Modeling Strategic Decision-Making on Networks with Context-Aware Agents

by

Anna Hirschmann

B.S., University of Colorado Boulder 2023

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Master of Science

Department of Applied Mathematics

2025

Committee Members:

Nancy Rodriguez, Chair

Zachary Kilpatrick

Hongjing Lu

Hirschmann, Anna (M.S., Applied Mathematics)

Modeling Strategic Decision-Making on Networks

with Context-Aware Agents

Thesis directed by Prof. Nancy Rodriguez

Mathematical models allow for deep analysis of observed phenomena by providing a framework to simulate reasonable outcomes when designing appropriately scaled experiments is not feasible. In our work, we develop several models to study the decision-making processes and overall success of individuals on a structured network when given a task and asked to come to consensus. Based on experiments, we investigate four cases: two different network structures (homogeneously mixed and spatially embedded) and two tasks of varying complexity. Importantly, we focus on effectively modeling the use of both contextual information (i.e., learning gained by interacting with others directly about a topic) and background information (i.e., personal biases about a topic brought into a learning situation) to make selections. A robust model for these experiments can give us insight into the ways that people combine social information and past experience to make decisions in groups. Implications of modeling these scenarios are a deeper understanding of how information is exchanged in online spaces and how regulations may be targeted directly at harmful discourse.

Dedication

To my parents, who always believed in me and supported me, even when I was difficult. May this paper signify my entrance into adulthood; you did it!

To my cat, who although she has no interest whatsoever in the content of this paper, kept me going in the final stretch of my degree.

Acknowledgements

Thank you to the UCLA team: Hunter Priniski, Bryce Linford, and Hongjing Lu for allowing me to work with you all on modeling your data! Thanks to Jeff Brantingham for always being responsive to new ideas and for your support. Thank you to Nancy and Zack for all your help and your patience through this process, I know I wasn't easy.

This research was supported in part by AFOSR MURI grant #FA9550-22-1-0380.

Contents

Chapter

Tables

Table

Figures

Figure

Chapter 1

Introduction

1.1 Motivation

The complexity of human interaction has grown immensely since the advent of the internet: direct contact no longer requires physical proximity. With the growing popularity of social media and other online communication platforms, complex interactions can take place instantaneously, unencumbered by physical distance limitations. The rapid spread of information can be wielded as a tool for organizing positive change, extending movements that would otherwise require geographical proximity [5], but can also be used to spread misinformation, conspiracy theories, and hate speech [12]. Modern social media allows users to re-share content, adding tags, hashtags, or new information, which can alter the meaning of the original post. It is difficult to study information spread in the age of the internet, because interactions are difficult to track [16] and the speed of online and offline information spread can lead to rapid, untraceable, and iterative reproduction of narratives, embedded with personal points of view [2].

In a time when polarization is rampant and online content moderation standards are changing rapidly [17] we want to understand how misinformation spreads and how people's biases can impact what they share and amplify on the internet, as well as how these biases change when new information is gained. With a deeper understanding of decision-making processes that lead to consensus, we can create better standards for the content that is allowed to be shared on various platforms, which has the potential to have salient offline impacts. A notable example of this effect is the protests that took place after George Floyd was murdered in Minnesota in spring 2020 [13].

Importantly, we want to understand and replicate decision-making processes and- ultimately- consensus formation on a complex network of interacting individuals with varying exposure to relevant narrative information combined with their personal biases. Implications of this understanding include a more concrete conceptualization of the transmission of information, including how network structure plays a role in consensus, which may give insight into how opinions form on modern communication platforms. Likewise, we hope to gain insight into the interplay between individual biases (background information) and the group consciousness gained from interaction (contextual information) when humans make choices in a given context.

To investigate these behaviors, our group designed several experiments that involved gathering information about participant behavior relative to a controlled narrative. The experiments had a three-phase structure whereby participants read and summarize an informative passage about an event, partake in a networked interaction with others in real time, and finally create new summaries for the passage. Experimental interaction tasks varied in their relevance to the narrative that all participants had read, which allowed for analysis of human behavior as it depends on environmental context. Importantly, we were able to compare the networked behavior of individuals in these different conditions and integrate both individual background (pre-interaction biases) and contextual (interaction) knowledge into a decision-making model for these exchanges, which is robust to different content conditions. The development of this model is the primary focus of this thesis. We begin with a brief discussion of the experimental procedure and results that motivated the model.

1.2 Experiment Design

The formal experiments that inspired the models expounded here were designed by a team of psychologists at the University of California, Los Angeles. To develop the experiments, we started with the work of Centola and Baronchelli (2015) [3], which investigated the emergence of conventions on networks given a cognitively simple task (naming an unfamiliar face). This inspired the control condition for our experiments, which we complemented with a more difficult task to discern the role of background knowledge and biases in the decision-making process. In this section, we summarize the procedures that are also laid out in [11].

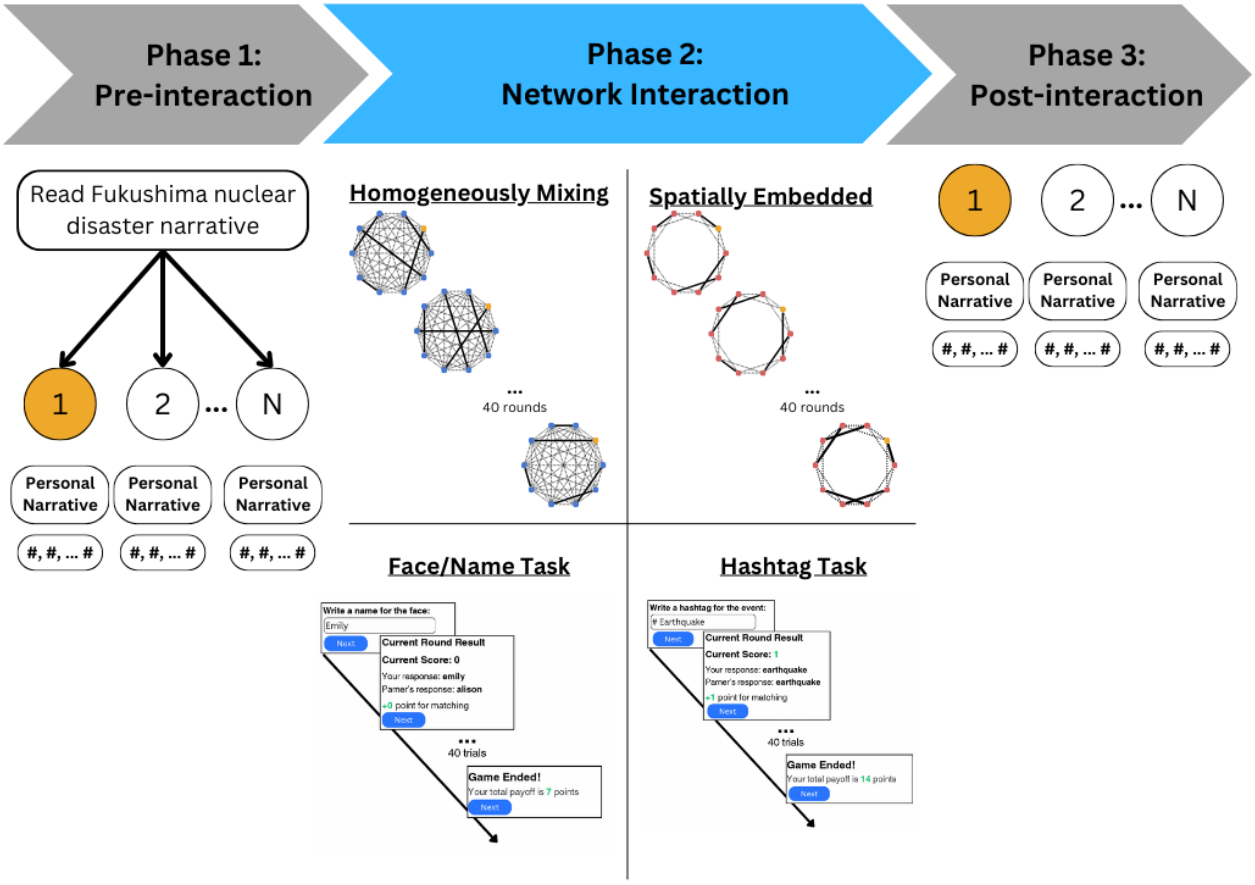

Figure 1.1: In **Phase 1**, all participants read about the Fukushima nuclear disaster and generated a personal narrative and 10 hashtags. This was followed by **Phase 2** interaction where participants were assigned to a particular task and network, and exchanged information in real time. Finally, in **Phase 3**, participants were asked to create new hashtags and narratives as in Phase 1.

Each experiment involved three phases: Phase 1 requires an independent report from all subjects; Phase 2 allows repeated updating of reports after a sequence of interactions between random pairs of subjects; Phase 3 repeats Phase 1. A diagram of the experimental procedure is illustrated in Figure 1.1.

In more detail: **Phase 1** (the pre-interaction phase) had participants read a short narrative on the Fukushima nuclear disaster (Appendix A) and were asked to create a "tweet" (140 words or less) and 10 hashtags to describe what they had read [11]. The network interaction phase, **Phase 2**, assigned each participant to one of six network structure conditions and presented them with a task. Let N denote the

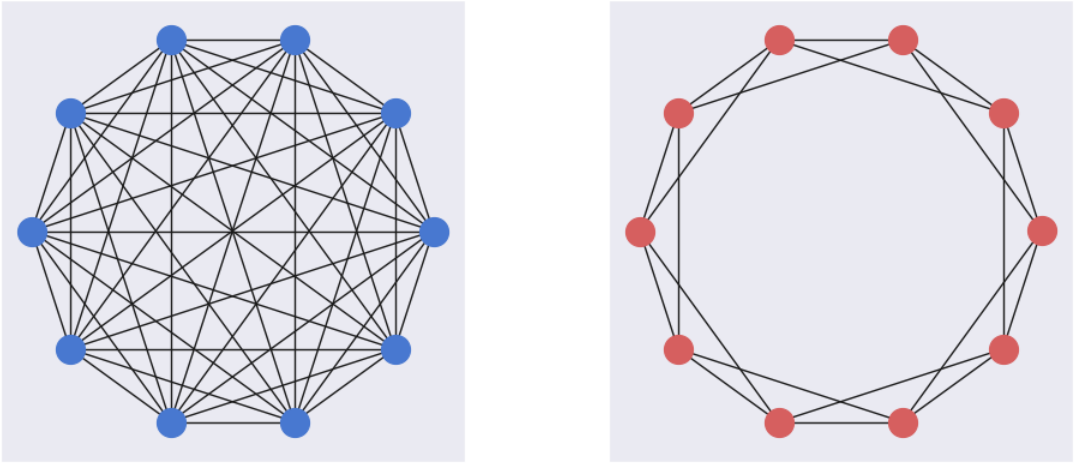

(a) Homogeneously Mixed Network

(b) Spatially Embedded Network

Figure 1.2: The homogeneously mixed network is a complete graph, so each node has degree $N-1$ and the diameter of the network is one. In the spatially embedded network, each node is linked to its four nearest neighbors regardless of N, so the diameter increases with N.

size of the network. The experimental conditions differed in the size of the network ($N = 20, 50$, or 100), and the structure of the network (homogeneous mixing or spatially embedded). The two network structures are shown in Figure 1.2, and the different tasks will be explained in detail below. In **Phase 3**, the post-interaction phase, participants were asked to recreate what they had done in Phase 1, generating 10 hashtags and a "tweet" for the nuclear disaster narrative.

This experimental design opened multiple avenues for investigation. For example, examining how Phase 2 interactions influence changes in individual outputs between Phases 1 and 3 and gaining a deeper understanding of the dynamics of networked interactions. The different features of the network structures lead to differences in information spread, which inform our understanding of online information exchange.

In all experiments, the Phase 2 interaction consisted of 40 trials where information was exchanged between randomized pairs of individuals on the network. In each experiment, the goal presented to par-

ticipants was to come to consensus with the rest of the network in their response to the given task, based on their one-on-one interactions with individuals they were connected to. In each trial, an individual was randomly paired with exactly one other person to whom they were connected via an edge on the network. The members of each pair proceeded to exchange responses and, if the responses matched, each individual scored a point. At the end of the experiment, there was a small financial payout for individuals based on the total number of points scored. The participants did not know the structure of the network or any additional information about their partners (e.g. if they had interacted previously) at any time during the experiment.

To understand how the narrative information from Phase 1 impacted decision-making, participants were assigned to one of two tasks with varying complexity. The simpler experiment asked participants "Write a name for the face" when presented with the image of an unfamiliar female [3]. This experiment (hereafter referred to as the "face/name task") was chosen for its simplicity and the fact that it is not related to the Phase 1 narrative. On the other hand, the more complex "hashtag task" asked participants to "Write a hashtag for the event," where the event in question is the Fukushima nuclear disaster from the Phase 1 narrative.

This hashtag task is directly related to what participants read in Phase 1 of the experiment, making the task more complex because it involves synthesizing information to provide a valuable answer. The Phase 3 responses were studied with careful attention to the difference in gained knowledge between participants in the face/name task and those in the hashtag task, with particular focus on the impacts of information exchange on the Phase 3 responses of hashtag task participants.

The disaster narrative was designed using causal language, which gave a semantic framework for studying how causal understanding and narrative construction changed between Phases 1 and 3 for the different groups according to their Phase 2 interactions [11]. As a supplementary investigation, our model zooms in on the Phase 2 interactions and the different dynamics between groups of individuals based on task and network structure.

1.3 Thesis Overview

Having outlined the experimental procedures in Chapter 2, we will provide an analysis of the results of the experimental runs, including an overview of group-level dynamics and individual decision behaviors. The dynamics of observed human interactions informed the development of a pair of probabilistic models. The first prioritizes network-level behaviors, while the second provides a deeper look into individual dynamics. We introduce and assess both models in Chapter 4 after exploring some existing models in Chapter 3. Finally, we suggest directions for future work in the context of content-influenced human interactions, including the extension of classical voter models to fit this context [14].

Chapter 2

Experimental Results & Analysis

A total of 26 experiments were conducted with 1,040 participants (1,002 of whom were active during the interaction phase). Table 2.1 shows the breakdown of these experiments by size, structure, and task. The experiments allowed us to identify several differences in the ways individuals participated in the networked interaction as a specific result of the task and the network structure they were assigned. One notable difference between tasks was the shape of the initial response distribution. All first responses for all trials were accumulated by task. These empirically estimated prior distributions indicated a higher level of "conditioning" for those participating in the hashtag task. This is a direct result of the specific causal language encoded in the disaster narrative read by all participants, which has immediate relevance to the hashtag task. Since all participants entered this task with fresh background information, there was less variability in the first individual responses compared to the control task. Visualizations of the top 25 trial-one responses

Table 2.1: Number of experimental runs for each network size and structure, according to the assigned task. The total number of participants assigned to each task is also indicated.

Task	**Network Structure**	**20 participants**	**50 participants**	**100 participants**	**Total Experiments**	**Total Participants**
Face/Name	Homogeneously Mixing	3	3	0	6	**420**
	Spatially Embedded	3	3	0	6	
Hashtag	Homogeneously Mixing	3	3	1	7	**620***
	Spatially Embedded	3	3	1	7	

*Only 582 of the participants assigned to the hashtag task took part in the network interaction portion of the experiment.

over all network experiments are illustrated in Figure 2.1 by task. One will notice that the most popular first response in the hashtag task is "nucleardisaster", which was said by over 70 different participants across all networks. This is over 2.5 times as many people who said "Emily" as their first response during the face/name task (this was the most popular first response in this task, said by 28 people). Note that even taking into account the larger number of participants in the hashtag experiments (see Table 2.1), this discrepancy in frequency of the most popular initial response is not entirely accounted for.

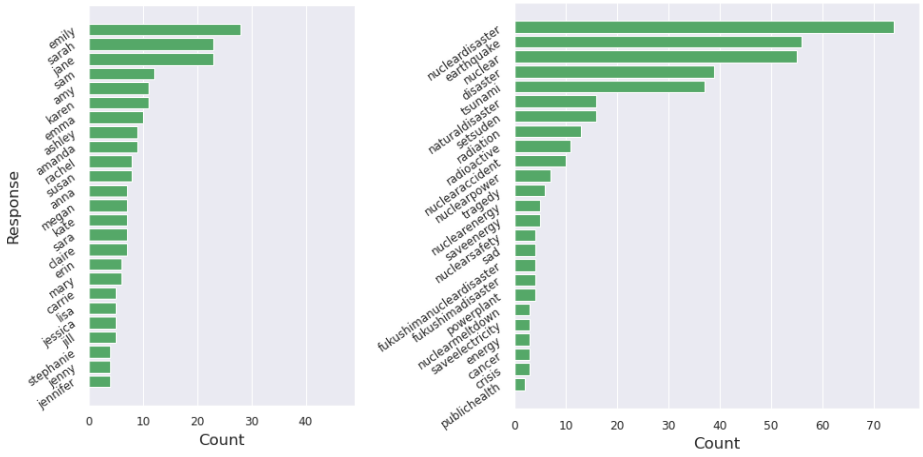

(a) Face/Name Task Prior (b) Hashtag Task Prior

Figure 2.1: The top 25 initial responses by task. This includes all responses given on trial one of the task, independent of network structure or size. Note that priors are both long-tailed, but only the top 25 responses are included. (a) Illustrates the Face/Name task prior. (b) Illustrates the Hashtag task prior.

2.1 Experimental Results: Network-Level

To quantify the coordination of responses over time, we used the Shannon entropy:

$$S_t = -\sum_{x \in R_t} p(x) \log_2(p(x)), \tag{2.1}$$

where t is the timestep, R_t is the set of responses on the network at time t, and $p(x)$ is the proportion of individuals responding x. For each experimental run, **exp**, we have a time series of entropy values,

$\mathbf{S_{exp}} = [S_1, S_2, \ldots, S_{40}]$. Since S_t quantifies the disorder (or variability) in responses at time t, lower values denote more coordination or order. Note that larger networks inherently have larger entropy, so it was necessary to scale S_t by the maximum possible entropy for the network size:

$$S_{max,N} = -\log_2\left(\frac{1}{N}\right). \tag{2.2}$$

This normalizes the entropy to have range $[0, 1]$. In Figure 2.2, we compare the average normalized entropies over time for the four different combinations of network structure and assigned task.

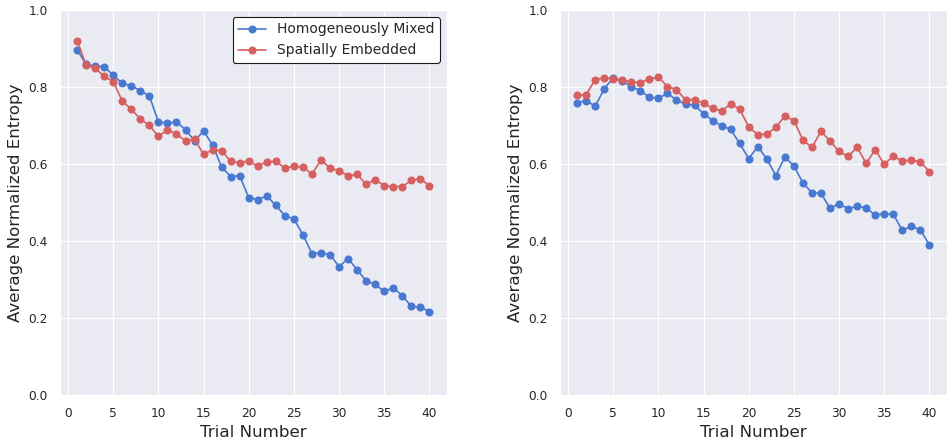

(a) Face/Name Task

(b) Hashtag Task

Figure 2.2: Average experimental entropy as a function of trial and normalized by network size. (a) Average experimental entropy of the face/name task. (b) Average experimental entropy of the hashtag task.

Note that networks with the hashtag task start with lower entropy relative to the face/name task. This reaffirms what we have noticed about the shape of the prior distributions illustrated in Figure 2.1, that more people shared initial responses in the hashtag task. Interestingly, this is followed by a period of increased entropy, which indicates that participants are exploring the space of possible responses [11]. Note also that in homogeneously mixed networks (regardless of task), the entropy decays more on average over the 40 trials than in spatially embedded networks. This is a consequence of the structure: while local coordination occurs in spatial networks, information cannot travel across the network quickly because the diameter is large (and

increases with N), so global coordination is less likely to happen than in a homogeneously mixed network with diameter one [11]. The entropy in spatial networks tends to flatten as participants agree with their neighbors, but have no exposure to responses in other parts of the network.

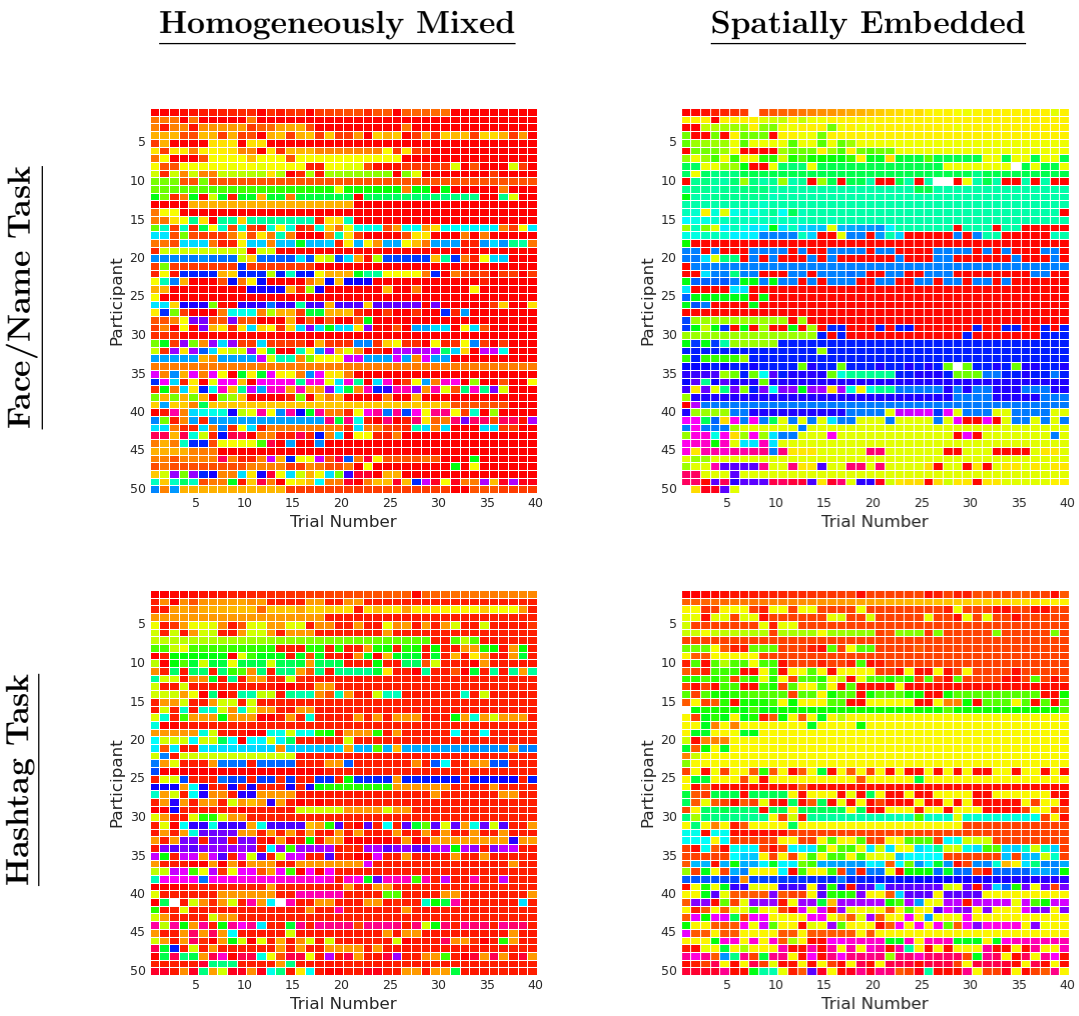

Figure 2.3: One of the $N = 50$ runs for each network structure and task combination. Rows are participants, and columns are trials. The different colors represent varying responses on the network.

Local versus global coordination is best visualized by plotting the responses of individuals over time, which illustrates the distribution of responses over the geometry of the network. See Figure 2.3, which illustrates the evolution of responses for the four experimental setups. In the figures, each column represents a given trial, and the rows represent a participant. A different color is assigned to a different name or

hashtag, depending on the task.

In the right column of Figure 2.3 (homogeneously mixed networks), we see dominant colors spread over the network, stabilizing at several different localizations around the same time. In contrast, the left column examples display a banded structure, illustrating the emergent local dominance of different responses. It is important to note that fully connected (or even just dense) networks are a better approximation of the types of online interaction we see today: digital infrastructure enables the rapid spread quickly and without spatial limitations [5].

2.2 Experimental Results: Individual-Level

While network-level statistics, such as entropy and overall response behavior, are informative of population response behavior at a high level, we also investigated how individuals behaved on each network. In particular, we know from the literature that an individual's decision is likely influenced by a combination of prior background and the contextual information gained through the experiment [6, 8, 18]. This motivated us to develop a "decision type" framework, which identifies an individual's behavior on a given trial as lying in one of four categories based on the varied levels of reliance on background and contextual knowledge: brand new (BN), earlier context (EC), repeat self (RS), and repeat partner (RP). The first category, BN, represents a complete reliance on background information – previous knowledge. A BN response is characterized by the fact that the individual has never said or heard that response during this phase of the experiment. Note that by definition, each subject produces a BN response in trial one. The other three strategies represent reliance on contextual information gained by interacting with others on the network. Denote the current trial by T, RS and RP represent an individual repeating either their response or their partner's response, respectively, from trial $T-1$. In contrast, EC represents a subject repeating something they either heard or said in some trial before $T-1$. It is important to note that the *responses* of the participants are a proxy for determining their decision types. Thus, in the case that an individual scored a point on trial $T-1$, the RS and RP strategies equivalently explain the response provided on trial T. In this case, we allocated the individuals randomly to one of the two decision types. The other decision types are mutually exclusive, which allows for a complex analysis of how individual and population-wide decision-type strategies impact

the group's ability to coordinate responses over time. In Figure 2.4, we observe how the responses of the individuals in the experiments translated into this categorization of decision types. Each decision type is color-coded: BN in red, EC in green, RS in yellow, and RP in blue. Note that there are cases where certain individuals do not respond; these instances are marked in gray.

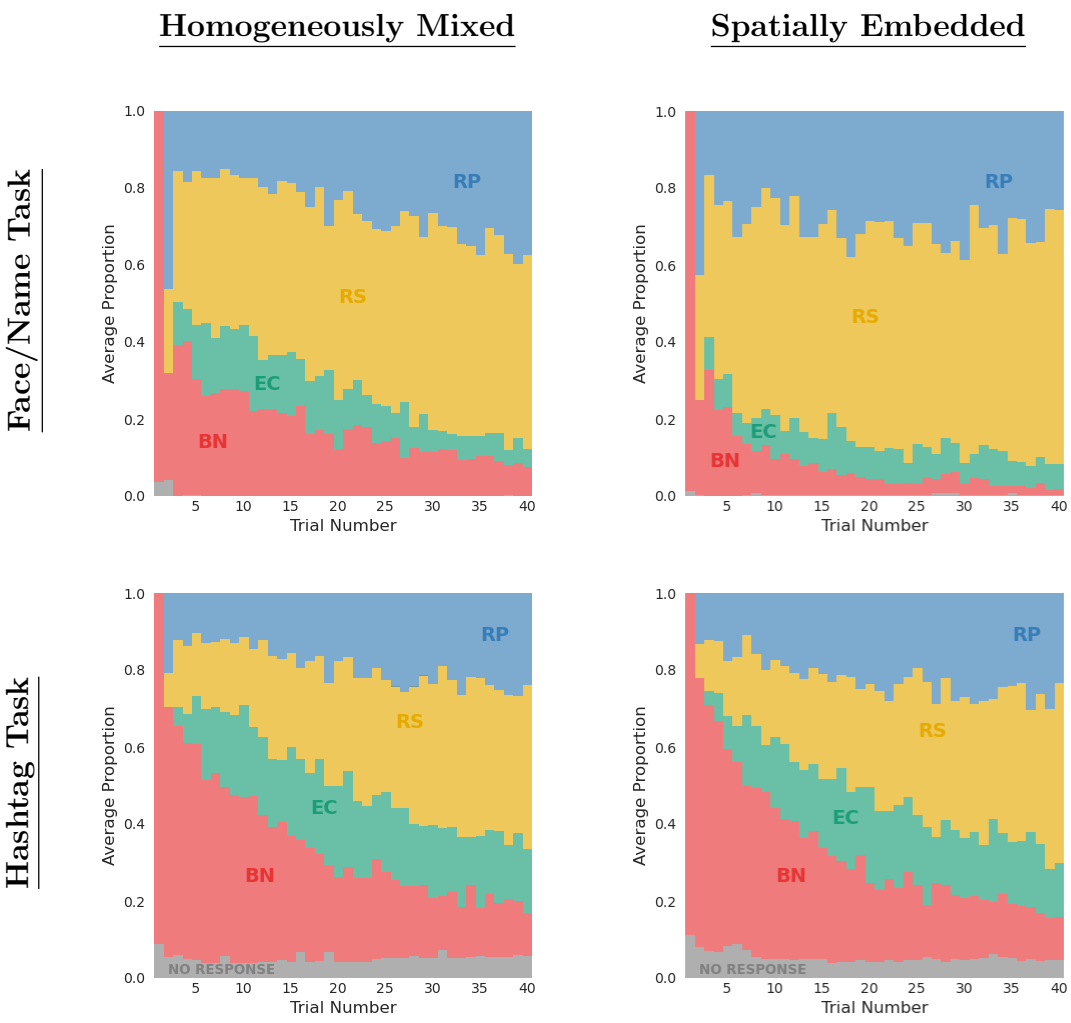

Figure 2.4: For each of the structure and task combinations, a stacked bar chart shows the average proportion of the networks using each decision-type at each trial number.

The most important difference between the tasks is the trend in the use of the BN strategy. Figure 2.4 illustrates that participants in the hashtag task rely more on the BN strategy and for longer than participants in the face/name task. This trend is likely due to an exploratory period for participants given the hashtag

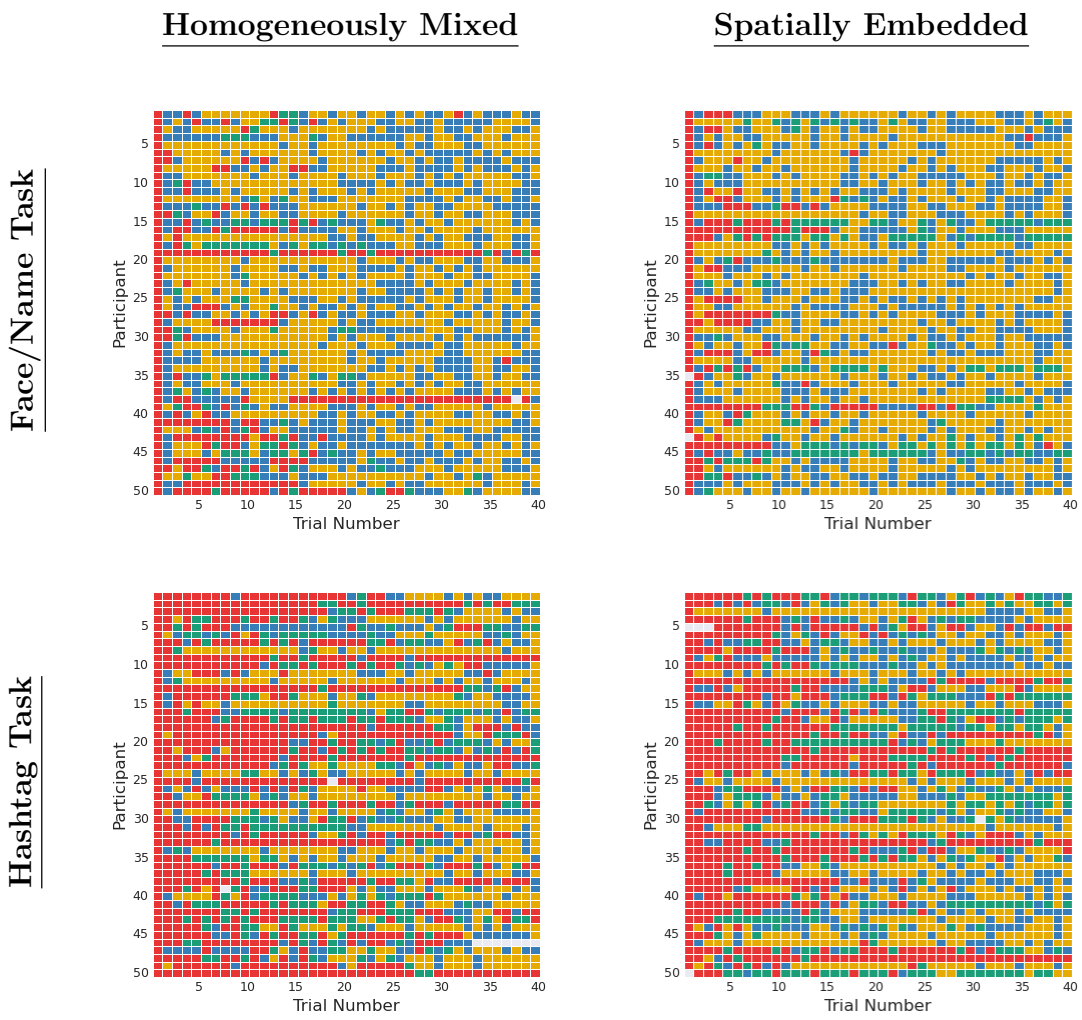

Figure 2.5: Individual participant decision types as they evolve with trial number across different network, task-type conditions. Columns represent trials, and rows represent participants. Each decision type is color-coded. In all cases, there is an abandonment of the BN decision type in favor of RS and RP.

task, where uncertainty leads to a more rigorous investigation of the response space [11].

Another way to visualize the decision types during the experiments is by displaying a grid of trial numbers against participants, and coloring each square according to the decision types used, as in Figure 2.5. In all four cases, we see abandonment of BN and EC decision types in favor of RS and RP, which indicates that learning is happening over time, regardless of task and structure. Importantly, we see that use of the BN (red) strategy in the face/name cases wanes much quicker than in the hashtag cases, which reinforces

our understanding that the increased complexity encourages an exploratory period where participants do not commit quickly to a limited number of responses.

In contrast to responses and entropy, decision types do not vary much according to network structure, which is a consequence of the fact that individuals in the experiments were unaware of the network structure during these interactions. Notice that decision types (particularly reliance on background information, which produces BN responses) are different between the tasks: there is a sharper decrease in the use of non-contextual information to make decisions in the face/name task. Since the hashtag task is more demanding [11], this may result from individuals combining different sources of information to create hashtags that are partially context-based but otherwise novel: an expansion of the response space. For example, in trial T, participants with the responses "nuclear" and "disaster" may be paired, and in trial $T + 1$, one or both may respond with "nucleardisaster." Unless they had been exposed to that response verbatim previously during the network interaction, their hashtag would be considered brand new and thus grouped into the use of background information. The face/name task is less susceptible to such situations because of its design. This illustrates the increased variability of individual group members' decision strategies in the hashtag task, which may follow from the higher demand of the task [11].

2.3 Experimental Results: Overview

Differences in human behavior due to varied tasks and network structure effects were present in the experiments. We can see that, as predicted, homogeneously mixed networks (a proxy for online communication) spread information faster than spatially embedded networks (a proxy for offline communication). The tasks to which participants were assigned impacted their behavior, which illuminates the importance of context awareness when modeling human decision-making, and was a driver for model development.

Chapter 3

Model Background

Game-theoretic models provide a natural framework for modeling our experiments, but they have some notable limitations. A well-known game theory framework is the replicator equation [4], which has been used to model the evolution of finitely many strategies in a population, favoring those with more comparative success. These models rely on strategies themselves being directly profitable, whereas in our experiments, this is not the case. Rather, the goal is group coordination, which requires cooperation between individuals in their *responses*, but different distributions of *strategies* across individuals may or may not be profitable depending on the network's makeup as a whole. Of course, some strategies may locally increase coordination, but it may even be true that individuals generating brand-new responses could drive the network as a whole toward coordination more rapidly. Because we quantify the agreement of the network using the Shannon entropy, which is a representation of the quantity of different responses and their popularity over time, it makes sense that the best strategies for consensus are to repeat either yourself or your partner from the immediately previous trial. If each individual on a network were to do this, the entropy would decrease at each timestep by definition, as no one would be *adding* responses to the network that were not there on the previous trial. Of course, a strict decrease requires a mix of these strategies: if the entire network committed to one or the other, the entropy would remain constant. Existing models are also heavily reliant on the idea that individuals make rational choices [1]. In our case, logic dictates that given the ultimate goal of the task, people would quickly learn to make profitable choices (i.e., RS or RP). However, our analysis of the experimental results shows that this is not the case. Subjects appear to continue to draw brand new and earlier context responses long into a task block.

Another potential avenue for investigation based on our four-strategy parsing of the experimental data is to develop a model based on existing models for bandit problems [10]. At first glance, our decision-type designation seems well-suited to modeling through this lens, though the time-dependence of reliance on different strategies implies that existing framework for these problems may not be able to capture differential behavior of individuals based on context. The complexity of our experiments, particularly with respect to the irrationality of participants and the time-dependence of their strategic learning, drove us to look beyond existing models and develop our own framework. Importantly, our new framework seeks to account for task-dependent differences in the prior information of participants, which would ideally be generalizable to a variety of different contextual situations. Since our experiment design was developed from Centola & Baronchelli (2015), we began our simulation efforts with the model developed in their paper [3].

3.1 The Centola & Baronchelli (CB) Model

In the Centola & Baronchelli (CB) model, there is a one-way transfer of information within each pair of agents: one person is the "speaker" and they share their response with the "hearer." This is an important deviation from the two-way information exchange that occurs in the experiments. Each participant has a "vocabulary" that stores the contextual information they refer to throughout the experiment. If the hearer already has the speaker's response in their vocabulary, both agents' vocabularies are overwritten by that single response, which becomes both of their responses for the next trial. If not, the speaker's response is added to the vocabulary of the hearer, and both agents randomly select their next response from their existing vocabulary. Pseudocode for this model is below:

Algorithm 1 Centola & Baronchelli (CB) Decision-Making Algorithm

```
for each round t in number of rounds do
    for each pair (agent1, agent2) in pairings do
        Randomly assign speaker and hearer roles to agent1 and agent2
        if speaker's current response is in the hearer's vocabulary then
            speaker's vocabulary ← {speaker's response}
            hearer's vocabulary ← {speaker's response}
        else
            Add speaker's response to hearer's vocabulary
        agent1's new response ← random selection from agent1's vocabulary
        agent2's new response ← random selection from agent2's vocabulary
```

The simplicity of this model is immediately appealing, but we do note that the interactions between the simulated agents do not appropriately reflect the two-way exchange of information that occurs between live agents. In addition, it quickly becomes clear that this model does not allow for the continued use of different strategies; importantly, the ability of agents to use brand new and earlier context strategies is severely limited because an agent's vocabulary only includes words they or a partner have said before. This model, aligned with existing game theory, promotes the idea that agents make rational choices, which we have seen is not the case in real life.

3.1.1 CB Simulations

We ran simulations to investigate the efficacy of the CB model (Algorithm 1) and compared it to our experimental data. The first metric we used for this comparison was Shannon entropy (Equation 2.1).

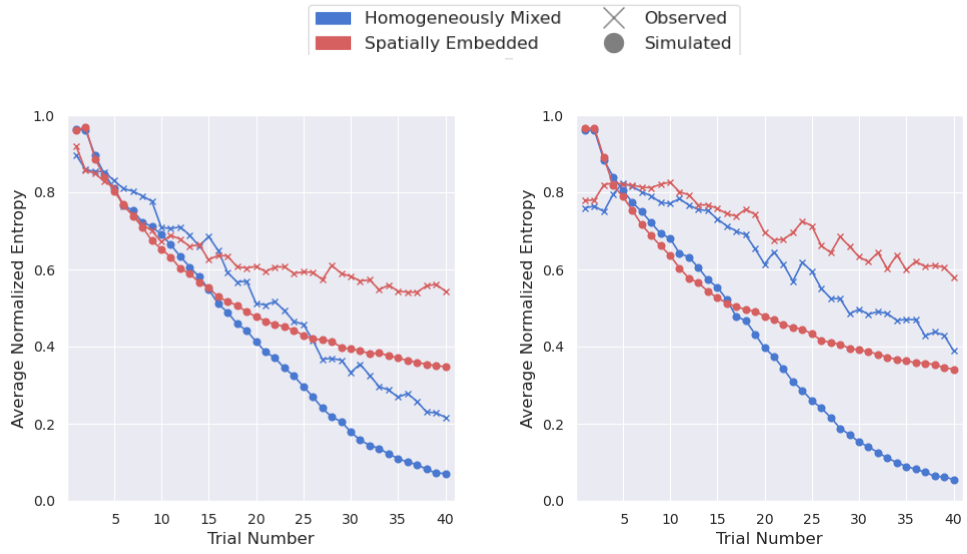

(a) Face/Name Task

(b) Hashtag Task

Figure 3.1: Observed average entropy and simulated average entropy for (a) the face/name task and (b) the hashtag task, both separated according to network structure. For each case, 40 simulations were run: 20 for each $N = 20$ and $N = 50$. In all four cases, the CB decision algorithm causes the network to perform better than our experiments, a consequence of the assumption of rational agents. Also note that the simulation is not task-sensitive: the behavior is the same for both tasks despite the differences we see in the experiments.

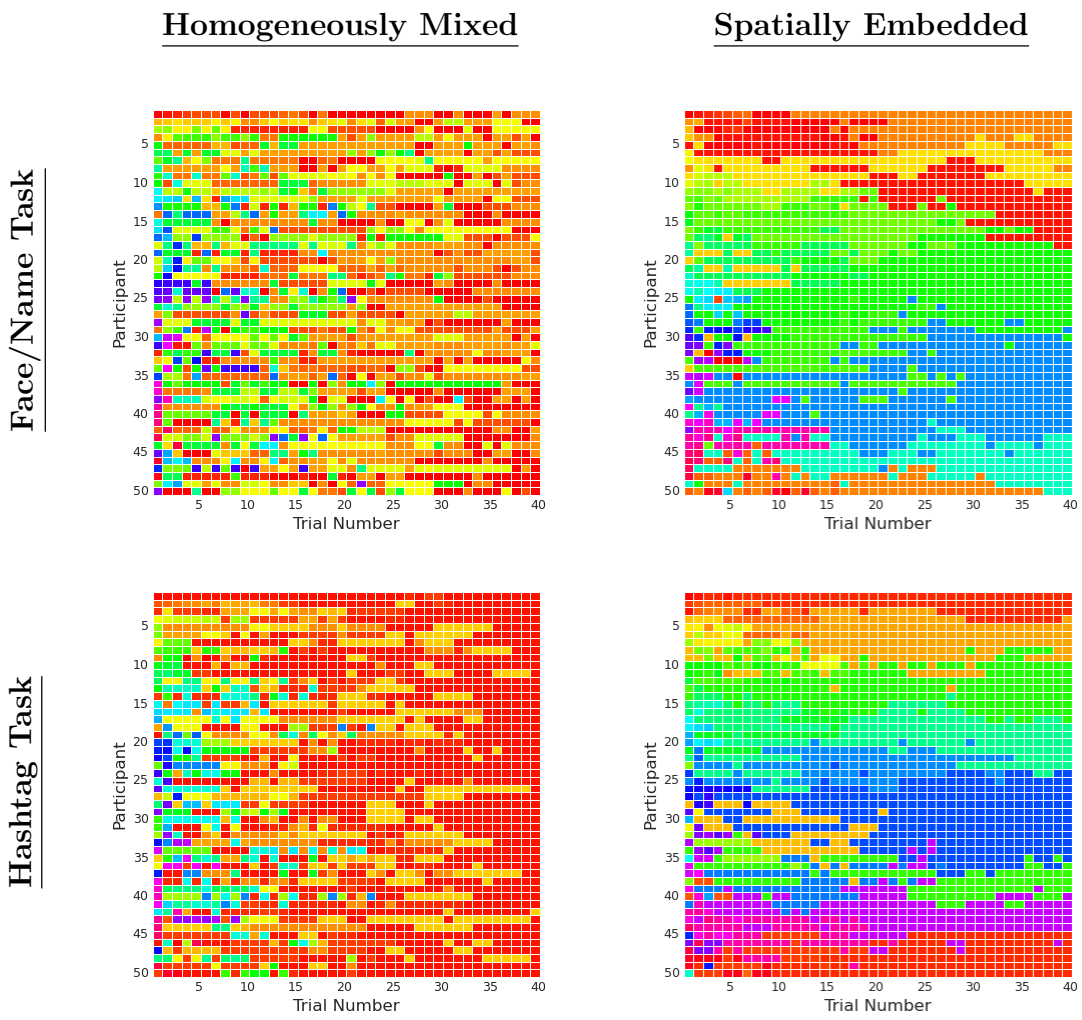

Figure 3.2: For each structure and task combination, a colormap shows how the *responses* change on the network over time for a selected $N = 50$ simulation. Each row represents a participant and each column is a trial, with cells colored according to the response given. We see qualitative features similar to the experiments: the spatially embedded cases (right column) exhibit banding behavior indicating local coordination, while the homogeneously mixed networks quickly see a limited number of dominant responses.

Figure 3.1 compares the simulated and experimental entropy for all of the experimental conditions. In all four cases, the simulation entropy decreases more on average, which implies more agreement on the networks. This is due to the assumption of rational agents, as promoted by classical game theory [4]. Likewise, we see that there is little difference in the behavior of the model based on the task, indicating that

this model is not affected by the specific background information participants bring into the interaction. Most importantly, the CB algorithm is unable to account for the difference in starting entropy between the two tasks, which is also demonstrated by the shape of the prior distributions in Figure 2.1.

Response colormaps are shown in Figure 3.2, which further illustrate the lack of robustness to task content, as we see the same qualitative features emerge for both face/name and hashtag task cases. The local coordination for the spatially embedded networks is evidenced by the localized bands of color, which emerge quickly. The qualitative features of Figure 3.2 can be compared to those of Figure 2.3 to see that the banding behavior emerges faster and is more stable in the simulated case than the experimental case, which is another consequence of over-rational agents.

As we did when analyzing the experiments, we then shifted to studying the behavior of individuals with respect to the decision types we identified: BN, EC, RS, and RP. Decision types were determined retroactively, as with the experimental data, based on the *responses* generated by agents throughout the trials, as well as the pairs interacting during the simulation and their responses.

For each combination of task and structure, there are 4 vectors which represent the average proportion of each decision type at each trial (one for each strategy). In the simulated case, this average is taken over 40 simulation runs: 20 with each $N = 20$ and $N = 50$. There are corresponding observed vectors, which contain the average proportion of each decision type at each time, taken over all of the *experiments* for that structure/task condition. This is also the height of the relevant colored bars in Figure 2.4. So, for each simulated vector $\mathbf{S}$ and corresponding observed vector $\mathbf{O}$, the element-wise distance is calculated $\mathbf{D}_i = |\mathbf{O}_i - \mathbf{S}_i|$. It is these vectors $\mathbf{D}$ that are plotted against trial number in Figure 3.3.

Upon inspection, we see high separation between the observed and simulated average decision-type proportion vectors, particularly in the case of the hashtag task. Notably, we see a significant difference in the proportion of the network referring to background information (using the BN strategy). As a note, someone who is the "speaker" at a given timestep *cannot* employ the RP strategy. Thus, in the case the **if** statement in line 4 of Algorithm 1 is entered, their strategy is RS by default. This indicates that there will be a much higher proportion of individuals using the RS strategy in the CB simulations than in our observed data, which explains the high distances between RS proportion vectors indicated by Figure 3.3.

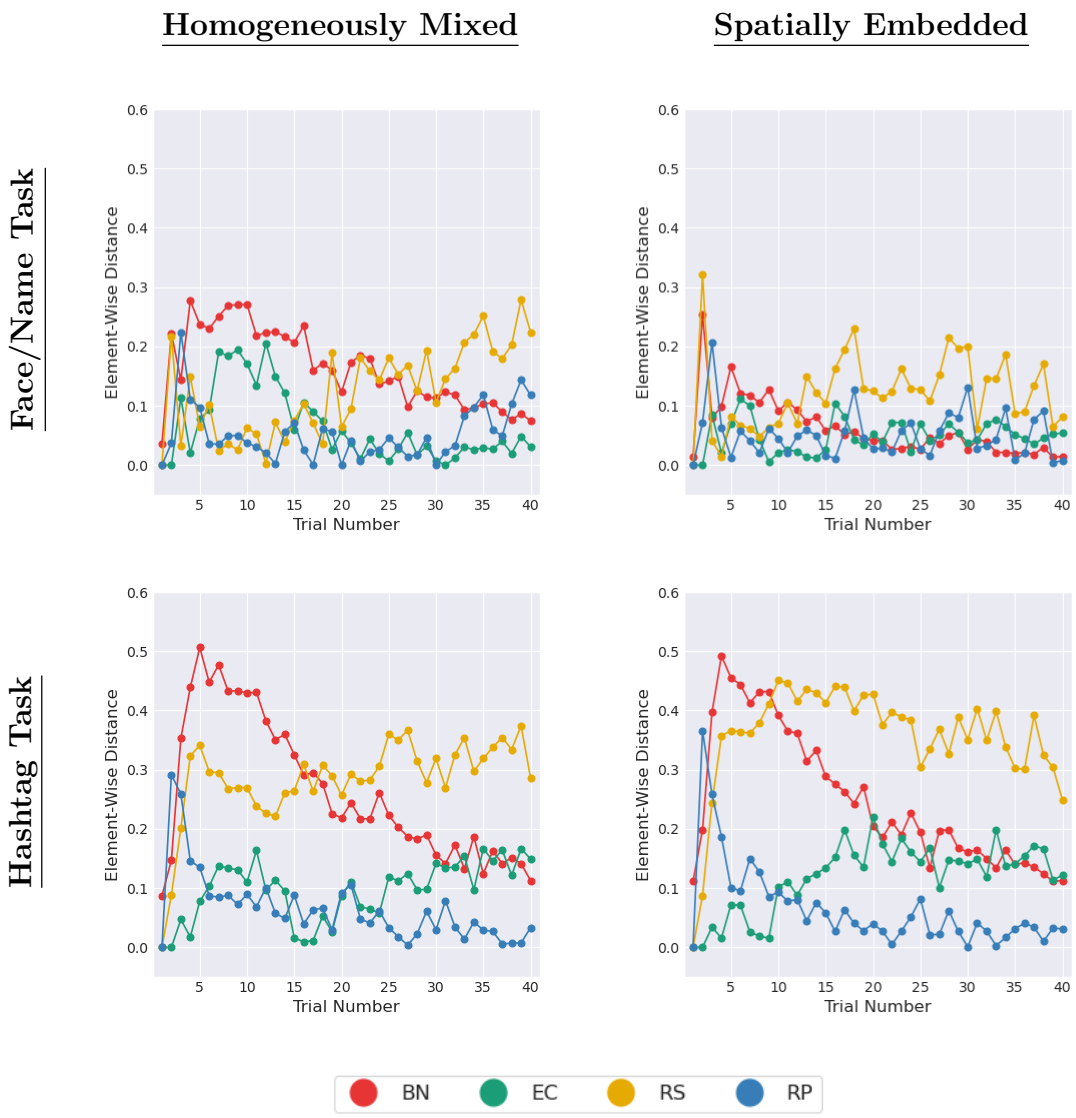

Figure 3.3: For each structure and task combination, line graphs show the vectors **D** which illustrate the element-wise distance between proportions of networks engaging in each decision-type over time in the observed and simulated (CB) cases. We see particularly high discrepancies between BN and RS strategies in both of the hashtag task cases.

To investigate *how* decision strategies are distributed over space and time in the simulations, we refer to Figure 3.4, which shows a color-map of the *decision-types* for a single $N = 50$ simulation for each structure/content pair.

In addition to the higher use of the RS strategy as a result of the design of the CB algorithm, the discrepancies in use of the BN strategy can similarly be explained: once a participant has been the "hearer"

it is impossible for them to use the BN strategy ever again because their vocabulary is nonempty. This phenomenon is much less obvious in the case of the face/name task since participants do not rely on the BN strategy very much in this task anyway, so the difference is far less pronounced. This is also why the RS

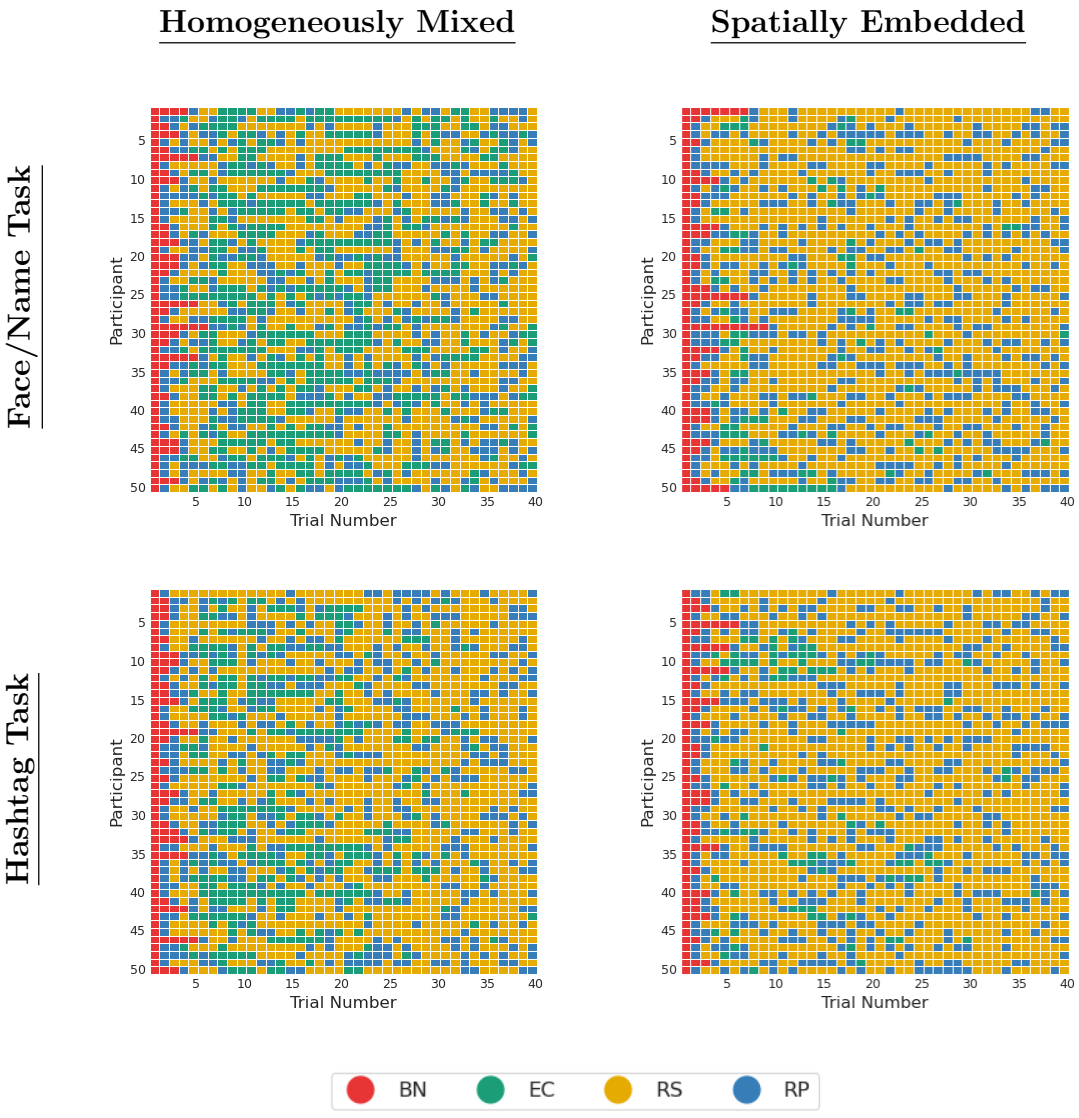

Figure 3.4: For each structure and task combination, a color-map shows which *strategies* are being used by agents through the course of the CB simulation. In all four cases, we see that BN use is constricted to the very early trials and RS dominates.

differences are less severe in the face/name case: participants already use this strategy more frequently for the duration of the experiment than hashtag task participants.

Using the CB algorithm we note that performance of simulated networks is reasonable and large-scale dynamics (entropy decrease, structure effects) are replicated well. This model performs particularly well for the face/name task, which is what provided the inspiration for its development [3]. With respect to the hashtag task, the CB algorithm does not capture the nuance of task context, one of the main goals for the models we develop. We also note that the CB algorithm is not suited toward replicating decision-type behavior as observed in the experiments, because it was not designed with this framework in mind.

Ultimately, the CB simulation results gave us a number of directions to pursue with our modeling efforts in the case of multiple tasks. The hashtag task is particularly difficult to model due to its complexity, which presented an interesting challenge for our modeling efforts.

Chapter 4

The Models: Building Context-Aware Agents

The results of network simulations using the Centola & Baronchelli (2015) decision strategy gave us a starting point for our model. We aimed to create an algorithm that, given a network structure and task:

- incorporates empirical prior information to account for entropy differences between the tasks,
- accurately models entropy decrease over time, particularly in the case of the hashtag task, and
- is true to the general proportion of decision types on the networks over time.

Our model development was initially motivated by a desire to successfully reproduce the main differences in network-level results, most importantly with respect to entropy decrease. As differential entropy decreases based on network structure are embedded in the proper implementation of interactions, we were most focused on making our model sensitive to the prior information agents brought into the interactions. We focused on successfully recreating the network-level trends of the observed agents without overfitting to the data, which drove us to develop a sequential decision-making structure based on the four decision types identified earlier. To study the performance of the model relative to the outputs of the experiments as well as the control model discussed previously, simulations were run over both homogeneously mixed and spatially embedded networks with various sizes ($N = 20$ or 50) and with both priors, which were passed as inputs into the model and used to generate initial responses for the simulated agents. The use of empirical prior information in the simulated networks allowed for the replication of behaviors that occurred as a result of biases within the different tasks. Each agent built their context throughout the simulation, which was updated as they shared information. A flowchart for the modeling framework is given in Figure 4.1.

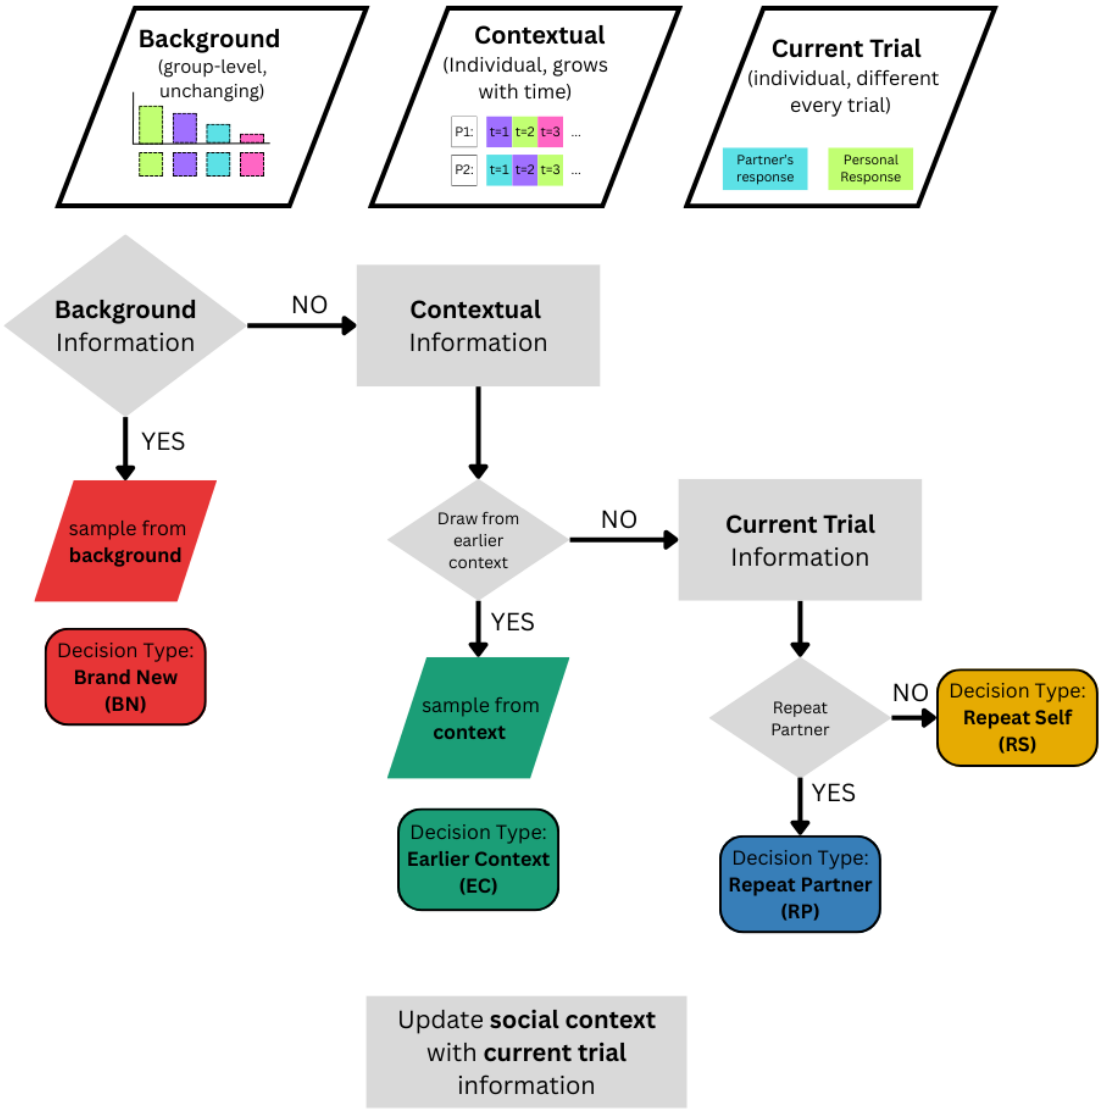

Figure 4.1: The model structure is a sequential decision-making path where agents consider a variety of possible decision-types and their probabilities of using each are determined by a combination of parameter values, perceived utility of a given strategy, and stochasticity. Once the *decision strategy* has been selected, a participant's response is generated by selecting from the appropriate pool of information (BN implies a random selection from the task prior, EC is a selection from personal memory which grows over time, and RS and RP imply specific responses).

4.1 Context-Aware Collective Learning (CACL)

To integrate the use of the Brand New strategy, we parameterized the first decision point using a network-level **learning parameter**, α. Lower α value leads individuals to switch to a strategy driven by context, as given by Equation (4.1a). The probabilities of individual i using each strategy at a timestep T are given by:

$$\begin{cases} P_{i,T}(BN) = \dfrac{\alpha}{T + \alpha + C_i}, & \text{(4.1a)} \\ P_{i,T}(EC) = (1 - P_{i,T}(BN)) \cdot \dfrac{S_{i,EC} + 3}{C_i + 20}, & \text{(4.1b)} \\ P_{i,T}(RP) = (1 - P_{i,T}(BN) - P_{i,T}(EC)) \cdot \dfrac{S_{i,RP} + 1}{S_{i,RP} + S_{i,RS} + 2}, & \text{(4.1c)} \\ P_{i,T}(RS) = (1 - P_{i,T}(BN) - P_{i,T}(EC) - P_{i,T}(RP)), & \text{(4.1d)} \end{cases}$$

where $S_{i,EC}$, $S_{i,RP}$, and $S_{i,RS}$ represent the number of points scored by individual i using the strategies EC, RP, and RS, respectively, and $C_i = S_{i,EC} + S_{i,RP} + S_{i,RS}$ is the total number of context points through trial $T - 1$ weighted by the trial number in which they were scored (e.g. a point scored using the EC decision strategy in trial 20 adds 20 to $S_{i,EC}$, which gives more weight to strategies that have paid off *recently*). Note that the entire network shares a single learning parameter α, but participants' probabilities of using any given strategy become individualized as they score points during the simulation. This balances the use of learned information while allowing for different use of background information at the network level. The 3/20 in Equation (4.1b) was empirically determined based on the overall frequency with which the EC strategy was used in the experiments.

As this collective learning model is structured, there is no upper limit on α, but as $\alpha \to \infty$, the probability of any person using the brand new strategy quickly goes to 0. Importantly, we wanted to ensure that observed differences in coordination behavior by network structure were a result of the topography, not accidentally influenced by the model. Individuals in the different tasks learn differently as implied by their varied use of the BN strategy. In order to ensure that network effects were a result of structure alone, we chose a single learning parameter for each task by minimizing the Euclidean distance between simulated and

experimental entropies on average for both network structures (see Figure 4.2).

Algorithm 2 shows the decision-making process for this simulation, referred to as the Context-Aware Collective Learning (CACL) algorithm.

Algorithm 2 Context-Aware Collective Learning (CACL) Decision-Making Algorithm

for each round T in number of rounds **do**
  for each pair (agent1, agent2) in pairings **do**
    if agent1's response = agent2's response **then**
      Update each agent's scores (and relevant decision-type scores)
    for each agent i **do**
      decision type is BN with probability $\frac{\alpha}{\alpha+T+C_i}$
      if decision type is BN **then**
        agent's new response ← random sample from experiment prior
      else
        decision type is EC with probability $\frac{S_{i,EC}+3}{C_i+20}$
        if decision type is EC **then**
          agent's new response ← random sample from agent's context
        else
          decision type is RP with probability $\frac{S_{i,RP}+1}{S_{i,RP}+S_{i,RS}+2}$
          if decision type is RP **then**
            agent's new choice ← partner's current choice
          else ▷ decision type is RS
            agent's new choice ← agent's current choice
    if agent1's response == agent2's response **then**
      for each agent **do**
        Add $50 \cdot T$ instances of partner's response to their context
        ▷ agents favor responses that have scored them points
    else
      for each agent **do**
        Add T instances of partner's response to their context

4.1.1 Fitting Parameters

Varying $\alpha \in [0, 10]$ and computing the average distance between the entropy of experimental subjects and model populations (Figure 4.2) demonstrates the dependence of task structure on this measure. Ultimately, we found two α values to model subject behavior, minimizing this entropy distance, which depended on the task structure: one for each task (prior distribution). Based on the results of this parameter sweep, simulations for the face/name experiment had a group-level learning parameter of $\alpha = 0.4$, in contrast to the hashtag simulations with a parameter value of $\alpha = 4.9$. As α is designed, this implies that the probability of using the BN strategy will, in general, be higher in the hashtag task than in the face/name task, which promises to more accurately reflect the decision strategies of the subjects. These parameter values are used

in all simulation results for this version of the model discussed hereafter.

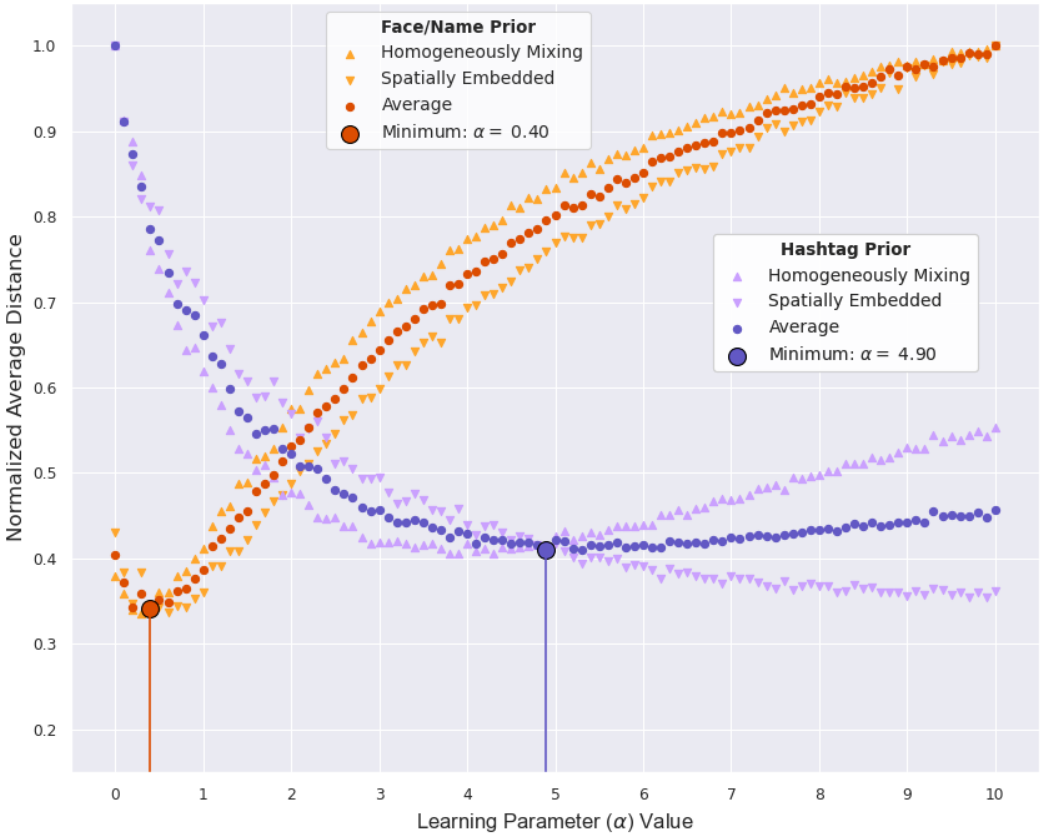

Figure 4.2: For each α, the Euclidean distance between the average experimental entropy and the simulated entropy is shown. The average is taken over 500 simulations. Colors represent different prior (task) conditions, and all distances are normalized so the maximum distance for each task is one. We see that the dependence of the distance on α is shaped by the task and network structures.

4.1.2 CACL Simulations

The CACL model is designed to more faithfully represent how subjects work to coordinate responses and reduce entropy. In Figure 4.3, we see a comparison between the experimental and CACL entropies using our optimized parameter for each task. We note that when compared to the entropy reduction in the CB

model, the CACL algorithm is able to better reflect the different behavior in all four conditions, particularly with respect to the task.

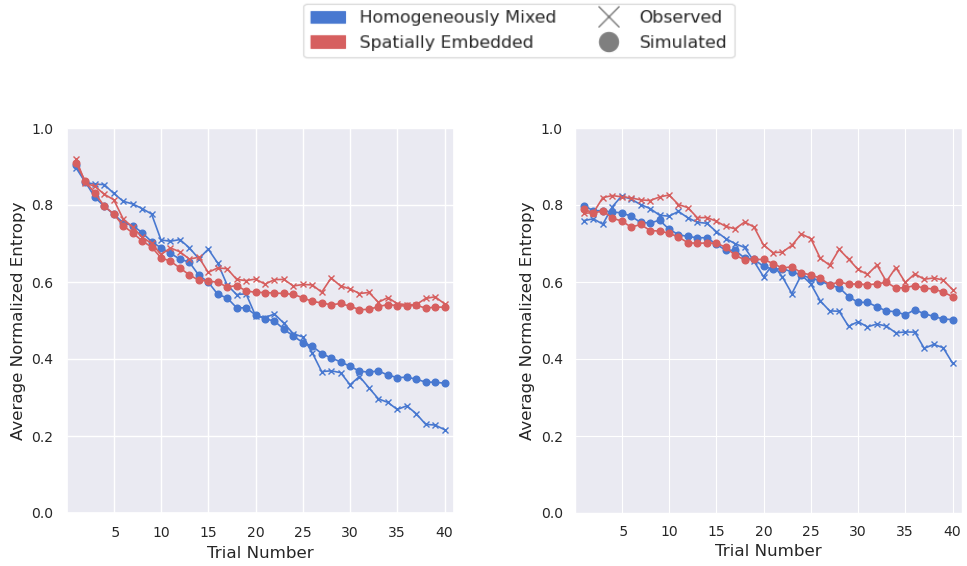

(a) Face/Name Task

(b) Hashtag Task

Figure 4.3: Observed average entropy and simulated average entropy for (a) the face/name task and (b) the hashtag task, both separated according to network structure. For each case, 40 simulations were run: 20 for each $N = 20$ and $N = 50$. The CACL model produces average network entropy decreases that are closer to the observed results in all four cases than those produced by the CB model (Figure 3.1). Importantly, this algorithm *is task-sensitive*, replicating the differences in behavior between the tasks seen in the experiments.

In general, the CACL model is more representative of the empirical data, but less sensitive to the network structure than the Centola & Baronchelli model. In Figure 4.4, we see an example of the distribution of responses on the network over time. Particularly in the hashtag task, the CACL model's slow learning rate limits the population's ability to coordinate responses to reflect the behavior of the experimental participants.

In the face/name case, we can see the impact of network structure on the choices of individuals across the population and time. These effects fail to emerge fully in the hashtag case due to the network-level slow rate of learning given by our optimized parameter, but this is observed in the experimental runs as well.

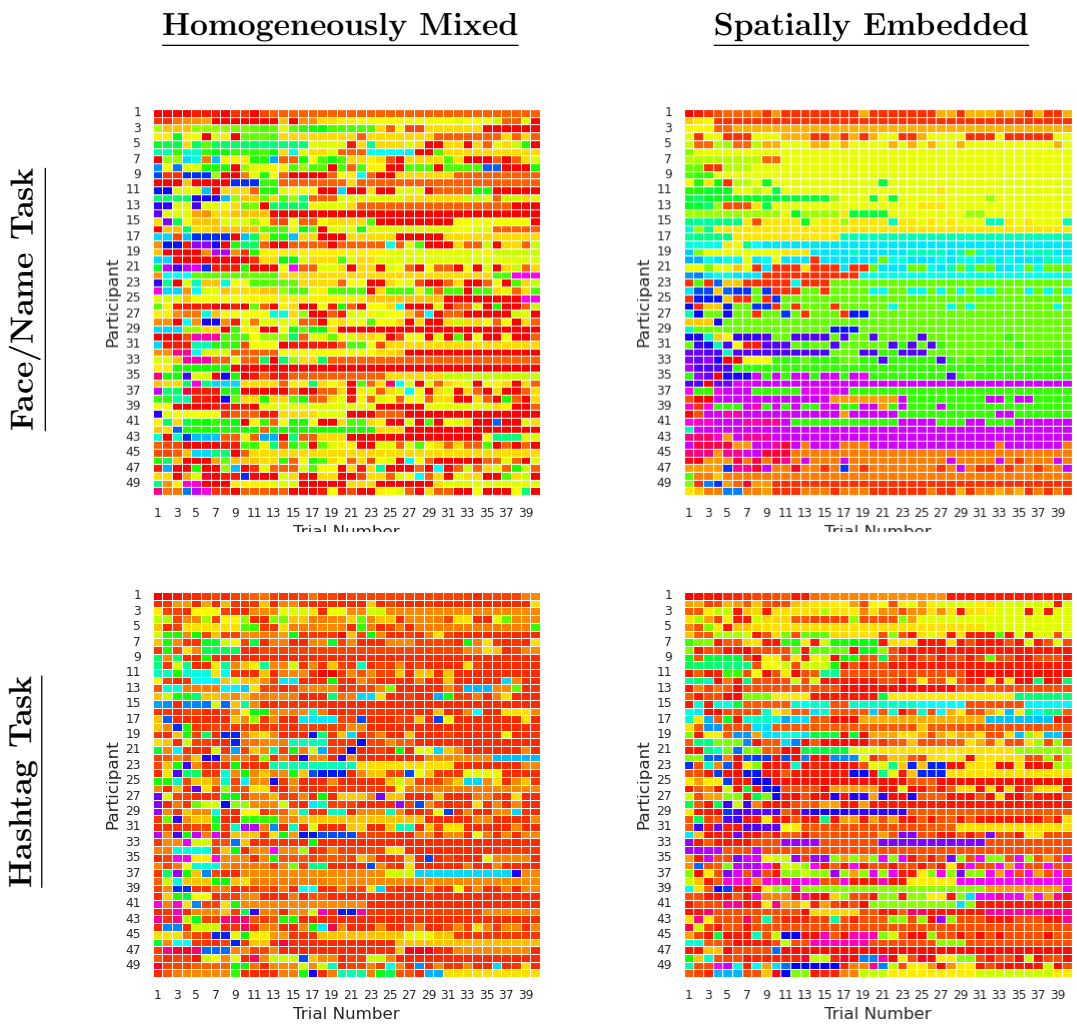

Figure 4.4: For each structure and task combination, a color-map shows how the *responses* change on the network over time for a selected $N = 50$ simulation. Each row represents a participant and each column is a trial, with cells colored according to the response given. Particularly in the spatially embedded case we see the striped features of the experimental runs emerging.

Both Figures 4.5 and 4.6 give insight into the decision-type behavior of agents in the CACL simulations. With the addition of the learning parameter, we expect to see a closer approximation of the behavior of experimental participants as α directly impacts the task-specific frequency with which individuals use the BN strategy (which was previously noted to be the most important decision-type difference between the tasks and also the most neglected by the CB algorithm).

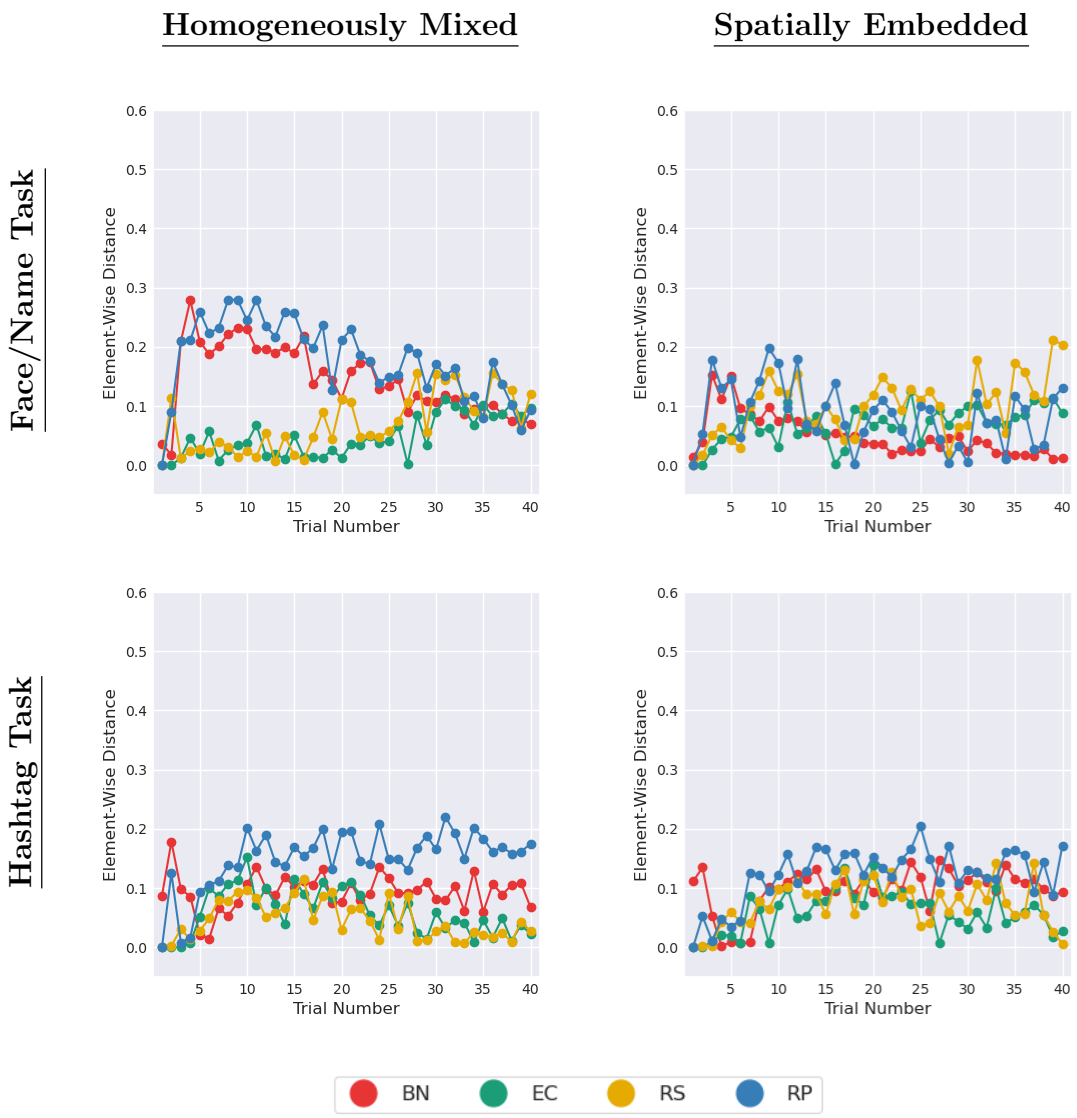

Figure 4.5: For each structure and task combination, line graphs show the vectors **D** which illustrate the element-wise distance between proportions of networks engaging in each decision-type over time in the observed and simulated (CACL) cases. In general, we see that the CACL model performs better than the CB model with respect to decision-type behavior. Recall that in the CB case, particularly for the hashtag task, use of the BN strategy was severely underestimated while use of the RS strategy was overestimated.

Notably, our algorithm does produce the group-level decision-type behavior seen in the experimental case by facilitating the exploratory period seen for the hashtag task. Specifically, the parameterization of the

model by each task and the specific role of α in the decision structure reproduced the early dominance of the BN strategy in the hashtag task. In all four structure/task conditions, the proportion of the network using each decision strategy on average is much closer to the data when using the CACL simulation as opposed to the CB version (as seen by comparing Figures 3.3 and 4.5).

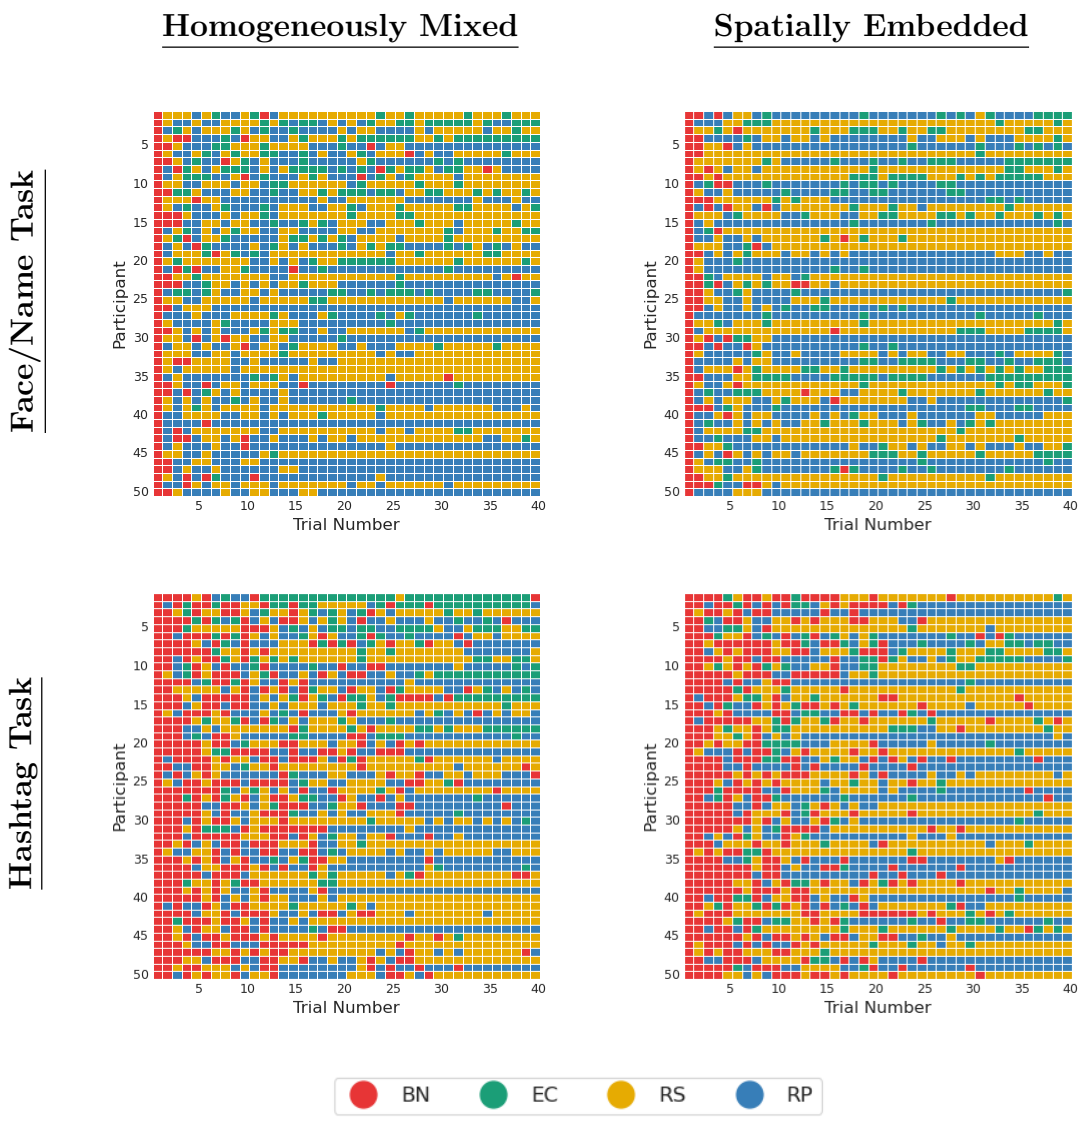

Figure 4.6: For each structure and task combination, a color-map shows which *strategies* are being used by agents through the course of the CACL simulation. We see that participants commit to specific strategies but the BN strategy is more popular in the hashtag task, which is true to the experimental data.

4.2 Context-Aware Personal Learning (CAPL)

While the CACL model represented clear gains in our ability to recreate observed differences, particularly between the tasks, the rigidity of a network-level learning parameter was unable to capture the differences in learning between *individuals* on the networks. For instance, though the proportions of decision strategies at use in the networks are truer to those observed in general, we do not see the emergence of stubborn non-learning individuals as can be seen by the horizontal red lines in Figure 2.3, which even emerge in the homogeneously mixed face case. This indicates participants who– in direct opposition to the rationality assumed by game theory– have decided to switch their response every trial as a way to introduce chaos into the system. This is a feature that is not well-modeled by the network-level learning of the CACL model. For this reason, we propose an updated algorithm (which still adheres to the model flow of Figure 4.1) that individualizes the learning process in order to account for these non-learning individuals. In the new algorithm, the decision of agent i at time T is governed as follows:

$$\begin{cases} P_{i,T}(BN) = \dfrac{S_{i,BN} + \frac{\alpha_i}{1-\alpha_i}}{T + S_{i,BN} + C_i + \frac{\alpha_i}{1-\alpha_i} + \frac{1-\alpha_i}{\alpha_i}}, & \text{(4.2a)} \\ P_{i,T}(EC) = (1 - P_{i,T}(BN)) \cdot \dfrac{S_{i,EC} + \frac{\beta_i}{1-\beta_i}}{C_i + \frac{\beta_i}{1-\beta_i} + \frac{1-\beta_i}{\beta_i}}, & \text{(4.2b)} \\ P_{i,T}(RP) = (1 - P_{i,T}(BN) - P_{i,T}(EC)) \cdot \dfrac{S_{i,RP} + 1}{S_{i,RP} + S_{i,RS} + 2}, & \text{(4.2c)} \\ P_{i,T}(RS) = (1 - P_{i,T}(BN) - P_{i,T}(EC) - P_{i,T}(RP)) & \text{(4.2d)} \end{cases}$$

where $S_{i,BN}$ is the number of points scored by agent i using the BN strategy, and α_i and β_i are agent-specific parameters which encode the likelihood of i using the BN and EC strategies, respectively. We note that α_i and β_i are both restricted to the domain: $(0, 1)$.

As encoded, α_i controls the likelihood of individual i using the brand new strategy. As $\alpha_i \to 0$, the participant becomes less likely to ever use the BN strategy. At the other end of the domain, as $\alpha_i \to 1$, participant i uses this strategy with a probability proportional to $\frac{1}{T}$, so use of the brand new strategy necessarily decreases over time as participants are exposed to more information on the network.

In contrast, β_i controls individual i's propensity to use the EC strategy, whose probability does not change over time and is described by the formula Eq. (4.2b). Thus, in this case, as $\beta_i \to 0$, agent i is biased against the EC strategy (always prefers RS or RP once reaching this decision juncture). On the other hand, as $\beta_i \to 1$, agent i always prefers EC over self or neighbor replicating decision strategies.

Note that at an early stage of our model development there was also a γ_i parameter governing the RP/RS decision-point such that $P_{i,T}(RP) = (1 - P_{i,T}(BN) - P_{i,T}(EC)) \cdot \frac{S_{i,RP}+\frac{\gamma_i}{1-\gamma_i}}{S_{i,RP}+S_{i,RS}+\frac{\gamma_i}{1-\gamma_i}+\frac{1-\gamma_i}{\gamma_i}}$. We compared the three-parameter version (with γ_i) to the two-parameter version (without) using the Bayesian Information Criterion for model selection, and every time the two-parameter model performed better. This indicated that the addition of γ_i did not have significant enough impacts on the modeling results to justify the additional complexity of the model introduced by the parameter. Also note that letting $\gamma_i = \frac{1}{2}$ in the above equation—indicating no inherent preference for either RS or RP over the other—gives 4.2c.

In context, this gives the Context-Aware Personal Learning decision-making algorithm:

Algorithm 3 Context-Aware Personal Learning (CAPL) Decision-Making Algorithm

for each round T in number of rounds **do**
  for each pair (agent1, agent2) in pairings **do**
    if agent1's response = agent2's response **then**
      Update each agent's scores (and relevant decision-type scores)
    for each agent i **do**
      decision type is BN with probability $\frac{S_{i,BN}+\frac{\alpha_i}{1-\alpha_i}}{T+S_{i,BN}+C_i+\frac{\alpha_i}{1-\alpha_i}+\frac{1-\alpha_i}{\alpha_i}}$
      if decision type is BN **then**
        agent's new response ← random sample from experiment prior
      else
        decision type is EC with probability $\frac{S_{i,EC}+\frac{\beta_i}{1-\beta_i}}{C_i+\frac{\beta_i}{1-\beta_i}+\frac{1-\beta_i}{\beta_i}}$
        if decision type is EC **then**
          agent's new response ← random sample from agent's context
        else
          decision type is RP with probability $\frac{S_{i,RP}+1}{S_{i,RP}+S_{i,RS}+2}$
          if decision type is RP **then**
            agent's new choice ← partner's current choice
          else ▷ decision type is RS
            agent's new choice ← agent's current choice
    if agent1's response == agent2's response **then**
      for each agent **do**
        Add $50 \cdot T$ instances of partner's response to their context
        ▷ agents favor responses that have scored them points
    else
      for each agent **do**
        Add T instances of partner's response to their context

4.2.1 Fitting Parameters

To find pairs of parameters (α_i, β_i) that replicate individuals' behaviors, we returned to the experimental results and used Maximum A Posteriori Estimation. We considered 41 possible values for each parameter:

$$\boldsymbol{\alpha} = \boldsymbol{\beta} = [0.0025, 0.025, 0.050, 0.075, 0.100, \ldots, 0.900, 0.925, 0.950, 0.975, 0.9975]$$

so that the parameter values were all within the appropriate domain (allowing parameter values to be equal to 0 or 1 produces division-by-zero errors, accounted for by the inclusion of the first and last elements of the parameter vectors, designed to get close to those values but not equal).

Each participant i, for $i = 1 : 1,002$, has a vector of decision types $\mathbf{DT^{(i)}}$ which are retroactively determined based on their responses and information about the information they exchanged and with whom over the course of the experiment, all of which was collected data during the experiment runs. Recall that in the case that two partnered individuals said the same thing on trial $T-1$ if either of them repeated that response on trial T their decision-type was randomly picked from RS and RP as both of those decision types would produce that response. We also have information about the points scored by individual i and with what strategy $\left(\mathbf{S_{BN}^{(i)}}, \mathbf{S_{EC}^{(i)}}, \mathbf{S_{RS}^{(i)}}, \mathbf{S_{RP}^{(i)}}\right)$. Also define $\mathbf{C^{(i)}} = \mathbf{S_{EC}^{(i)}} + \mathbf{S_{RS}^{(i)}} + \mathbf{S_{RP}^{(i)}}$. Assuming that participant i responded on each trial, all of these vectors are of length 40. For each combination of parameter values and each participant i, we computed a value proportional to the probability that a given pair of parameters $(\boldsymbol{\alpha}_j, \boldsymbol{\beta}_k)$ would give the appropriate decision types $\mathbf{DT^{(i)}}$. Note that using the Maximum A Posteriori (MAP) estimate for each i, we are in search of

$$\underset{\alpha,\beta}{\arg\max}\ P\left(\mathbf{DT^{(i)}} \mid \alpha, \beta, \mathbf{T}, \mathbf{C^{(i)}}, \mathbf{S_{BN}^{(i)}}, \mathbf{S_{EC}^{(i)}}, \mathbf{S_{RS}^{(i)}}, \mathbf{S_{RP}^{(i)}}\right) \tag{4.3}$$

because we assume that all possible values of α and β are equally likely. Because we know $\mathbf{T}$ (the vector of trial numbers), $\mathbf{C^{(i)}}$, $\mathbf{S_{BN}^{(i)}}$, $\mathbf{S_{EC}^{(i)}}$, $\mathbf{S_{RS}^{(i)}}$, and $\mathbf{S_{RP}^{(i)}}$, we can assume the independence of the decision types over time. If j is the trial number, define $P_j = P\left((\mathbf{DT^{(i)}})_j \mid \alpha, \beta, (\mathbf{S_{BN}^{(i)}})_j, (\mathbf{S_{EC}^{(i)}})_j, (\mathbf{S_{RS}^{(i)}})_j, (\mathbf{S_{RP}^{(i)}})_j, (\mathbf{C^{(i)}})_j, \mathbf{T}_j\right)$ and equation 4.3 becomes:

$$\underset{\alpha,\beta}{\arg\max}\ \prod_{j=1}^{40} P_j \tag{4.4}$$

We note that:

$$P_j = \begin{cases} P(BN)_j = \dfrac{(\mathbf{S_{BN}^{(i)}})_j + \frac{\alpha}{1-\alpha}}{\mathbf{T}_j + (\mathbf{S_{BN}^{(i)}})_j + (\mathbf{C^{(i)}})_j + \frac{\alpha}{1-\alpha} + \frac{1-\alpha}{\alpha}}, & \text{if } (\mathbf{DT^{(i)}})_j = \text{BN}, \\ P(EC)_j = (1 - P(BN)_j) \cdot \dfrac{(\mathbf{S_{EC}^{(i)}})_j + \frac{\beta}{1-\beta}}{(\mathbf{C^{(i)}})_j + \frac{\beta}{1-\beta} + \frac{1-\beta}{\beta}}, & \text{if } (\mathbf{DT^{(i)}})_j = \text{EC}, \\ P(RP)_j = (1 - P(BN)_j - P(EC)_j) \cdot \dfrac{(\mathbf{S_{RP}^{(i)}})_j + 1}{(\mathbf{S_{RP}^{(i)}})_j + (\mathbf{S_{RS}^{(i)}})_j + 2}, & \text{if } (\mathbf{DT^{(i)}})_j = \text{RP}, \\ P(RS)_j = (1 - P(BN)_j - P(EC)_j - P(RP)_j), & \text{if } (\mathbf{DT^{(i)}})_j = \text{RS}. \end{cases}$$

For each individual and combination of parameter values, we computed this 40-term product of P_j's, and then found the maximum to determine which pair of α, β is most likely. Examples of some posterior distributions for specific individuals are shown in Figure 4.7.

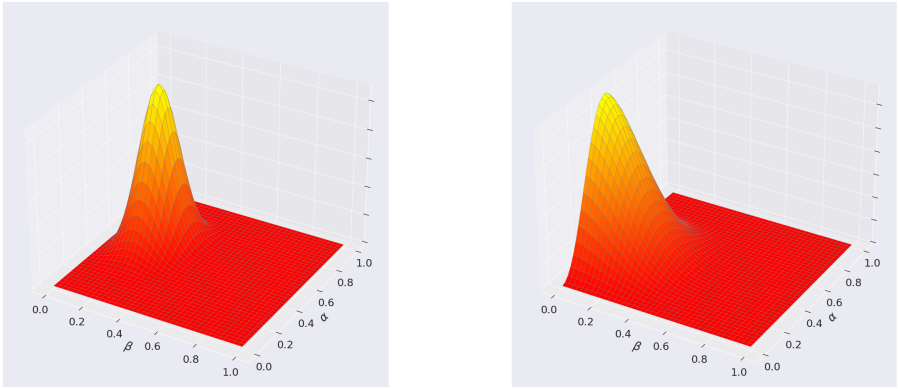

(a) Individual 1

$\alpha = 0.625$, $\beta = 0.2$

(b) Individual 2

$\alpha = 0.4$, $\beta = 0.0025$

Figure 4.7: Posterior distribution over α, β grid for two different experimental participants. Height is proportional to the probability that (α, β) produced the decision types of these individuals over all 40 trials, and the optimal parameter pairs for each individual are indicated.

After finding the best parameter pair for each of the individuals, we were able to analyze the distributions of the parameter values among the live participants using kernel density estimate (KDE) plots. Figure

4.8 shows such a plot including all of the individuals.

While looking at the distribution of parameters for all individuals gives us an idea of reasonable parameter pairs to approximate observed human behavior, it does not provide insight into the impact of network, task, or individual characteristics that may impact a person's behavior and, thus, influence which parameter values best model their decisions. To study which factors have the most salient impacts on the parameters, we have divided the dataset according to several characteristics, including task, network structure, number of total points scored by an individual, and their "contribution" to network consensus (defined below).

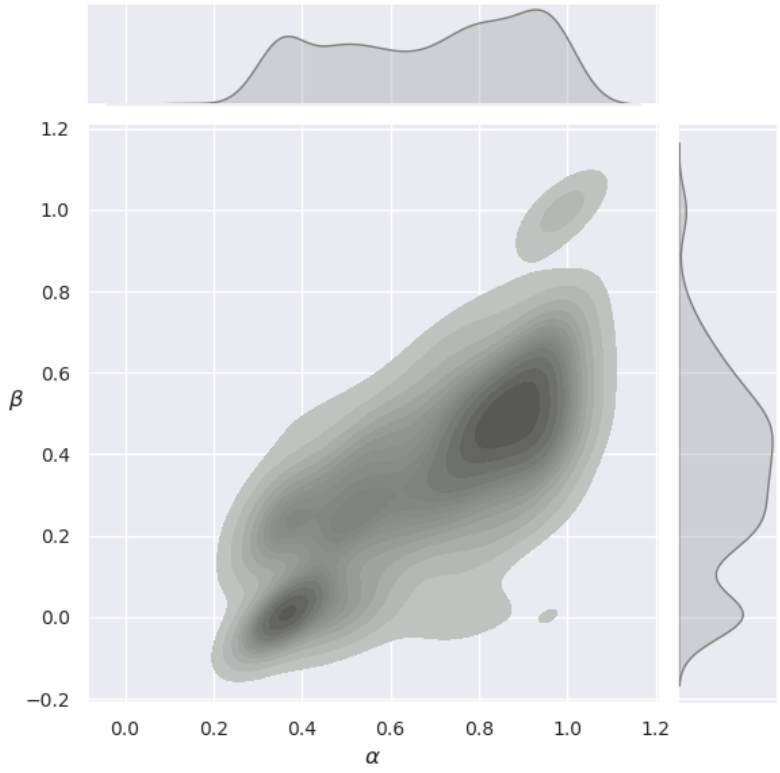

Figure 4.8: Density of participants at different points on the $\alpha\beta$ plane. All participants, regardless of task or network structure, are included. Darker colors represent higher densities, and the one-dimensional KDE plots for each variable are included at the top and right. We notice two areas of high concentration: one near $(0.3, 0.0)$ and the other near $(0.9, 0.5)$. Respectively, these pockets represent quick learners (those who begin using RS/RP dominantly early on) and slower learners who continue to use the BN and EC strategies.

The first clear way to split the data to see qualitative differences in the behavior of individuals was by task, as shown in Figure 4.9. As expected, we see a general tendency for individuals with the face/name task (Figure 4.9a) to have lower α and β values, indicating that they exhibited higher bias against the BN and EC strategies and, thus, relied on more immediate information (from the previous trial).

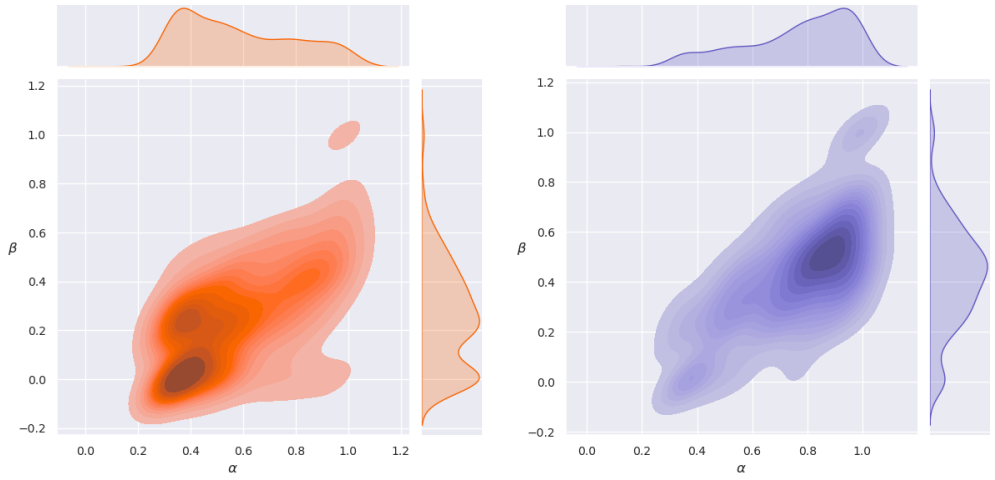

(a) Face/Name Task

(b) Hashtag Task

Figure 4.9: Density of participants at different points on the $\alpha\beta$ plane for (a) the face/name task and (b) the hashtag task. Notice that the highest concentration of face/name participants is near low values of both α and β, indicating quicker learning that the hashtag participants, who are concentrated near higher values of both parameters and thus demonstrate higher use of both the BN and EC strategies.

The observed longer exploratory period in the hashtag task, indicated by increased use of the BN strategy, is reflected in the KDE plots by task. Face/name participants tend to have lower values of both parameters, meaning they quickly shift their strategies to RS and RP (which are the most effective at lowering entropy). On the other hand, hashtag participants demonstrate higher values for both parameters. Specifically the concentration of α near 1 indicates that the decrease in use of the strategy can best be modeled for many participants as a direct relationship with time (driven by the T in the denominator of Equation (4.2a)).

On the other hand, when we split the data according to structure, we expect to see little-to-no difference in the variable distributions between these two groups assuming that decision-type behavior is not swayed by network structure because it is unknown to the participants. This is not exactly the case: in the spatially embedded case (4.10b), we see two peak concentrations, indicating that there is a potential difference in behavior between individuals based on the network structure. Investigating this further, we can see that there is a statistically significant difference (see Table 4.1) between the structure distributions within the face task, but not within the hashtag task. All of these distributions are shown in Figure 4.11.

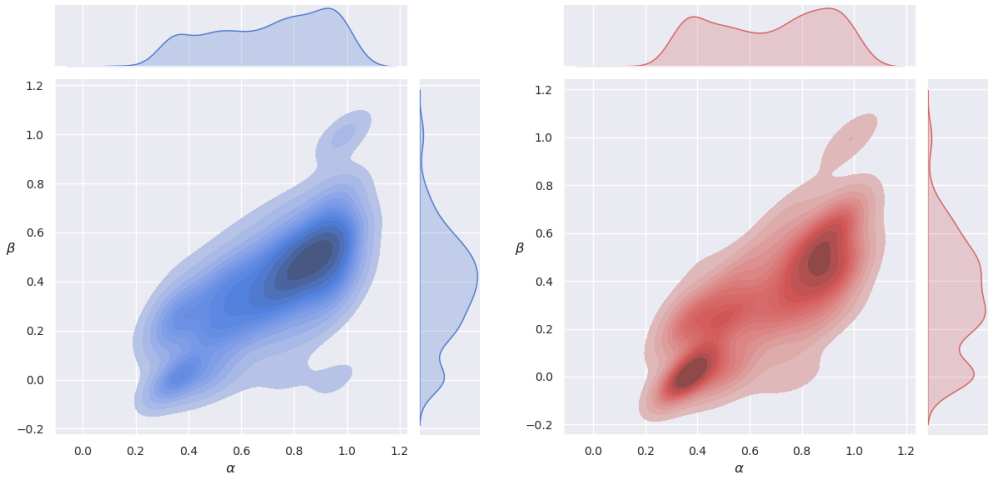

(a) Homogeneously Mixed

(b) Spatially Embedded

Figure 4.10: Density of participants at different points on the $\alpha\beta$ plane, separated by network structure. We see a concentration of quicker learners in the spatially embedded case that is not similarly reflected in the homogeneously mixed case. This is likely a result of the neighborhood size of individuals: participants in spatially embedded networks switch to RS/RP strategies as their exposure to new information is limited.

We note that the spatially embedded networks contain quicker learners, which is a result of their small neighborhood size and limited exposure to new information. This contradicts our assumption that network size does not impact learning behavior because it is unknown: although participants do not know the structure, spatially embedded networks restrict the amount of incoming information and result in localized agreement, which reduces the likelihood with which participants will introduce new responses to the network.

We also see that this difference is more pronounced in the face/name task than in the hashtag task (as shown in Figure 4.11) which is a combination of the inherent quicker learning of the less complex task with the reduction in information exchange based on the spatially embedded structure.

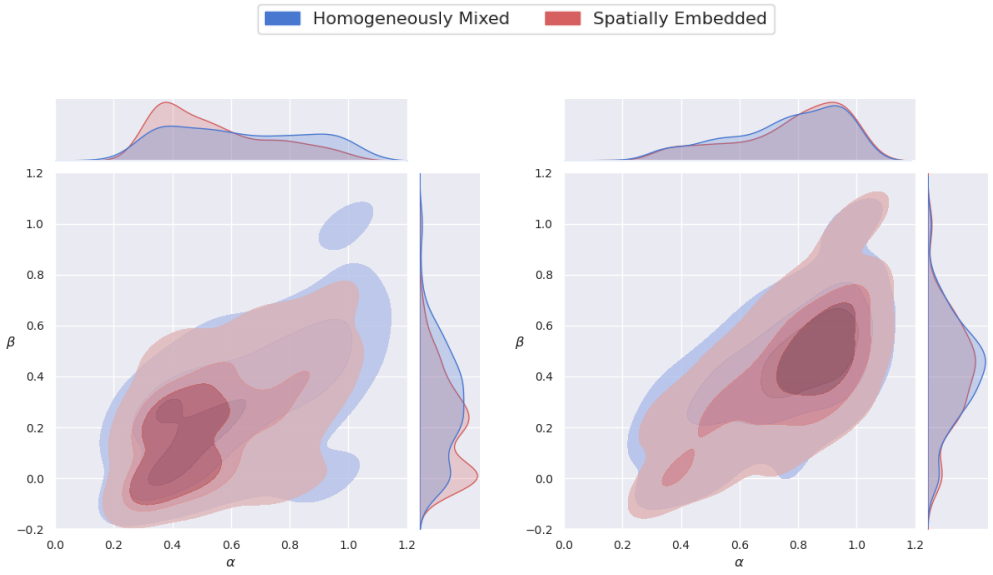

(a) Face/Name Task (b) Hashtag Task

Figure 4.11: Density of participants at different points on the $\alpha\beta$ plane, separated according to their assigned task, and then further split according to structure to see a side-by-side comparison of these subsets. Statistically significant differences between the homogeneously mixed and spatially embedded networks that had the face task can be seen qualitatively, while the lack of any significant difference between the groups in the hashtag case is also clear.

In addition to the natural comparison of individuals with different tasks and network structures, we also divided the group according to their agreeability in the network using two different metrics: points scored and contribution toward consensus. A participant's contribution toward consensus is defined

$$A_i = \frac{1}{20} \sum_{t=20}^{40} p(x_{t,i}), \tag{4.5}$$

or the average proportion of individuals on the rest of the network saying the same response as i for the last

half of the experiment. In contrast to total points scored, which only quantifies how many times a subject happened to agree with their partner, this metric takes into account the larger trends toward consensus on the network. Naturally, in spatially embedded networks, participants tend to score more points, but homogeneously mixed participants often have higher contribution (see Figure 4.14).

The KDE plots comparing individuals who scored a lot of points (at least 10 out of 40 possible) with those who did not is in Figure 4.12 and a similar plot comparing those with high contribution to consensus (agreement with at least 25% of the rest of the network on average for the last half of the experiment) to those with lower contribution is in Figure 4.13. Taking both of these metrics as proxies for agent success in the activity we see, as we predicted, that more successful agents (Figures 4.12a and 4.13a) tend to exhibit lower values of both parameters than their less successful counterparts.

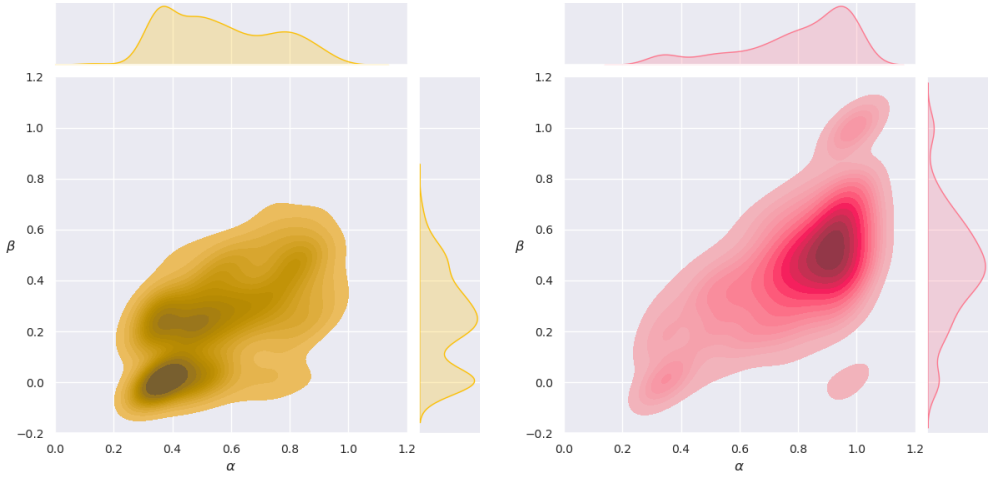

(a) At least 10 points scored

(b) Fewer than 10 points scored

Figure 4.12: Density of participants at different points on the $\alpha\beta$ plane, separated according to the number of points scored during the entire experiment. Because the length of each experiment is 40 trials, participants are able to score up to 40 points during the experiment (assuming they agreed with their partner on every single trial). Participants who scored at least a quarter of the possible points available (a) tend, as expected, to be quicker learners than those who scored fewer than 10 points (b).

As noted, participants on spatially embedded networks tend to score more points than those on homogeneously mixed networks because there is a limited amount of conflicting information they are exposed to during the course of the experiment.

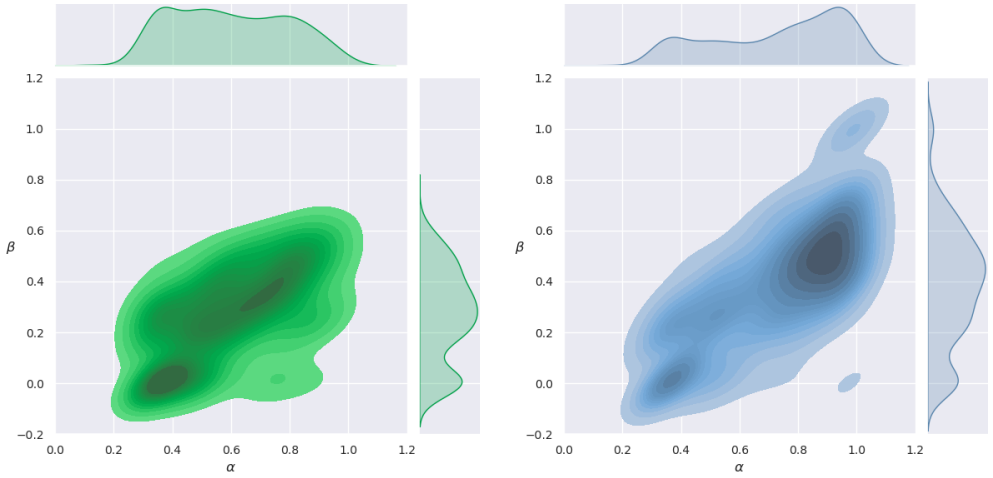

(a) $A_i \geq 0.25$ (b) $A_i < 0.25$

Figure 4.13: Density of participants at different points on the $\alpha\beta$ plane, separated by their contribution to consensus as defined in Equation (4.5). While individuals who are more agreeable have lower parameter values in general, we see a high density of people emerging near $(0.7, 0.4)$. This is the case because the EC strategy does not hinder contribution to consensus (though it may directly influence the number of points scored). We also note that contribution is based only on agreement over the second half of the experiment, which is why higher α values can produce agreeable people.

Table 4.1: Energy distance and p-values for comparing the distributions of the indicated subsets of individuals from the live experiments.

Comparison	**Energy Distance**	**P-value**
Face vs. Hashtag	0.1391	0.00000
Homogeneous vs. Spatial	0.0082	0.00125
Homogeneous Face vs. Spatial Face	0.0385	0.00000
Homogeneous Hashtag vs. Spatial Hashtag	0.0021	0.44903
Homogeneous Face vs. Homogeneous Hashtag	0.0763	0.00000
Spatial Face vs. Spatial Hashtag	0.2228	0.00000
High Pts vs. Low Points	0.2217	0.00000
High Contribution vs. Low Contribution	0.0680	0.00000

The energy distance and significance between the subsets of individuals discussed in Figures 4.9—4.13 are included in Table 4.1. Likewise, the relationship between points scored and contribution to consensus split according to network structure and task can be seen in Figure 4.14.

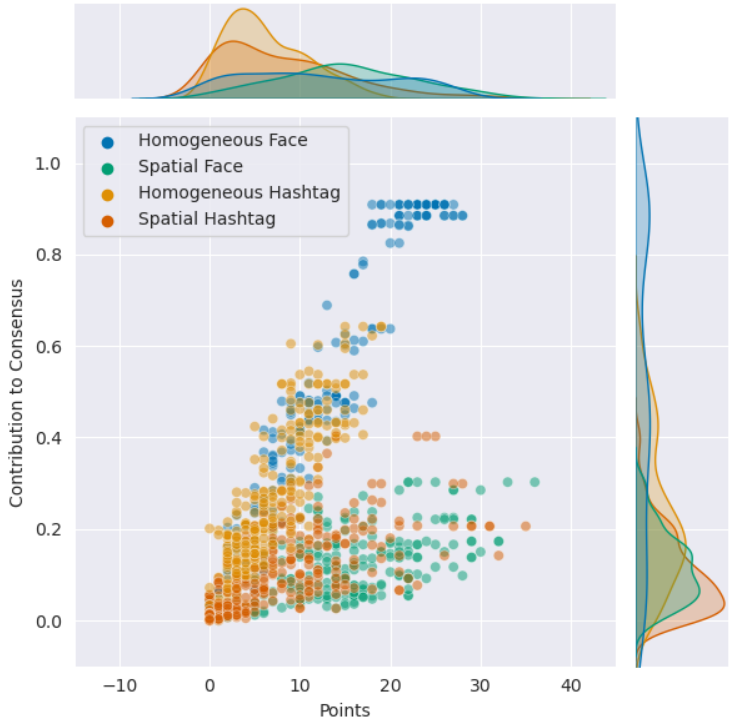

Figure 4.14: Number of points scored plotted against contribution to consensus and colored according to structure/task condition. Spatial participants tend to score more points, while homogeneous participants tend to have higher contribution to global consensus, a direct result of the two network structures. Also note that face/name task participants tend to do better than hashtag task participants with respect to these metrics, a consequence of the reduced complexity of the task.

Ultimately, the most significant distances between parameter distributions occurred between those who participated in the face/name task and those who participated in the hashtag task (see Table 4.1). For this reason, the simulations that utilized the CAPL decision-making algorithm involved random draws from these separate distributions for runs on the different priors. Importantly, the network structure did not (for either task) result in a draw from a separate distribution, even though we noticed a statistically significant difference between the homogeneously mixed face/name and the spatially embedded face/name participants.

This was to avoid overfitting to the observed data, and supported by the low energy distance between the two distributions. We note that because weighted points scored impact how participants utilize certain strategies *in addition* to their parameter values, this difference should emerge as a result of the structure, rather than something necessary to impose via the parameter values.

4.2.2 CAPL Simulations

As we did for the other models, we used the CAPL decision-making structure to run a series of simulations for all four experimental cases. Preliminary results of these simulations can be seen below, including entropy and response behavior in Figures 4.15 and 4.16, respectively.

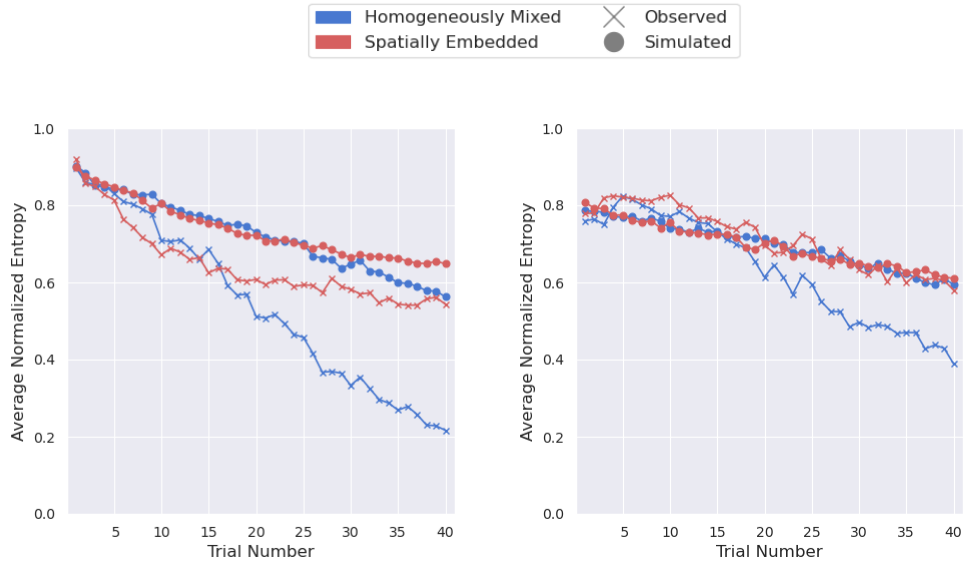

(a) Face/Name Task

(b) Hashtag Task

Figure 4.15: Observed average entropy and simulated average entropy for (a) the face/name task and (b) the hashtag task, both separated according to network structure. For each case, 40 simulations were run: 20 for each $N = 20$ and $N = 50$. The CAPL model performs relatively poorly with respect to entropy decrease, not replicating the difference based on network structure well. Since this model was optimized according to decision-type behavior while the CACL model parameters were based on entropy, we would expect the CACL model to do better with regard to entropy decrease.

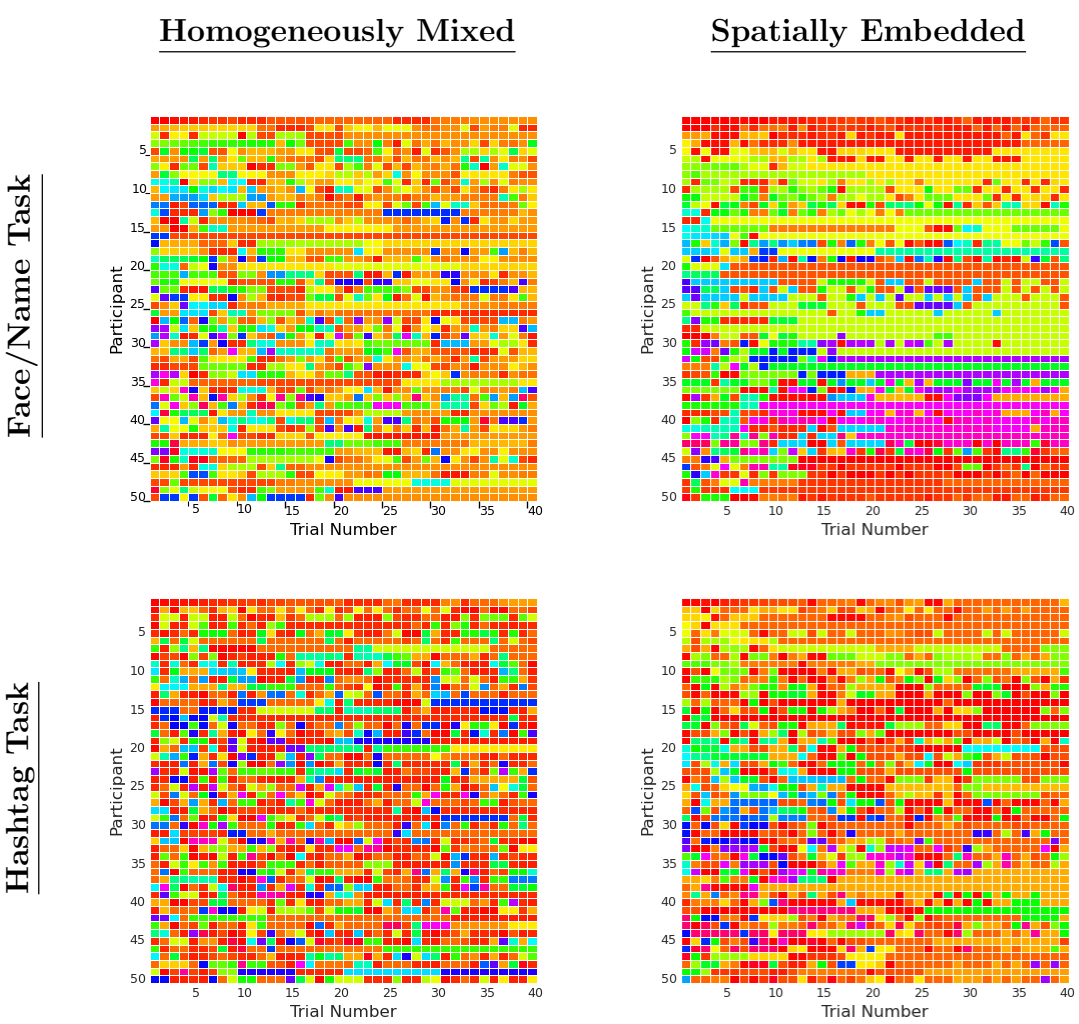

Figure 4.16: For each structure and task combination, a color-map shows how the *responses* change on the network over time for a selected $N = 50$ simulation. Each row represents a participant and each column is a trial, with cells colored according to the response given. This model produces less of the qualitative features from the experiments than either the CACL or CB model.

At the network level, we note that the preliminary results of these simulations are not as promising as the results of simulations that relied on collective learning. Note that, like the CACL model, the CAPL model is context-aware, and is thus able to recreate the difference in starting entropy between the two tasks by generating the initial responses of simulated agents from the empirical prior distributions. It is clear, however, that entropy decrease over time is not well facilitated by this model. Similarly, with respect to Figure 4.16, we note that the overall distribution of responses in homogeneously mixed networks does not

begin to exhibit the same coordination seen in the experimental data or either of the previously discussed models. The localized coordination in the spatially embedded cases is apparent, but to a limited extent. This can be attributed to the fact that the CAPL parameter fitting was based on optimizing *individual decision-types*, in contrast to the CACL case where α values were chosen to best fit observed entropy trends.

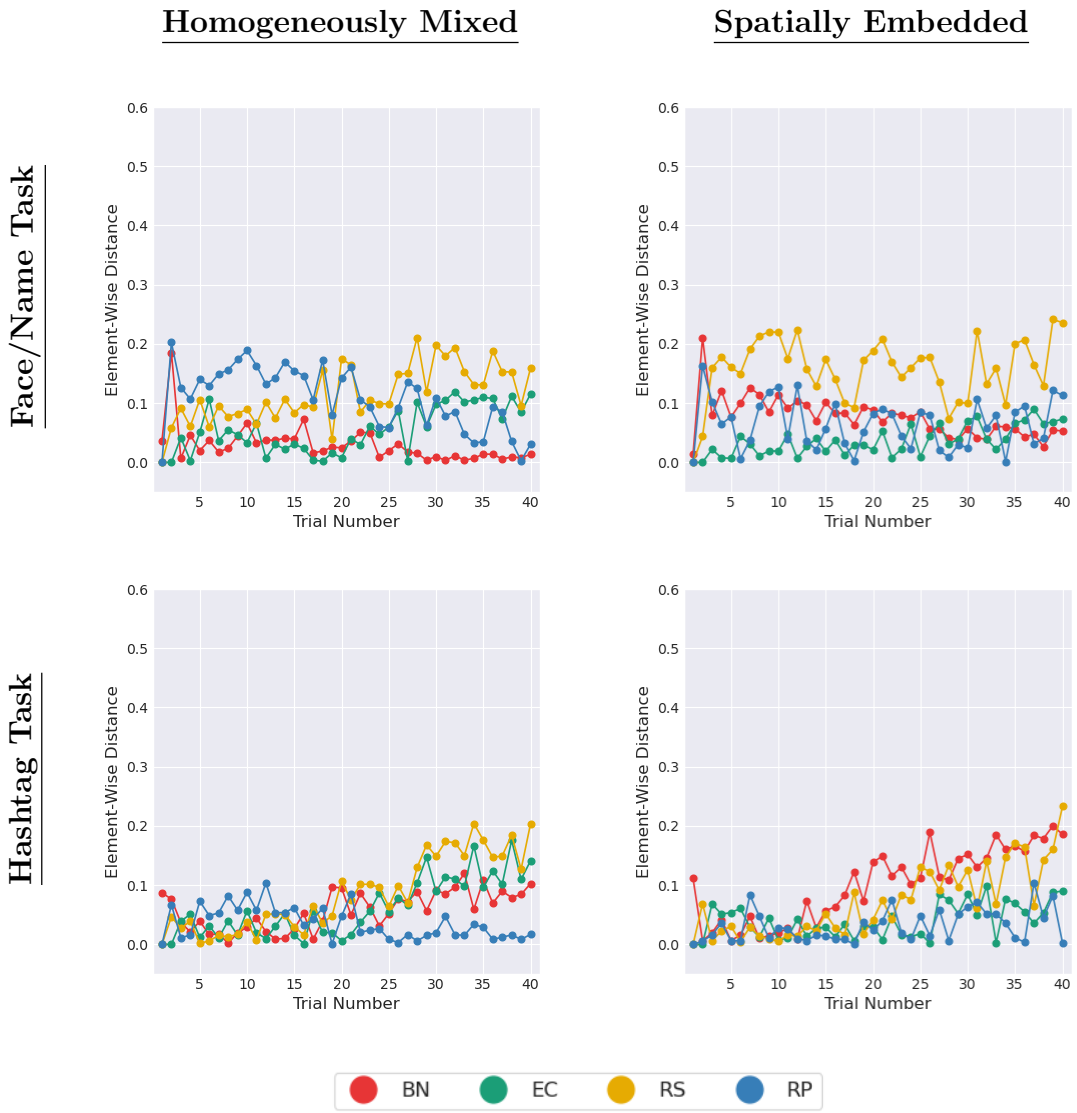

Figure 4.17: For each structure and task combination, line graphs show the vectors **D** which illustrate the element-wise distance between proportions of networks engaging in each decision-type over time in the observed and simulated (CAPL) cases. In general, we see that the CAPL model performs better than either the CB or the CACL model with respect to decision-type behavior, as can be expected.

Figures 4.17 and 4.18 show the decision-type behavior of CAPL simulations. In Figure 4.17 we see that the behavior of the simulations with respect to the proportion of the networks using each decision type at a given trial number is truer to the observed behavior. In addition, the CAPL model is able to capture

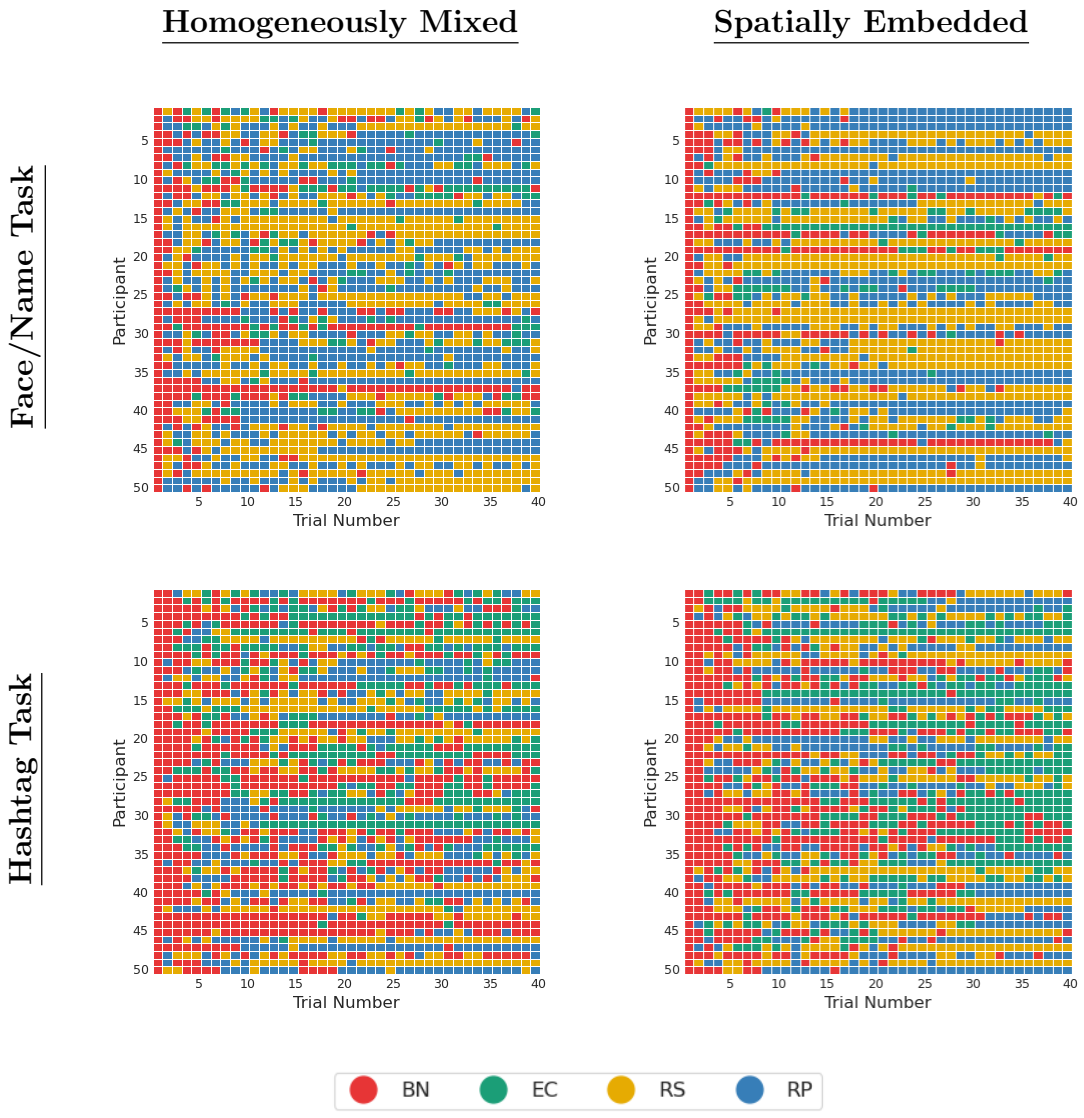

Figure 4.18: For each structure and task combination, a color-map shows which *strategies* are being used by agents through the course of the CAPL simulation. Note that the stubbornness of individuals in the experimental case as noted above is appropriately reflected by the CAPL simulation.

the presence of stubborn, non-learning individuals in all of the different cases, which better encapsulates the complexity of how these different decision types are distributed across the network (as seen in Figure 4.18).

4.3 Context-Aware Modeling Overview

Our efforts to develop a task-aware model were successful and provided a number of insights with respect to participant behavior given a task. Through collective-learning, we were able to produce global behaviors similar to those of the experiments, particularly with respect to propensity toward agreement over time in the four different experimental cases. To improve the performance of the context-aware model with regard to individual use of decision strategies, we integrated personalized learning. While the individual-learning case produced greater nuance in the strategies used by different participants, some of the global dynamics became less visible. This result indicates that a hybrid context-aware model (one which involves network-level and personal tuning) may provide further insight into the interaction between these two scales of dynamics.

Chapter 5

Conclusion

5.1 Discussion & Implications

Our investigation revealed several characteristics of networked human interactions and has yielded promising preliminary results for simulations of such behavior. The experiments are built on the findings of Centola and Baronchelli (2015) concerning the face/name task [3], and other conclusions can be drawn from the addition of the hashtag task. Notably, our modeling efforts extend those of Centola and Baronchelli, particularly concerning decision-type behavior, which has been neglected in the existing literature. Both of our decision-making algorithms were able to more closely replicate the decision-types as witnessed on the experimental networks, which indicates a promising future for the simulation of interaction in the case of varying task-relevant background information and individual decision rigidity. Implications of such modeling capability include the opportunity to understand the formation of echo chambers [15], particularly in online interaction networks with limited regulation.

Importantly, the decision-type structure we developed reveals differences between behavior in the two tasks, which gives vital information about human content reproduction in the age of polarized information exchange. The ability to model decision-making in a content-sensitive way is imperative as we try to understand how online interactions breed offline phenomena.

While our models proved more accurate with respect to individual behavior than existing literature, we did not see the same trends toward consensus formation in all of our simulations as we did in the experimental results. This indicates that perhaps there are other, more complex, drivers of decision-making

when presented with such tasks. For instance, our observations of the experimental data indicate that people with the hashtag task were more exploratory than in the face/name task [11] which is indicated by the slower decrease in use of the BN strategy in the hashtag cases and also by the less prominent decrease in entropy over time. These differences again underscore the importance of content-sensitivity in such tasks, which has natural extensions to real-world information exchange.

A couple of reasons for this seem plausible and should be investigated, including that in cognitively more complex tasks (i.e., the hashtag task), participants may develop a stronger belief that some responses are correct because there is a perceived intellectual assessment built into the task. In these cases, participants may be more likely to gather and combine information, building upon their understanding of such content rather than abandoning a response in favor of conforming to someone else, given a financial incentive for agreement. Faster entropy decrease in the face/name tasks can be explained by a large-scale abandonment of information that is foreign to the network in favor of more immediately profitable strategies (in particular, repeating yourself and repeating your partner). Importantly, we see a statistically significant difference between CAPL parameter distributions for participants in the face/name task based on the network structure, even though the network structure is concealed from participants. A potential implication of this is that interaction structures become more perceptible when group-conscious decisions are made faster because there is quicker agreement among the neighbors of a particular node, but the neighborhood in the homogeneously mixed network is much larger $(N-1)$ than in the spatially embedded case (4). For this reason, we see lower mean values of both α and β for the spatially embedded face/name case than the homogeneously mixed case, as the amount of new information participants are gaining more quickly decreases. This effect is only magnified by the above argument that there is no intellectual value encoded in the responses people give in the face/name task, so they are more susceptible to conformity.

Ultimately, these understandings imply that the problem of modeling contentious human interactions on large online networks is difficult not only because of network topology, but also because of the way decisions are formed when people develop stronger personal connections to opinions. We do gain an understanding that informational content is important, which provides possible directions for future experiments that may be able to shed more light on information exchange.

5.2 Directions for Future Work

Additional insight can still be gained from these experiments and simulations that will give us a deeper understanding of human information exchange. Future work should focus on improving simulations and formalizing models that can be extended to other contexts. The simplicity of the binary choice voter models outlined by Redner (2019) provides an appealing theoretical basis for the understanding of content-specific opinion formation rather [14]. Building on this work, we can hope to gain a theoretical understanding of how network opinions evolve and the important drivers behind that evolution. Importantly, starting with theory can provide insight into how these simpler models could be used as a tool for understanding how the tradeoff between being *correct* and being *agreeable* manifests in large group interactions and online.

Based on the results of the context-aware models we have developed we expect that a hybrid model (perhaps a localized learning model or a model where parameter values are time-sensitive) may be more able to capture all of the important features of the experimental results and thus easier to extend to additional contexts.

We argue that in the hashtag case, a perceived desire to be correct drove longer exploratory periods and less immediate consensus formation as networked individuals engaged in an intellectual information exchange. In contrast, face/name task participants engaged in much more agreeable decision strategies, and ultimately, those networks were quicker to converge to fewer responses. As political parties become more polarized and intellectual discourse wanes [7, 9], it would be interesting to study how these ideas play out given a task that involves an issue of high political visibility that people may have opinions on already (for example, green energy or abortion issues).

In addition, future experiments and models should involve more complex network structures that are able to replicate more closely the ways information is shared online. For example, we note that information exchange online is often not bidirectional and "influencer" nodes should be included that broadcast information in one direction. Similarly, it would be appropriate to include weights for the trustworthiness of information coming from different sources and how much people tend to espouse content from those sources.

Understanding and replicating human decision-making and beliefs is important in an era where com-

plex interactions are developing and polarizing content is emerging as central to our discourse. Through modeling efforts, we can provide a theoretical basis for understanding how misinformation and negative ideas can spread and effective efforts for intervention on a large scale in an internet space that has become widely unregulated [17].

Bibliography

[1] G. Askari, M. E. Gordji, and C. Park, The behavioral model and game theory, Nature, (2019).

[2] D. Boyd, S. Golder, and G. Lotan, Tweet, Tweet, Retweet: Conversational Aspects of Retweeting on Twitter, in 2010 43rd Hawaii International Conference on System Sciences, 2010, pp. 1–10.

[3] D. Centola and A. Baronchelli, The spontaneous emergence of conventions: An experimental study of cultural evolution, Proceedings of the National Academy of Sciences, 112 (2015), pp. 1989–1994.

[4] R. Cressman and Y. Tao, The replicator equation and other game dynamics, PNAS, (2014).

[5] P. Dawson, Hashtag narrative: Emergent storytelling and affective publics in the digital age, International Journal of Cultural Studies, 23 (2020), pp. 968–983.

[6] G. Gigerenzer and W. Gaissmaier, Heuristic decision making, Annual Review of Psychology, 62 (2011), pp. 451–482.

[7] M. Hannon, Public discourse and its problems, Politics, Philosophy, & Economics, 22 (2023), pp. 336–356.

[8] G. Klein, Naturalistic decision making, Human Factors, 50 (2008), pp. 456–460.

[9] M. Leoputre, Democratic speech in divided times, Oxford University Press, 2021.

[10] H. Park and M. K. S. Faradonbeh, Efficient algorithms for learning to control bandits with unobserved contexts, IFAC Papers Online, 55 (2022), pp. 383–388.

[11] J. Priniski, B. Linford, A. Hirschmann, S. Krishna, F. Morstatter, N. Rodriguez, J. Brantingham, and H. Lu, Cognitive complexity and network structure shape belief dynamics in networked groups, submitted to Nature Human Behavior, (2025).

[12] J. H. Priniski, M. McClay, and K. J. Holyoak, Rise of qanon: A mental model of good and evil stews in an echochamber, arXiv preprint arXiv:2105.04632, (2021).

[13] J. H. Priniski, N. Mokhberian, B. Harandizadeh, F. Morstatter, K. Lerman, H. Lu, and P. J. Brantingham, Mapping moral valence of tweets following the killing of george floyd, arXiv preprint arXiv:2104.09578, (2021).

[14] S. Redner, Reality-inspired voter models: A mini-review, Phisique, 20 (2019), pp. 275–292.

[15] K. Sasahara, W. Chen, H. Peng, G. L. Ciampaglia, A. Flammini, and F. Menczer, Social influence and unfollowing accelerate the emergence of echo chambers, Journal of Computational Social Science, 4 (2021), pp. 381–402.

[16] T. R. Tangherlini, S. Shahsavari, B. Shahbazi, E. Ebrahimzadeh, and V. Roychowdhury, An automated pipeline for the discovery of conspiracy and conspiracy theory narrative frameworks: Bridgegate, Pizzagate and storytelling on the web, PLOS ONE, 15 (2020).

[17] S. A. Thompson and K. Conger, Meet the next fact-checker, debunker, and moderator: You, The New York Times, (2025). Available at: `https://www.nytimes.com/2025/01/07/technology/meta-facebook-content-moderation.html` (Accessed: February 24, 2025).

[18] A. Tversky and D. Kahneman, Judgment under uncertainty: Heuristics and biases, Science, 185 (1974), pp. 1124–1131.

Appendix A

Fukushima Nuclear Disaster Narrative

The Fukushima Nuclear Disaster was a 2011 nuclear accident at the Daiichi Nuclear Power Plant in Fukushima, Japan. The cause of the nuclear disaster was the Tōhoku earthquake on March 11, 2011, the most powerful earthquake ever recorded in Japan. The earthquake triggered a tsunami with waves up to 130 feet tall, with 45 foot tall waves causing direct damage to the nuclear power plant. The damage inflicted dramatic harm both locally and globally.

The damage caused radioactive isotopes in reactor coolant to discharge into the sea, therefore Japanese authorities quickly implemented a 100-foot exclusion zone around the power plant. Large quantities of radioactive particles were found shortly after throughout the Pacific Ocean and reached the California coast.

The exclusion zone resulted in the displacement of approximately 156,000 people in years to follow. Independent commissions continue to recognize that affected residents are still struggling and facing grave concerns. Indeed, a WHO report predicts that infant girls exposed to the radiation are 70% more likely to develop thyroid cancer.

The resulting energy shortage inspired media campaigns to encourage Japanese households and businesses to cut back on electrical usage, which led to the national movement *Setsuden* ("saving electricity"). The movement caused a dramatic decrease in the country's energy consumption during the crisis and later inspired the Japanese government to pass a battery of policies focused on reducing the energy consumption of large companies and households.