Ideas Worth Spreading: A Free Energy Proposal for Cumulative Cultural Dynamics

Kastel, N.; Hesp, C.

DOI
10.1007/978-3-030-93736-2_55

Publication date
2021

Document Version
Submitted manuscript

Published in
Machine Learning and Principles and Practice of Knowledge Discovery in Databases

Ideas worth Spreading: A Free Energy Proposal for Cumulative Cultural Dynamics

Abstract. While there is a fast growing body of theoretical work on characterizing cumulative culture, quantifiable models underlining its dynamics remain scarce. This paper provides an active-inference formalization and accompanying simulations of cumulative culture in three steps: Firstly, we cast cultural transmission as a bi-directional process of communication that induces a generalized synchrony (operationalized as a particular convergence) between the internal states of interlocutors. Secondly, we cast cumulative culture as the emergence of accumulated modifications to cultural beliefs from the local efforts of agents to converge on a shared narrative.

Keywords: Active Inference, Generalized Synchrony, Communication, Cumulative Culture, Cultural Dynamics.

1 Introduction

While there is a fast growing body of theoretical and empirical literature on the processes of cultural evolution, quantitative models that are able to integrate different approaches and insights from multiple disciplines into unified, quantifiable interpretations of theory and empirical data are in rapidly growing demand (Creanza, Kolodny & Feldman, 2017). This is particularly true for the mechanisms of social transmission, which have been especially reviewed under theoretical models while mathematical models for cultural transmission remain scarce in this field. A prominent stream of quantitative models for cultural development are inspired by epidemiology, and convert models used for predicting the spread of a virus to formalize the spread of an idea (Bettencourt, Cintrón-Arias, Kaiser & Castillo-Chávez, 2006).

While the comparison of an idea to a virus has its benefits from a structural perspective, it implies the controversial notion that an idea is simply copied during its transmission through cultural exchange between individuals. This notion is not only intuitively insufficient for a realistic portrayal of communication dynamics, but also conflicts with established theoretical models of transmission on these same grounds.

Current literature in cultural psychology indicates that rather than being simply duplicated during transmission, cultural beliefs and practices are modified through the active interpretation of each individual (Kashima et al. 2019). Another example for the discrepancy between quantitative epidemiology models for transmission and theory is taken from the psychology of communication. Research in this field suggests that communication is conditioned upon a mutual shared reality (Echterhoff, Higgins & Levine, 2009), or “common ground” (Clark & Brennan, 1991) between interlocutors. According to these theories, not only does cultural information change during communication, but (contradictory to the one-sided transmission of cultural information from “transmitter” to “receiver” that is implicit in epidemiological models) both interlocutors are active participants in generating this change. “Grounding” theories suggest that communication involves more than simply formulating a message and sending it off, but requires the mutual belief that what is being said will be understood by all parties.

Crucially, the notion that cultural information resists alterations during its transmission conflicts with a fundamental and particularly distinguished theory of cultural transmission: cumulative culture (Dunstone & Caldwell, 2018; Stout & Hecht, 2017). This approach to cultural evolution reflects the idea that cultural traits are gradually modified through transmission such that adaptive modifications accumulate over historical time (Dean, Vale, Laland, Flynn & Kendal, 2014). This theory operates from a basic assumption that transmission of cultural information naturally involves its modification in a way that fundamentally conflicts with the depiction of transmission under a disease spread formalisation.

The cumulative conceptualisation of modifications to cultural information is prominent in the literature and may be the most representative of genuine complexities underlying cultural dynamics. However, this triumph entails perhaps an inevitable downfall in that such a complex depiction of culture has proven exceptionally challenging to model in quantitative accounts (Buskell, Enquist, & Jansson, 2019). This paper provides an active-inference based quantitative account of cumulative culture as an accumulation of changes to cultural information over multiple transmissions.

2 Method

An emerging conclusion from the literature is that the term "transmission" for describing the spread of cultural information seems impoverished, as it leaves out the retention of cultural information. As implied by active inference and theoretical models of communication, the acquisition of cultural beliefs is as fundamental to the understanding of cultural information spread as their transmission. For this reason, we will henceforth be referring to what is known in the literature as cultural transmission as communication, or more technically- the local dynamics of cumulative culture.

2.1 Simulating the Local Dynamics of Communication

In our model, cultural transmission is cast as the mutual attunement of actively inferring agents to each other's internal belief states. This builds on a recent formalisation of communication as active inference (Friston & Frith, 2015) which resolves the problem of hermeneutics, (i.e., provides a model for the way in which people are able to understand each other rather precisely despite lacking direct access to each other's internal representations of meaning) by appealing to the notion of generalised synchrony as signaling the emergence of a shared narrative to which both interlocutors refer to. In active inference, this shared narrative is attained through the minimisation of uncertainty, or (variational) free energy when both communicating parties employ sufficiently similar generative models. We build on this to suggest that having sufficiently similar generative models allows communicating agents to recombine distinct representations of a belief (expressed as generative models) into one synchronised, shared model of the world. When we simulate the belief-updating dynamics between interacting agents, the cultural reproduction of a particular idea takes the form of a specific convergence between their respective generative models.

Under this theory, the elementary unit of heritable information takes the form of an internal belief state, held by an agent with a certain probability. When we simulate the belief-updating dynamics between interacting agents, a reproduced cultural belief is carried by the minds (or generative models) of both interlocutors as a site of cultural selection, where it may be further reproduced through the same process. Our simulations of communication involve two active inference agents with distinct generative models and belief claims that engage in communication over a hundred time steps.

2.2 Simulating the Global Dynamics of Cumulative Culture

Cultural beliefs and practices spread within a society through communication, a process which we have referred to as the local dynamics of cumulative culture. This description is appropriate because the accumulated outcomes of each (local) dyadic interaction collectively determine the degree to which an idea is prevalent in a culture. Moving from local communication dynamics to a degree to which an idea is prevalent in a cumulative culture is what we will refer to as the global dynamics of cumulative culture.

In our simulations of a cumulative culture, 50 active inference agents simultaneously engage in local dyadic communication as shown in our first simulation, such that 25 couples are engaged in conversation at every given time step. At the first time step, all agents have relatively similar belief states- referred to as the status quo. When we introduce an agent holding a divergent belief state to that of the status quo in the population, it propagates through it via pseudo-random engagements of agents in dialogue. In a simulated world of actively inferring agents, their individual mental (generative) models are slightly modified with every interlocutor they encounter, as their distinct representations converge to a shared narrative (Constant, Ramstead, Veissière, & Friston, 2019). The attunement of interlocutor's to each other's generative models on the microscale thus translates over time and with multiple encounters into collective free energy minimisation on the macroscale.

3 A Generative Model of Communication

In our simulations, agents attempt to convince each other of a cultural belief by utilising generative models that operate with local information only. For the establishment of such generative models, we will formulate a partially observed Markov decision process (MDP), where beliefs take the form of discrete probability distributions (for more details on the technical basis for MDP'S under an active inference framework, see Hesp 2019).

Under the formalism of a partially observed Markov decision process, active inference entails a particular structure. Typically, variables such as agent's hidden states (x, s), observable outcomes (o) and action policies (u) are defined, alongside parameters (representing matrices of categorical probability distributions).

3.1 Perceptual Inference

The first level of this generative model aims to capture how agents process belief claims they are introduced to through conversation with other agents. The perception of others' beliefs (regarded in active inference as evidence) requires prior beliefs(represented as likelihood mapping A1 about how hidden states (s1) generate sensory outcomes (o). Specifically, our agents predict the likelihood of perceiving evidence toward a particular expressed belief, given that this belief is "the actual state of the world". Parameterising an agent's perception of an interlocutor's expression of belief in terms of precision values can be simply understood as variability in agents' general sensitivity to model evidence. High precisions here correspond to high responsiveness to evidence for a hidden state and low precisions to low responsiveness to evidence. Precisions for each agent were generated from a continuous gamma distribution which is skewed in favour of high sensitivity to evidence on a population level (See figure 1: Perception).

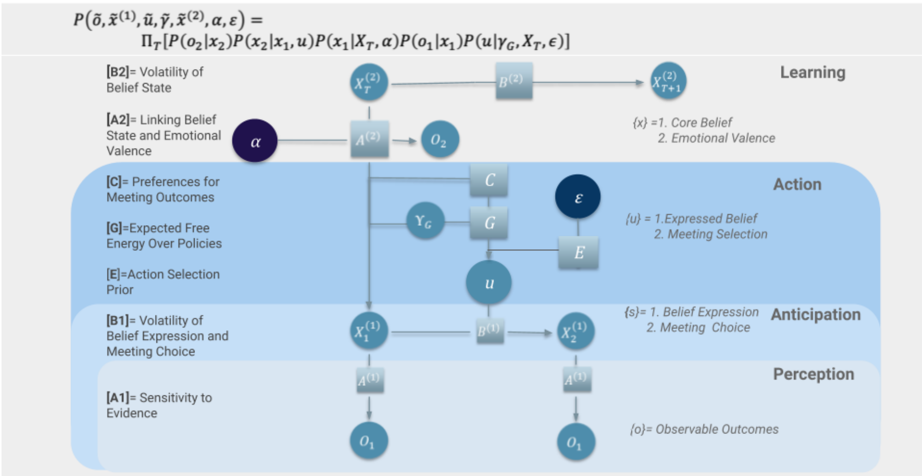

Fig. 1. A generative model of communication. Variables are visualised as circles, parameters as squares and concentration parameters as dark blue circles. Visualised on a horizontal line from left to right-states evolve in time. Visualised on a vertical line from bottom to top- parameters build to a hierarchical structure that is in alignment with cognitive functions. Parameters are described to the left of the generative model and variables are described on the right.

Updating of core belief based on beliefs expressed by self and another agent after each meeting (detailed descriptions of the computations involved in perceptual inference can be found under appendix):

$$Q\left(x_{core}^{(2)}\right) = \sigma\left(\ln x_{core}^{(2)} + \gamma_{A,self}^{(2)} \ln o_{expr,self} + \gamma_{A,other}^{(2)} \ln o_{expr,other}\right) \quad (1)$$

3.2 Anticipation

At this level, our generative model specifies agents' beliefs about how hidden states (detailed in appendix A2) evolve over time. State transition probabilities [B1] define a particular value for the volatility of an agent's meeting selection (s2) and belief expression (s1) [B1]. For each agent, this precision parameter is sampled from a gamma distribution, determining the a priori probability of changing state, relative to maintaining a current state. Note that belief states themselves are defined on the continuous range <0, 1> (i.e., as a probability distribution on a binary state), such that multiplication tends to result in a continuous decay of confidence over time in the absence of new evidence (where the rate of decay is inversely proportional to the precision on B) (See figure 1: Anticipation).

3.3 Action

After perceiving and anticipating hidden belief states in the world, our agents carry out deliberate actions biased towards the minimum of the expected free energy given each action (a lower level generative model for action is detailed in appendix A4 and A5). At each time point, a policy (U) is chosen out of a set of possible sequences for action. In our simulations, two types of actions are allowed: selecting an agent to meet at each given time point (u2) and selecting a specific belief to express in conversation (u1). The first allowable action holds 50 possible outcomes (one for each agent in the simulation) while the second is expressed on the range <0,1>, where the extremes correspond to complete confidence in denying or supporting the belief claim, respectively. Each policy under the G matrix specifies a particular combination of action outcomes weighted by its expected negative free energy value and a free energy minimising policy is chosen (See figure 1: Action).

Voluntary Meeting Selection. While the choice of interlocutor is predetermined in a dyad, our multi-agent simulations required some sophistication in formulating the underlying process behind agents' selection for a conversational partner (s3) at each of the hundred time points. Building on previous work on active inference navigation and planning (Kaplan & Friston, 2018), agents' meeting selection in our model is represented as a preferred location on a grid, where each cell on the grid represents a possible agent to meet (Appendix).

We demonstrate the feasibility of incorporating empirical cultural data within an active inference model by incorporating (1) confirmation bias through state-dependent preferences [C], biasing meeting selection through the risk component of expected free energy (G) and (2) novelty seeking through the ambiguity component of expected free energy. The first form of bias reflects the widely observed phenomenon in psychology research that people's choices tend to be biased towards confirming their current beliefs (Nickerson, 1998). The second form of bias reflects the extent to which agents are driven by the minimisation of ambiguity about the beliefs of other agents, driving them towards seeking out agents they have not met yet.

3.4 Perceptual Learning

On this level agents anticipate how core belief states (specified in appendix A1) might change over time [B2] (figure 2.3). This is the highest level of cognitive control, where agents experience learning as a high cognitive function (higher level generative model is detailed in appendix A3). By talking with other simulated agents and observing their emotional and belief states, our agents learn associations between EV and beliefs via a high level likelihood mapping [A2], (updated via concentration parameter α). The Updating of core belief, based on beliefs expressed by other agents, is detailed in appendix A7. This learning is important because it provides our agents with certainty regarding the emotional value they can expect from holding the alternative belief to the status quo, which has low precision at the beginning of the simulation (before the population is introduced to an agent proclaiming this belief). The prior P(A) for this likelihood mapping is specified in terms of a Dirichlet distribution (Appendix).

4 Results

4.1 Local Dynamics of Coupled Communication

In nature, generalised synchrony emerges from a specific coupling between the internal states of dissipative chaotic systems (Pikovsky, Kurths, Rosenblum & Kurths, 2003). In active inference communication, agents are coupled in a bidirectional action-perception cycle in which they can be described as coupled dynamical systems (Friston & Frith 2015). Specifically, our model defines perceptual inference as the coupling parameter linking the internal states of interlocutors.
Also understood as sensitivity to model evidence (A1), perceptual inference is a direct and explicit form of coupling that occurs over the span of a single dialogue such that it modulates agents' convergence of internal belief states during conversation (Fig.2).
Our results indicate that without sufficiently high precisions on sensitivity to model evidence, agents' ability to listen and attune to the belief expression of their partner is limited to the extent that they are responsive to sensory evidence from their environment.

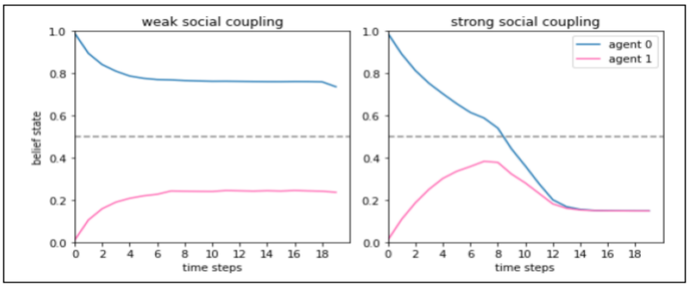

Fig.2: The [A1] parameter (sensitivity to model evidence) modulates the level of social coupling between agents in dialogue. **(Left)** When precision on sensitivity to model evidence is low (for both agents) their internal states are very weakly coupled, which results in each agent sticking to their own belief. **(Right)** When both agents have high sensitivity to model evidence, their beliefs converge into a shared representation of an idea that inhabits both of their generative models.

To get a sense of the implications of these simulations, it is important to make explicit the way in which they tie in to previous work on active inference communication. In 2015, Friston & Frith provided evidence for the notion that generalised synchrony becomes altogether unattainable when agents do not possess sufficiently similar generative models. Our model goes beyond this to provide evidence for the idea that only when generalised synchrony is attainable (i.e., when interlocutors possess sufficiently similar generative models), communication underlies a convergence between their belief states. Our simulations should therefore be understood as taking generalised synchrony for granted while providing evidence for the premise that the level to which agents' beliefs converge (i.e., the level of synchrony between their internal states) is modulated by their sensitivity to model evidence [A1].

4.2 Global Dynamics of Cumulative Culture

Our simulations of a cumulative culture should be understood as capturing the dynamics of a culture that is the sum (or-accumulation) of modifications to cultural beliefs and practices over time (Fig.3). While the local dynamics simulated in the previous section represent a single modification to cultural information (as a convergence between distinct belief states held by individual agents), these simulations accumulate these modifications and expose their emerging dynamics within the population. The fundamental achievement of these results is therefore their methodologically consistent and novel depiction of cumulative culture under a quantitative and measurable framework (namely, active inference).

We explain the communicative isolation observed in our simulations (Fig.3) as a self organised separation between groups of agents when they hold intractably divergent beliefs, such that communicative isolation best ensures local and collective free energy minimization. In other words, when an intractable divergent belief propagates within a homogenous population, communicative isolation between incongruent groups emerges as a strategy to minimize expected free energy, while the same strategy homogenizes the belief states of agents within congruent groups.

The above simulations also show how changes to parameters that determine levels of confirmation bias [C] and novelty seeking [G] affect the segregation within the population into groups of agents holding either status quo congruent beliefs or the alternative belief. When novelty seeking is upregulated, the population evolves such that the majority of agents end up subscribing to the alternative belief. However, when confirmation bias is upregulated, the majority of agents end up subscribing to the status quo. What these results indicate is that novelty seeking on a local level stimulates the population as a whole toward the adoption of a belief that is divergent from the status quo. This happens when novelty seeking individuals, which are open and willing to meet with agents of unknown beliefs, are intrinsically encouraged by their own curiosity to engage with a divergent belief. Once such agents become gradually more favoring of this belief they start to popularize it to the rest of the population. If the population is however, populated by a vast majority of agents biased toward meeting individuals with confirming beliefs, they do not engage with an alternative belief and it does not get popularized.

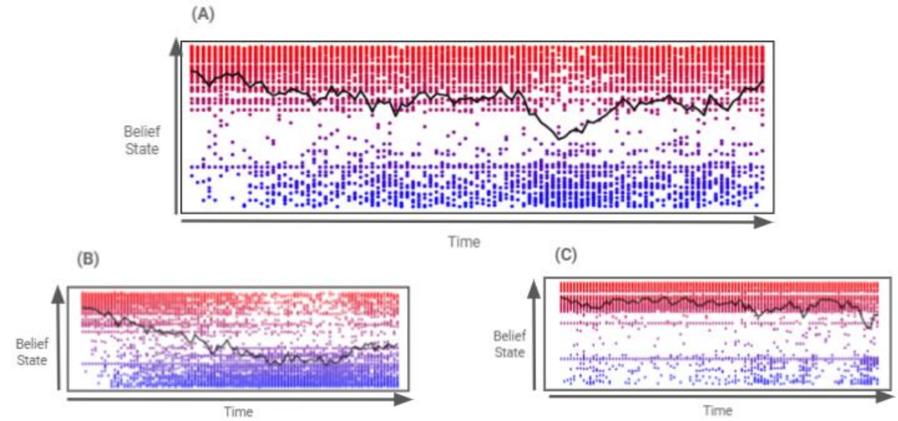

Fig 3: simulations of the spread of each agent's belief state (y) across time (x). 5o agents were used in this simulation and each of the 100 time steps represents the reproduced belief state outcomes of a particular combination of agents in dialogue. **(A)** Simulation of a Cumulative Culture. When a divergent belief state (blue) is introduced to the status quo population (red) at the first time step, it spreads through it via pseudo-random engagements of agents in dialogue that cumulatively change the belief structure within the population. Most notably, the introduction of a divergent belief seems to split the population into two subgroups: those holding a belief state that approximates the new divergent belief, and those holding an approximate status quo belief. This effect is modulated by agents' individual strategies for choosing which interlocutors to engage in conversation with (s3). **(B)** When novelty seeking is high in the population (above 10% of agents present high novelty seeking), the population is divided in favour of the divergent belief state, with more agents eventually holding this belief than the status quo. **(C)** When confirmation bias is high in the population (above 90% of agents present high confirmation bias) the population is divided in favour of the status quo belief, with more agents holding to this belief than the new and divergent belief.

5 Conclusion

In this paper, we employed an active inference model to tackle the complex task of formulating the dynamics underlying cumulative culture. Under this account, transmission is cast as a bidirectional process of communication that induces a generalised synchrony between the internal (belief) states of agents holding sufficiently similar generative models. Generalised synchrony is operationalised in our model as a particular convergence between the internal states of interlocutors, which is shown to be largely modulated by sensitivity to model evidence [A1].

When we simulate a population of agents that simultaneously engage in the converging dynamics of communication over time, cumulative culture emerges as the collective behavior brought about by these local modifications to cultural beliefs and practices. When a divergent belief is introduced to the status quo, it spreads within the population and brings about a collective behaviour that seems to be characterised by a divide between different belief groups. The level to which the status quo population defects to the divergent belief is mediated by local psychological strategies of confirmation bias and novelty seeking.

References

1. Buskell, A., Enquist, M., & Jansson, F. . A systems approach to cultural evolution. Palgrave Communications, 5(1), 1-15 (2019).
2. Bettencourt, L. M., Cintrón-Arias, A., Kaiser, D. I., & Castillo-Chávez, C. The power of a good idea: Quantitative modeling of the spread of ideas from epidemiological models. Physica A: Statistical Mechanics and its Applications, 364, 513-536 (2006).
3. Clark, H. H., & Brennan, S. E. Grounding in communication (1991). .
4. Constant, A., Ramstead, M. J., Veissière, S. P., & Friston, K. Regimes of expectations: An active inference model of social conformity and human decision making. Frontiers in psychology, 10, 679 (2019).
5. Creanza, N., Kolodny, O., & Feldman, M. W. Cultural evolutionary theory: How culture evolves and why it matters. Proceedings of the National Academy of Sciences, 114(30), 7782-7789 (2017).
6. Dean, L. G., Vale, G. L., Laland, K. N., Flynn, E., & Kendal, R. L. Human cumulative culture: a comparative perspective. Biological Reviews, 89(2), 284-301(2014).
7. Dunstone, J., & Caldwell, C. A. Cumulative culture and explicit metacognition: A review of theories, evidence and key predictions. Palgrave Communications, 4(1), 1-11 (2018).
8. Echterhoff, G., Higgins, E. T., & Levine, J. M. . Shared reality: Experiencing commonality with others' inner states about the world. Perspectives on Psychological Science, 4(5), 496-521 (2009).
9. Friston, K., & Frith, C. A duet for one. Consciousness and cognition, 36, 390-405 (2015).
10. Friston, K. J., & Frith, C. D. Active inference, communication and hermeneutics. cortex, 68, 129-143 (2015).
11. Hesp, C., Ramstead, M., Constant, A., Badcock, P., Kirchhoff, M., & Friston, K. . A multi-scale view of the emergent complexity of life: A free-energy proposal. In Evolution, Development and Complexity (pp. 195-227). Springer, Cham (2019) .
12. Kaplan, R., & Friston, K. J. Planning and navigation as active inference. Biological cybernetics, 112(4), 323-343 (2018).
13. Kashima, Y., Bain, P. G., & Perfors, A. The psychology of cultural dynamics: What is it, what do we know, and what is yet to be known?. Annual Review of Psychology, 70, 499-529 (2019)..

14. Pikovsky, A., Kurths, J., Rosenblum, M., & Kurths, J. Synchronization: a universal concept in nonlinear sciences (No. 12). Cambridge university press (2003).
15. Stout, D., & Hecht, E. E. Evolutionary neuroscience of cumulative culture. Proceedings of the National Academy of Sciences, 114(30), 7861-7868 (2017).

Appendix A Generative Model Architecture, Factors and Parameters

A.1 Higher level hidden state factors:

$$x^{(2)} : \{x^{(2)}_{core},$$
[core beliefs of self and others about a particular claim, across days]
$$x^{(2)}_{mem},$$
[memory of having visited each agent]
$$x^{(2)}_{habit}\ \}$$
[habits of self, across days]

A.2 Lower level hidden state factors (specify events on a given 'day')

$$x^{(1)} : \{\ x^{(1)}_{loc},$$
[self location, where each agent has a unique 'home' location]
$$x^{(1)}_{belief},$$
[beliefs of self and others about a particular claim]
$$x^{(1)}_{visit},$$
[beliefs about having visited each agent]
$$x^{(1)}_{sat}\ \}$$
[satisfaction of self and others]

A.3 Higher level generative model:

$$x^{(1)}_{belief} = A^{(2)}_{belief} x^{(2)}_{core}$$
[core beliefs specify prior expectations for beliefs on the lower level]

$$x^{(1)}_{sat} = A^{(2)}_{sat} x^{(2)}_{core}$$
[core beliefs specify satisfaction states for the lower level]

$$x^{(1)}_{visit} = A^{(2)}_{mem} x^{(2)}_{mem}$$
[memory specifies beliefs about having visited each agent on the lower level]
$$E_{expr} = A^{(2)}_{expr} x^{(2)}_{habit}$$
[habits of self specify prior tendency for belief expression]

$$x^{(2)}_{T+1} = B^{(2)} x^{(2)}_{T}$$
[higher-level states decay over time: gradual forgetting]

A.4 Lower level generative model for action:

Action model for meeting selection:

In our simulations, We have incorporated psychological biases in agents' preferences for meeting similar (i.e belief compatible) or unknown agents. Note that while agents biased toward confirming beliefs would tend toward individuals with similar beliefs to their own, novelty seekers would not look for the opposite of this (i.e look for individuals of divergent beliefs to their own), but rather have a preference for individuals of yet unknown beliefs.

In active inference, action selection is guided by the expected free energy [G], which entails maximising the expected benefit or utility of an action (known as pragmatic value), while also maximising the potential information gain of future actions by reducing uncertainty about the causes of valuable outcomes (known as epistemic value). These constraints to action selection could be interpreted as formalising the exploration-exploitation trade-off in learning systems. Epistemic value (exploration) refers to the benefit related to searching over a sample space in order to get a better estimation of promising areas that will maximise pragmatic value (exploitation). Active-inference agents would therefore maximise epistemic value until information gain is low, after which the maximisation of pragmatic value and exploitation are assured (Friston, Rigoli, Ognibene, Mathys, Fitzgerald & Pezzulo, 2015).

In our model, agents' choice in meeting interlocutors with known and similar beliefs versus those with unknown beliefs can be cast in terms of a tradeoff between pragmatic and epistemic value. On the one hand, a confirmation bias emerges from the maximisation of expected utility, increasing synchronisation between interlocutors' internal models, thus allowing for the emergence of shared expectations (Hesp et al., 2019). On the other hand, novelty seeking emerges from the maximisation of information gain, allowing for the exploration of the sample space. Also understood as intrinsically motivated curious behaviour (Friston, Lin, Frith, Pezzulo, Hobson & Ondobaka, 2017), maximisation of epistemic value allows individuals to better predict the consequences of their actions (e.g.,, which agent to meet) through greater certainty about the hidden states of their environment (e.g., the beliefs of other agents).

From the point of view of agents in our simulations, increasing pragmatic value translates into selecting to meet interlocutors with similar beliefs, while increasing epistemic value translates into selecting agents whose beliefs are unknown or highly uncertain (This way, a meeting increases information gain). From this point of view, it is clear the two values constrain each other and maximizing both simultaneously is partially (but not entirely) paradoxical. While maximising pragmatic value requires agents to choose to meet with an interlocutor they know is similar to them, maximising epistemic value necessitates they meet with one they do not know at all.

$$P(u_{loc}) = \sigma\left(-\gamma_{G,loc}G_{loc} + \gamma_{E,loc}E_{loc}\right)$$
$$G_{loc} = o_{u,belief} \cdot (\ln\ o_{u,belief} - C_{belief})\ + H \cdot x_{u,2,visit}$$
$$x_{u,2}^{(1)} = B_u^{(1)}x_1^{(1)}$$
$$o_{u,belief} = A_{belief}^{(1)}x_{u,2}^{(1)}$$
$$C_{belief} = \ln\ \left(A_C^{(2)}x_{core}^{(2)}\right)$$

$if\ x_{visit,j} = 1:$

[equals 1 if agent visited a particular agent j]

$H_j = 0$

[ambiguity is zero if agent visited this agent j already] $else:$

$H_j = 0.1$

[ambiguity is non-zero if agent has not visited agent j yet]

A6. Generative process:

Generative process for meeting selection:

$u_{loc} \sim P(u_{loc})$ *[actual meeting u_{loc} is sampled from meeting selection prior $P(u_{loc})$]*

Generative process for belief expression and EV (satisfaction) of each agent:

At a high level of cognitive control, agents incorporate a series of processes underlying the selection of a particular belief for expression (u2). Other than the partial reliance on a low level habitual factor [E], this action involves multiple higher order considerations.

First, an agent considers their core belief state (x), and the way this state apriori maps on to one of two discrete emotional valence states (s2) via an initial likelihood mapping [A2] Emotional Valence (EV) is defined as the extent to which an emotion is positive or negative (Feldman Barrett & Russell, 1999), such that agents' core beliefs are apriori associated with either positive emotional valence or negative emotional valence (with some probability). As a minimal form of vicarious learning, the initial mapping is further updated based on associations agents observe between their interlocutors' expressed belief state and EV value The initial mapping therefore involves minimal precision for the expected EV for belief 2, since agents are first introduced to this belief (and associated EV) during the simulations. For this reason, the initial likelihood mapping between states is updated throughout our simulation via a crucial concentration parameter (α).

EV states are generated from core belief states, using a (learnable) likelihood mapping:

$$x_{sat}^{(1)} = A_{sat}^{(2)} x_{core}^{(2)}$$

Confidence of belief expression is generated using a Gamma distribution, where the rate parameter expris the Bayesian model average of (+,-)values associated with high and low satisfaction:

$$P\left(\gamma_{expr}\right) \approx \Gamma\left(1, \beta_{expr}\right)$$
$$\beta_{expr} = \beta^{(+,-)} \cdot x_{sat}^{(1)}, \qquad \beta^{(+,-)} = [0.25, 2.0]$$

The expression of beliefs is guided by current core beliefs (scaled with satisfaction-dependent expr) and by habitual belief expression Eexpr(scaled with a fixed parameter E,expr):

$$P\left(u_{expr} | \gamma_{expr}\right) = \sigma\left(-\gamma_{expr} \ln x_{core}^{(2)} + \gamma_{E,expr} E_{expr}\right)$$

The intrinsically stochastic and itinerant nature of the generative process of communication is modeled by using a two-dimensional Dirichlet distribution to generate observed expressions on
the range [0,1], where each agent's belief expression prior Puexpr|expr is used to specify their concentration parameters (multiplied by 12 to reduce variance):

$$o_{expr} = Dir(12 u_{expr})$$

Generative process for emotional valence expressed by each agent:

$$o_{sat} = A_{sat}^{(1)} x_{sat}^{(1)}$$

[satisfaction observed by interaction partner corresponds to actual satisfaction]

The EV state predicted is then used to generate an action confidence value (γ) such that positive EV generates high confidence in a certain expression of the belief state (u1) and negative EV generates low confidence values. Higher confidence values produce higher precision on the expected free energy (G) for one's belief expressed in the current conversation.

A7. Perception:

Updating beliefs about the other agent's belief based on their expression:

$$Q\left(x^{(1)}_{belief}\right) = o_{expr}$$

Updating of core belief based on beliefs expressed by other agents:

$$Q\left(x^{(2)}_{core}\right) = \sigma\left(\ln ln\, x^{(2)}_{core} + \gamma^{(2)}_{A,self} \ln ln\, o_{expr,self} + \gamma^{(2)}_{A,other} \ln ln\, o_{expr,other}\right)$$

A8. Learning:

Habit learning for meeting selection:

$$P\,(E_{loc}) = Dir\,(e_{loc})$$
$$Q\,(E_{loc}) = Dir\,(e_{loc} + 0.05u_{loc})$$

Habit learning for belief expression:

$$P\,(E_{expr}) = Dir\,(e_{expr})$$
$$Q\,(E_{expr}) = Dir\,(e_{expr} + 0.1o_{expr})$$

Perceptual learning for the mapping between satisfaction and core beliefs, based on the expressions of other agents:

$$P\left(A^{(2)}_{sat}\right) = Dir\left(a^{(2)}_{sat}\right)$$

$$Q\left(A^{(2)}_{sat}\right) = Dir\left(a^{(2)}_{sat} + \gamma^{(2)}_{A} o_{expr} \ln ln\, x^{(1)}_{sat}\right)$$

A9. Initialisation of parameters for each agent:

$$\gamma^{(2)}_{A,belief} \sim \Gamma\,(5, 6)$$

[regulates the integration of beliefs of other agents in one's own core belief]

$$\gamma^{(2)}_{A,sat} \sim \Gamma\,(10, 1)$$

[regulates learning rate of mappings between satisfaction and core belief, based on observed correspondences in other agents]

$$\gamma_{G,loc} \sim \Gamma\,(1, 1)$$

[regulates reliance on action model in selecting agent to meet]

$$\gamma_{E,loc} \sim \Gamma\,(1, 1)$$

[regulates reliance on habitual prior in selecting agent to meet]

$$\gamma_{E,expr} \sim N\left(\frac{\gamma_{E,loc}}{10}, \frac{\gamma_{E,loc}}{200}\right)$$

[regulates reliance on habitual prior in expressing action, which correlates with $\gamma_{E,loc}$*]*

$$\gamma^{(2)}_{B,core} \sim \Gamma\,(4, .5)$$

[regulates stability of core beliefs across days]

$$\gamma^{(2)}_{B,habits} \sim \Gamma\,(.5, 1)$$

[regulates stability of expression habits across days]

$$B^{(2)}_{0} = [[.75, .25]\,, [.25, .75]]$$

[specifies baseline transition probabilities]

$$B^{(2)} = \sigma\left(\gamma^{(2)}_{B} \ln\, B^{(2)}_{0}\right)$$

[corrects $B^{(2)}_{0}$ *using the agent-specific* $\gamma^{(2)}_{B}$ *values]*

Agents with relatively weak confirmation bias:

$$A^{(2)}_{C} \sim Dir(6, 4)$$

[induces weak reliance on core beliefs for specifying lower-level preferences]

Agents with relatively strong confirmation bias:

$$A^{(2)}_{C,1} \sim Dir\,(999, 1)$$

[induces strong reliance on core beliefs for specifying lower-level preferences]