CellPress
OPEN ACCESS

Review

Formalising social representation to explain psychiatric symptoms

Joseph M. Barnby ⓘ ,[1,*] Peter Dayan,[2,3] and Vaughan Bell[4,5]

Recent work in social cognition has moved beyond a focus on how people process social rewards to examine how healthy people represent other agents and how this is altered in psychiatric disorders. However, formal modelling of social representation has not kept pace with these changes, impeding our understanding of how core aspects of social cognition function, and fail, in psychopathology. Here, we suggest that belief-based computational models provide a basis for an integrated sociocognitive approach to psychiatry, with the potential to address important but unexamined pathologies of social representation, such as maladaptive schemas and illusory social agents.

Formalising social representation

The field of cognitive neuroscience has taken important steps toward uncovering the mechanisms by which humans build representations of our environment. Integrated experimental and theoretical work suggests that humans encode, organise, recall, and generalise knowledge within cognitive structures to allow for flexible behaviour [1–4]. There is now ample evidence that social information is structured in a similar way [5]; relationship structures, mental states, traits of individuals, and social environments can be mapped as conceptual representations that are recoverable from experimental investigations [6], observational studies [7], and neural activation [8,9]. One notable feature of social environments is that they are highly interactive: humans react, cooperate, compete, and conflict with others who have their own beliefs, desires, and traits. Therefore, humans need methods for dynamically modelling other agents [10] and, by extension, theorists working in this domain require computational models that can adequately represent this state space. However, while many elements involved in goal-directed nonsocial information processing have been captured through computational models that centre on model-based learning and prospective reasoning [11,12], the use of formal computational models to conceptualise **social representation** (see Glossary) are relatively newer [13–16].

Computational models allow integration of biological and psychological determinants of decision making to provide refined explanations of mental disorder that cut across multiple levels of explanation [17]. Indeed, changes in how humans encode and predict basic value information in the environment have been used to explain biases related to delusion formation [18,19] and hallucinations [20,21], where overweighted and rigid **prior** beliefs coupled with sensory uncertainty generate aberrant perceptions and false beliefs. Similarly, **reinforcement learning (RL)** models have been useful when characterising major depressive disorder [22] and anxiety [23]; **artificial neural networks** have been proposed as models of aberrant perception [24]; and **attractor state models** have been used to explain probabilistic reasoning in delusions [25]. Psychiatric disorders frequently, and sometimes fundamentally, involve alterations to social perception and behaviour [26,27], meaning that computational models of social interaction in psychiatric disorder, although still nascent, are an exciting development in cognitive science [28–30].

Highlights

Bayesian-belief models formalise rich state spaces and use full posterior probability distributions that can characterise and extend sociocognitive theories; they also offer an opportunity to test psychiatric theories that implicate changes in the structural representations of self and others.

Tasks that involve a range of joint outcomes for the participant and their partner(s) interrogate critical aspects of complex social processing, and better reflect how we come to know other minds.

Computational models of social representations provide a test-bed to examine how uncertainty, social decision making, and learning processes contribute to core psychiatric phenomenology, such as social schemas and the production of illusory social agents in psychosis.

[1]Social Computation and Cognitive Representation Lab, Department of Psychology, Royal Holloway, University of London, Egham TW20 0EX, UK
[2]Max Planck Institute for Biological Cybernetics, Tübingen, 72076, Germany
[3]University of Tübingen, Tübingen, 72074, Germany
[4]Clinical, Educational, and Health Psychology, University College London, London WC1E 7HB, UK
[5]South London and Maudsley NHS Foundation Trust, London SE5 8AZ, UK

*Correspondence:
Joseph.barnby@rhul.ac.uk
(J.M. Barnby).

In this review, we provide a framework for examining how belief-based generative models can be used to specify and test hypotheses about how humans (which we refer to as 'selves' or 'self') infer the presence of, and represent, other agents ('others' or 'other') during social decision making. We first consider normal function and then look at how these models have been, or could be, applied to understand psychiatric disorders. This review focuses on approximate **Bayesian models**, which consider the uncertainties that individuals have about their beliefs, and how these uncertainties affect both inference and learning. Bayesian models offer a versatile and clear framework for representing individual and interactive task spaces [16,31,32]. Since reasoning about others involves considering a self's own model of the other in a specific context, it would be incomplete to use generic (e.g., neural network) models to explain interactive learning. Even if generic models are used in a model-free manner to distil the interaction, model-based methods would still provide essential understanding.

We collectively refer to formal theoretical models that make claims about the probabilistic structure of self and other mental states as computational social representation models. We develop our ideas from foundational work that outlines the importance of rich computational techniques to explain general economic [13] and social psychological phenomena [15,17,33] and the role of social representation in social cognition and psychiatric disorders [34–38].

Social representations and psychiatric disorder

Inherent aspects of social interaction (uncertainty, contextual reward, and social dependencies) lend themselves to structured, probabilistic models, with Bayesian formulations modelling how uncertainty about others' intentions can be integrated into social decision making. 'Structured probabilistic model' refers to a generative hypothesis that links prior probability distributions, such as the probability that a self values equality or the self's assumptions about an other's preferences for equality, to value functions, such as the probability that the self will be prosocial in this moment or an other will act prosocially. Variance within each distribution allows for the quantification of uncertainty over the experimenter's parameter approximation of a self. It also allows experimenters to model a self's own epistemic uncertainty of their characterisation of themselves and their environment (Figure 1; reviewed in [14]; also see [39,40]). It is often the case that these two forms of uncertainty may be hard to distinguish using practical experimental procedures and, thus, robust fitting procedures should be used. These dimensions provide multiple abstract representations of value components of self and others (e.g., uncertainty and valence) that reward-learning models of decision making exploit to generate appropriate value signals [41]. Iterative game theoretic tasks (Box 1; reviewed in [33]) have offered particularly fruitful paradigms for modelling social interactions in which a self's internal representations of the other agent(s) are required for drawing task-relevant inferences.

Here, we review formal theories that are intended to model generative processes in social cognition. We focus on models that characterise the self and other(s) using rich **state spaces** and consider their use for probing psychiatric disorder (for a summary see Table 1), drawing from and extending previous work [14,33]. We adopt model classes that embody different temporal, recursive, and causal structural features (for a visual summary see Figure 2, Key figure): shallow depth-of-mentalisation (sDoM; including a self's beliefs, and a self's beliefs about the other given past actions), hierarchical DoM (hDoM; self-beliefs about how the self's actions will influence the other's actions at different levels of recursion), and group mentalisation (how the self may hold beliefs about multiple others). We note here that, while RL formulations of social interaction have been developed, may fit within our taxonomy, and make important contributions [42–46], these are not the focus of our review. Nevertheless, we hope that our taxonomy provides a useful framing for wider work in this area.

Glossary

Artificial neural network: collection of inter-related nodes, inspired loosely by biological neurons. Artificial nodes connect to one another using inputs and outputs, where inputs may be raw data (e.g., pixel colour within an image) or the weighted output of another node. Nodes generate probabilistic outcomes using mathematical functions (e.g., SoftMax) to transmit information.

Attractor state model: attractor states are cast as hubs generated by the dynamic interplay of units within neural networks; over time, small network units generate the ability to influence a larger proportion of units within a system. In psychological systems, this can be framed as the influence of nodes within a dynamic network.

Bayesian model: probabilistic model that involves the interaction of priors and likelihoods combined to form **posterior** beliefs about the chance of an event occurring. Beliefs in these models are not necessarily propositional attitudes, as defined in the philosophy of personal belief, but rather 'effective' beliefs. In other words, beliefs are probability distributions over *a priori* qualitative dimensions necessary for performing approximate probabilistic inference.

Computational model: set of mathematical equations that determine the flow of information (e.g., beliefs → actions) given their structure and interdependency. A computational model can generate the actions of an agent given values that determine the solution of each equation.
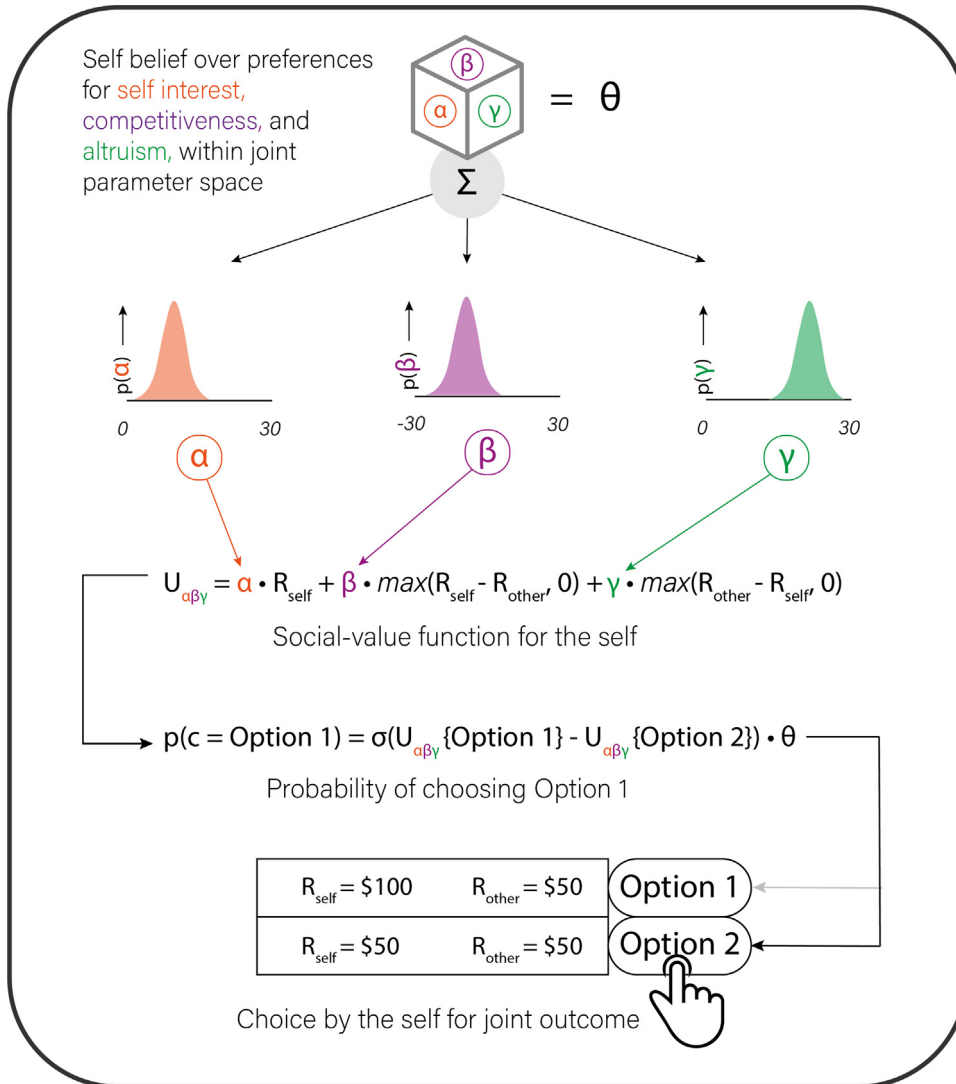
Likelihood: probability of observations given a set of prior beliefs, such as the probability a prosocial other made a competitive decision to split some money: $p(D_{competitive}|Prosocial)$.

Posterior: product of the likelihood and the initial belief(s), normalised by the marginalised sum of the product to form a new belief, such as the probability an other is prosocial given their recent competitive decision:

$$p(Prosocial|D_{competitive}) = \frac{p(D_{competitive}|Prosocial) \times p(Prosocial)}{\sum p(D_{competitive}|Prosocial) \times p(Prosocial)}$$

Prior: initial belief(s) before any evidence has been considered, such as the starting probability an other is prosocial: $p(Prosocial)$.

Reinforcement learning (RL): associative formulation that determines the interface between an agent's beliefs and actions. Associative RL models

Trends in Cognitive Sciences

Figure 1. How rich representations of the self translate into joint social action. Social representation models can hold rich components that represent self and/or an other(s) social values. By placing probability distributions over state space, social representation models can make predictions about how a self may act in each social environment, where outcomes determine joint rewards for the self and an other. In this example, drawing from a Fehr–Schmidt (FS) model [130], we demonstrate how self-beliefs (θ) that contain the subjective self-valuation of absolute self-reward (α), and relative rewards of inequality aversion (β) and altruism (γ) may form an action probability when combined with the difference in utility of choosing Option 1 or Option 2 across the entire state space. In terms of implementation, joint values of self-belief (probability distributions over state space) and state space (values with a range accommodating variables of psychological interest) are both encoded as separate 3D matrices. Self-beliefs along each dimension (α, β, or γ) can be calculated by summing (marginalising) along each axis. Importantly, the FS equation is only itself part of the explanation; it specifies the additional model components that provide the causal model structure around the function. Our illustration depicts internal distributions placed over the FS equation; rather than a fitting procedure, it allows experimenters to estimate the uncertainty a self may hold over their own social preferences. In conjunction with other model elements (e.g., social recursion in hierarchical depth-of-mentalisation or bias terms in shallow depth-of-mentalisation) and probabilistic distributions, the FS equation can become part of an explanatory theory of cognitive processes. σ, sigmoid function, c, choice; therefore p(c = Option 1), probability of choosing Option 1.

typically describe updates about the inferred value (V) of different stimuli over time (t), and use the difference between expected and actual outcomes (δ) to drive learning about environments:

$$V^{t+1} = V^t + \alpha\delta$$

$$\delta = V^t - reward$$

α : *learning rate*

Social representation: mental information-bearing structures that organise knowledge about the internal states of others, and the relationship of others to the self; they involve the dynamic interaction and transitions between traits, states, and actions [34]. **State space:** span of independent latent values (S) that determine action probabilities and outcomes (O): $p(O^t|S^t)$. This provides a basis for an agent to form beliefs within the context at hand. State spaces may be continuous (e.g., 0–30) or discrete (0,1) and contain transition structures that determine when an agent moves from one state to another: $p(S^{t+1}|S^t)$.

Box 1. Tasks to assess social agent representation

To test social decision making, experimental frameworks must accommodate social-regarding, multiperson decisions; that is, decisions that involve outcomes for the self and the interacting other(s). This allows the experimenter to capture more ecologically valid ways of how selves come to know other minds in everyday life encounters [114,115] rather than only relying on private, 'spectator' states, which may have a minor emphasis on how humans come to know others [116].

Microeconomic games have been used to create carefully controlled socially regarding situations, in which a self's actions influence themselves and other(s), which allow for cooperative and competitive outcomes and have been cited as particularly useful tools for probing interpersonal distortions in psychiatric disorders [13]. Game theoretic concepts operationalised within socially regarding tasks can offer a substrate for mathematically transparent maps from objective value utility to subjective outcomes.

Iterative game theoretic tasks allow the experimenter to study how selves make sequences of social-regarding decisions over time. For example, in one version of the trust task we describe in the main text, selves act as investors and, in each round, send a sum of money from an endowment to an other. The experimenter triples the money as it transfers to the other, who then decides, as a dictator, how much to send back to the self. The roles could also be reversed. Optimising outcomes in the trust task generally requires selves to build and exploit a predictive model of how the other might feel and act in response to the self's choices and, therefore, the behaviour of the self is tied to inferential processes and downstream strategic considerations (what might the other feel and then do if I send less money? [74]). Some iterative game-theoretic tasks have separated choice behaviour and social inference from strategic goals. For example, using a repeated Dictator game, selves are asked to infer the true intentions of their other after each other decision (*cf.* [56]). Despite joint outcomes, the other is unable to view the self's attributions, which, therefore, are observed free from strategic repercussions.

Repeated tasks have also been useful to track social belief change over time, for example in social-value orientation (SVO) tasks (*cf.* [65]), in which alternation between self-guided and other-guided decision making allows experimenters to estimate how a self's individual-level beliefs influence their estimations of a social other's SVO.

Shallow depth-of-mentalisation

sDoM models allow experimenters to formalise hypotheses about the beliefs of the self (their motivations, mood, and intentions), and the beliefs the self holds about an interaction other's internal state, conditioned on their past behaviour (Figure 2A [14]). sDoM models are shallow in the sense of not including the sort of deeper social recursion and cognitive hierarchy that we consider later. These models may be important in cases in which hierarchical social planning is restricted [47] or when selves are simply observing others, rather than interacting with them (in a game-theoretic sense).

One role of sDoM models is to probe inferences about the internal representations of external agents from observation alone. For example, building on descriptive accounts of rational action planning [48], experimenters tested how selves predict the action path of an other's movements around a 2D grid (a 'maze-world') given the cost of action and action variability [49–51]. Likewise, prior work [52] used a gambling task that required selves to infer emotions of an other based on the integration of facial cues, utterances, and financial results. sDoM models have also been used to investigate the approval rate of individuals within a group following feedback, and the subsequent impact on self-esteem [53]. Given the basic disruptions in mental state inferences as observed in schizophrenia [54], these models provide a formal platform to assess when greater consistency in an other's actions ceases to be an inferred reliable predictor of their hidden motives, and the conditions under which models that explain the link between inference and action may begin to fail.

Socially regarding paradigms (Box 1) take the step of considering interaction between self and others. These permit investigation of the effect of joint outcomes on action selection and beliefs, and formally address how social context influences cognitive dynamics, such as learning and uncertainty. One example is the use of a repeated Trust Game and a slot machine paradigm to study generalised anxiety disorder [55]. In the repeated Trust Game, one participant is an investor

Table 1. Summary of models of social agent representation

Class[a]	Type[b]	Ostensive purpose	Refs
sDoM	Social regarding	Likelihood of harmful intent and self-interest attributions being emitted in response to an other's decisions in a repeated Dictator game.	[57,58]
	Nonsocial regarding	Estimate the degree to which humans use a Bayesian inverse planning process to predict the actions of an intentional, rational agent.	[49–51]
	Social regarding	Approximate the magnitude and uncertainty of a self's trust beliefs about interacting others given the investment of the other during an iterative trust task.	[55]
	Social regarding	Approximate the changes to belief about self-esteem in response to social evaluation.	[53]
	Non-social regarding	Integration of facial and utterance trust-cues on emotional inference.	[52]
sDoM; multi-phase design	Social regarding	Approximate the social-value preferences for the absolute and relative payoffs of a self, and a self's belief about the preferences of an other along the same dimensions.	[65]
	Non-social regarding	Estimate the change in monetary preferences of the self following exposure to an other using a Delegated Interpersonal Discounting Task [62].	[60]
hDoM	Social regarding	Estimate the mentalisation depth of selves. The model allows recursive estimation of a self's strategy, and the strategy of their other within a 'centipede' game.	[70,71]
		Estimate the mentalisation depth of selves. The model allows recursive estimation of a self's strategy, and the strategy of their other within a 'Hide and Seek' game.	[82,83]
		Estimate utterances of a speaker to signal utility of semantic meaning for listener, and listener's model of speaker's intentions.	[72,73]
		Generate choices of a self in an iterative trust task, conditioned on beliefs of self and their beliefs about their other.	[47,74,75];
		Allow a recursive estimation of a self's strategy, and recursive strategy of their other, given the evolving conditions of the Stag Hunt Game.	[79–81]
Group mentalisation	Nonsocial regarding	Estimate the likelihood of differential group membership of several others during a choice preference task.	[87,90]
	Social regarding	Estimate individual and group utility of investment, and higher-order probability of other group members free-riding.	[91,92]

[a]Classes of model concern the overall generative purpose.
[b]Type of model describes whether the purpose of the model is to detect social regarding (the actions of the participant and/or the partner influence joint outcomes) or nonsocial-regarding outcomes (actions of the participant and/or partner do not influence joint outcomes).

who may choose to send a proportion of an endowment to another participant, who acts as the trustee. The experimenter scales the investment by a multiplicative factor (e.g., three); the trustee can then choose how much to send back to the investor. The game is played over multiple rounds, with the investor and trustee adopting the same roles each time. Within this task, it has been reported that selves more heavily weighted social losses compared with nonsocial losses, and positive first impressions lasted throughout social, but not nonsocial contexts when updating beliefs about the trustworthiness of the other [55]. Selves who scored above the clinical cut-off on a generalised anxiety disorder scale showed a selective reduction in learning about exploitative partners and negative events, and a failure to enhance learning as social uncertainty increased. This result goes against the conventional Bayesian expectation that learning should increase because of epistemic uncertainty (if a model is not yet reliable, more sampling is needed), suggesting a general insensitivity to uncertainty. Results also contrast with RL models [23], which suggest that, in nonsocial contexts, negative (rather than positive) outcomes are overvalued in those with high anxiety. Nevertheless, specific task demands may preclude comparison between parameter estimates derived from either study, and care must be taken when generalising parameter estimates into new contexts.

Social-regarding paradigms have also been used to understand better the coupling of policy uncertainty with priors and **likelihoods** over mental state attributions during interactions in an iterative Dictator Game [56–58]. These models provide formalisations that encompass earlier

Key figure

Taxonomy of computational social representation

Figure 2. (A) Schematic of shallow depth-of-mentalisation (DoM) models with recursive depth ($k \leq 0$). In this example, the experimenters first learn about the magnitude and uncertainty of a self's beliefs about the intentions of the general population. The self then approximates the likelihood about the intentions of the other given the other's decisions. (B) Shallow DoM models can be used in inter-temporal tasks to assess changes to the representation of the self and other social agents. The experimenter's beliefs over the preferences of the self and the self's beliefs about the other can provide a basic scientific tool to test how the self and other are statistically represented, and what may perturb these representations. In this example, the experimenters first learn about the magnitude and uncertainty of a self's beliefs. The experimenters can then use the self's predictions about the other to build a possibly approximate model of the self's beliefs about the preferences of the other. (C) Hierarchical DoM models ($k > 0$) allow selves to make recursive inferences in their model of an other, and the model an other may hold about the self, up to level k (in one flavour, assuming that the other is level $k-1$; in another flavour, that the other is level $-1...k-1$ or $0...k-1$, with a prior probability for each such level). (D) Group mentalisation models are affiliated with the inferences and social behaviours of the group. This may be nonsocially regarding (e.g., group classification based on an other's private preferences) or socially regarding (e.g., the probability of defection in a Public Goods Game based on the prior history of the group).

ideas about the predictive mind in social cognition [35], such as the flexibility over attributions (trait), probabilities over the current intentional motives of the other (state), and subsequent attributional change (action). As a result, these models allow core computational theory to be extended to characterise an environment that allows for joint social dependencies, and to probe clinically relevant phenomena (e.g., multiple mental state attributions) while retaining core computational facets (e.g., learning or uncertainty).

Multiphase task structures allow simultaneous approximation of multiple representations and can assess how the self's beliefs change following social observation, or how the self suppresses updating of irrelevant representations to use context-appropriate beliefs (Figure 2B). To note, these designs probe the neutrality of self after social observation; being exposed to others may alter self-preferences [59] or effect inferences about the self [60]. Inter-temporal discounting tasks [61–63] estimate whether a self prefers a smaller monetary reward now or a larger one at a later time point (phase 1), how a self learns about an interaction other's preferences (phase 2), and whether exposure to an interaction other changes the self's preferences (phase 3). By estimating belief distributions about each agent [60], experimenters can approximate the causal influence of exposure to an other on the beliefs of the self (e.g., being exposed to a patient other can inspire greater patience in the self) and the way this depends on the self's initial uncertainty about themselves (and, putatively, the assumed similarity between the two).

Multiphase designs have been extended to social regarding contexts, testing how self social-value preferences integrate into beliefs about the social-value orientation (SVO) of an other, and how paranoia may distort this process. This extends prior assessments of how social influence may govern a self's SVO [64]. Using models that estimate the preference for relative and absolute joint pay-offs (Figure 1 [65]), experimenters can estimate how uncertainties and social orientation of self-preferences may disrupt the process of learning about the social values of an other along multiple, contextually relevant dimensions (i.e., relative versus absolute pay-off preferences). These multiphase configurations allow testing of key psychiatric phenomena, such as the misattunement hypothesis [66], with results suggesting that discrepancies in self-other representations make individuals more prone to developing paranoid explanations of social events [65]. By extending prior imaging work using multiphase paradigms [61–63], it will be useful to probe whether computational biases as a function of paranoia are reflected in reduced functional modelling of others, whether therapeutic outcomes may be marked by increased functional modelling of others, and, importantly, whether this is reflected in positive outcomes.

Hierarchical depth-of-mentalisation

hDoM models capture the recursive inference of value functions between two players in strategic planning. hDoM are used to formalise how agents model first, second, or k^{th}-order nested beliefs (e.g., the self's model of the other's model of the self's model of the other, etc; Figure 2C) [67–69]. This can then branch off over the orthogonal dimension of multiple steps of forward planning, when self and other interact over several rounds [47,70–73]. Therefore, such models formalise the processes underlying social dependencies between self and other [14].

For example, in the Trust Game, one major factor governing the interaction is inequality aversion. If the trustee can persuade the investor that they are inequality averse, then the investor has a reason to trust that it is safe to invest. Recursive reasoning about inequity aversion in the trust task proceeds as follows: a level 0 investor would consider their actions in light of their beliefs about the trustee (based on the trustee's history of decisions; e.g., [55]); a level 1 Investor would also consider the beliefs the trustee holds about the investor and, therefore, how the investor's own actions may influence the trustee's subsequent actions in this and future rounds; a level 2

investor considers their actions in light of a model that encompasses the belief in the trustee's beliefs about the investor's beliefs, and so on.

Recursive modelling was used in a ten-round version of the repeated Trust Game [74,75]. This extends local behavioural assessments of coaxing by an other [76]. Experimenters tested hypotheses about the level of such modelling and the degree of forward planning in which trustees or investors indulge as they make their choices. It has been found, for instance, that people with borderline personality disorder (BPD), when acting as trustees, have shallower recursive depth compared with healthy volunteers [74]. Later versions have characterised additional factors, including a form of risk aversion of the investor, operationalised as their reluctance to invest because of the danger of a poor return by the trustee; and (state) irritation, which accumulates according to (trait) irritability in the light of unsatisfactory behaviour of the partner and leads to a breakdown of recursive reasoning, forward planning, and cooperation [77,78]. People with BPD are more irritable as trustees.

The Stag Hunt Game is a microeconomic game of cooperation: two players (as hunters) can capture a stag if they both opt to do so and will both be rewarded with many points. However, either or both players could instead opt individually to capture a rabbit, which is worth fewer points. If one opts for the stag, the other for the rabbit, then only the latter gets any points. The Stag Hunt Game has been modelled similarly to the Trust Game to approximate the forward planning of selves [79], and how forward planning may be implemented in the brain [80]. Importantly, this model has been used to test those with autism spectrum disorder (ASD). Experimenters found that, while controls were best fit by a recursive model that considered the strategy of the other, selves with ASD were best fit by a fixed model, which contained significantly less recursive depth, and were less flexible (i.e., were more tied to past strategies), despite being no different to controls in their cooperation rate overall [81]. This finding converges with later work using a hDoM model [82] during a 'hide and seek' game [83]; those with ASD were less flexible in their strategy compared with neurotypical individuals and displayed shallower mentalisation, although this was only a disadvantage when paired with a social versus nonsocial other.

The convergence of shallower forward-planning during social paradigms in individuals with either BPD or ASD is intriguing given their respective symptom profiles: ASD reflects an undersensitivity to social cues, whereas BPD reflects an oversensitivity. Here, computational modelling has the potential to disentangle how each may show similar behavioural phenotypes, but with distinct underlying mechanisms. For example, the introduction of irritability [77,78] makes it possible to assess whether reduced forward-planning performance in those with BPD can be explained by exaggerated irritation rather than by a deficit in forward planning per se. Likewise, it may be that, under different social contexts (e.g., cooperation versus competition), forward planning becomes restricted or not used to inform decision making (e.g., [81]) and may highlight where ASD and BPD diverge. Notably, when prosocial behaviour is sufficiently predictive, selves may not need to develop sophisticated models of their others [79]; a form of equilibrium can be achieved through simple prosocial social-value biases, which may be less computationally costly [84]. Therefore, those with ASD may have difficulties forming complex, sophisticated models of others, rather than a deficit of basic associative processes. Indeed, there are normative contextual predictions about forward planning that can be used to test divergence: higher levels of strategic planning are optimal under competitive conditions [70,71,82] and lower levels of sophistication are optimal under cooperative conditions [85].

Group mentalisation
Although most models aim to understand social interaction on a one-to-one level, understanding and strategising based on group characteristics, such as collective beliefs, actions, and group

identity, are also a key part of navigating the social world. Generative processes hypothesised to govern social categorisation have been formalised using nonsocial-regarding preference tasks (Figure 2D [86]). In one such example [87], the self is required to infer the group membership of others by observing their pattern of preferences, which include socially benign characteristics (e.g., movie choices). In another, participants learn characteristics that more strongly determine action and judgement (e.g., political/moral beliefs) [88]. Models of social structure learning [87–90] explain the statistical process through which selves assign others to social clusters, and how the influence of group members may distort social categorisation (e.g., similarity to a member of the group).

The social-regarding beliefs a self holds about a group have been modelled using the Public Goods Game (PGG), in which selves secretly decide how much money to contribute to the group. The total amount is multiplied by a factor of i (between 1 and the total number of players) and is then distributed evenly among all players, including any who declined to contribute themselves. Modelling of the PGG has suggested selves statistically represent a group as a single entity [91]. This has been extended to use a blend of individual utility and group utility functions, estimating the joint utility of a self's own investment in the group, and their belief that others in the group would free-ride [92]. Deeper forward planning was not found to predict behaviour over and above shallow mentalisation; the model captured beliefs about the selves' immediate reward versus the long-term group reward.

To our knowledge, such Group Mentalisation models have yet to be applied to psychiatric phenomena, despite the likely importance of group-level cognition in a range of psychopathological states. For example, in game theoretic tasks, perceived group cohesion increases paranoid attributions of intent toward groups of opponents [93]; in the Cyberball task, those at clinical high risk for psychosis show higher sensitivity to attributing paranoia when socially excluded [94]; using a PGG, youths show high levels of cooperation but at the expense of being vulnerable to exploitation [95].

Future directions for formal theories of social representation in psychiatry

Any disorder that impacts cognition might lead to social disability. However, we suggest two components of psychopathology where computational social representation models could make a particular contribution. This is owing to their specifically representational nature and their evidenced role in distress and impairment.

The first is explaining the role of maladaptive schemas in psychiatric disorder. Schemas are high-level cognitive structures comprising memories, models, and action plans that relate to the self, the world, and others, and which are developed during childhood and are updated through experience [30,38]. With respect to theories of social representation [35], schemas may be considered a trait that has downstream influence on state and action space. The development of maladaptive schemas is strongly linked to adverse childhood experiences [96] and there is good evidence that they are causal factors in depression, anxiety, and eating disorders [97]. Importantly, along with their effect in nonsocial domains, schemas predominantly involve representations of self and others, and these representations are considered key to their causal role in biasing perceptions and actions [35,36]. Although typically conceptualised in descriptive terms, the social representation components of schemas could be captured through formal models that quantify biases in social-value functions. This would require identifying how other-related schema components relate to sensitivity to engage appropriate forward planning and inference across social contexts, learning in the face of uncertainty about others, and sensitivity to identify oneself with a positive group. This procedure would also benefit from longitudinal work to ascertain the impact of abusive or neglectful social environments on alterations to learning in childhood/adolescence, and the resultant

maladaptive social representations in adulthood. Previous work [30,98] has suggested ways in which therapeutic mechanisms of interest are translated into computationally tractable questions; using this foundational work, it is possible to make specific predictions around the change of schema in psychiatric disorder using established models (Box 2).

The second potentially fruitful future target for computational social representation approaches is the experience of illusory social agents in psychosis. Psychotic episodes, common in disorders such as schizophrenia and severe bipolar disorder, and occurring after a range of neuropsychiatric difficulties, are notable for the fact that those affected frequently experience interaction with illusory social agents in the form of communication with personified hallucinated voices [99] or have delusional beliefs that involve being targeted by illusory individuals, groups, or other unreal social agents [100]. Unlike problems such as social anxiety disorder, where social representations may be distorted but remained tethered to genuine individuals and their actions, in psychosis, the representations of social agents may be internally generated, seemingly *sui generis* in some cases or after becoming untethered from the genuine social agents with whom they are initially identified and instead begin to be perceived as being autonomous social actors [34]. These experiences are a significant driver of distress [101,102] and are now a target for therapeutic intervention [103]. The degree to which humans normally attribute agency in otherwise nonagentive contexts is a matter of continuing investigation (Box 3). We suggest that computational social representation models have the potential to formally model how illusory intentional agency develops, becomes untethered from the social environment, is structured, and drives disability. For example, testing whether deeper forward planning is engaged when outcomes are stochastic is a potential marker indicating that intentional agency is being inferred from coincidental outcomes. This frames illusory intentional agency as an identifiable state [37] or process, which may become more frequent under conditions of environmental or physiological stress [104] and may be probed through game theoretic paradigms.

A related consideration is the scope of amodal interactions, which have been focused on for characterising the issues that arise in real-world social dynamics. For example, analysis of eye-gaze and facial displays has been fruitful for uncovering differences in social perception in ASD [105] and ecological momentary assessment has offered a granular picture of social and

---

Box 2. Simulating a social exchange to test psychiatric theory

Theories of the effect of early trauma on later risk of developing psychopathology have variously suggested that it disrupts certainty about the sense of self [117], leads to overcertainty about threat from others [118], or leads to maladaptive schemas that can represent a range of pathological beliefs about the self and others [30,98].

Ways in which descriptive therapeutic questions may be translated into computational terms for specific hypothesis testing have been suggested [30]. This work highlights how experimenters can operationalise descriptive theories using computational models that place distinct theoretical, statistical bounds on the uncertainty around how a self may represent themselves and the intentions of a social other, and the interaction of both on future beliefs about the self. For example, experimenters can then test whether those who have been subject to early life trauma attribute greater uncertainties to the self (and self-regarding values), and/or attribute greater uncertainties to representations of others and their intentions. The relationship of changes to structural representations of the self, and subsequent neural and physiological function, may then become a more tractable question to answer [119].

To illustrate, we outline a model that defines the joint probability of self-valuing money for themselves, the probability of the self valuing equal monetary outcomes, and the self's beliefs about what the other believes along the same dimensions [65], in an interaction involving the risk of personal loss and joint payoff outcomes. The model can also be extended to assess how much a self's representations have changed after exposure to the other (see [60]) to test the impact of the interaction itself.

Simulations indicate how outcomes would operationalise verbal theories from the literature (Figure I). Disruptions to the uncertainty of a self's beliefs about their own social preferences leads to more stochastic outcomes when the self is the decider, and subsequently greater shifts in beliefs following exposure to an interaction other. By contrast, disruptions to the precision of a self's beliefs about the interaction other leads to poor predictive accuracy, suggesting an inability to successfully model others. Dynamic partner behaviour, based on the experimenter's estimate of self-preferences in phase 1, or partners created based on the behaviour of clinical populations (e.g., [76]), may also allow well-controlled predictions and efficient inference about participants as in a form of optimal experimental design. The failure of the model to capture predictions from verbal theories should drive revision of these hypotheses.
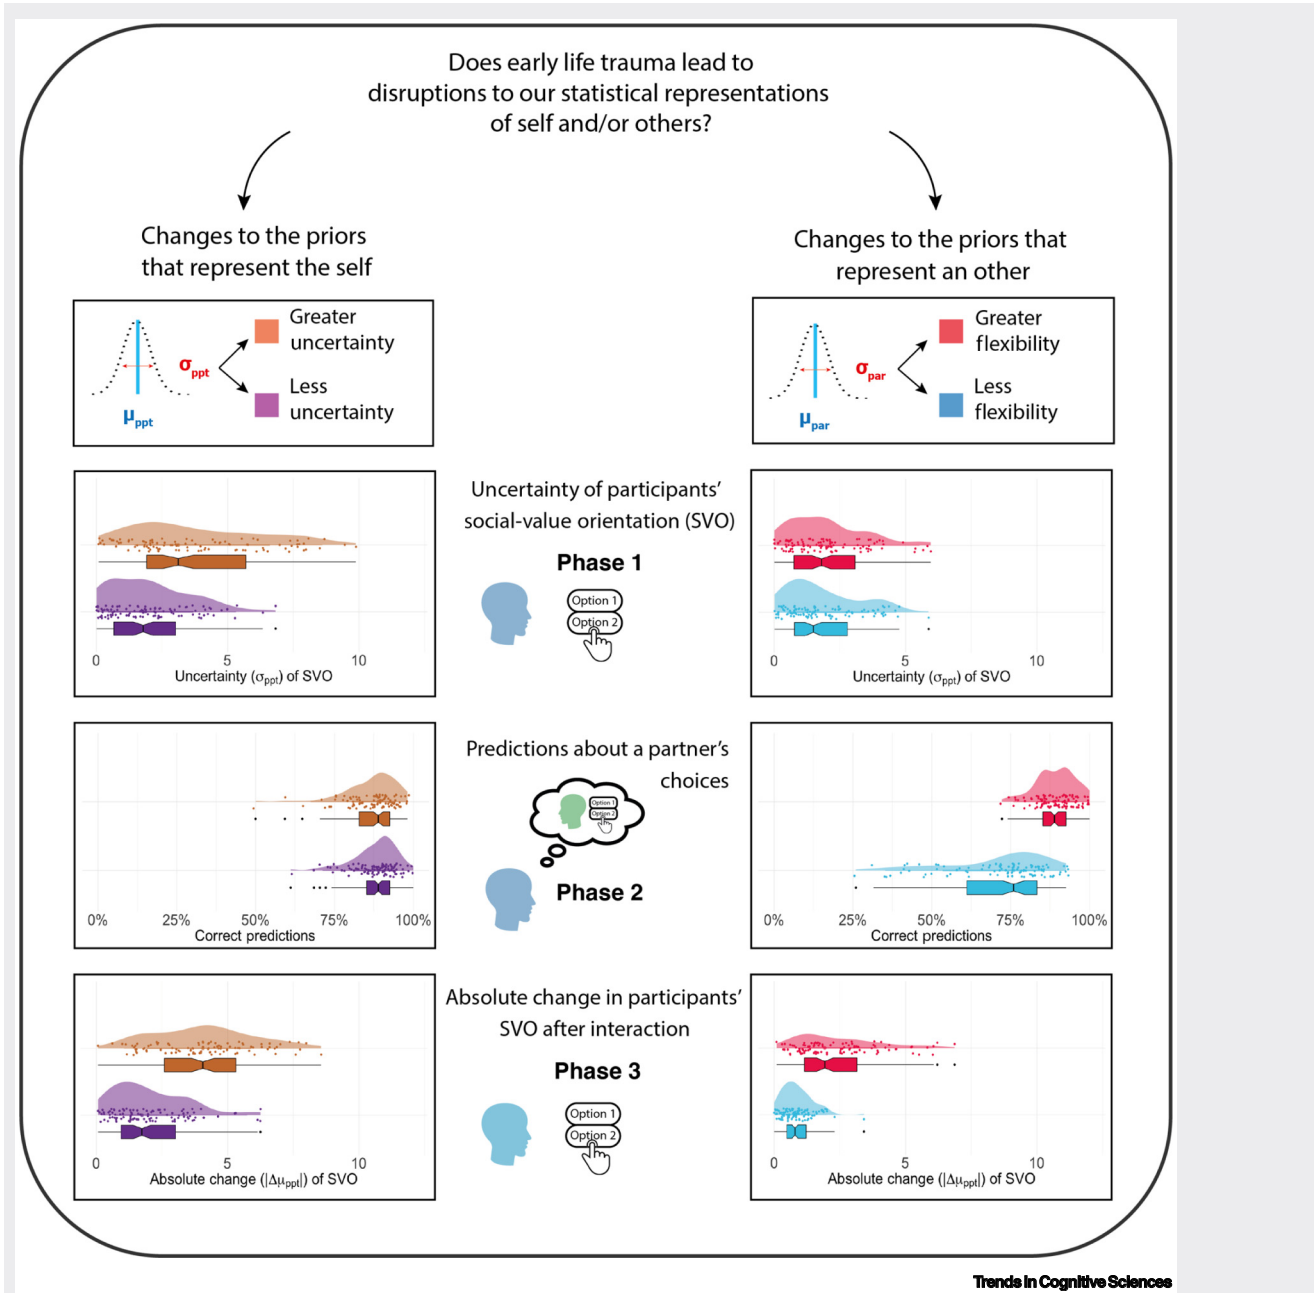
Figure I. Simulations to test changes in self-other representations. Simulating perturbations to the generative process can make firm predictions to formally assess competing hypotheses. Here, synthetic participants are playing an extended social-value orientation (SVO) task. Participants are matched with three different anonymous interaction partners over different phases; one per phase (see [44]). In phase 1, participants need to decide how to split money with their interaction partner over a series of trials; this is governed by a magnitude (μ_{ppt}) and uncertainty (σ_{ppt}) over self SVO. In phase 2, participants need to guess their new interaction partner's decisions as accurately as possible; this is governed by flexibility around a self's beliefs about an other's SVO (σ_{par}), where greater values represent a greater ability to step away from self SVO to accommodate the other's SVO. Phase 3 allows estimation of changes in SVO of participants following exposure to their interaction partner following the interaction of distributions in phases 1 and 2 (see [65]). Right column: uncertainty does not differ between groups in phase 1. In phase 2, the blue group find it much harder to learn about their partner (fewer % correct). In phase 3, the blue group's SVO is shown to be very rigid compared with that of the red group (lower absolute change in a self's SVO; $|\Delta\mu_{ppt}|$). Left column: in phase 1, there is larger uncertainty self-SVO in the orange group. However, both groups are as good as each other in predicting their partner in phase 2. In phase 3, due to higher uncertainty over their own SVO, the orange group is highly influenced by their interaction partner versus the purple group.

Box 3. Agency and intention in sociocognitive tasks

How are social cognitive processes involved in interactions with intentional social agents? In the laboratory, experimenters can compare conditions in which others believe that they are interacting with an other that is either a person or an algorithm [120] or test two or more people during interpersonal exchange [121]. In online studies, coordinating genuine social interaction is potentially more challenging. One method is to use 'live' social tasks that match participants in real time although this can require complex software to support this. Another method uses asynchronous 'ex-post matched' designs, in which selves respond to genuine decisions collected from others taking part at other times [122]. This is more convenient for experimenters but limits the degrees of freedom in the paradigm to ensure that the full range of responses is available for later matching.

However, interactive paradigms are important, because prior work has found that social cognitive processes may be differentially deployed when selves are, or believe they are, interacting with an intentional agent compared with an algorithm as other. This has led to the hypothesis that there may be a process of interpretation that leads humans to deploy a 'social' or 'nonsocial' framework that determines which cognitive resources are applied to a task [120,123].

A related issue is why humans frequently perceive the presence of intentional agents when there is none, or adopt the strategy of acting in accordance with the idea that an interaction may be intentional, even in a range of non-intentional situations. Indeed, humans attribute agency readily to nonsocial objects that appear to mimic a behavioural policy, for example, where participants attribute intentions to animated shapes [124,125]. Anthropologically, an approach to understanding the world focused on the role of unseen, intentional agents has been noted as a common factor across religions [126]. In psychopathology, the overperception of external social agency to the point of distress and disability is a key experience of psychosis [34] and has been found in studies that test the attribution of agency in the psychosis spectrum [127–129]. This raises the question of whether patients with psychosis systematically misdeploy a combination of agentive inference (i.e., misdeployment of DoM) and maladaptive social representations (e.g., high weight placed on inferences of harmful intent) in non-intentional contexts.

---

physiological sensitivities in psychosis [106]. Despite the utility of modelling approaches to capture alterations in psychiatric diagnoses, such as bipolar disorder, paranoia, and ASD cited previously, the framework in which we situate our formal models cannot capture all possible naturalistic elements of social interaction. Integrating formal models alongside real-world, naturalistic markers of social interaction may provide a more complete picture of social representation in psychiatric disorders and should be a future goal of the field.

Concluding remarks

There is a clear need to understand how we structure beliefs about ourselves and others in ways that are amenable to hypothesis-driven cognitive science. Game theoretic tasks can provide controlled paradigms to create socially regarding scenarios that encourage cooperation, conflict, and evolving social inference, thereby facilitating this approach. Crucially, creating simulations from a specified model can provide normative measures of social interaction to assess individual deviation (e.g., away from normative forward planning under different social contexts). This latter property allows for the testing of core descriptive clinical theory that captures rich phenomenology of interest in a shared social environment that can answer questions around the causal processes that go awry in psychiatric disorders.

Nevertheless, we note some potential challenges with the application of computational social representation. The first is the extent to which computational models distinguish between latent representations, which may be causally important but unobservable and even inaccessible to the affected individual, and overt, conscious representations that have a vivid role in mental life. Both are important features of psychopathology, and it is not clear how belief-based models, in principle, can distinguish between them. The issue might ultimately be an empirical one, with hypothesised representations needing to be investigated to test the extent to which they reflect implicit versus explicit social representations as experienced by the patient. Current models of metacognition may be particularly important in assessing this, because they provide a window onto the insight the self has about its own states [107]. It is an open question as to whether

Outstanding questions

How do mechanisms of basic value learning contribute to social representation? Creating carefully designed within-participant causal experiments can test how specific disruptions to basic value learning alter the development of novel social representations of interaction partners.

How stable are parameter and model estimates within computational models of social representation? Estimating the prior beliefs of self and other with a social representation model requires that experimenter's approximations of parameters are reliable, and that winning models are differentiable from other alternative models on recovery. Ensuring recoverability and robust fitting procedures is essential. Stable estimation then allows for the tracking of fluctuations of self-other representations during longitudinal measurements.

Do parameter estimates from multiple formal models of social representation cluster provide a 'social fingerprint' that defines intraindividual distortions to self and other representation? Considerations include the reliability of models and socially regarding tasks within participants, rigour in model fitting procedures, and the need to gather large, high-quality sample sizes to detect small effects.

How can computational models of social representation fit within current dynamic theories of mental disorder? Network and dynamic systems models of psychopathology suggest that disorders are self-maintaining networks of symptoms. Given that altered social representations are likely important in psychopathology, understanding how parameters within social representation models feedback and interact with each other is crucial.

When do we make agentive attributions in nonsocial contexts? The extent to which we perceive agency or explicitly use the 'intentional stance' to interpret nonsocial situations is not well understood and may question our assumption that nonsocial tasks are necessarily devoid of social processing.

metacognitive insensitivity to explicit representations of self may also lead to recursive distortions the self's implicit representations from the perspective of the other.

Second, aberrant social representations, such as maladaptive schemas and illusory social agents, in psychosis are known to evolve in content and accessibility over the lifespan [108,109]. Although this is not an inherent limitation of computational social representation models, it has not been substantially tested in psychiatry, because most studies have been restricted to brief laboratory-based tasks rather than being integrated into longitudinal cohort paradigms, where changes in social representations over time can be effectively studied [110].

In this review, we do not consider biological realisations. Bayesian models, and indeed computational models in general, are abstractions that make precise predictions and provide formal explanations of cognitive processes. How faithfully neural processing hews to the approximately Bayesian precepts that we have articulated is an important, but unanswered, question across many areas of cognitive science [111,112]. Nevertheless, computational models and *in silico* 'lesion studies' simulating psychiatric dysfunction may still have the capacity to constrain explanations at lower levels, implying their relevance for understanding biological mechanisms [113]).

Finally, the impact of task engagement and comprehension when estimating the approximated beliefs of participants may not be immediately apparent; this is especially true in cases where cognitive impairment is common (e.g., in schizophrenia) and may confound parameter estimates and task performance. Careful consideration must be made within task design to ensure factors, such as general cognitive function, attention, and comprehension, are controlled for, or model outcomes are compared across patient groups to demonstrate specificity of model predictions, experimental manipulations, and reliability (see Outstanding questions).

Formal models of social representation provide a tractable scientific tool to assess a range of social representations that have been cited as centrally important in myriad psychiatric disorders. We believe computational modelling of social representation can provide a precise, clear, and detailed framework with which to assess some of the most complex social phenomena in psychiatry, and that it may enrich prior work on the representation of minds [35] to explain higher level abstractions implemented in social neurocognitive processes.

Declaration of interests

None declared by authors.

References

1. Behrens, T.E. *et al.* (2018) What is a cognitive map? Organizing knowledge for flexible behavior. *Neuron* 100, 490–509
2. Epstein, R.A. *et al.* (2017) The cognitive map in humans: spatial navigation and beyond. *Nat. Neurosci.* 20, 1504–1513
3. Liu, Y. *et al.* (2022) Decoding cognition from spontaneous neural activity. *Nat. Rev. Neurosci.* 23, 204–214
4. Tolman, E.C. (1968) Cognitive maps in rats and men. *Psychol. Rev.* 55, 192
5. Schafer, M. and Schiller, D. (2018) Navigating social space. *Neuron* 100, 476–489
6. Conway, J.R. *et al.* (2020) Understanding how minds vary relates to skill in inferring mental states, personality, and intelligence. *J. Exp. Psychol. Gen.* 149, 1032
7. Thornton, M.A. *et al.* (2022) The 3d Mind Model characterizes how people understand mental states across modern and historical cultures. *Affect. Sci.* 3, 93–104
8. Peer, M. *et al.* (2021) Structuring knowledge with cognitive maps and cognitive graphs. *Trends Cogn. Sci.* 25, 37–54
9. Park, S.A. *et al.* (2021) Inferences on a multidimensional social hierarchy use a grid-like code. *Nat. Neurosci.* 24, 1292–1301
10. Schilbach, L. *et al.* (2013) Toward a second-person neuroscience 1. *Behav. Brain Sci.* 36, 393–414
11. Daw, N.D. *et al.* (2011) Model-based influences on humans' choices and striatal prediction errors. *Neuron* 69, 1204–1215
12. Keramati, M. *et al.* (2016) Adaptive integration of habits into depth-limited planning defines a habitual-goal-directed spectrum. *Proc. Natl. Acad. Sci. U. S. A.* 113, 12868–12873

13. Camerer, C.F. (2009) Behavioral game theory and the neural basis of strategic choice. In *Neuroeconomics* (Glimcher, P.W. *et al.*, eds), pp. 193–206, Academic Press
14. FeldmanHall, O. and Nassar, M.R. (2021) The computational challenge of social learning. *Trends Cogn. Sci.* 25, 1045–1057
15. Montague, P.R. *et al.* (2012) Computational psychiatry. *Trends Cogn. Sci.* 16, 72–80
16. Vélez, N. and Gweon, H. (2021) Learning from other minds: an optimistic critique of reinforcement learning models of social learning. *Curr. Opin. Behav. Sci.* 38, 110–115
17. Lee, D. (2013) Decision making: from neuroscience to psychiatry. *Neuron* 78, 233–248
18. Ashinoff, B.K. *et al.* (2022) Rethinking delusions: a selective review of delusion research through a computational lens. *Schizophr. Res.* 245, 23–41
19. Feeney, E.J. *et al.* (2017) Explaining delusions: reducing uncertainty through basic and computational neuroscience. *Schizophr. Bull.* 43, 263–272
20. Corlett, P.R. *et al.* (2019) Hallucinations and strong priors. *Trends Cogn. Sci.* 23, 114–127
21. Horga, G. and Abi-Dargham, A. (2019) An integrative framework for perceptual disturbances in psychosis. *Nat. Rev. Neurosci.* 20, 763–778
22. Chen, C. *et al.* (2015) Reinforcement learning in depression: A review of computational research. *Neurosci. Biobehav. Rev.* 55, 247–267
23. Pike, A.C. and Robinson, O.J. (2022) Reinforcement learning in patients with mood and anxiety disorders vs control individuals: a systematic review and meta-analysis. *JAMA Psychiatry* 79, 313–322
24. Keshavan, M.S. and Sudarshan, M. (2017) Deep dreaming, aberrant salience and psychosis: connecting the dots by artificial neural networks. *Schizophr. Res.* 188, 178–181
25. Adams, R.A. *et al.* (2018) Attractor-like dynamics in belief updating in schizophrenia. *J. Neurosci.* 38, 9471–9485
26. Hitchcock, P.F. *et al.* (2022) Computational psychiatry needs time and context. *Annu. Rev. Psychol.* 73, 243
27. Schilbach, L. (2016) The neuroscience of mimicry during social interactions. *Emot. Mimicry Social Context* 72, 72–89
28. Huys, Q.J. *et al.* (2011) Are computational models of any use to psychiatry? *Neural Netw.* 24, 544–551
29. Huys, Q.J. *et al.* (2016) Computational psychiatry as a bridge from neuroscience to clinical applications. *Nat. Neurosci.* 19, 404–413
30. Moutoussis, M. *et al.* (2018) Computation in psychotherapy, or how computational psychiatry can aid learning-based psychological therapies. *Comput. Psychiatry (Cambridge, Mass.)* 2, 50
31. Gershman, S.J. and Niv, Y. (2010) Learning latent structure: carving nature at its joints. *Curr. Opin. Neurobiol.* 20, 251–256
32. Tervo, D.G.R. *et al.* (2016) Toward the neural implementation of structure learning. *Curr. Opin. Neurobiol.* 37, 99–105
33. Rusch, T. *et al.* (2020) Theory of mind and decision science: towards a typology of tasks and computational models. *Neuropsychologia* 146, 107488
34. Bell, V. *et al.* (2017) Rethinking social cognition in light of psychosis: reciprocal implications for cognition and psychopathology. *Clin. Psychol. Sci.* 5, 537–550
35. Tamir, D.I. and Thornton, M.A. (2018) Modelling the predictive social mind. *Trends Cogn. Sci.* 22, 201–212
36. Tariq, A. *et al.* (2021) Relationship between early maladaptive schemas and anxiety in adolescence and young adulthood: a systematic review and meta-analysis. *J. Affect. Disord.* 295, 1462–1473
37. Tariq, A. *et al.* (2021) A meta-analysis of the relationship between early maladaptive schemas and depression in adolescence and young adulthood. *Psychol. Med.* 51, 1233–1248
38. Young, J.E. *et al.* (2003) *Schema Therapy: A Practitioner's Guide*, Guilford Press
39. Kreps, D.M. and Wilson, R. (1982) Reputation and imperfect information. *J. Econ. Theory* 27, 253–279
40. Fudenberg, D. and Levine, D.K. (1989) *Equilibrium Payoffs with Long-run and Short-run Players and Imperfect Public Information*, MIT Press
41. Ruff, C.C. and Fehr, E. (2014) The neurobiology of rewards and values in social decision making. *Nat. Rev. Neurosci.* 15, 549–562
42. Hampton, A.N. *et al.* (2008) Neural correlates of mentalizing-related computations during strategic interactions in humans. *Proc. Natl. Acad. Sci. U. S. A.* 105, 6741–6746
43. Suzuki, S. *et al.* (2015) Neural mechanisms underlying human consensus decision-making. *Neuron* 86, 591–602
44. Zhang, L. and Gläscher, J. (2020) A brain network supporting social influences in human decision-making. *Science. Advances* 6, eabb4159
45. Na, S. *et al.* (2022) Computational mechanisms underlying illusion of control in delusional individuals. *Schizophr. Res.* 245, 50–58
46. Na, S. *et al.* (2021) Humans use forward thinking to exploit social controllability. *eLife* 10, e64983
47. Hula, A. *et al.* (2015) Monte Carlo planning method estimates planning horizons during interactive social exchange. *PLoS Comput. Biol.* 11, e1004254
48. Gergely, G. and Csibra, G. (2003) Teleological reasoning in infancy: The naive theory of rational action. *Trends Cogn. Sci.* 7, 287–292
49. Baker, C.L. *et al.* (2009) Action understanding as inverse planning. *Cognition* 113, 329–349
50. Baker, C. *et al.* (2011) Bayesian theory of mind: modelling joint belief-desire attribution. *Proc. Annu. Meet. Cognit. Sci. Soc.* 33, 2469–2474
51. Baker, C.L. *et al.* (2017) Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nat. Hum. Behav.* 1, 1–10
52. Ong, D.C. *et al.* (2015) Affective cognition: exploring lay theories of emotion. *Cognition* 143, 141–162
53. Low, A.A.Y. *et al.* (2022) Self-esteem depends on beliefs about the rate of change of social approval. *Sci. Rep.* 12, 1–15
54. van der Weiden, A. *et al.* (2015) Self–other integration and distinction in schizophrenia: a theoretical analysis and a review of the evidence. *Neurosci. Biobehav. Rev.* 57, 220–237
55. Lamba, A. *et al.* (2020) Anxiety impedes adaptive social learning under uncertainty. *Psychol. Sci.* 31, 592–603
56. Barnby, J.M. *et al.* (2020) Paranoia, sensitization and social inference: findings from two large-scale, multi-round behavioural experiments. *R. Soc. Open Sci.* 7, 191525
57. Barnby, J.M. *et al.* (2020) Reduction in social learning and increased policy uncertainty about harmful intent is associated with pre-existing paranoid beliefs: evidence from modelling a modified serial dictator game. *PLoS Comput. Biol.* 16, e1008372
58. Barnby, J.M. *et al.* (2022) The computational relationship between reinforcement learning, social inference, and paranoia. *PLoS Comput. Biol.* 18, e1010326
59. Devaine, M. and Daunizeau, J. (2017) Learning about and from others' prudence, impatience or laziness: the computational bases of attitude alignment. *PLoS Comput. Biol.* 13, e1005422
60. Moutoussis, M. *et al.* (2016) How people use social information to find out what to want in the paradigmatic case of intertemporal preferences. *PLoS Comput. Biol.* 12, e1004965
61. Garvert, M.M. *et al.* (2015) Learning-induced plasticity in medial prefrontal cortex predicts preference malleability. *Neuron* 85, 418–428
62. Nicolle, A. *et al.* (2012) An agent independent axis for executed and modeled choice in medial prefrontal cortex. *Neuron* 75, 1114–1121
63. Thomas, L. *et al.* (2022) Contagion of temporal discounting value preferences in neurotypical and autistic adults. *J. Autism Dev. Disord.* 52, 700–713
64. Murphy, R.O. *et al.* (2011) Measuring social value orientation. *Judgm. Decis. Mak.* 6, 771–781
65. Barnby, J.M. *et al.* (2022) Knowing me, knowing you: interpersonal similarity improves predictive accuracy and reduces attributions of harmful intent. *Cognition* 225, 105098
66. Bolis, D. *et al.* (2017) Beyond autism: introducing the dialectical misattunement hypothesis and a Bayesian account of intersubjectivity. *Psychopathology* 50, 355–372
67. Costa-Gomes, M. *et al.* (2001) Cognition and behavior in normal-form games: an experimental study. *Econometrica* 69, 1193–1235
68. Camerer, C.F. *et al.* (2004) A cognitive hierarchy model of games. *Q. J. Econ.* 119, 861–898
69. Stahl, D.O. (1993) Evolution of smart n players. *Games Econ. Behav.* 5, 604–617

70. Doshi, P. *et al.* (2012) Modeling human recursive reasoning using empirically informed interactive partially observable Markov decision processes. *IEEE Trans. Syst. Man Cybern Part A Syst. Humans* 42, 1529–1542

71. Doshi, P. *et al.* (2014) Decision-theoretic planning in multi-agent settings with application to behavioral modeling. In *Plan, Activity, and Intent Recognition: Theory and Practice* (Sukthankar, G. *et al.*, eds), pp. 205–224, Elsevier

72. Goodman, N.D. and Frank, M.C. (2016) Pragmatic language interpretation as probabilistic inference. *Trends Cogn. Sci.* 20, 818–829

73. Scontras, G. *et al.* (2021) A practical introduction to the Rational Speech Act modeling framework. *arXiv* 2021 arXiv:2105.09867

74. Xiang, T. *et al.* (2012) Computational phenotyping of two-person interactions reveals differential neural response to depth-of-thought. *PLoS Comput. Biol.* 8, e1002841

75. Ray, D. *et al.* (2008) Bayesian model of behaviour in economic games. *Adv. Neural Inf. Proces. Syst.* 21, 1–8

76. King-Casas, B. *et al.* (2008) The rupture and repair of cooperation in borderline personality disorder. *Science* 321, 806–810

77. Hula, A. *et al.* (2018) A model of risk and mental state shifts during social interaction. *PLoS Comput. Biol.* 14, e1005935

78. Hula, A. *et al.* (2021) Multi-round trust game quantifies inter-individual differences in social exchange from adolescence to adulthood. *Comput. Psychiatry* 5, 102–118

79. Yoshida, W. *et al.* (2008) Game theory of mind. *PLoS Comput. Biol.* 4, e1000254

80. Yoshida, W. *et al.* (2010) Neural mechanisms of belief inference during cooperative games. *J. Neurosci.* 30, 10744–10751

81. Yoshida, W. *et al.* (2010) Cooperation and heterogeneity of the autistic mind. *J. Neurosci.* 30, 8815–8818

82. Devaine, M. *et al.* (2014) The social Bayesian brain: does mentalizing make a difference when we learn? *PLoS Comput. Biol.* 10, e1003992

83. D'Arc, B. *et al.* (2020) Social behavioural adaptation in Autism. *PLoS Comput. Biol.* 16, e1007700

84. Seymour, B. *et al.* (2009) Altruistic learning. *Front. Behav. Neurosci.* 23

85. Devaine, M. *et al.* (2014) Theory of mind: did evolution fool us? *PLoS ONE* 9, e87619

86. Gershman, S.J. and Cikara, M. (2020) Social-structure learning. *Curr. Dir. Psychol. Sci.* 29, 460–466

87. Lau, T. *et al.* (2018) Discovering social groups via latent structure learning. *J. Exp. Psychol. Gen.* 147, 1881

88. Lau, T. (2021) Reframing social categorization as latent structure learning for understanding political behaviour. *Philos. Trans. R. Soc. B* 376

89. Lau, T. *et al.* (2020) Social structure learning in human anterior insula. *eLife* 9, e53162

90. Gershman, S.J. *et al.* (2017) Learning the structure of social influence. *Cogn. Sci.* 41, 545–575

91. Khalvati, K. *et al.* (2019) Modeling other minds: Bayesian inference explains human choices in group decision-making. *Science Advances* 5, eaax8783

92. Park, S.A. *et al.* (2019) Neural computations underlying strategic social decision-making in groups. *Nat. Commun.* 10, 1–12

93. Greenburgh, A. *et al.* (2019) Paranoia and conspiracy: group cohesion increases harmful intent attribution in the Trust Game. *PeerJ* 7, e7403

94. Lincoln, T.M. *et al.* (2018) The relevance of emotion regulation in explaining why social exclusion triggers paranoia in individuals at clinical high risk of psychosis. *Schizophr. Bull.* 44, 757–767

95. Keil, J. *et al.* (2019) Getting less than their fair share: maltreated youth are hyper-cooperative yet vulnerable to exploitation in a public goods game. *Dev. Sci.* 22, e12765

96. Pilkington, P.D. *et al.* (2021) Adverse childhood experiences and early maladaptive schemas in adulthood: a systematic review and meta-analysis. *Clin. Psychol. Psychother.* 28, 569–584

97. Nicol, A. *et al.* (2020) The relationships between early maladaptive schemas and youth mental health: a systematic review. *Cogn. Ther. Res.* 44, 715–751

98. Moutoussis, M. *et al.* (2014) Bayesian inferences about the self (and others): a review. *Conscious. Cogn.* 25, 67–76

99. Alderson-Day, B. *et al.* (2021) Voice-hearing and personification: characterizing social qualities of auditory verbal hallucinations in early psychosis. *Schizophr. Bull.* 47, 228–236

100. Lemonde, A.C. *et al.* (2021) Delusional content at initial presentation to a catchment-based early intervention service for psychosis. *Br. J. Psychiatry* 218, 217–223

101. Paulik, G. (2012) The role of social schema in the experience of auditory hallucinations: a systematic review and a proposal for the inclusion of social schema in a cognitive behavioural model of voice hearing. *Clin. Psychol. Psychother.* 19, 459–472

102. Tsang, A. *et al.* (2021) The relationship between appraisals of voices (auditory verbal hallucinations) and distress in voice-hearers with schizophrenia-spectrum diagnoses: a meta-analytic review. *Schizophr. Res.* 230, 38–47

103. Dellazizzo, L. *et al.* (2022) A systematic review of relational-based therapies for the treatment of auditory hallucinations in patients with psychotic disorders. *Psychol. Med.* 52, 2001–2008

104. Barnby, J.M. *et al.* (2022) The Felt Presence experience: from cognition to the clinic. *PsyArXiv* Published online November 7, 2022. http://dx.doi.org/10.31234/osf.io/qykhd

105. Cañigueral, R. *et al.* (2021) Effects of being watched on eye gaze and facial displays of typical and autistic individuals during conversation. *Autism* 25, 210–226

106. Schlier, B. *et al.* (2019) Autonomic arousal during psychosis spectrum experiences: results from a high resolution ambulatory assessment study over the course of symptom on-and offset. *Schizophr. Res.* 212, 163–170

107. Seow, T.X. *et al.* (2021) How local and global metacognition shape mental health. *Biol. Psychiatry* 90, 436–446

108. Yttri, J.E. *et al.* (2022) Auditory verbal hallucinations in schizophrenia, part II: phenomenological qualities and evolution. *J. Nerv. Ment. Dis.* 10–1097

109. Wang, C.E. *et al.* (2010) Stability of dysfunctional attitudes and early maladaptive schemas: A 9-year follow-up study of clinically depressed subjects. *J. Behav. Ther. Exp. Psychiatry* 41, 389–396

110. Dayan, P. *et al.* (2015) Taming the shrewdness of neural function: methodological challenges in computational psychiatry. *Curr. Opin. Behav. Sci.* 5, 128–132

111. Gershman, S.J. and Uchida, N. (2019) Believing in dopamine. *Nat. Rev. Neurosci.* 20, 703–714

112. Lowet, A.S. *et al.* (2020) Distributional reinforcement learning in the brain. *Trends Neurosci.* 43, 980–997

113. Piray, P. and Daw, N.D. (2021) A model for learning based on the joint estimation of stochasticity and volatility. *Nat. Commun.* 12, 1–16

114. Reddy, V. (2008) *How Infants Know Minds*, Harvard University Press

115. Timmermans, B. *et al.* (2013) Social interaction builds the we-mode. Comment on: Gallotti and Frith (2013) Social Cognition in the We-mode. TiCS 17, 160–165. *Trends Cognit. Sci.* Published online May 21, 2013. www.cell.com/trends/cognitive-sciences/abstract/S1364-6613(13)00041-7#Comments

116. Schilbach, L. (2014) On the relationship of online and offline social cognition. *Front. Hum. Neurosci.* 8, 278

117. Hayward, L.E. *et al.* (2020) How might childhood adversity predict adult psychological distress? Applying the identity disruption model to understanding depression and anxiety disorders. *J. Affect. Disord.* 265, 112–119

118. LoPilato, A.M. *et al.* (2021) Associations between childhood adversity, cognitive schemas and attenuated psychotic symptoms. *Early Interven. Psychiatry* 15, 818–827

119. Koban, L. *et al.* (2021) The self in context: brain systems linking mental and physical health. *Nat. Rev. Neurosci.* 22, 309–322

120. Abu-Akel, A.M. *et al.* (2020) Re-imaging the intentional stance. *Proc. R. Soc. B* 287, 20200464

121. Bahrami, B. *et al.* (2012) Together, slowly but surely: the role of social interaction and feedback on the build-up of benefit in collective decision-making. *J. Exp. Psychol. Hum. Percept. Perform.* 38, 3

122. Raihani, N.J. *et al.* (2013) The effect of 1, 5 and 10 stakes in an online dictator game. *PLoS ONE* 8, e73131

123. Takahashi, H. *et al.* (2014) Different impressions of other agents obtained through social interaction uniquely modulate dorsal and ventral pathway activities in the social human brain. *Cortex* 58, 289–300

124. Abell, F. *et al.* (2000) Do triangles play tricks? Attribution of mental states to animated shapes in normal and abnormal development. *Cogn. Dev.* 15, 1–16

125. Castelli, F. *et al.* (2002) Autism, Asperger syndrome and brain mechanisms for the attribution of mental states to animated shapes. *Brain* 125, 1839–1849

126. Boyer, P. (2003) Religious thought and behaviour as by-products of brain function. *Trends Cogn. Sci.* 7, 119–124

127. Lisøy, R.S. *et al.* (2022) Seeing minds–a signal detection study of agency attribution along the autism-psychosis continuum. *Cognit. Neuropsychiatry* 17, 356–372

128. Okruszek, Ł. *et al.* (2018) Biological motion sensitivity, but not interpersonal predictive coding is impaired in schizophrenia. *J. Abnorm. Psychol.* 127, 305–313

129. Murphy, P. *et al.* (2018) The paranoia as defence model of persecutory delusions: a systematic review and meta-analysis. *Lancet Psychiatry* 5, 913–929

130. Fehr, E. and Schmidt, K.M. (1999) A theory of fairness, competition, and cooperation. *Q. J. Econ.* 114, 817–868