Advances in Complex Systems

ATTITUDES, IDEOLOGIES AND SELF-ORGANISATION: INFORMATION LOAD MINIMISATION IN MULTI-AGENT DECISION MAKING

KIRSTY KITTO

School of Information Systems, Queensland University of Technology
2 George Street, Brisbane, 4000, Australia.
kirsty.kitto@qut.edu.au

FABIO BOSCHETTI

Marine Research, Commonwealth Scientific and Industrial Research Organisation (CSIRO)
Private Bag 5, Wembley, 6913, Australia
School of Earth and Geographical Sciences, The University of Western Australia, 35 Stirling Highway, Crawley 6009, Australia
Fabio.Boschetti@csiro.au

Received (received date)
Revised (revised date)

Sophisticated models of human social behaviour are fast becoming highly desirable in an increasingly complex and interrelated world. Here, we propose that rather than taking established theories from the physical sciences and naively mapping them into the social world, the advanced concepts and theories of social psychology should be taken as a starting point, and used to develop a new modelling methodology. In order to illustrate how such an approach might be carried out, we attempt to model the *low elaboration* attitude changes of a society of agents in an evolving social context. We propose a geometric model of an agent in context, where individual agent attitudes are seen to self-organise to form ideologies, which then serve to guide further agent-based attitude changes. A computational implementation of the model is shown to exhibit a number of interesting phenomena, including a tendency for a measure of the entropy in the system to decrease, and a potential for externally guiding a population of agents towards a new desired ideology.

Keywords: context; attitudes; information; quantum decision theory; self-organisation

1. The Challenges of Modelling Social Systems

How are we to model human social interactions? The sheer complexity of social systems has historically left them beyond the reach of analytical descriptions and numerical simulations, however, recent years have seen a shift towards this frontier of analysis. With an ever growing supply of social data, increasing computational power, and a far more sophisticated understanding of human cognition, advanced models of social systems are fast becoming feasible. The importance of this frontier is also increasingly apparent; as our species grapples with complex problems such as climate change, the global financial crisis, and the forecasted age of scarcity, all of which have an inherently social theme, we must find appropriate models and methods that can be integrated with existing physical, financial and biological approaches. This integration must happen in both directions however, and those attempting to mathematically and computationally simulate social systems must pay attention to the lessons that have already been learned in fields such as psychology and sociology.

For example, consider the challenges facing policy makers as they attempt to adapt to and mitigate for the effects of climate change. While it is essential that we have sound models of the bio-physical processes underlying this phenomenon, it is equally important that we understand how humans will respond to the information presented to them about this global problem in a wide variety of fora. Thus climate models; policy initiatives; media campaigns; and word of mouth will all affect the long term behaviour of the people in the system. A set of climate models that somehow incorporated the opinions and attitudes, and hence likely decisions and actions, of humans, would significantly aid policy makers as they attempt to guide our society through a period of intense change and readjustment. Thus, while it is essential that we understand the physical processes underlying climate change, we ignore the attitudes and opinions of the world population at our peril. Furthermore, changing the attitudes of a population is likely to prove one of the more controllable modifications that might be attempted by the governments of the world, as long as they can be provided with adequate models of the cognitive processes involved. It is still highly challenging to predict the outcomes of direct interventions in the earth's climate, but if the attitudes of the global population can be shifted towards a scenario where it desires to act in such a way that a low emissions outcome can be obtained then direct intervention in the climate system may prove unnecessary; human behaviour is an obvious intervention point, and so a capability to dynamically model the changing attitudes and opinions of the humans in the system is fast becoming highly desirable.

Social psychology provides a wide range of results and techniques that have been used to understand human social behaviour [7, 50, 1, 28]. It focuses on social actions and interrelations, and looks at the way in which personality, values, and cognitive states affect, and are affected by, social structure and culture. As such, it provides a wide range of data, theories and initial models that complex systems scientists can make use of if they wish to start modelling social processes. However, many of the established theories of this field are somewhat ill defined, making them difficult to use in quantitative modelling. Here we would like to suggest that rather than directly applying techniques from fields such as Statistical Mechanics [18] and Network Theory [2], which tends to generate models which oversimplify their assumptions about human cognitive states, a far richer set of models can be created if effort is made to incorporate the lessons that have already been learned within psychology about human cognition and behaviour. Indeed, making such an effort is likely to result in models that gain a wider acceptance in psychology generally, as they will contain core concepts and features of the field that are not at present well treated analytically. Thus, this paper is intended to be a 'proof of concept' contribution towards the generation of a genuine interaction between modellers and social psychologists.

To this end, we shall introduce a mathematical model of attitude change, which takes key results from social psychology as its starting point. Section 1.1 introduces the notion of attitudes, and attitude change, and so traces out the scope of effects that analytical and computational models must navigate before they can be considered as psychologically plausible. In particular, we shall emphasise the manner in which a *social context* can drive the behaviour of human agents. We then introduce a geometric approach in sections 2–4, which is capable of capturing many of the key aspects of this socially adaptive behaviour. Our model takes *Quantum Decision Theory* (QDT) [17, 16] as its starting point, due to its implicit capacity to represent the effect of context upon a decision. This theory is has been shown capable of providing a unified explanation for many of the so called 'violations' of rational decision theory that are exhibited by individual humans. We shall first introduce this QDT, and then extend its theoretical basis to the case of multiple decision making agents. We shall then present the results of a computational implementation of the model in section 5, discussing the manner in which this model is capable of exhibiting behaviour typical of Guided Self-Organisation (GSO) in section 6. Thus, we shall show that it is possible to formalise the manner in which groups of agents can be guided by the information presented to them, which is in turn the result of a self-organising process that arises between the agents themselves. We shall start with a consideration of the notion of an attitude.

1.1. *Attitude Changes and the Social Context of an Agent*

The social behaviour of individuals is frequently driven by their internally held *attitudes*, a construct essential to the field of social psychology [3, 1, 28]. For example, privately held attitudes play a critical role in people's personal choices about their health, education, social groups, and housing, as well as the importance they attribute to national issues such as the environment, immigration and state security. They help to determine a wide variety of highly consequential outcomes at a social scale, and are a fundamental construct in the field [46]. Indeed, even when writing the first *Handbook of Social Psychology* in 1935, Allport claimed that the attitude construct was *the* most indispensable concept in the field [3]. This fundamental status makes attitudes a prime candidate for mathematical and computational analysis; is it possible to formalise the way in which the attitudes of a society of agents changes and evolves?

Attitudes are highly contextual, and this makes them extremely difficult to model formally. How will a given person think about 'global warming' vs 'climate change'? What if their daughter has just had her house flooded? Or if they are about to make a very large tax payment that includes a carbon component? People's attitudes are not static immutable objects, but change in response to persuasion [51], and the attempt to maintain cognitive consistency [20]. We often express different attitudes and opinions in accordance with the social scenario that we find ourselves to be in [4, 10], and it is frequently the case that an explicitly expressed attitude is quite different from an internally held one [31].

The review article by Petty & Wegener [46] discusses the historical process that led to the current understanding of attitude change in psychology. It draws attention to the manner in which the contextuality of the process of attitude change in itself made the construction of a theory highly challenging. Thus, while an extensive collection of empirical results had been published by the 1970's, very little conceptual coherence had been established in the varying psychological accounts of the process. The same variable (e.g., source credibility, or mood) was frequently found to have a different effect in different scenarios, and could sometimes even produce the same persuasion outcome via different processes in different situations [46]. The late 1970's then saw the emergence of two attempts at unifying and general accounts of attitude change; the Elaboration Likelihood Model (ELM) [45]; and the Heuristic-Systematic Model (HSM) [19]. Both of these models utilised a dual-process approach, taking mental effort as a key switching variable.[a] This development meant that the many early models initially proposed to explain the varying empirical results could be brought into a general framework; they were not competing or contradictory but operated in different circumstances. To this day, the framework created within these models is considered as the standard when it comes to the modelling of attitude change.

In essence, both models posit that some processes of attitude change require relatively high amounts of mental effort, resulting from situations where individuals are motivated to pay attention to a message, or have the cognitive capacities to consider it carefully. In these *high effort* or *high elaboration* processes, people's attitudes will be determined by an effortful examination of all relevant information, and so changing them will expend high amounts of cognitive energy. In contrast, other *low effort* or *low elaboration* processes of persuasion require relatively little mental consideration by the persuadee, resulting in attitudes determined by factors like emotions, 'gut feeling', liking, and reference to authority. Large swings in attitude can be produced via either process. However, the changes induced by the high mental effort processes are postulated to be more persistent, resistant to counter-persuasion, and predictive of behaviour in the long term than low effort attitude changes.

The difference between these two processes has a number of implications for public policy. In an era of high-frequency press reporting periods (i.e. the 24 hour

[a]Despite a slight difference in emphasis (the ELM high effort route arises from the cognitive processing of a message, whereas the HSM emphasises the "effort exerted in comprehending message content, not effort exerted in cognitive responding or thinking about message content" [25]), both models are highly similar, and they can generally accommodate the same empirical results.

news cycle) we have entered a climate where low effort attitudes appear to predominate [49, 57], and the transitional nature of this process could be seen to result in the apparent increase in undecided or swinging voters in the modern age. This in turn has led to dramatic shifts in public opinion about issues such as climate change, immigration etc. that often catch policy makers by surprise; how are they to predict these often illogical and highly emotional attitudes?

There are few analytical models capable of describing the dynamics of low elaboration attitude change. One computational implementation of attitude change has been created using the framework of the ELM [43]. This model was used to force a more accurate specification of the ELM, which was still largely heuristic, and so emphasised the advantage of simulations for the purpose of theory formalisation. This same approach has since been used to model social phenomena such as environmental campaigns [42] and so provides a very interesting first step towards general models of attitude change in context. However, we see a key weakness with this approach, lying in its treatment of low elaboration processes. Specifically, while there are many variables working together in this model, leading to nonlinear effects and indeterminacies that are hard to predict, there is no uncertainty in the model itself; an agent will always respond in the same manner to a situation that is identical, and we do not think that this approach is in keeping with the genuine contextuality of low elaboration attitude change. Although Mosler & Martens do recognise this determinism as a weakness and intend to implement random generators that drive individual behaviour with a well defined variance [42], this in itself raises an interesting question as to what kind of uncertainty is appropriate in a social simulation.

1.2. *Cognitive Uncertainty in Social Decision Making*

Does the uncertainty in a simulation of social behaviour arise from a lack of knowledge on the part of the modeller, or does it result from an undecided agent? Thus, is uncertainty internal to an agent and dependent upon their cognitive state, or is it something that arises from the essentially external process of observing that agent? It is likely that there will be scenarios where both forms of uncertainty arise in social modelling, but we are left wondering if current models are capable of treating both forms of uncertainty well.

More specifically, current modelling approaches tend to assume that agents have a well defined but epistemologically unknowable state; as modellers we "know too little" [14] about that state, and this is the cause of the uncertainty most commonly incorporated into models of the system. However, people are frequently *genuinely undecided* about issues and courses of action to follow; they have yet to make up their minds and so their state is in some sense undefined. Philosophically, this difference is quite profound. An agent who has already formed an attitude towards a social issue (which we admittedly might not know about) may exhibit very different behaviour from one who has not considered their response to that same issue. Indeed, the second agent may, from a given initial state, respond very differently to a question, survey, piece of information etc. depending upon how it is framed [58, 14, 60, 17]. This *contextuality* of a social agent is not something that is well treated by current analytical approaches to social modelling, but it seems likely that this lack will have the largest impact upon models of low elaboration attitude change. Thus, while an agent who has engaged in a high elaboration process will commonly settle upon a specific attitude that will be difficult to change (i.e. they will be decided), an agent who has reached the same attitude through a low elaboration process will not prove to be so firmly decided. Indeed, it seems likely that if they were then subjected to a new social context then they may well prove to have a very different attitude, and this effect is frequently observed in psychology [4, 10]. A modelling problem presents; if low elaboration thought processes are becoming more common in our society, and are highly contextual, then how are we to analyse their effects upon that society? Similarly, we might ask how the genuine uncertainty of an agent can be represented mathematically to incorporate contextually dependent social decisions? This paper is a first step towards answering questions such as these.

For this reason, the focus of this paper will be upon low elaboration attitude change. We shall attempt to capture the manner in which the social context of an agent affects the attitudes that they have towards a given issue, and how this in turn affects their eventual decisions and actions with regards to that issue.

An example will help to clarify the concepts just discussed, and will serve to illustrate the coming discussion. Let us consider an agent, called Alice, who, from a particular *cognitive state*, representing her current attitude towards some issue with a social component, must decide whether or not to act in some way. She might be answering a question, she might be voting for a particular politician, perhaps she has to work out if she should immunise her child, or drive to work. In order to maintain generality in the model that follows we shall term all of these different decisions as *actions.* However, Alice has not yet made her decision, and how she eventually does choose to act will depend upon both her own attitudes (implicit and explicit), and on the attitudes of those that surround her. Thus, an extra set of factors will affect our agent, and henceforth these will be referred to as the *context* of the agent. Note that any agent with the same initial cognitive state may choose a different course of action if they find themselves in a different context, and this uncertainty should lie in the mind of the agent, not of the modeller, a situation that we feel reflects the true uncertainty of human decision making (and its modelling). Finally, we note that this example also draws attention to the recursiveness of attitudes; the actions of Alice will feed back into the context of other agents in the system, so changing the context of every agent's decisions.

We shall now attempt to incorporate this behaviour into a conceptually simple, and yet psychologically plausible, *geometric* model of an agent's attitudes as they are affected by their social context.

2. An Agent in Context

We shall begin with a consideration of an agent A, whose cognitive state is represented as a vector $|A\rangle$ in Dirac form.[b] Note that A may not have direct access to their cognitive state (i.e. A may not be aware of this state for reasons of context to be explained below). This starting point is the basis of a new *Quantum Decision Theory* (QDT), which has been widely explored and described in a set of recent works (e.g. [47, 17, 16]), and provides a unified explanations for many of the so called 'violations' of rational decision theory that are exhibited by humans. Here, we extend the initial set of QDT models by assuming that A refers one particular agent who is considering a set of social *issues*, as represented by the cognitive state $|A\rangle$. This space in which $|A\rangle$ resides could be very high dimensional, or it might be considerably smaller. Its general structure will depend upon the nature of the issue under consideration and will change if a different issue is being considered by the agent. Thus, for example, A might be considering the issue of climate change, or which candidate to vote for in an election, or how to school their child. The model that we construct shall consider one issue alone, thus all of the agents within a system will be making decisions within one vector space, although this may encompass a very complex issue (e.g. how an agent intends to vote in a coming election given their attitude towards climate change).[c]

For the issue that is currently under consideration, if A has decided to act then we shall denote this state of action using the symbol $|1\rangle$, to represent a situation where it is *true* that they have *chosen to act* (in contrast to a state of inaction which we denote as $|0\rangle$). However, a decision to act (or not) depends on the context in which it is made; we are immediately faced with the dilemma that our social agent cannot be described as making a decision without reference to a context. Thus, we must specify that *within a given context*, termed p say, our agent will have a certain probability of acting, and note that a change in context might change this probability. Thus, the decisions to *act* or *not to act* in the context p are represented within a slightly expanded notation: $|1_p\rangle$, $|0_p\rangle$.

In what follows, we propose that a *geometric model* of an agent in a changeable context provides a minimal framework capable of dynamically modelling low elaboration attitude change. This model takes inspiration from quantum theory (QT), which has a probability structure that is markedly different from that of classical

[b]Dirac notation was invented as a shorthand for quantum physics [33]. It explicitly allows us to represent a vector a using a *ket*, $|a\rangle$, with the transpose given as a *bra* $\langle a|$. This allows for an immediate recognition of the inner product between two vectors $\langle a|b\rangle$ (a *bra-ket*) and of the outer product $|a\rangle\langle b|$. We use it here to make explicit the difference between an agent A and their cognitive state $|A\rangle$, a distinction that will become important when the effects of social context are discussed.

[c]We note that it is possible for different issues to have overlapping vector spaces, although we do not anticipate that they will all overlap completely, which can itself lead to *interference effects* in quantum physics [33], a point that we shall not discuss here, but which has been used in the psychological literature to explain a number of apparent 'violations of rational decision theory' [17].

probability theory [33]. For example, classical probability theory generally assumes that a system has some specified state which measurements then ascertain (with uncertainty arising due to a lack of knowledge). However, as was discussed above, much of the uncertainty involved in social modelling is of a different form; it is ontological rather than epistemological, which means that very different responses can be obtained from the same agent if they are asked the same question in a different way, or even in the same way. QT provides a very natural formalism for describing such a state of affairs. Indeed, quantum systems behave in a markedly similar manner, and this is the motivation underlying the geometric nature of our model.

Geometry provides a very natural way of incorporating the importance of context into the representation of the current state of a agent, via the Pythagoras' theorem. In what follows, we shall represent both the cognitive state of an agent, and that of their context, explicitly. This is achieved in our model through its use of a vector in a Hilbert space to represent the cognitive state of the agent, rather than that of a point in a configuration space.

Hilbert spaces are vector spaces that can be of a real or complex form. They must have an inner product defined, and must form a complete metric space with respect to the distance function induced by the inner product [33]. Using this framework, we can immediately see that the set of states $\{|0_p\rangle, |1_p\rangle\}$ are orthogonal, and so can be taken to define an orthonormal basis of the 2D subspace representing the agent's decision of whether or not to act. This subspace lies within the higher dimensional space representing the agent's complete cognitive state, and so may not be spanned by the 2D decision subspace. Indeed, it is quite possible that a higher dimensional decision subspace will be required for some actions (when, for example an agent might be choosing between three orthogonal or mutually exclusive alternatives). While the current simplification to 2D might seem inappropriate, spectral theory [33] allows for the representation of states lying in these higher dimensional spaces as a sum of orthonormal projection operators (i.e. a sum of decisions to act or not to act). This suggests that far more complex decisions can be represented within this more simplistic formalism, and future work will turn to this extension. Despite this potential future complexity, recognising that $\{|0_p\rangle, |1_p\rangle\}$ is a basis means that the inner product (denoted in Dirac notation as $\langle .|.\rangle$) returns 0 or 1. That is: $\langle 0_p|0_p\rangle = \langle 1_p|1_p\rangle = 1$ and $\langle 1_p|0_p\rangle = \langle 0_p|1_p\rangle = 0$. Our orthonormal basis thus represents the set of 'act' or 'not act' decisions to be made by our agent in the context p. We note that in this case orthogonality is entirely appropriate as an agent cannot both 'act' and 'not act' at the same time.

With this added formalism, it is now possible to model the cognitive state of an undecided agent. A is certain to do something. This implies that the probability of A acting *or* not acting that is equal to 1 (as is standard). As the length of the cognitive state $|A\rangle$ corresponds to A's total probability of acting, we are left with one obvious candidate [33] for the representation of the cognitive state of our agent,

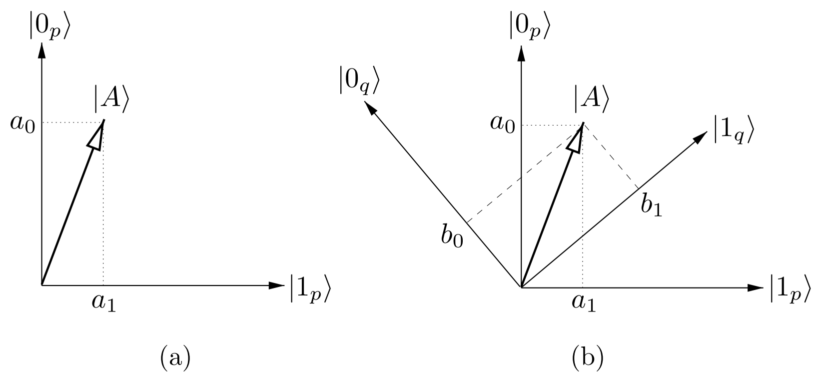

Fig. 1. An agent attempts to decide whether or not to act. (a) Their probability of action is proportional to the length squared of the projection of their state onto the axes labelled $|0_p\rangle$ (no action) and $|1_p\rangle$ (action); (b) The changing context of a decision. The probability of the agent acting changes between the two depicted contexts, which can immediately be seen by the different lengths of the projections from the state $|A\rangle$ onto the two different 'act' axes $|1_p\rangle$ and $|1_q\rangle$.

defined with respect to the context p:

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

a situation that is illustrated in Figure 1(a). With reference to Figure 1(a), we see that in the context p our agent is genuinely undecided. However, in order to make this point clear, we must now define a notion of *measurement*.

When a person responds to a survey they are undergoing a process that is highly similar to a *quantum measurement* [33], and the same can be said of all actions as they were defined above. The decision to act (or not) entails the measurement of a state of an agent, but this very act of measurement may itself affect the decision to act (precisely as it does in the quantum scenario). For example, consider the manner in which the framing of a decision in a positive or negative light can lead to risk averse or risk taking behaviour [58]. Such results suggest that the act of measurement, which necessarily defines the context of the system, can itself influence the outcomes that are obtained. The current geometric formulation can easily incorporate such effects.

Measurement of the state (1) is defined in this approach with respect to a projection operator V, where

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

Thus, the basis vectors $\{|0_p\rangle, |1_p\rangle\}$ *define the current context* p of our agent, which in turn affects their decisions about whether or not to perform an action during the process of measurement. This effect is reflected in the probability that A will act in a given context p, which is given by

$$
\begin{aligned}
P =& \langle A|V_1|A\rangle & (3)\\
=& \langle A|1_p\rangle\langle 1_p|A\rangle & (4)\\
=& \big(a_0^*\langle 0_p|1_p\rangle + a_1^*\langle 1_p|1_p\rangle\big) \times \big(a_0\langle 1_p|0_p\rangle + a_1\langle 1_p|1_p\rangle\big) & (5)\\
=& |a_1|^2 & (6)
\end{aligned}
$$

and similarly, their probability of inaction is given by $|a_0|^2$. Note that this probability arises not due to our lack of knowledge about how the agent intends to act, but from a genuinely undecided agent.

Perhaps the most important feature of this new model arises from a consideration of context itself; it is not just a label. We can immediately develop a far richer notion of context by asking: what would happen if the context changed? QT provides us with a particularly elegant mechanism for dealing with this scenario via a change of basis. Consider figure 1(b), which is an elaboration of figure 1(a), and represents the changing probabilities of action that arise in the case of two different contexts, p and q. With reference to figure 1(b) we can quickly see that while our agent is highly likely to act in context q, this is not the case in context p, where A is much less likely to act (since by examination of the figure we can see that while $|a_0| > |a_1|$ in context p, $|b_1| > |b_0|$ in context q).

We shall now extend this simple QDT model to the description of a society of agents, each making decisions to act (or not) within a social context.

3. A Multi-Agent Model

The simple 2D model introduced above can be naturally extended across a set of multiple agents which we shall call a *society*, all of whom are concurrently considering an issue. For the purposes of initial implementation, as well as representation of the model itself, we shall from now on, assume that all decisions and frames are represented within the same 2D subspace. However, it is important to recognise that the general framework introduced via this simplified 2D model can be generalised to a much higher dimensional scenario, and that such a future extension will be necessary for a complete model of low elaboration attitude attitude change.

In Figure 2 we have drawn a collection of agents, $\{|A\rangle, |B\rangle, |C\rangle \dots\}$, where each individual X is described with a cognitive state $|X\rangle$ which is expected to change in time. Each agent is attempting to make a decision within some set of social contexts, and these are represented by the bases (or axes) that surround each agent. Two potential social contexts (blue $\{|0_b\rangle, |1_b\rangle\}$ and red $\{|0_r\rangle, |1_r\rangle\}$) have been drawn in this figure, and a simple geometric application of the Pythagorean theorem shows the manner in which the probability of each agent's decision to act can substantially change with reference to a different social context.

Figure 2 has depicted agents $|A\rangle$ and $|B\rangle$ with very similar states, while agent $|C\rangle$ appears to have a difference in opinion. As the probabilities of action for these three different agents are extracted by taking a projection of their state onto an

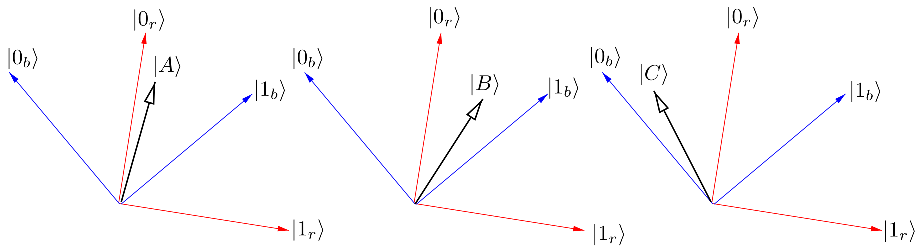

Fig. 2. A set of agents $\{|A\rangle, |B\rangle, |C\rangle \ldots\}$ all making a choice to act (or not) within a set of social contexts. Each has a different cognitive state (e.g. set of attitudes), which can be measured with respect to one of two different social groups or framings of the problem. These global contexts are represented by the red and the blue axes (represented by the basis states $\{|0_r\rangle, |1_r\rangle\}$ and $\{|0_b\rangle, |1_b\rangle\}$ respectively.

axis labelled with $|1\rangle$, this simple example shows a marked difference between the probable actions of the three agents in the illustrated social contexts. Thus, while $|C\rangle$ is unlikely to act in either context, $|A\rangle$ is slightly more likely to act than not in the blue frame, and $|B\rangle$ shows a strong propensity to act in the blue frame.

Note that the same red and the blue social contexts are drawn for each agent in the system. We designate such frames as *global*, as they are shared by all agents in the society, however, it is important to recognise that as agents can only make a decision with respect to one context at a time, they can only consider an issue using one global frame at a time. This leads us to introduce the notion of an *ideology*; in the model that we develop each global frame is taken to represent a specific ideology within a society. Thus, global frames are meant to represent the varying ideologies that parts of a society will use to understand an issue, e.g. liberalism vs conservatism; communist vs socialist vs capitalist vs anarchist; pro-euthanasia vs anti etc.

We have yet to discuss the manner in which an ideology comes to exist, or how it might change in time. It seems natural to anticipate that an ideology could potentially be understood differently by every agent who adheres to it. Thus, there are many different forms of republicanism and liberalism, and these different understandings of the global ideology must be considered in a model of social decision making. This leads us to introduce the notion of a *local* framing of an issue, which represents each agent's personal understanding of that issue. The local frames of the individuals in a society might be similar to a global understanding, or they might differ substantially, depending upon the agent and how they think about the world. Local frames might arise from a wide range of both external and internal factors, such as the socioeconomic status of an agent, their educational background, race etc. and so are likely to be highly complex, and multidependent variables, however, as a first approximation, we shall model them as another basis in the two dimensional vector space already introduced for the states and global frame.

This is obviously an unsatisfactory simplification from the perspective of social psychology, there is no guarantee that these three constructs will arise in the same vector space. Even the assumption of a two dimensional space is unrealistic. However, this assumption had to be taken in order to develop a first order simple model. Future implementations of this model will investigate a more complex understanding of the complex interplay between states and frames.

Despite these potential conceptual complexities, this simplifying step of defining frames as bases in a two dimensional vector space allows for a straightforward designation of global frames as resulting from an aggregation function[d] applied to the local frames of every agent who somehow identifies with that ideology.

Agents can make decisions to act within either their local or the global context. This is taken to represent the manner in which, while we frequently make internal or private decisions (as represented by the local frame), we must sometimes cast our choices within a societal domain (as represented by the global frame) when for example, we must vote in a general election.

We claim that this framework provides an opportunity to model low elaboration processes of attitude change nontrivially, due to its explicit recognition of the context in which an agent makes a decision. The geometric approach allows for the probability of an agent acting to vary over the full range $(0, 1)$ in response to the range of angles that can be taken by the cognitive state of the agent within the Hilbert space that represents the issue currently under consideration. Thus, within this model, an agent A's decision to act or not depends not just upon their internally held state. Rather, it depends on two interdependent factors:

(1) The current cognitive state of the agent, $|A\rangle$.
(2) The social context p of the agent A, as represented by a global or local frame.

Since the social context of an agent has to fundamentally arise from the attitudes of every other agent in the system, we can quickly see that these two factors will recursively interact through time, and that both the cognitive states and the different framings of the issue will evolve in time. We shall now start to more fully formalise the intuitions of this model.

3.1. *Uncertainty and Cognitive Dissonance*

We start with a consideration of the *uncertainty* that an agent experiences about how they are likely to act within a given context. An agent whose cognitive state lies close to the axes representing their current frame will be more certain about their likely future actions than one whose cognitive state lies between those axes (i.e. has the cognitive state forms a 45° angle between choosing to act and choosing not to act in the frame p). This leads us to introduce a measure of this uncertainty and *binary entropy* provides a suitable formalisation. Defined as the entropy of a

[d] In what follows we shall use clustering, however, we anticipate that there are many potential aggregation functions, and that different ones will prove necessary for different issues.

Bernoulli trial (e.g. a two-outcome random variable such as a coin toss), with a probability of success given by P, it is specified as:

$$H_b(P) \equiv -P \log_2 P - (1-P)\log_2(1-P), \tag{7}$$

which is the function depicted in Figure 3 that takes its minimum values at $P = 0$ and $P = 1$, and its maximum at $P = 1/2$.

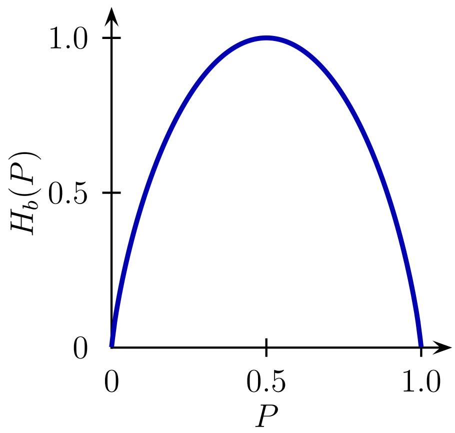

Fig. 3. Binary entropy function $H_b(P)$.

Referring to Figure 1(a), we can rewrite the binary entropy (7) for our agent within the context p using a set of geometric variables

$$H_b(P) = -|a_1|^2 . \log_2(|a_1|^2) - |a_0|^2 . \log_2(|a_0|^2) \tag{8}$$

$$= -\cos^2\theta . \log_2(\cos^2\theta) - \sin^2\theta . \log_2(\sin^2\theta) \tag{9}$$

where θ is the angle between the $|1_p\rangle$ basis state and the state of the agent $|A\rangle$. Rewriting (7) in this manner makes obvious the way in which the entropy of the agent will change if either (a) the agent undergoes a change in state, or (b) finds themselves in a changed context. How might these two situations arise?

Psychologically, an agent who has made a decision is likely to feel a certain amount of *cognitive dissonance* [20] as their internal cognitive state will not be aligned with their decision (unless their cognitive state was already aligned with the relevant frame from which they are currently considering an issue). This means that they will feel a certain amount of psychological discomfort, which will drive them to alter their view of the world to fit with their decision. They can do this in the current model by adjusting either their cognitive state, or their local framing of the issue, to more accurately reflect their decision, and both moves will result in a decrease in the binary entropy associated with their current state. Note that as global framings result from an aggregation function performed over the local frames, it will not be possible for the agent to directly adjust the global frame that they currently subscribe to.

The literature on cognitive dissonance suggests that some people are more comfortable with such nonalignment than others. For example, some agents will feel far less comfortable with uncertainty than others, and so be more affected by dissonance [53, 54]. Thus the *personalities* of an agent will play a key role in how this adjustment occurs. We shall now formalise this intuition with the addition of two key personality parameters in the model.

3.2. *The Personality of an Agent*

Our model considers the personality of an agent as an essential factor in its dynamics. Thus, agents are driven by an attempt to navigate two different drives for cognitive consistency:

(1) A desire for internal cognitive consistency. This results in a drive to align their cognitive state with the local frame within which they are currently considering an issue. Such a drive results in a process where agents act to reduce the binary entropy associated with their current cognitive state, either with a shift in attitude (i.e. by adjusting their cognitive state), or with a shift in their local understanding of the issue (i.e. by shifting their local frame).
(2) A desire to 'fit in' with the society and its current norms. This desire is expressed by a pull of agent's local frames towards the current global frame (or ideology) to which they belong, which serves to reframe the agent's understanding of the issue.

These two drives may prove to compete with one another in the mind of the agent, and indeed, they might have a different pull for agents of different personality types (e.g. a 'conformist' agent vs a highly 'individualistic' one). Defining Θ as the angle between the agent's current state $|A\rangle$ and the decision to act in the global frame to which they currently belong (defined here as the closest global axis to their current state), and taking θ to perform a similar function in their local frame, we define each agent's individual entropy, with reference to both frames, as:

$$H(|A\rangle, \theta, \Theta) = w_i(A)H_b(P(\theta)) + w_s(A)H_b(P(\Theta)) \tag{10}$$

where the weights $w_i(A)$ and $w_s(A)$ refer to agent A's need for internal consistency and social conformity respectively. These weights can be set to range over a population of agents, indicating a rough parameterisation of a society's personality make-up. We note here that, with an increase in the dimensionality of the space that models the attitudes of the population, it becomes possible to anticipate the manner in which a number of the standard personality scales and theories [21, 1, 24, 55] might start to be incorporated into this model.

We note that neither of the two terms in (10) are guaranteed to minimise; each agent will have to balance their desire to fit in with their preference for internal cognitive consistency. We further anticipate that how an agent chooses to achieve this balance will depend upon their personality type (as defined by the same personality weights). Thus, while a process of minimisation could be expected to occur as the agents adjust their attitudes to bring them more into line with their personalities, it is unlikely that an actual minimal value will be obtained. However, it becomes interesting to ask how (10) will decrease as the agents adjust their understanding of an issue in time.

4. Agents and Contexts Evolving In Time

Over time, agents will choose to act within either their own local framing of an issue, or within the society's (as is represented by whichever global frame they identify with most strongly from the society's set of ideologies). Thus, for example, an agent might locally express their opinion about climate change to a group of friends (a local framing), or they might be expected to vote in an election where the core item of debate has become whether or not a carbon tax should be implemented (which would entail a global frame). The proportion of global to local decisions is likely to depend upon the social scenario that is under consideration, however, once a frame has been chosen, an agent will then make a decision to act (or not) in that frame, which will result in a compulsion to update their cognitive state (and possibly their frame) as a result of their personal discomfort with cognitive dissonance. This section will describe the details of this time evolution process in more detail.

4.1. *A Choice of Frame*

Agents act in either their local frame, or in one of the global frames belonging to the society. We expect that the choice of frame will depend upon the psychological make-up of the agent, which leaves a number of plausible options within the current model. We shall list some obvious contenders here, but caution the reader that many more are possible even within the current simple model. Indeed, one sensible parameterisation of the model would include a collection of agents acting according to a mixture of any one of the following strategies.

4.1.1. *Local Frames Dominate*

Given that local frames represent an agent's specific understanding of a particular issue that has arisen in a society, it is possible that decisions will be dominated by these local frames. This could be justified by considering those scenarios where the individuals in a society are more frequently called upon to express their opinions or to commit to actions on a largely local basis (such as deciding to compost their foodscraps at home, or not to immunise their child etc.). If this dynamics is considered most relevant to the model at hand, then it would be reasonable to weight the choice of frame process towards decisions within the local frame.

Indeed, it is possible to imagine a model where agents *only* make their decisions within their individual local framings of a problem *except for when* all agents in a society make a decision within the current global frames. This singular global decision could represent a general election, but we expect that agents are likely to be called upon to make a global decision more frequently than would be anticipated by such a instantiation of the model.

4.1.2. *Conformity Probabilistically Decides the Frame*

Alternatively, a person who has a strong desire to 'fit in' with their society could be biased towards making decisions in the global frame and thus care more about conformity, perhaps at the expense of their own internal consistency (note that both conformity and consistency parameters can be high). Situations where conformity dominates could be used to model results such as the by now famous Asch conformity experiments [4], where subjects gave obviously incorrect responses in a trial in order to fit in with the responses given by a confederate group who had been instructed to lie.

Such a scenario can be modelled by taking the conformity parameter $w_s(A)$, and using it to weight a frame choice function $f(A) = w_s(A)/w_s(\max)$, where $w_s(\max)$ is the maximally allowed conformity value (which is not necessarily obtained in any given run). As $0 < f(A) < 1$, this can quite easily work to probabilistically assign the decision of agent A to either the relevant global frame (with probability $f(A)$) or their local frame (with probability $1 - f(A)$). Thus, agents who have a high conformity will be more likely to make decisions in the global frame that pertains to a given social issue.

4.1.3. *A Full Personality Model*

Both of the above two choice models are extreme, and somewhat naive. We anticipate that a more sophisticated set of choice of frame algorithms could eventually be required. We will not explore this full model in any further detail here, as the resulting added complexity does not seem to be necessary at this stage in our development of the model. The interested reader is encouraged to consult http://www.per.marine.csiro.au/staff/Fabio.Boschetti/quantumPeople.html for a listing of the frame choice procedures that have been implemented so far. As we will show in section 5.3, a number of very interesting dynamics can be obtained even with a very simple instantiation of the model.

4.2. *An Evolving Agent State*

Once an agent has chosen (not necessarily consciously) the frame in which they will make their decision, the probability of them acting will be decided according to the standard QDT described in section 2. Thus, the probability of action can be extracted using equation (3). However, once a decision to act or not has been made, the agent will find that their internal cognitive state will not correspond with the decision that they have made (as is represented by the relevant frame). Cognitive dissonance will therefore force agents to update their states and/or their local framings of the issue that they are considering, and this will in turn affect the ideologies of the society, so changing the social context of all agents.

At present, we update agent states and local frames slightly differently according to the frame in which the decision was initially made.

4.2.1. *Local Decisions*

If the decision was in the local frame, then only the cognitive state of the agent is updated (within the local frame). Thus, an agent who has chosen to act within a certain framing of a problem will shift their state towards the decision ('yes' or 'no') that they made in that context. The size of this shift is defined as dependent upon two factors: (1) the personality profile of the agent (given in this case as w_i, as it represents the desire of an agent to align their cognitive state with their local frame); (2) the angle θ. Writing θ_0 for the angle between the agent's state and the $|0_p\rangle$ axis, and θ_1 for the angle between their state and the $|1_p\rangle$ axis, the new angle between the agent's state and the frame will become:

$$\text{if } A \text{ decides} \begin{cases} \text{to act: } \theta_1(|A\rangle_{t+1}, w(A)) = \theta_1(|A_t\rangle) \times w(A) \\ \text{not to act: } \theta_0(|A\rangle_{t+1}, w(A)) = \theta_0(|A_t\rangle) \times w(A) \end{cases} \tag{11}$$

where $w(A)$ depends upon the comfort of A with holding an attitude that is dissonant from their decision. Thus, for this update process $w(A) = w_i(A)$. Agents who decide to act will thus experience a rotation of their cognitive state by a certain distance dependent upon their personality towards the $|1_p\rangle$ axis (recall that θ is the distance between the $|1_p\rangle$ axis and the current state of the agent $|A\rangle$), and agents who decide not to act will experience a rotation of their cognitive state in the opposite direction.

4.2.2. *Global Decisions*

If the decision was made in the global frame, then both the cognitive state of the agent and their local frame are updated (with reference to their global frame). Thus, in addition to the update of the cognitive state that is represented in equation (11), the local frame of the agent will shift towards the global axis that represents the decision made by the agent. The amount by which the local frame shifts is given by an equivalent version of equation (11), thus the new angle between the local frame and the global frame is given by (11), but with $w(A) = w_s(A)$.

4.2.3. *Differential Responses*

This model exhibits naturally differential responses according to the personality of the agents. Thus, those individuals who are comfortable with dissonance will likely be able to maintain attitudes that do not conform to their actions, and will be more likely to respond to the same question differently if it was asked twice in a row. In contrast, those agents who prefer a consistent cognitive state will experience significant swings in attitude as a result of actions that they choose to take. This means that the consistency variable for each agent ($w_i(A)$) will play an important role in the dynamics of this model, as over time, agents will change their state in an attempt to gain the cognitive consistency that their personality mandates. Furthermore, this suggests that if the personality variables are fixed at the same value for all agents then the dynamics of the model will be much simpler, and we shall indeed show that this is the case in section 5.3.

4.3. *An Entropy-Minimising Global Social Context*

Global frames are defined in this model through a simple clustering approach (see section 5.1 for the details of this process). Thus, for the issue under consideration, the local frames of the agents in the system are analysed and categorised into a set of ideologies which are represented by a set of global frames. This means that as the local frames adjust throughout a run, the global frames can also move, and so the agents will find themselves in a fluid situation where the ideologies of the society adjust and shift. Conceptually, this means that low elaboration ideologies are neither pre-defined nor defined by the work of an intellectual leader, rather, they arise from a representation of what is common among the views of all agents who share similar attitudes. In tracking the emergence of these ideologies, it is possible to ask questions about the entropy of the society as a whole, and how it might evolve in time.

We define this entropy by considering the summed binary entropy of all N agents:

$$H_{b_{Tot}} = \sum_{i=1}^{N} H(|i\rangle, \theta_i, \Theta_i) \tag{12}$$

$$= \sum_{i=1}^{N} [w_i(i) H_b(P(|i\rangle, \theta)) + w_s(i) H_b(P(|i\rangle, \Theta))] \tag{13}$$

which we propose should spontaneously minimise over time according to the social make up of the system. Thus, if a society is composed of a large number of individuals who have a conformist make-up then it makes sense to expect that this entropy function would decrease over time as the agents in that society (1) seek to align with the opinions of one other, and (2) to align their cognitive states with their local frames, or understandings of an issue.

Intuitively, an absolute minimum of (12) is possible for any situation where all agents are polarised onto the relevant basis states. That is, if all agents are 'decided' (either to act or not to act) then the total binary entropy of the system will be a minimum. However, this completely aligned global understanding of a problem seems quite implausible from a psychological perspective. Frequently, it is possible for the members of a society to understand a problem very differently, or to frame it in a number of different ways [14, 30], and this will contribute to the social context of all agents. Indeed, as was discussed above in section 3, it is frequently the case that a society understands an issue from a small number of broadly definable perspectives, and minimisation is likely to prove very difficult to achieve in such scenarios. However, over time, we expect the agents to self-organise towards a situation where they are highly aligned within groups. This ongoing process will be measured by the total entropy of the system (12), which can be expected to decrease when alignment increases.

The next section will describe a simple implementation of this model and its resulting dynamics over time. In particular, we shall use (12) to investigate the behaviour of the system, as its agents attempt to navigate the frequently conflicting demands of social cohesion and internal consistency.

5. Implementation and Preliminary Results

A proof of concept model has been implemented in MATLAB, which allows for an investigation of the timewise behaviour of this new agent based modelling paradigm. In section 5.1 we discuss implementation specific choices that were made, before moving onto a specification of the algorithm as it was implemented in section 5.2, and then a discussion in section 5.3–6 of some early results that have been obtained from this very simple model. We direct the interested reader towards the actual MATLAB script, made available at http://www.per.marine.csiro.au/staff/Fabio.Boschetti/quantumPeople.html for more details than room permits here. Similarly, we encourage the reader to use this script to gain a more detailed feel for the time dynamics of this model.

5.1. *Implementation*

Implementing the model described in sections 3–4 requires a number of specific choices that are not necessarily fundamental to its dynamics, and could be changed in the future for specific scenarios, or more generally. In this section we discuss some of these choices, and point to their potential future extensions.

Firstly, the problem of finding global frames is non-trivial. We have chosen clustering for the purposes of this paper, but this choice will no doubt depend upon the specifics of the system to be modelled. Indeed, recent work on the theory of judgement aggregation [41] suggests that this process can be very complex, and highly dependent upon the social scenario under consideration. As a first step towards investigating this highly complex issue, we have implemented a k-means algorithm, modified for the specific needs of the geometrical representation used in this work.

The dimensionality of the system that has been implemented is of the lowest possible form. Thus, the cognitive states of the agents, and the frames in which they are making their decisions, are both represented on a single 2D plane. There is no reason beyond simplicity for this choice. Indeed, there is every reason to expect that attitudes should be represented on a much higher dimensional space, and this will be investigated in the future. However, this choice does allow for a straightforward visualisation of the system dynamics in the current forum.

The 2D nature of the model presented means that it is obviously symmetric. Thus, agents who are at precisely 180° to one another will exhibit the same probabilities of action in the one global frame. For this reason, the current implementation is restricted to 180°. This has the advantage of simplifying the interpretation of the visualisation used. However, this choice is not the most general, and future more sophisticated implementations will most likely require a return to 360° in higher dimensional (and probably complex number based) Hilbert spaces. In particular, with an extension to a complex Hilbert space, two states at 180° are not the same, which can result in *interference effects* [33] that are peculiar to quantum theory. These have been used by a number of QDT approaches to model a number of violations of so-called rational decision theory [16, 47, 60]. The possibility of this extended implementation will be reserved for future work, as a number of interesting effects can be obtained even for the simple real-valued model that we have so far implemented.

Also, we note that the choice of personality distribution for a society of agents has a profound influence upon the dynamics of that society. When a random distribution was chosen, the resultant behaviour was essentially unpredictable, although broad patterns emerge in the dynamics. The current implementation allows for changing the distribution of personalities in the society, and hence the resulting dynamics. This point will become more clear in the sections that follow. At present, agents are initialised with randomly assigned states (ranging over the full 180° of possibilities) and frames. Personality variables range from [0–1], and can be either assigned randomly or initialised with a fixed distribution of coherence and consistency as is desired.

Finally, as was discussed in section 4.1, there are a number of possibilities for designating in which frame an agent is likely to act. We have implemented two schemes at present:

(1) A 'Weighted' model where the probability of an agent acting in the global frame is proportional to their conformity, $w_s(A)$, and the probability of them acting in their local frame is proportional to their consistency, $w_i(A)$.
(2) An 'Ideology' model where the probability of an agent acting in the global frame is equal to their conformity, $w_s(A)$, and the probability of them acting in their local frame is $1 - w_s(A)$.

However, we remind the reader that it is easy to conceive of many other possibilities. Future implementations will seek to extend the decision model as other socially relevant scenarios are identified in the psychological literature.

5.2. *The Algorithm*

Figure 4 shows the basic pseudocode of the model. For anyone interested in the full implementation details, the code itself can be downloaded as a MATLAB script at http://www.per.marine.csiro.au/staff/Fabio.Boschetti/quantumPeople.html .

```
Number of global frames = G
Number of agents = N
For i=1..N
  Assign coherence & consistency variables
    If RandomPersonality = 0 then conformity = 0.5 and consistency = 0.5
    If RandomPersonality = 1 then consistency & conformity range over [0-1]
  Assign cognitive states & local frames randomly (angle ranges [0-180] degrees)
For each timestep
  Find the position of the global frames (use k-means)
  For each agent
    Calculate which global frame the agent belongs to
      (This is global frame the smallest angle away from cognitive state)
    Probabilistically choose to act or not within appropriate frame
    If acting in local frame then update cognitive state
    If acting in global frame then update cognitive state and local frame
    Calculate entropy of the agent
  Calculate total entropy of system
```

Fig. 4. Basic pseudocode for the algorithm that was implemented (see sections 3–5.1 for details).

5.3. *Results*

Throughout this section, we shall utilise the following convention in all figures shown: cognitive states are represented using black lines, global frames are represented by the large dots above the cognitive states, and local $|1\rangle$ frames are the small black spots that range above these (note that the cognitive states and local frames are not necessarily in line with one another for any given agent).

5.3.1. *Behaviour with Random Personality Variables*

Fig. 5 depicts a typical run of the implemented model with 100 agents and 100 time steps, for a random assignment of the personality variables. The random configuration of agent personality variables means that each agent A has consistency ranging $0 \leq w_i(A) \leq 1$ and conformity similarly ranging $0 \leq w_s(A) \leq 1$. Nine time points

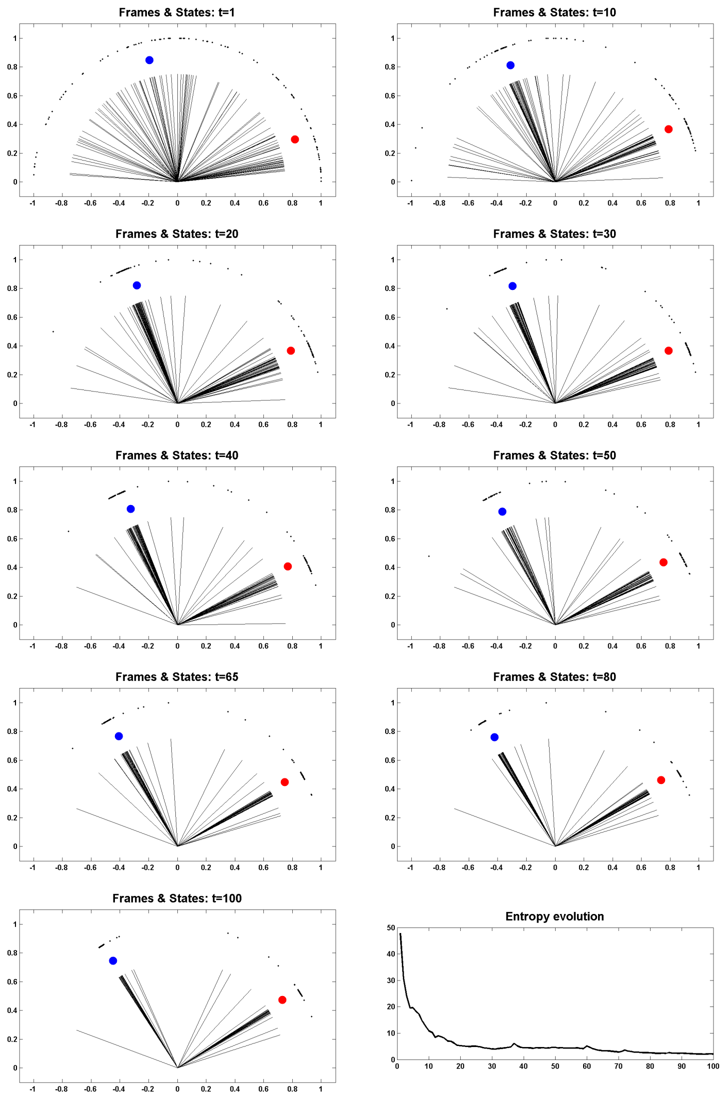

Fig. 5. A typical time evolution pattern for a system of 100 agents over 100 iterations, and its associated total entropy signature (obtained with reference to (12)). Agent's cognitive states are represented using black lines, global frames by the large dots above the cognitive states, and local $|1\rangle$ frames are the small black spots. The total entropy of the system decreases as agents evolve towards ideologies.

are depicted, which show the manner in which, while initially all agents are randomly distributed throughout all available attitudes, they quickly self-organise to a scenario where many agents are clustered around the two markedly stable global frame. While the basic structure persists for the remainder of the run, we note that the situation remains fluid, with local frames and agents still capable of exhibiting substantial shifts in position over the remainder of the simulation. The evolution of the total entropy (10) is also depicted, and shows the expected tendency to decrease throughout the run.

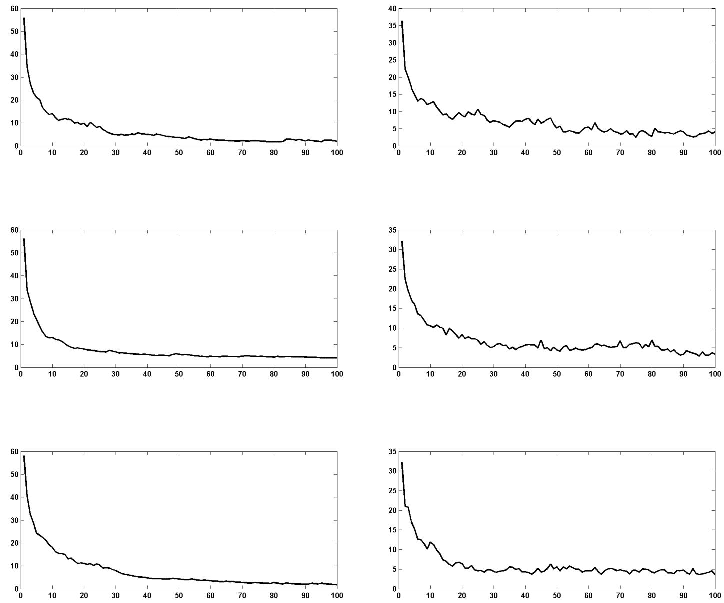

(a) (b)

Fig. 6. Standard entropy evolution signatures for (a) 1, and (b) 3 global frames, as specified by the number of clusters, k, designated in the clustering algorithm.

This pattern of general entropy decrease appears to be generally ubiquitous throughout the model. Figure 6 shows a set of typical entropy signatures for the case of randomly assigned personality variables, in Figure 6(a) for 1, and then in Figure 6(b) for 3 global frames. In both cases we see the entropy generally tending to minimise as agents align with the various frames. However, their different personality profiles keep the agents from reaching a truly stable situation, as there are many competing requirements across the society. This scenario is very different when the personality variables are fixed in a more stable configuration.

5.3.2. *Behaviour when Personality is Fixed*

When runs are implemented with non-random assignment of the personality variables (w_i and w_s) then the behaviour of this model becomes far more deterministic. In Figure 7 we see a run of the model which started with the same random seed

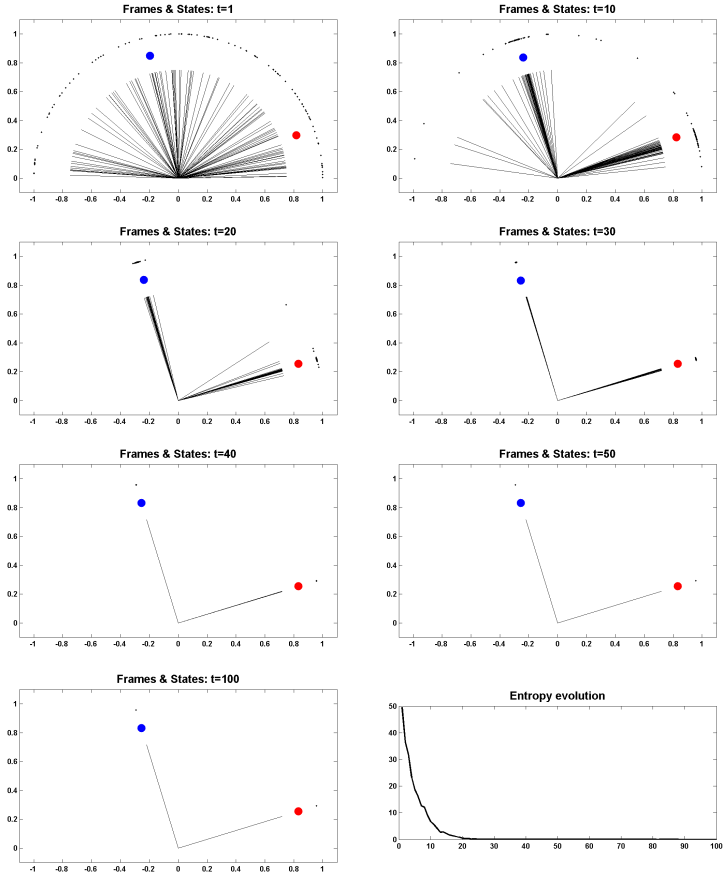

Fig. 7. When $w_i = w_s = 0.5$ for the same initial conditions as was used in Figure 5, the behaviour of the system quickly becomes much more stable than for the randomly assigned personality variables scenario, with the entropy quickly falling to zero.

as that depicted in Figure 5, but for a population which had its personality variables all set equally to $w_i = w_s = 0.5$. Here, we see that the system quickly becomes highly uniform, with all agents eventually settling into one of the two global frames, which themselves settle down to a distance of 90° apart. The associated entropy signature quickly becomes zero.

This feature of the model has been more fully explored in a recent paper [36], which examines more fully the effects that personality variables can have in this model.

5.4. *The Impact of a New Issue*

Figure 8 takes the same scenario as was shown in Figure 5 up until step 50 when the model has settled into a strongly polarised setting, characterised by two, almost opposite views. This could represent, as an example, the agents' attitude towards

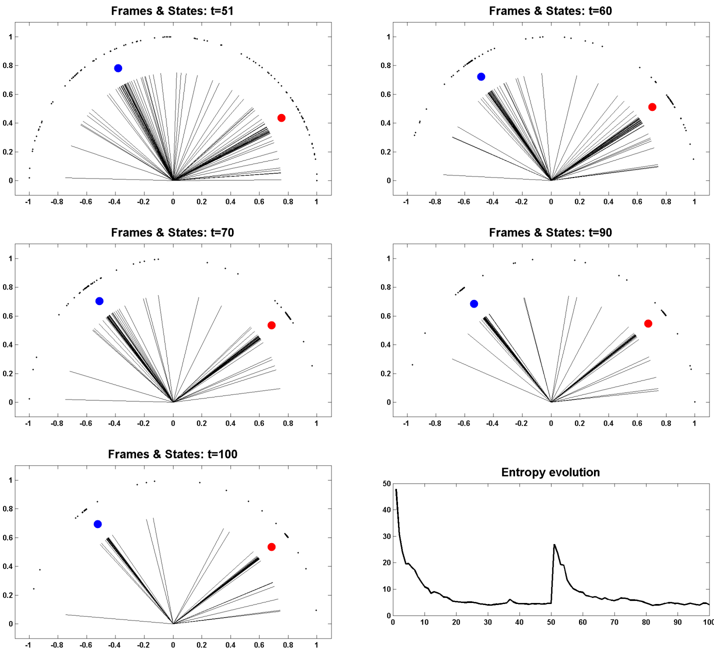

Fig. 8. The same initial dynamics as were represented in Figure 5 up until step 51, when we model the appearance of a new issue by resetting all of the local frames. Note that the extant global frames quickly pull agents towards the previous issue, and the system settles down into stability around this original setting. The entropy signature of this run shows a marked increase at step 51, followed by a stabilisation.

illegal immigration. At iteration 51, we re-initialised the local frames of all agents, a move that is intended to model the sudden appearance of a new issue over which agents need to take a decision. Such a scenario might correspond to the consideration of a newly proposed mitigation initiative aimed to combat climate change. Because the issue is novel, and thus initially unrelated to previous issue considered by the society, we model the agents' understanding of it as random. However, as we see in Figure 8, the new local frames have little impact on the global dynamics. The existence on two polarised global frame soon 'attracts' the discourse about the new issue (climate change) and casts it within the pre-existing polarisation. Thus, rather than the new issue redefining the polarisation, it is the current polarisation which redefines the new issue by forcing the local frames to align (stochastically) with the established polarisation. The two issues become tightly coupled, with agents remaining in their original tendencies towards action or inaction. This models fairly well the dynamics of the discourse on issues including climate change in several countries [32, 34, 39, 40].

6. Ideologies as Emergent Features with ‘Causal’ Properties

It is interesting to analyse the role of the global frames, and their relation to the local frames, within the framework of complex system science. The global frames can be identified as ‘emergent structures’ according to several definitions. Besides their trivial identification in terms of pattern formation, the global frames can be understood as an example of intrinsic emergence [22] in the sense that they provide information processing capabilities to the agents in the system. Notice that in our model there is no communication among agents; each agent understands what other agents do and what the shared understanding of the problem is only via the global frames, which become the avenue for internal information processing. Without the global frames, no internal organisation among the agents in the system would be possible.

The global frames can also be identified as emergent structures according to the Efficiency of Prediction view of emergence [52]. According to this view, the global frames identify the level of analysis at which it is most efficient to describe the system; if we want to understand a society’s perception of a problem, then we could either analyse all individual local frames, at a considerable processing cost, or we could reach a similar understanding by just considering the global frames (i.e. the shared ideologies of all agents in the system). This understanding would quite possibly miss many subtleties, but it would also come at a considerably reduced computation cost.

As a more stringent requirement, emergent structures may be required to display some sort of independent causal power, that is a causal power which does not reduce to the underlying components that serve to form that structure [11]. In our model, this causal power is represented by the influence of the global frames upon each agent’s state and local frame; there is an obvious feedback loop between the local frames which, via the clustering algorithm, determine the global frames, and the global frames which, once identified, influence the agents’ states and their local frames by affecting their position (a similar feedback loop occurs between local frames and agents states).

It is this causal power which provides an avenue for intervention in the system via the global frames. Figure 9(a) shows the same system as was depicted in Figure 5 after 50 iterations, but with a new intervention performed at step 51, when we perturb one of the frames by imposing a rotation. This may be seen as an external intervention to redefine what an ideology represents; it could be a party redefining its values or making a new statement through the provision of new information that redefines a currently held position on an issue. Later timesteps show the system

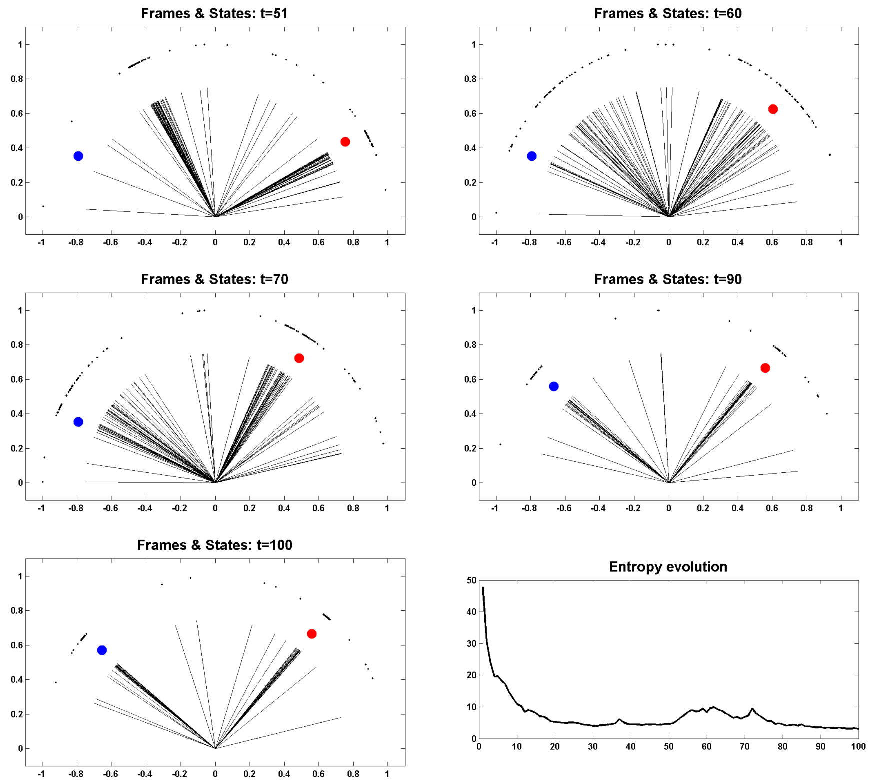

Fig. 9. Example of causal properties in the global frames for the scenario initially depicted in Figure 5. After 50 iterations the system reaches a stable state with 2 global frames and all agents roughly aligned with their local frames as well as to one of the global frames. At iteration 51 an external perturbation rotates the blue frame 40 degrees anticlockwise. We then see the system re-organising. After 100 iterations the system reaches a new, different stable state. The entropy panel shows the time evolution of the total entropy measure, displaying the effect of the intervention at iteration 51.

responding to the perturbation and trying to reorganise itself before it converges to a final stable state. Finally, the associated total entropy signature is also shown, and highlights the intervention point with a surge in global entropy which then settles down as the new stable state is realised.

By slightly paraphrasing the definition of self-organisation given in [48], Figure 9 provides an example of the self-organised nature of the system as well as an avenue for guiding the process. Reinterpreting equations (29) and (30) from that paper, we can see that a considerable re-organisation in the system (a re-organisation which involves considerable information processing by all agents in the system) has been obtained by a single external action. This action involves a much reduced information processing effort than would have been required were the system not organised, in which case the same action would have had to be carried out on each agent subscribing to the perturbed ideology. In other words, rather than having to convince a large population of agents to shift their local framing of an issue, all that would be needed in such a scenario would be a change in the ideology that they individually subscribe to. Naturally, this provides an avenue for guiding the system towards desirable final configurations, a scenario that could be represented by the externally driven shift of frame discussed here. Such a shift of frame could be taken to represent a global change in a society's understanding of an issue. For example: a political party might decide to adjust its policy settings with a new understanding of what constitutes a refugee; a highly publicised act of terrorism might take place leading to a new public perception of safety and an associated need for heightened state security; or perhaps a local and quite specific case of corruption could be widely publicised for the purposes of changing a political narrative. Thus, the current tendency towards political 'spin' and close management of public relations and perceptions suggests that this point has already been implicitly recognised by politicians and the media [57]. However, it must be recognised that many such attempts to guide the attitudes and opinions of the public often end in failure.

7. Testing and Validation

This model is very much in its early phases, and while the theory presented here is general well beyond our simple 2-D computational implementation, the current formulation of this theory can at best only be considered a toy model at present. However, at this point we shall briefly consider the potential viability of this model when it comes to the testing of its predictions and the associated problem of validation. We shall start with a consideration of the personality variables that were added to the model in section 3.2. Do they have any psychological validity?

7.1. *A Mapping to Cultural Theory*

Adding parameters to a model is something that should be considered carefully. Indeed, almost any behaviour can be modelled with the addition of enough parameters to what was fundamentally an inadequate theoretical construct. This point gives us reason to pause. Is there any room for validation or comparison with existing psychological datasets in this model? We have found one potential mapping to a psychologically justified theory, which could allow for a realistic initialisation of the model as it gains in sophistication, and this section will briefly explore this potential avenue of testing and future validation.

Cultural Theory [24] has been used to understand differences in the perception of environmental risk, and the associated management strategies that are likely to be preferred by an individual [55]. It explains why people perceive dangers differently and focus on particular threats at the expense of others, which provides useful insights into divergent community responses to contested risks like climate change [13]. Shared values and beliefs within different cultural groups lead to the attribution of blame to different institutions thought to violate the socially accepted standards of each group. Selective attention towards specific issues within different cultural groups is therefore thought to represent cultural biases, which are termed worldviews. Thus, "Cultural theory is based on the axiom that what matters most to people is their relationships with other people and other people's relationships with them." p5–6 [59], and this allows the theory to define four basic worldviews according to a person's answers to two fundamental questions: Who am I?; and, What shall I do? [59] The identity question is answered differently according to whether individuals either (a) belong to a strong group which makes decisions for all members (e.g. a collective), or (b) only weakly identify with a group and that their choices conform only with their own norms (e.g. individualism). The strength by which an individual identifies with these two possible responses is classified on a continuum, termed the *group* scale. The action question is answered by considering whether the individual is bound by few restrictions (e.g. considers themselves a free spirit able to act how they desire), or is tightly constrained and does not feel free to act according to their desires. This *grid* scale is similarly represented on a continuum.

The overlap between these two scales allows for the definition of four basic cultural types [59] (see figure 10):

Hierarchical worldviews are typified by strong groups with numerous prescriptions or rules (high group, high grid).

Egalitarian worldviews consist of strong groups with few prescriptions (high group, low grid).

Fatalistic worldviews are espoused by people who only weakly identify with groups but strong prescriptions that are considered to be imposed upon them from the outside (low group, high grid).

Individualistic worldviews have few group identifications and few prescriptions (low group, low grid).

These four types display good predictive properties when applied to the perceptions that an individual might have about many different social issues, including politics (e.g. liberal vs conservative), economics (e.g. free markets vs regulated), and environmental management (e.g. exploitative vs protective), as generally people who espouse a particular worldview for one issue usually share common environmental and political attitudes.

We shall not consider Cultural Theory in any more detail here, but the interested reader is encouraged to consult any of the above mentioned references. Instead, we will now point to a potential mapping of the grid and group notions onto the consistency and conformity variables that we defined in section 3.2. We propose that an individual who identifies strongly with their group can immediately be understood as having a high conformity value, while those with a low conformity value will lie on the low group end of the spectrum. While this mapping of the conformity parameter to the group scale is likely to prove uncontroversial, the grid mapping is less obvious. However, it seems likely that an individual who considers

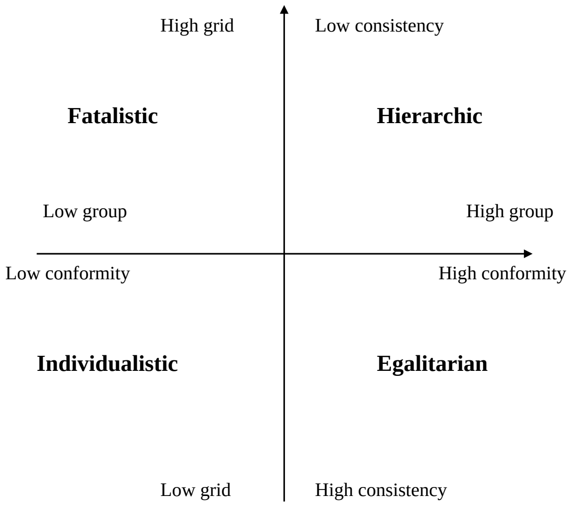

Fig. 10. The grid-group model of worldviews along with a potential mapping to the consistency-conformity variables defined in this model.

it important that they be true to themselves (high consistency) will exhibit low grid characteristics, while low consistency could be seen to correspond to high grid values. This hypothesis recognises the manner in which the desire of Alice to be true to herself will generally overrule her recognition of social rules and norms, and that this will lead to her feeling that her actions should not be constrained by these rules. This hypothesis is less straightforward, and likely to prove incorrect in a fully psychological valid model, however, it provides an intriguing first approximation that allows us to correlate the model with data that is currently being obtained by large scale surveys of attitude change and social dynamics [13, 38, 37], and so start to test its validity.

While it must be admitted that the mapping proposed here is quite simple, it does open up a number of modelling possibilities, and suggests that this approach can indeed both be seeded with psychologically plausible parameterisations, and then tested. While we anticipate that the initial mapping proposed here will most likely require extension and refinement, especially beyond its current 2D form, the straightforward manner in which the new parameters introduced in this model can be mapped to psychologically plausible constructs is encouraging. We reserve a more detailed treatment of this mapping to future work.

7.2. *The General Validation of Social Models*

When it comes to any type of model the obvious question is whether the results should be considered a 'reliable' representation of real dynamics. This issue is particularly difficult to resolve for models of social behaviour, and it is beyond the scope of this paper to address this at depth. Instead, we summarise here the three main challenges that must be considered in order to assess the reliability of the results generated by a model of social behaviour:

(1) Whether we believe such models have predictive power or can at best just explain observed behaviours [5, 44, 8, 9, 6, 15].
(2) Whether model prediction can be validated even in principle, given that social processes can be affected by expectations and thus by existing knowledge about the predictions [23].
(3) Whether the theory behind the model is itself validated and predictive [26].

On the first point, we believe that explanation does entail the ability to predict, and that models of social behaviour should not be exempted from this responsibility, although we need to be clear about what we define as prediction [12, 56]. The second point is discussed at length in the philosophy of science literature under the term 'theory absorption' [23]. In principle the method we propose could be used to model this very issue: the prediction of the model could be represented as a new context and thus a new global frame which can in turns affect the agents' future behaviour. The third point is discussed in Elsenbroich [26], who claims that we cannot expect a model to provide better prediction/validation results than the theory it implements. Thus, at this point it seems more reasonable to ask whether some of the processes described in the cognitive science literature, which are relevant to decision making in a social context, are captured by our model.

As we discussed in section 7.1, there is a clear potential for the parameters used in this model to be mapped into existing psychological models, and with a sophisticated seeding of a population according to a measured set of personality variables we believe that this model could indeed be expected to capture many essential characteristics of low elaboration attitude change. Indeed, as extended time series of attitude change are collected for a variety of social issues (such as the ongoing CSIRO Annual Survey of Australian Attitudes to Climate Change [38, 37]), we anticipate that a model correctly seeded with the a distribution of personality types similar to the Australian population should be capable of at the very least capturing some of the behaviour exhibited by those datasets. If it was also capable of prediction, then the first two challenges in the list above would have been met by this model. However, in its present simple form, the manner in which the model captures the socially motivated interventions shown in Figures 8 and 9 leaves us with reason to believe that it is very much capable of at least starting to test the underlying theories of social psychology. Similarly, there are a number of other debates in psychology that have generally remained beyond the realms of scientific investigation due to the ill-defined nature of the arguments involved. Indeed, just as Mosler et al. [43] used their computational implementation of the ELM to explore the characteristics of a key social model and thus make it more specific, we anticipate that our model will provide a similar resource for those wishing to truly test the predictions of social psychology. For example, the *person-situation* debate [27] was a controversy in personality psychology that arose over the question of whether the person *or* the situation is more influential in determining the behaviour of a person. This debate has now essentially been resolved, with most researchers understanding that both the personality *and* the situation will influence a person's behaviour. This is precisely the scenario that is modelled by our approach, and so we anticipate that our model could be used to make the relevant variables of both personality and situation more explicit. Such a move would allow for a testing of the underlying psychological theory, so answering questions that have hitherto been difficult to rigorously explore due to their generally ill-defined nature.

Perhaps the best hope for agent based models of social behaviour lies in their ability to point towards the necessary and sufficient parameters required in an adequate social theory. Thus, we find ourselves agreeing with Mosler et al. [43]; the validity of social models lies much more in their ability to make specific those concepts from psychology that are currently not well modelled, and so likely to be poorly understood. Indeed, the need for us to make a number of specific choices throughout our implementation (as was discussed in section 5.1) draws attention to this fact.

8. Future Directions

An obvious set of possibilities for extension in this model will all require more realistic social modelling. Firstly, a spatial implementation, where conformity can only be satisfied over a defined semi-local network seems desirable. At the moment conformists strive to reach a scenario where they agree with all agents in their relevant global frame. Future implementations will seek to provide the model with a social network style dependency for this agreement. Thus, in this model, social clades or subgroups could spontaneously emerge with a well defined spatial (or communication based) boundary; who you talk to matters in the social world. Similarly the different implementations of choosing the frame in which a decision is made (section 4.1) is another obvious candidate for future investigation and extension. Perhaps most importantly, the use of a single plane for the current model places a highly unrealistic assumption upon the model, and its extension to a higher dimensional model that incorporates the numerous results from social psychology about underlying attitude and personality variables is a high priority.

Another set of possibilities arise from the quantum inspired nature of this model. Indeed, a highly developed set of quantum inspired models of human decision making are coming to the fore, all of which utilise marked *interference effects* to explain the apparent inconsistencies violations of standard probability theory that are exhibited by humans every day as we make our decisions in a complex world [17, 60, 35]. These models all seek to explain the decisions of a single human agent, and it has been the purpose of this paper to present a model that 'scales them up' into a social context, however, we have refrained from a consideration of the manner in which the different decisions to act that an individual agent makes might interfere with one another and so affect the dynamics of the society. This intriguing possibility also presents a ripe opportunity for future investigation.

Indeed, we feel that with the more realistic enhancements described above, this model might be capable of shedding genuine light upon the complex dynamics of social evolution and re-adjustment. Ultimately, we believe that a more realistic future extension of this model could be used to predict how likely a society is to undergo such phenomena as attitudinal phase transitions, or re-framings of issues and debates. Thus, policy makers could track the evolution of attitudes in a society and identify scenarios where the society was entering phases of instability. They may even be able to predict situations where the likely hijacking of opinions and attitudes by interested parties was becoming a significant possibility. Needless to say, there are many applications of such tools beyond a policy setting.

9. Conclusions

This paper has presented a model of human decision making that is based upon the highly complex notion of a low elaboration attitude. Social psychology has developed a rich set of empirical results that shed light upon this concept, however, many of the more mathematical models of attitude change are naive, relying upon objective states and properties which suggest that attitudes are held regardless of the context in which humans find themselves. This is a problematic assumption, as humans are cognitive misers [29], and our opinions are apt to change in response to many contextual factors; frequently our minds are only 'made up' at the point where we are forced to make a choice. The model presented here attempts to approach the problem of modelling social decision making from this perspective.

We have presented the notion of an agent, who, while they possess a definite cognitive state or attitude $|A\rangle$, may make very different decisions depending upon the context in which they find themselves. This is represented probabilistically, using a geometric approach to uncertainty which captures the notion of a decision in context. However, the since the context of an attitude can be understood both externally (as some sort of aggregate over the opinions of an entire society) and internally (as the individual framings of a social issue that an individual adopts), this led us to define both global and local framings of an issue. We then showed how cognitive dissonance can be understood to drive individuals towards changing both their cognitive state, and their local understanding of an issue after they make a decision to act (or not). Personality was a key variable in defining the dynamics of the individual agents in the system, and this was in turn used to motivate a model of the state update of agents over time, with different members of the population subjected to different update regimes depending upon how uncomfortable they are with internal inconsistency and social non-conformity. A measure of the entropy of a society was introduced, which was shown to reduce in a simple computational implementation of the theoretical model. We argued that the ongoing update of agent cognitive states and local frames will in turn result in the emergence of stable global ideologies in a society, resultant from a process of self-organisation.

More generally, beyond the likely advantages of a genuinely new approach to social modelling, we believe that our model provides an unlikely avenue of communication between natural scientists and engineers on one side and social scientists on the other. Effective communication between these groups is often often lost in this arena, insurmountably challenged by a crucial difference in their divergent approaches to knowledge; while natural scientists and engineers are trained to think that there is 'a truth' which needs to be discovered, social scientists tend to believe truth is a metal construct and thus contextual. In the first case uncertainty arises from not knowing the truth (knowing too little), in the second about choosing which truth to accept (knowing too differently), an understanding that at first glance appears not amenable to rigorous formal analysis, and is therefore often refuted by natural scientists for this reason alone. We believe that the geometric representation employed in this paper allows natural scientists and engineers to model, and thus more easily accept, the views that social scientists hold so dear. Similarly, the adoption of our proposed framework may in some cases provide social scientists with some confidence that important aspects of social theory can be considered within quantitative models, so making them relevant to the real world problems that they are seeking to address.

Overall, we feel that this proposed new class of model offers a promising avenue for future research. It allows for the sophisticated modelling of humans working within the many frames and contexts that affect their decisions and choices. Such models are likely to prove essential for including the actual dynamics of human decision making in the complex and often contradictory world that we inhabit as a society, and so an honest exploration of their possibilities could open up a new frontier of mathematical and computational analysis. This alone makes them worthy of attention.

Acknowledgments

KK was supported by an Australian Research Council Post Doctoral Fellowship (DP1094974) while completing this work. The authors would like to acknowledge many helpful discussions with Peter Bruza and Iain Walker during the creation of this model, as well as the assistance of the referees, whose comments served to markedly improve this paper.

References

[1] Ajzen, I., *Attitudes, Personality and Behaviour*, 2nd edn. (Open University Press, 2005).

[2] Albert, R. and Barabási, A.-L., Statistical mechanics of complex networks, *Reviews of Modern Physics* **74** (2002) 47–97.

[3] Allport, G. W., Attitudes, in *Handbook of social psychology*, ed. Murchison, C. (Clark University Press, Worcester, MA, 1935), pp. 798–884.

[4] Asch, S. E., Studies of independence and conformity: A minority of one against a unanimous majority, *Psychological Monographs* **70** (1956) 1–70.
[5] Ascher, W., Beyond accuracy, *International Journal of Forecasting* **5** (1989) 469–484.
[6] Ascher, W., The ambiguous nature of forecasts in project evaluation: Diagnosing the over-optimism of rate-of-return analysis, *International Journal of Forecasting* **9** (1993) 109–115.
[7] Augoustinos, M., Walker, I., and Donaghue, N., *Social Cognition*, 2nd edn. (Sage, London, 2006).
[8] Beven, K., Towards a coherent philosophy for environmental modelling, *Proceedings of the Royal Society A* **458** (2002) 2465–2484.
[9] Beven, K., A manifesto for the equifinality thesis, *Journal of Hydrology* **320** (2006) 18–36.
[10] Bond, R. and Smith, P. B., Culture and Conformity: A Meta-Analysis of Studies Using Asch's (1952b, 1956) Line Judgment Task, *Psychological Bulletin* **119** (1996) 111–137.
[11] Boschetti, F. and Gray, R., A Turing test for Emergence, in *Advances in Applied Self-organizing Systems*, ed. Prokopenko, M. (Springer-Verlag, London, UK, 2007), pp. 349–364.
[12] Boschetti, F., Grigg, N., and Enting, I., Modelling = conditional prediction, *Ecological Complexity* **8** (2011) 86–91.
[13] Boschetti, F., Richert, C., Walker, I., Price, J., and Dutra, L., Assessing attitudes and cognitive styles of stakeholders in environmental projects involving computer modelling, *Ecological Modelling* (2012).
[14] Brugnach, M., Dewulf, A., Pahl-Wostl, C., and Taillieu, T., Toward a relational concept of uncertainty: about knowing too little, knowing too differently, and accepting not to know, *Ecology and Society* **13** (2008) 30.
[15] Brunner, R., Predictions and policy decisions, *Technological Forecasting and Social Change* **62** (1999) 73–86.
[16] Busemeyer, J. and Bruza, P., *Quantum Models of Cognition and Decision* (Cambridge University Press, 2012).
[17] Busemeyer, J. R., Pothos, E., Franco, R., and Trueblood, J., A Quantum Theoretical Explanation for Probability Judgment Errors, *Psychological Review* **118** (2011) 193–218.
[18] Castellano, C., Fortunato, S., and Loreto, V., Statistical physics of social dynamics, *Rev. Mod. Phys.* **81** (2009) 591–646.
[19] Chaiken, S. M., The heuristic model of persuasion, in *Social influence: The Ontario symposium*, eds. Zanna, P., Olson, J. M., and Herman, C. P., Vol. 5 (Erlbaum, Hillsdale, NJ, 1987), pp. 3–39.
[20] Cooper, J., *Cognitive dissonance: 50 years of a classic theory* (Sage, 2007).
[21] Costa, P. T. and McCrae, R. R., Normal personality assessment in clinical practice: The neo personality inventory., *Psychological Assessment* **4** (1992) 5–13.
[22] Crutchfield, J., The Calculi of Emergence: Computation, Dynamics and Induction, *Physica D* **75** (1994) 11–54.
[23] Dacey, R., Theory absorption and the testability of economic theory, *Zeitschr. f. Nationalökonomie* **36** (1976) 247–267.
[24] Douglas, M., *Risk acceptability according to the social sciences* (Russell Sage Foundation, New York, 1985).
[25] Eagly, A. H., Wood, W., and Chaiken, S., Causal inferences about communicators and their effect on opinion change, *Journal of Personality and Social Psychology* **36** (1978) 424–435.

[26] Elsenbroich, C., Explanation in agent-based modelling: Functions, causality or mechanisms?, *Journal of Artificial Societies and Social Simulation* **15** (2012).

[27] Epstein, S. and O'Brien, E. J., The person-situation debate in historical and current perspective, *Psychological Bulletin* **98** (1985) 513–537.

[28] Fazio, R. and Petty, R. (eds.), *Attitudes: Their Structure, Function, and Consequences* (Psychology Press, 2008).

[29] Fiske, S. and Taylor, S., *Social cognition*, 2nd edn. (McGraw-Hill, New York, 1991).

[30] Gray, B., Strong opposition: frame-based resistance to collaboration, *Journal of Community and Applied Social Psychology* **14** (2004) 166–176.

[31] Greenwald, A. G. and Banaji, M. R., Implicit social cognition: Attitudes, self-esteem, and stereotypes, *Psychological Review* **102** (1995) 4–27.

[32] Heath, Y. and Gifford, R., Free-market ideology and environmental degradation: The case of belief in global climate change, *Environment and Behavior* **38** (2006) 48–71.

[33] Isham, C. J., *Lectures on Quantum Theory* (Imperial College Press, London, 1995).

[34] Kahan, D. M., Braman, D., Slovic, P., Gastil, J., and Cohen, G. L., The second national risk and culture study: Making sense of - and making progress in - the American culture war of fact, Technical report, Yale Law School (2007), public Law Working Paper No. 154; GWU Law School Public Law Research Paper No. 370; Harvard Law School Program on Risk Regulation Research Paper No. 08-26. Available at SSRN: http://ssrn.com/abstract=1017189 or http://dx.doi.org/10.2139/ssrn.1017189.

[35] Khrennikov, A. Y., *Ubiquitous Quantum Structure: From Psychology to Finance* (Springer, 2010).

[36] Kitto, K. and Boschetti, F., The effects of personality in a social context. In Press *35th Annual Cognitive Science Society Conference (CogSci2013)* (2013).

[37] Leviston, Z. and Walker, I., Beliefs and denials about climate change: An australian perspective, *Ecopsychology* .

[38] Leviston, Z. and Walker, I., Second annual survey of australian attitudes to climate change: Interim report, Technical report, The Commonwealth Scientific and Industrial Research Organisation (CSIRO) (2011).

[39] Lewandowsky, S., Popular consensus: Climate change is set to continue, *Psychological Science* **22** (2011) 460–463.

[40] Lewandowsky, W., S. abd Stritzke, Oberauer, K., and Morales, M., Memory for fact, fiction, and misinformation, *Psychological Science* **16** (2005) 190–195.

[41] List, C., The theory of judgment aggregation: An introductory review, *Synthese* **187** (2012) 179–207.

[42] Mosler, H.-J. and Martens, T., Designing environmental campaigns by using agent-based simulations: Strategies for changing environmental attitudes, *Journal of Environmental Management* **88** (2008) 805 – 816.

[43] Mosler, H.-J., Schwarz, K., Ammann, F., and Gutscher, H., Computer Simulation as a Method of Further Developing a Theory: Simulating the Elaboration Likelihood Model, *Personality and Social Psychology Review* **5** (2001) 201–215.

[44] Oreskes, N., Why believe a computer? models, measures, and meaning in the natural world, in *The Earth Around Us: Maintaining a Livable Planet*, ed. Schneiderman, J. (W.H. Freeman and Co, San Francisco, 2000), pp. 70–82.

[45] Petty, R. E. and Cacioppo, J. T., *Communication and Persuasion: Central and Peripheral Routes to Attitude Change* (Springer-Verlag, New York, 1986).

[46] Petty, R. E. and Wegener, D. T., Attitude change: Multiple roles for persuasion variables, in *The handbook of social psychology*, eds. Gilbert, D., Fiske, S., and Lindzey, G. (McGraw-Hill, 1998), pp. 323–390.

[47] Pothos, E. M. and Busemeyer, J. R., A quantum probability explanation for violations

of 'rational' decision theory, *Proceedings of the Royal Society B* **276** (2009) 2171–2178.

[48] Prokopenko, M., Boschetti, F., and Ryan, A. J., An information-theoretic primer on complexity, self-organization, and emergence, *Complexity* **15** (2008) 1–18.

[49] Rosenberg, H. and Feldman, C. S., *No Time To Think: The Menace of Media Speed and the 24-hour News Cycle* (Continuum, London, 2008).

[50] Schneider, F., Gruman, J., and Coutts, L. (eds.), *Applied Social Psychology: Understanding and Addressing Social and Practical Problems*, 2nd edn. (Sage, 2012).

[51] Seiter, R. H. and Gass, J. S., *Persuasion, social influence, and compliance gaining*, 4th edn. (Allyn & Bacon, Boston, 2010).

[52] Shalizi, C. and Crutchfield, J., Computational mechanics: Pattern and prediction, structure and simplicity, *Journal of Statistical Physics* **104** (2001) 819–881.

[53] Sorrentino, R. and Roney, C., *Uncertain Mind: Individual Differences in Facing the Unknown* (Psychology Press, 2000).

[54] Sorrentino, R. M. and Hewitt, E. C., The uncertainty-reducing properties of achievement tasks revisited, *Journal of Personality and Social Psychology* **47** (1984) 884–899.

[55] Steg, L. and Sievers, I., Cultural theory and individual perceptions of environmental risks, *Environment and Behavior* **32** (2000) 250–269.

[56] Symons, J. and Boschetti, F., How computational models predict the behavior of complex systems, *Foundations of Science* (2012).

[57] Tanner, L., *Sideshow: Dumbing Down Democracy* (Scribe Publications, Australia, 2011).

[58] Tversky, A. and Kahneman, D., The framing of decisions and the psychology of choice, *Science* **211** (1981) 453–458.

[59] Wildavsky, A., Choosing preferences by constructing institutions: a cultural theory of preference formation, *The American political science review* **81** (1987) 3–22.

[60] Yukalov, V. I. and Sornette, D., Decision theory with prospect interference and entanglement, *Theory and Decision* (2010) 1–46.