Mental disorders emerge from the gut microbiome gradually transitioning from a symbiotic to a parasitic state: a 4E cognition approach to computational psychiatry

Mayank Srivastava (mayanks23@iiserb.ac.in)[1,2] & Pragati Nayak[1,2] & Mukesh Makwana[2] & Sanjay Chandrasekharan[2]

[1]Department of Data Science and Engineering, Indian Institute of Science Education and Research, Bhopal, India
[2]The Learning Sciences Research Group, Homi Bhabha Centre for Science Education,
Tata Institute of Fundamental Research, Mumbai, India

Abstract

Recent studies suggest that changes in the gut microbiome, as well as social isolation, are closely correlated with mental disorders. We argue that these correlations are intertwined, and the way these interlinkages lead to mental disorders could be productively explored using computational models of agent-environment interactions. To illustrate this approach, we present a reinforcement learning model of social isolation, and discuss how an extension of this model – based on social stress-related changes to gut states – could account for the role of the gut microbiome in the emergence of mental disorders.

Keywords: Mental disorders; 4E cognition; Social isolation; Gut microbiome; Gut-brain axis; Computational psychiatry

The pearl is the oyster's autobiography.
– Federico Fellini

Introduction

4E cognition approaches – which consider cognition as emerging through agents' extended, embedded, enactive, and embodied interactions with the environment – have led to mechanism accounts of higher-order cognition processes such as language (Glenberg & Gallese, 2012) and making behavior (Chandrasekharan et al., 2025). However, such accounts have focused on agent-environment interactions leading to 'normal' cognition. Abnormal behavior – particularly mental disorders – have not been studied much by 4E approaches, though philosophers have explored this area (Núñez de Prado-Gordillo & López-Silva, 2025). Computational psychiatry models based on the 4E approach focus mostly on predictive processing (Friston, 2023). Dominant research threads in computational psychiatry are based on information processing models of the brain, which promote biochemical interventions. These do not align with models of mental disorders as emergent socio-psycho-bio processes.

In contrast to this brain emphasis, recent studies show close correlations between states of the gut microbiome and mental disorders, including schizophrenia (Zheng et al., 2019) , bipolar disorder (Evans et al., 2017) , and anxiety/depression (Peirce & Alviña, 2019). This correlation also exists for alcohol and substance dependencies (Jerlhag, 2019). For instance, the transfer of the microbiome of alcohol-dependent mice to genetically modified mice (with no gut microbiome) leads to the latter group exhibiting alcohol dependence behavior (Xiao et al., 2018). Brain states that arise from social interaction, such as anxiety and stress, can also change the gut microbiome, as these neural states change gastric acid and pH levels (Kessing et al., 2015; Mahl & Brody, 1954). A key stress factor that contributes significantly to mental disorders is social isolation, particularly in refugee populations (Blackmore et al., 2020; Giacco et al., 2018).

As the gut microbiome changes constantly in relation to diet and other agent-environmental interactions (Alcock et al., 2014), computational psychiatry models based on 4E cognition could provide mechanism accounts of gut-brain interlinkages. As a starting point to model mental disorders as interactive processes, we present a computational model that explores the emergence of social isolation through agent-environment interactions. Building on this model's results, we propose *a theoretical account* of how the emergence of social isolation and related social stress could change the gut microbiome, and how this transition could lead to mental disorders. This model – *the key contribution of the paper* – is based on the 'Porcupine's Dilemma' (PD) thought experiment proposed by Schopenhauer, which postulates that in winters porcupines move close to each other to maintain body heat. However, as they move closer, their quills start pricking others, causing pain. This makes the porcupines move apart, which makes them cold, making them come together again, leading to pain. This oscillatory behavior continues until an equilibrium between temperature and pain is reached. By analogy, the thought experiment suggests that people seek out social relationships, but turn back when faced with rejection. After many approach-avoid oscillations, they settle on relationships at a distance, or none. Since this thought experiment is based purely on agent-environment dynamics, with no appeal to psychological states, it provides a useful *initial toy model* to explore the gradual and dynamic emergence of social isolation in biological agents, including humans. We implemented this toy model (extending (Chandrasekharan & Stewart, 2007)), to examine: 1) whether the thought experiment is accurate, and 2) what factors mediate the emergence of social isolation. We discuss the results, and outline how ongoing work extends this toy model, to understand the relationship between social isolation, the gut microbiome, and mental illness.

Our multi-agent simulation study used PD as an analogy for human social behavior. We first implemented the model using a fixed rule (FR)-based architecture. Once this system stabilized, we moved to a reinforcement learning (RL)-based architecture. Both systems had the same changing environment. The agents had the same sensors, physical limits, and

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

ways of interacting. Given our 4E cognition focus, our discussion is primarily based on the RL results. However, we present the study results side-by-side, as this design allows comparison between hand-crafted behavioral heuristics and learned policies under identical environmental conditions.

Model Design

The environment is represented as a 2D toroidal grid, for spatial continuity of agent movements without boundary effects. Ambient temperature is a global environmental temperature (range $\in [-20, 40]$°C). Each grid cell has a temperature value that evolves over time, due to ambient temperature and heat from agents. Agents have a triangular body with a pointed head indicating heading. The body occupies three grid cells (head, mid-body and rear) along the heading direction. Agents move in a continuous space. At each time step, agents feel the nearby temperature, check discomfort caused by other agents, choose an action (using fixed rules or RL) and then update their own state. Each agent maintains a body temperature (internal heat level), pain level, cold tolerance CT (how strongly an agent is affected by ambient temperatures, range: [0,1]), and pain tolerance PT (modulates sensitivity to physical discomfort caused by nearby agents, range: [0,1]). Agents move at a constant speed in the heading direction. The temperature at every grid cell slowly shifts at each time step, towards the ambient temperature. For a grid cell with temperature T_t and ambient temperature T_a, the decay is defined as: $T_{t+1} = T_a + (T_t - T_a) \cdot \lambda$, where $\lambda = 0.8$ controls the rate of thermal dissipation. The heat dissipated by agents thus stay for a short time, but slowly fades away. Heat is distributed continuously (bilinear interpolation), allowing smooth spatial diffusion. Agents have more quills on the back and fewer quills near the head. Quills cause (direction-dependent) pain when the quills of one agent come into contact with the body of another agent and vice-versa. For example, if agent A's head (with few quills) approaches agent B's back (with dense quills), agent A experiences greater pain, while agent B will experience little pain. Pain is induced when the shortest toroidal distance between agents falls below a quill interaction threshold: $d < r_b + r_q$; r_b is body radius, r_q is effective quill reach (depends on the number and size of quills; directionally modulated, using the relative angle θ between the agent's heading and the contact direction). Directional modulation is defined as: $m(\theta) = \frac{1-\cos(\theta)}{2}$. Pain increment is computed as: $\Delta p = (1 - \tau_p) \cdot D \cdot O \cdot S$, τ_p is PT, D is how far another agent comes close to the quills, O is whether the contact is full body or just with quills, and S is (quill number) x (quill size). At each time step, pain is clipped to $[0, 1]$. Both FR and RL models explored how agent behavior changed for low ($\in [0, 0.3]$), medium ($\in (0.3, 0.7)$) and high ($\in [0.7, 1]$) PT.

Fixed Rules (FR)-based System

In the first architecture (Fig 1), agents act according to fixed rules, which allow them to stay within comfortable temperature, avoiding pain. The simulation starts with random movement of agents (pain = 0 and comfortable body heat $\in [35, 39]$). At each time step, the agent checks if the body temperature is within the comfortable range. If yes, agents move randomly, not interacting with other agents. Agents gain heat by: $T_i(t+1) = T_i(t) + \alpha\big(T_{\text{local}} - T_i(t)\big) + \beta\, n_i + \Delta T_i^{\text{contact}}$, where $T_i(t)$ is the body temperature of agent i at time t, T_{local} is the local environment temperature sensed by the agent, n_i is the number of nearby agents, $\Delta T_i^{\text{contact}}$ is direct contact heat transfer with agent j and $\Delta T_i^{\text{contact}} = \pm 0.1\, T_i(t)$, where the sign is negative if $T_i(t) > T_j(t)$ and positive if $T_i(t) < T_j(t)$, α controls heat exchange with the environment as $\alpha = 0.1\,\tau_i$, and β represents heat gained from neighboring agents ($\beta = 0.8\,(1 + \tau_i)$, where τ_i denotes the CT of agent.) If body heat is below the comfortable range, agents will move towards other agents, to gain warmth. Other agents' quills will then cause pain. If pain becomes too high, the agent moves away. The agent then again interacts with other agents. The agent stops when its body heat range is within the comfortable range and it feels minimal pain. Pain is normalized by tolerance to form the pain ratio: $\rho_t = \frac{p_t}{\tau_p}$ (p_t is current pain, τ_p is PT). The probability of the agent's stopping movement is computed by: $P_{\text{stop}} = \frac{1}{1+e^{2(\rho_t - 0.7)}}$ (0.7 represents a critical normalized pain threshold; at $\rho_t = 0.7$, Pstop ≈ 0.5, there is equal likelihood of stopping or continuing movement). This value defines the transition point between tolerance-dominated and avoidance-dominated behavior. Agent will stop moving when other agents are in comfortable body heat range and $P_{\text{stop}} > 0.5$.

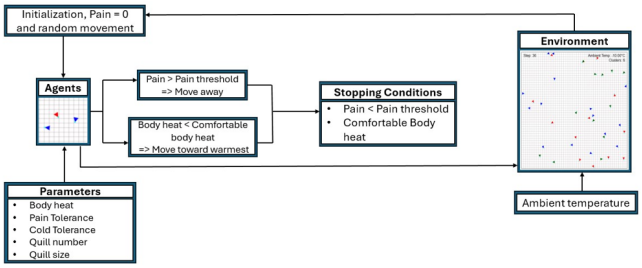

Figure 1: Rule–based agent–environment interaction

Reinforcement Learning (RL)-based System

In the RL architecture (Fig 2), agents learn movement behavior by interacting with the environment, rather than following fixed rules. Agents learn from past experiences through rewards, and adjust their actions accordingly. They do not react only to current pain, but evaluate the trade-off between discomfort and future benefits. RL was based on tabular Q-learning. It provides a simple, model-free, and interpretable framework suited to our discrete state–action space. The ϵ-greedy system provides a balance between exploration and exploitation. Each agent kept track of thermal mismatch (difference between the surrounding temperature and own body temperature) and current pain level. Thermal mismatch was mapped to 25 discrete bins, spanning the full environmental temperature range, (−20 to 40°C). Pain was discretized

into 25 bins (interval [0, 1]). The total state space is (25 × 25 = 625) discrete states. The action set is defined as A = $\{a_0, a_1, a_2, a_3, a_4\}$, where a_0 is left rotation, a_1 is right rotation, a_2 is forward movement, a_3 is backward movement, and a_4 is maintaining the current direction without translational movement. Each agent has a Q-table $Q(s, a)$ (dimensions $25 \times 25 \times 5$). It stores an estimate of the expected cumulative reward when action a in state s. Agents select actions using an ϵ-greedy policy, balancing exploration and exploitation. The Q-value update rule is:

$$Q(s,a) \leftarrow Q(s,a) + \alpha \left[r + \gamma \max_{a'} Q(s',a') - Q(s,a) \right] \quad (1)$$

α is the learning rate, γ the discount factor, r the reward, and s' the next state. The parameter values are in Table 1. The learning rate decays over time until convergence. The exploration rate ϵ decays exponentially for further exploitation. The reward function combines thermal comfort and pain minimization. Pain reward is defined as:

$$R_{\text{pain}} = -k_p \cdot p \cdot (1 - \tau_p) \quad (2)$$

where p is current pain level, τ_p is PT, k_p is pain scaling coefficient (set to 5, to balance the influence of pain and thermal comfort within the reward function).

The temperature reward is calculated based on the deviation from the comfortable body heat range:

$$R_{\text{temp}} = \begin{cases} k_t \left(1 - \dfrac{|T - T_m|}{\Delta T}\right), & T \in [T_{\min}, T_{\max}] \\ -c \cdot |T - T_m| \cdot (1 - \tau_c), & \text{otherwise} \end{cases} \quad (3)$$

where T_m is the midpoint of the comfortable body heat range, ΔT is half of its width, τ_c is CT, K_t is thermal comfort scaling coefficient (5), and c is cold penalty coefficient (2).

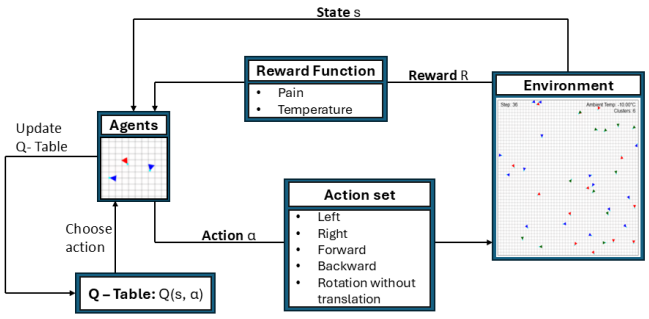

Figure 2: Q-learning–based agent–environment interaction

Total reward is: sum of the pain and thermal components.

$$R = R_{\text{pain}} + R_{\text{temp}} \quad (4)$$

This design allows agents to maintain thermal comfort while avoiding pain. The implementation (developed using HTML, CSS, and JavaScript) is open-source, available at: https://github.com/Mayank-Srivastava-21/Porcupine-Sim. Web interface: https://porcupine-simulation.netlify.app/.

Table 1: Parameters of the reinforcement learning system

Parameter	Value
Learning rate (α)	0.1
Learning rate decay	1×10^{-4}
Discount factor (γ)	0.8
Initial exploration rate (ϵ_0)	0.3
Minimum exploration rate ($\epsilon_{\min}$)	0.01
Exploration decay	0.99995

Study Design

We explored the social behavior of agents (both FR and RL) under different conditions, particularly how agents adjust their distance, form groups, and organize under different temperature and pain constraints. Each experiment was run for 50 iterations and the average of all iterations was taken. All experiments were conducted with six different agent population sizes (9, 18, 27, 36, 45 and 54). The simulation configurations for each experiment is shown in Table 2. Quill size was $\in$ [0.5, 1] and number was $\in$ [8, 18]. Experiment 1 explored how the distance between agents changed across simulation time steps. Agents had random PT and CT. Ambient temperature was set to -10. Experiment 2 explored how agents reacted when the ambient temperature of the environment was changed. Agents had 0.1 CT. PT was uniformly distributed, i.e. equally divided between agents. For example, if we take 9 agents, and the range of the PT was from [0,1], then PT(PT) was divided uniformly as 0.1 PT to any 3 agents, 0.5 PT to any other 3 agents and 1 PT to rest of the agents. The ambient temperature was set as -20, -10, 0, 10, 20, and 30. Experiment 3 explored how the agents behaved under low PT, middle PT and high PT. Each agent had 0.1 CT. PT was uniformly distributed. The ambient temperature was set to -10. Experiment 4 explored how the behavior of agents changed with different quill numbers. Agents had 0.1 CT and uniformly distributed PT. The quill number of agents was varied from [8,18] and [18,36]. Experiment 5 explored the overlapping behavior of agents. Agents had varying CT and PT. Particularly, we examined at what values the benefit of warmth outweighs the cost of pain, and agents thus came too close to each other.

Table 2: Simulation configuration

Parameter	Value
No. of iterations per run	50
Simulation time steps	200000
Grid size	50 X 50 cells
Cell size	10 pixels

Results

All line plots show average data, with shaded regions representing 95% confidence intervals As shown in Figure 3, in experiment 1, during the early and middle stages of the simulation, agents move closer to stay warm, but move apart

to avoid the pain from getting too close to other agents. This back-and-forth movement forms an oscillating pattern, as suggested by the thought experiment. The oscillations became smaller as the simulation advanced, because agents adjusted their positions to gain warmth with minimal pain. Over time, the average distance settled at a low, steady value, as agents formed clusters. Figure 3 shows results for 9 agents. A similar oscillatory pattern graph was observed with 18, 27, 36, 45, and 54 agents.

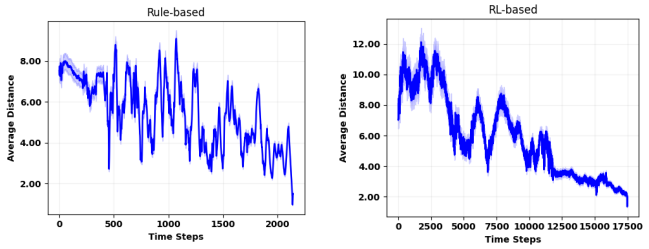

(a) **(b)**

Figure 3: The plots shows the change in average distance between agents over time

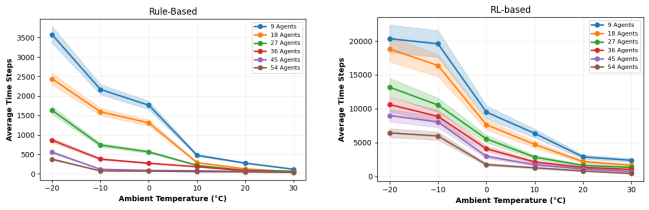

(a) **(b)**

Figure 4: The plots shows how the effect of ambient temperature on clustering time for rule-based and Q-learning agents

In experiment 2, as shown in Figure 4, the average time steps decreased for FR and RL systems when the ambient temperature increased. In colder environments, agents took longer to cluster, as they experienced a stronger thermal stress, and balanced the need for warmth with pain avoidance. As the environment became warmer, agents reached thermal comfort more easily, and settled into clusters faster. FR agents consistently clustered much faster than RL agents, across all temperatures, as agents could immediately respond to temperature cues. RL agents required time to explore actions and learn through rewards before converging.

In experiment 3, the behavior of agents changed as the number of agents in the environment increased. Figure 5(a) shows the average time taken by agents to form stable clusters as the number of agents increased. Clustering time decreased for both systems when more agents were present, because higher number of agents increased the likelihood of interactions, leading to easier cluster formation. FR agents clustered much faster than RL agents, across all population sizes. Figure 5(b) shows the average number of clusters formed as the number of agents increased. The number of clusters increased with population size for both FR and RL systems. RL agents formed more clusters than FR agents. RL agents also tended to maintain more distributed group structures, whereas FR agents aggregated into fewer, denser clusters.

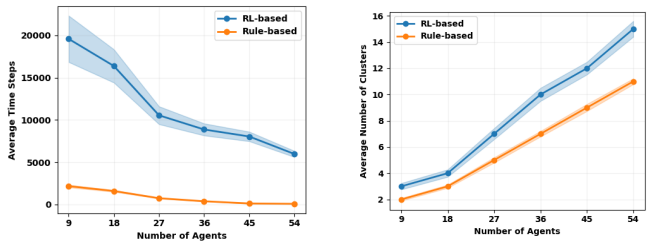

(a) **(b)**

Figure 5: The plots show how increasing the number of agents affects the (a) average time required to form stable clusters (b) the average number of clusters formed

Figure 6, shows how Mean Nearest-Neighbor distance (the average distance between an agent and all agents within the cluster) changed as the number of agents increased, for different levels of PT. The average value of nearest-neighbor distance was taken for different levels of PT. *Agents with low PT maintained the largest distances from their neighbors within their cluster, for both FR and RL systems. These agents avoided close contact.* Agents with high PT cluster tightly, resulting in smallest nearest-neighbor distances. Agents with medium PT showed intermediate spacing. The curves for low, medium, and high PT moved closer together as the number of agents increased.

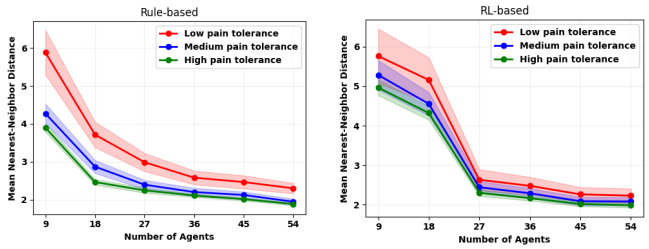

(a) **(b)**

Figure 6: The plot shows the mean Nearest Neighbor Distance when Number of agents increases

Figure 7, shows how the average cluster membership ratio (number of time steps the agent is in a cluster divided by the total number of time steps) changes against the number of agents, for different levels of PT. For both FR and RL systems, the cluster membership ratio increased as the number of agents grew. Agents with low PT showed the lowest cluster membership when the population was small, as these agents avoid close contact and thus remain more isolated. Agents with medium and high PT form clusters more readily, leading to higher membership ratios even at low population sizes. At high densities, most agents stayed clustered almost all the time, because the lower ambient temperature led them to seek

warmth, and the crowded environment forced them to stay together. RL led to more efficient and robust cluster formation, as learning agents were able to adapt their movement strategies in a way that allowed them to remain in clusters despite pain, whereas FR agents could not do this well because they relied strictly on fixed avoidance rules.

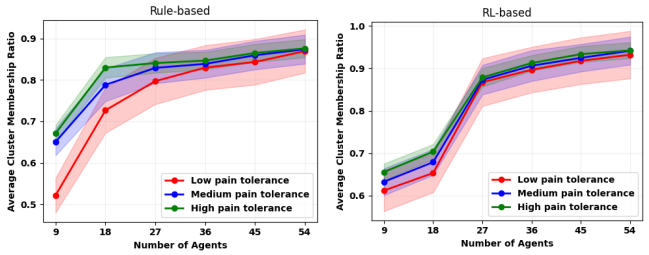

(a) **(b)**

Figure 7: The plot shows the average cluster membership ratio when number of agents increases for different PT levels

Experiment 5 (Figure 8(a)) showed that when quill numbers were higher (quills cover the whole body), agents took more time to cluster, compared to fewer quills. More quills generated stronger pain in other agents, which discouraged close contact, and slowed down aggregation. Agents thus needed more time steps to arrive at the optimal distance. This effect was strong for RL agents, who took longer times to cluster than FR agents. Figure 8(b), shows the average number of clusters formed. As the number of agents increased, the number of clusters increased, for all cases. Agents with more quills distributed themselves across a larger number of smaller clusters, rather than forming a few dense groups.

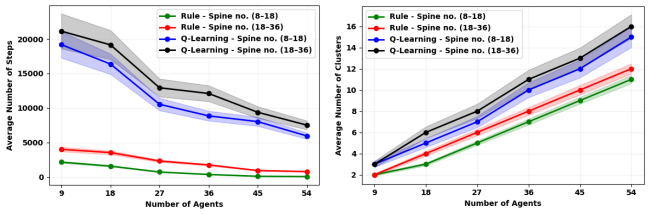

(a) **(b)**

Figure 8: The plot shows the effect of quill length on clustering dynamics for rule-based and Q-learning agents

Interestingly, experiment 5 showed that some agents overlap when their PT is high ($p \gtrsim 0.7$) and CT is low ($c \lesssim 0.2$). High PT agents have more ability to bear pain. Low CT agents cannot bear low temperatures. When these agents interact with other agents, they opt to suffer the pain, moving too close to other agents, resulting in overlapping.

An RL model with a shared Q-table – analogous to social learning – showed similar results across all conditions.

From social isolation to parasitic microbiomes

The main results from running the PD toy model – as an analogy for social isolation in humans – are the following:

1. Social isolation is a dynamic process that emerges over time, after many oscillations in terms of social distance, as suggested by PD. This process is mediated by a number of factors, such as the ability to inflict pain (quill numbers), the ability to suffer pain (PT), and the need for society (CT).

2. Social isolation depends significantly on the ability of agents to tolerate pain. Agents with low PT are more likely to be socially isolated. In the human context, this indicates that social isolation emerges over time in individuals with a low ability to handle rejection.

3. Social isolation is higher in populations with many agents who can inflict pain.

4. Social isolation is lower in dense populations.

To develop our model further, we extend the results from this minimal model to the human case. We acknowledge that this extension not fully warranted. However, we are encouraged by the fit between our results and empirical findings. For instance, social isolation is found to be higher among refugees – compared to migrants – in rich countries (Blackmore et al., 2020). Extending our results, a possible explanation could be: refugees typically are *forced to* immigrate – from dense populations, where they are less isolated – to less-populated rich countries, where social interaction possibilities are much lower. When interactions do occur, there is a high rejection rate (as interactions levels are low, compared to their home countries) – a painful experience for which they are not prepared. Migrants, on the other hand, are *motivated to* move, and many find employment in the countries they move to. They are thus more prepared for painful rejections, and they also face lower rejection rates, due to their employment status.

How can this toy-model of social isolation be extended to account for the emergence of mental disorders? To explore this, we are currently implementing a model of alcohol addiction, where dependence on alcohol gradually emerges through two factors – changes in different gut microbe populations, and RL. The microbial populations change slowly in response to intake of alcohol, through a crowding-out process. In a group that socially drinks alcohol every evening, microbes that can survive in alcohol environments crowd-out others, leading to a gut biome with lower diversity. Over time, the alcohol-friendly microbes seek to survive by producing chemicals that generate cravings for alcohol in the evening (Alcock et al., 2014). These biochemicals orient the brain towards choices that include alcohol intake (Zhang et al., 2025). Within this population, a fraction become dependent on alcohol (via RL), developing behavior such as drinking alone. This social transition – considered a starting point of alcohol dependence – is partly driven by epigenetic factors, and partly the 'kick level' generated by alcohol in each agent. Agents who get drunk with limited intake do not move to such risky social behavior, while agents who require more intake to get drunk move to socially-isolated drinking. This is because more alcohol intake leads to higher crowding-out by microbes that can survive alcohol, leading to very low microbiome diversity (a marker of alcohol dependence) and thus higher levels of craving.

In this model, changes in the microbiome lead to social isolation, which accentuates alcohol dependence. Mental disorders (without substance dependence) could follow a similar trajectory. In such cases, the change in the gut state would be driven by gastric acid and pH level changes, in response to stress generated by social situations. *One* possible pathway for this gut-state change is via histamine (Ito, 2000). For instance: [social stress *promotes* heightened alertness, *which promotes* higher histamine production in the brain, *which promotes* rising production of histamine in the gut, *which promotes* higher stomach acid secretion]. This pathway is possible because histamine mediates *both* alertness (in the brain), and acid production (in the gut). A related pathway is via norepinephrine, a critical gut-brain messenger. Such interlinked processes can change acid and pH levels during social stress, leading to microbiome changes over time. Depending on the composition of the biome, this transition could eventually lead to the proactive generation of craving biochemicals by the microbes – to maintain the right acid and pH levels that allow them to survive (Trevelline & Kohl, 2022). Modulated by epigenetic and psycho-social factors, this change could reorient the brain towards states that accentuate social stress – such as depressive and anxious states. These states generate behavioral loops that allow the less-diverse microbiome to persist and flourish, which reorients the system even further.

Such a pattern is similar to the way parasites control host behavior (Yong, 2016). The above theoretical account thus suggests that stressful social interactions can – over time – transition the gut microbiome to an *interlocking configuration with the brain*, where the biome *turns parasitic* and alters behavior. This symbiotic-to-parasitic transition is similar to the immune system transitioning to a state where it attacks the body (autoimmunity). The parasitic transition of the biome is driven by the *interlocking of socio-psycho-bio processes*, which changes behavior and leads to mental disorders.

An important implication of this socio-psycho-bio account is that addictions and mental disorders have a gut infection aspect, which can persist even after neural-level interventions. Such persistent infections could explain why purely neurochemical approaches do not fully address mental disorders, particularly why such behaviors tend to return when medication is stopped. The 4E account we have outlined suggests that addressing mental disorders require moving beyond neurochemistry – to develop multi-pronged interventions, specifically ones that seek to *unravel the emergent interlocking of socio-psycho-bio processes* that lead to parasitic microbiomes.

References

Alcock, J., Maley, C. C., & Aktipis, A. A. (2014). Is eating behavior manipulated by the gastrointestinal microbiota? evolutionary pressures and potential mechanisms. *BioEssays*, *36*(10), 940–949.

Blackmore, R., Boyle, J. A., Fazel, M., Ranasinha, S., Gray, K. M., & Gibson-Helm, M. (2020). The prevalence of mental illness in refugees and asylum seekers: A systematic review and meta-analysis. *PLoS Medicine*, *17*(9).

Chandrasekharan, S., Sinha, R., & Date, G. (2025). Making cognition: A mechanism account of the way humans develop the ability to build and alter their environments. *Synthese*, *206*(1), 7.

Chandrasekharan, S., & Stewart, T. C. (2007). The origin of epistemic structures and proto-representations. *Adaptive Behavior*, *15*(3), 329–353.

Evans, S. J., Bassis, C. M., Hein, R., Assari, S., Flowers, S. A., & McInnis, M. G. (2017). The gut microbiome composition associates with bipolar disorder and illness severity. *Journal of Psychiatric Research*, *87*, 23–29.

Friston, K. (2023). Computational psychiatry: From synapses to sentience. *Molecular Psychiatry*, *28*(1), 256–268.

Giacco, D., Laxhman, N., & Priebe, S. (2018). Prevalence of and risk factors for mental disorders in refugees. *Seminars in Cell and Developmental Biology*, *77*, 144–152.

Glenberg, A. M., & Gallese, V. (2012). Action-based language: A theory of language acquisition, comprehension, and production. *Cortex*, *48*(7), 905–922.

Ito, C. (2000). The role of brain histamine in acute and chronic stresses. *Biomed. and Pharmacotherapy*, *54*(5), 263–267.

Jerlhag, E. (2019). Gut–brain axis and addictive disorders: A review with focus on alcohol and drugs of abuse. *Pharmacology and Therapeutics*, *196*, 1–14.

Kessing, B. F., Bredenoord, A. J., Saleh, C. M., & Smout, A. J. (2015). Effects of anxiety and depression in patients with gastroesophageal reflux disease. *Clinical Gastroenterology and Hepatology*, *13*(6), 1089–1095.

Mahl, G. F., & Brody, E. B. (1954). Chronic anxiety symptomatology, experimental stress, and HCl secretion. *AMA Archives of Neurology and Psychiatry*, *71*(3), 314–325.

Núñez de Prado-Gordillo, M., & López-Silva, P. (2025). Making sense of the 4E cognition turn in mental health research. *Philosophy, Psychiatry, and Psychology*, *32*(2), 131–150.

Peirce, J. M., & Alviña, K. (2019). The role of inflammation and the gut microbiome in depression and anxiety. *Journal of Neuroscience Research*, *97*(10), 1223–1241.

Trevelline, B. K., & Kohl, K. D. (2022). The gut microbiome influences host diet selection behavior. *Proceedings of the National Academy of Sciences*, *119*(17), e2117537119.

Xiao, H.-W., Ge, C., Feng, G.-X., Li, Y., Luo, D., & Fan, S.-J. (2018). Gut microbiota modulates alcohol withdrawal-induced anxiety in mice. *Toxicology Letters*, *287*, 23–30.

Yong, E. (2016). *I contain multitudes: The microbes within us and a grander view of life*. Random House.

Zhang, T., Wang, W., Li, J., Ye, X., Wang, Z., Cui, S., Shen, S., Liang, X., Chen, Y. Q., & Zhu, S. (2025). Free fatty acid receptor 4 modulates dietary sugar preference via the gut microbiota. *Nature Microbiology*, *10*(2), 348–361.

Zheng, P., Zeng, B., Liu, M., et al. (2019). The gut microbiome from patients with schizophrenia modulates the glutamate–glutamine–gaba cycle and schizophrenia-relevant behaviors in mice. *Science Advances*, *5*(2).