winning an election debate.

We shall now turn to a consideration of how two related questions might interact in this model.

4 A RELATED QUESTION

In Figure 1(b) we saw that participants in the trial who had been exposed to a biased worm exhibited a different distribution of judgements about both the winner of the debate, and their likely future voting behaviour. In this section we shall consider the manner in which related questions can be expected to interfere with one another. First however, we must consider two different types of related question: (1) a *compatible* related question; and (2) an *incompatible* related question.

We define compatible questions as those which can be represented in one frame (or basis). As described by Busemeyer *et al.* (2011), when the feature sets used to represent a frame overlap then it is possible to consider the two questions simultaneously. Thus, two compatible questions can both be asked of a person without any expected interference between them. In the case of our electoral worm, it seems likely that the answers to the two questions: who won the debate? and; who will you vote for in the coming election? can be represented in one geometric space listing all four response options, since the two bases are compatible (as they are both the same). In this situation we can choose to ignore other factors as 'irrelevant', at least as a first approximation.

In contrast, if we now decide to concurrently consider the public debate about human induced climate change (and how policy might affect voting behaviour), then the problem of representation becomes much more complex. There will be many respondents who do not normally consider climate change when making their electoral decisions, but if we ask them questions about climate change *before* asking them questions about who they will vote for, then we will need to consider the manner in which these two questions might interfere with one another. It seems likely that in this case, the set of responses that might be given by a subject about their belief in the need for action on climate change (e.g. yes, maybe, no) will influence their intentions to vote in the coming election (for, e.g. Brown, Clegg, Cameron, Don't know), potentially changing their response. As these two questions must be represented using two different frames, it will be necessary to combine them into a larger (12-dimensional) vector space in order to consider them concurrently, but responding to one question will collapse a subject's state onto the relevant axis, which will in turn *interfere* with their future answer to the next question (Busemeyer *et al.*, 2011). To do this we note that interference effects are typified by a violation of the Law of Total Probability (LTP), which states that if $\{B_1, B_2, \dots\}$ is a finite measurable sample space, whose union is the entire probability space, then for any A in the same probability space, $P(A) = \sum_n P(A \cap B_n) = \sum_n P(A|B_n)P(B_n)$. In a situation where LTP is violated, quantum theory provides an extra *interference* term which describes the manner in which two different decisions influence one another in the mind of their maker. While this term can be motivated theoretically (Khrennikov, 2010), it is possible to write this term down without reference to the full apparatus of QT. For the above case, if data could be found that showed a violation of LTP for the case where subjects exhibited a difference between their predicted voting patterns even when summed over all possible climate change responses, then we would have good reason to motivate a quantum inspired approach. Indeed such an approach could fit an interference term to such data which would allow us to continue with probabilistic reasoning (albeit quantum). Thus, the probability of a vote for Brown $P(B)$, given the three different climate responses $\{y = \text{'yes'}, n = \text{'no'}, m = \text{'maybe'}\}$ could be modelled as:

$$P(B) = P(B|y)P(y) + P(B|n)P(n) + P(B|m)P(m) + I(|A), context), \quad (9)$$

where $I(|A), context)$ is an interference term which depends upon both the subject's current state ($|A)$) and the measurement of their *local* frame (i.e. their context), which can range from -1 to +1.

While such a fit can most likely also be approximated using Bayesian techniques (Pearl, 1988) which would adjust the probability of a participant's vote according to a new prior probability, there are a number of reasons to pursue the quantum model as an alternative. Firstly a Bayesian approach does not give a clear indication of what causal mechanisms underlie this change in behaviour (Jones and Love, 2011), and so does not help us to understand the dynamics of the social system we are trying to model. While the quantum approach does not attempt to describe the actual cognitive behaviour that leads to the violation of the LTP, it does provide us with an understanding of the general mechanism underlying this behaviour (Busemeyer *et al.*, 2011) and so holds the possibility of deeper understanding. Secondly, a Bayesian approach assumes that a given subject is in a definite state. Conditioning may change that state into another definite state, but answering a question corresponds to the recording of a judgement corresponding to that definite state. In contrast, a quantum approach represents the cognitive state of an individual as indefinite (superposed), recognising that often we don't jump from state to state, but instead we have to make up our minds. Our judgement is constructed (made definite) when the question is asked (i.e. it is not simply a representation of our existing judgments). Finally, it is interesting to speculate about whether datasets might be envisaged which can only be well modelled using a quantum rather than a Bayesian approach. For example, the quantum approach leads us to ask whether there might be situations where due to a change in context the prior probabilities of Bayesian analysis *cease to exist*, or where *new unexpected variables arise*. This scenario poses problems for both standard probability theory as well as the Bayesian, since the probability spaces of classical probability theory do not appear to recognise the concept of a random variable that is only defined in some instances.¹ In contrast, this scenario is represented in the quantum formalism which explicitly recognises the manner in which a change in context can induce a new probability space. Future work will investigate the differences between these two representations in more detail, and search for data sets that can highlight their respective behaviours.

¹Although this concept of different probability spaces was recognised by Kolmogorov when he introduced the notion of a probability space (Khrennikov, 2010).

5 DISCUSSION AND FUTURE DIRECTIONS

It is important to appreciate that the model proposed in this paper offers a significant break from standard probabilistic reasoning. While contextual probability shares some features with conditional probability, the two are not the same. Conditional probabilities can only be compared if they are defined in an equivalent Kolmogorovian probability space. In contrast, the contextualised probability structures utilised in the current approach is defined in an explicitly represented context space, which specifies the variables that exist in that specific context (Khrennikov, 2010). Thus, the adaptive interaction that occurs between an agent and their social context are somewhat strangely treated by standard probability, which must assume that all opinions formed by an agent are compatible and defined in the one essentially universal space. This assumption is not required in a quantum-inspired approach, which allows for, among many other effects, non-commutative scenarios (where the order in which an agent makes their judgements matters), incompatible judgements, and interference between judgements (Isham, 1995).

Representation of the agents and of their contexts via a ‘natural’ choice of basis will be a key aspect of this model which must be carefully considered in future work. It is likely that much can be learned here through close work with social psychologists in terms of natural choices for basis states etc. (Augoustinos *et al.*, 2006). For now, we note that there are a number of likely candidates for the basis states, including for the model presented in this paper: political orientation; environmental concern; fiscal concern; conservative vs liberal views etc. However, it is important to realise that different representations of an agent’s internal state and their external context will likely prove important in different modelling scenarios; no one representation is likely to prove sufficient in the complex world of social modelling.

It is anticipated that this model will provide a powerful new modelling methodology for investigating political effects such as media bias, push polling, public opinion of policy, even phase transitions in public opinion etc. In adopting a technique that to some extent externalises the opinions and decisions of agents, we can start to understand the manner in which the social context of an agent can have a strong influence upon that agent. While such effects are well known in social psychology, our modelling technology is rarely up to the challenge of their representation, and this appears to have led to a certain amount of cynicism exhibited by social scientists surrounding the new models and techniques proposed by modellers. It is hoped that by treating the social context of human behaviour as a core effect, rather than the all-too-frequent add on, we might start towards a genuine dialogue between two fields that frequently ‘talk past’ one another.

ACKNOWLEDGEMENT

This project was supported by the Australian Research Council Discovery grant DP1094974. Thanks to Colin Davis for providing the authors with the raw dataset for the results discussed in Davis *et al.* (2011).

REFERENCES

- Augoustinos, M., I. Walker, and N. Donaghue (2006). *Social Cognition* (2 ed.). London: Sage.
- Brézillon, P. (1999, May). Context in problem solving. *Knowledge Engineering Review* 14, 47–80.
- Busemeyer, J. R., E. Pothos, R. Franco, and J. Trueblood (2011). A Quantum Theoretical Explanation for Probability Judgment Errors. *Psychological Review* 118(2), 193–218.
- Davis, C. J., J. S. Bowers, and A. Memon (2011). Social influence in televised election debates: A potential distortion of democracy. *PLoS ONE* 6(3), e18154.
- Isham, C. J. (1995). *Lectures on Quantum Theory*. London: Imperial College Press.
- Jones, M. and B. C. Love (2011). Bayesian Fundamentalism or Enlightenment? *Behavioral and Brain Sciences*. In Press.
- Khrennikov, A. Y. (2010). *Ubiquitous Quantum Structure: From Psychology to Finance*. Springer.
- Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems*. San Francisco, CA: Morgan Kaufmann.