Collaborative meaning-making: the emergence of novel languages in humans, machines, and human-machine interactions

Kouwenhoven, T.

2

Grounding and the Need for Structure

Language is a unique hallmark of humans, it is both learned and symbolic, which poses the problem of emergence: if neither form nor meaning is known, how can individuals communicate in the first place? The current study replicates work that investigates the emergence of signal forms and meanings and explores how Personal Need for Structure (PNS) of interacting partners can aid or hinder the emergence of communicative systems. We use an existing measure of personal need for structure to investigate its relationship with the emergence of such systems while participants play the embodied communication game (ECG). Similar to the original study, our work shows that a bootstrapping process and sufficient common ground are integral to the recognition of signalhood. Moreover, this process appears to be more successful for individuals who respond differently to a lack of structure compared to their interaction partners. Contrary to what is usually assumed, our results indicate that not only shared expectations and biases seem to matter in communicative tasks, but that diversity in biases of communication partners can also be beneficial for the emergence of new communication systems.

Originally published as: Kouwenhoven, T., de Kleijn, R.E., Raaijmakers, S.A., Verhoef, T.(2023). Need for Structure and the Emergence of Communication. In J. Culbertson, A. Perfors., H. Rabagliati. & V. Ramenzoni., editors, *Proceedings of the Annual Meeting of the Cognitive Science Society*, Volume 44, pages 549-555. Cognitive Science Society.

2.1 Introduction

Humans can share and accumulate knowledge through language, enabling them to pass this knowledge on to future generations. Communication through language can be formulated as the joint action that emerges when speakers and listeners perform actions in coordination (Clark, 1996), and uses signals that are both symbolic and learned. The emergence of signals is therefore a defining event in human cognitive evolution. However, the exact dynamics of language emergence—the settling of two individuals on an effective interchange through discrete, grounded symbols—is complex and not yet fully understood (Tylén et al., 2013; Scott-Phillips and Kirby, 2010). If form and meaning are unknown, one fundamental question concerns the cooperative process of agreeing on what form should refer to what meaning (Oliphant, 2002). This process has been studied quite extensively through laboratory experiments in which participants need to invent and negotiate novel signals to solve a communicative or cooperative task (Steels, 2006; Scott-Phillips and Kirby, 2010; Tylén et al., 2013). A general finding from such studies is that participants are able, through social coordination, to establish conventions and gradually develop a communication system. Consistently, researchers report on the importance of common ground and the reliance on shared biases and expectations between interacting partners on the road to success. However, building an entirely novel system of signals from scratch is not easy, and in such experiments, it is often the case that not all pairs manage to solve the game. Analyses tend to focus on the conventions established in successful games, which have generated many insights, but we propose that a focus on differences in coordination outcomes and properties of the individuals involved can help to understand these dynamics better. In this chapter, we show how sometimes diversity rather than alignment of initial cognitive biases and preferences of individuals might positively influence success in the social coordination of a shared language.

In essence, the emergence of signals can be formulated as a cooperation problem, where individuals have a common goal and need to figure out how to influence each other in an *initially unstructured* environment. It has been proposed that the emergence of language is influenced by human biases to prefer compressible, simple systems (Kemp and Regier, 2012; Kirby et al., 2015). Such a bias can, for example, drive the emergence of systematic structure over generations of transmissions (Kirby et al., 2015). Individuals have been found to differ in their *personal need for structure* (Neuberg and Newsom, 1993) which can affect problem-solving capabilities such as solving maths problems (Svecova and Pavlovicova, 2016) and learning a foreign language or text comprehension (Eva et al., 2014). As such, the social coordination of a shared language, which is initially unstructured, can potentially also be influenced by an individual's personal need for structure. We expect that PNS might also affect how individuals act in language emergence tasks, and investigate how a personal need for structure affects the evolution of a communication system that is created de novo.

Specifically, the experiment presented in this chapter was designed to study the relationship

between personal need for structure as measured by the PNS questionnaire (Neuberg and Newsom, 1993), its F1 and F2 sub-factors and the emergence of a communicative system while playing the Embodied Communication Game (Scott-Phillips et al., 2009), which is described in detail in the next section.

2.2 Background

The current study is based on an experiment designed by Scott-Phillips et al. (2009), who investigated the emergence of newly created communication systems when humans are not able to communicate verbally, or in any other conventional way. Participants played the ECG, a cooperative game, and the results revealed how signals acquire informative meaning without pre-defining a communication channel, roles of signaler and receiver, or a form space.

2.2.1 The Embodied Communication Game (ECG)

The ECG is a cooperative two-player game that consists of two 2×2 grid worlds, where players are embodied in the sense that they are given a physical form (a black square) to move around with. Each quadrant has one of four colours (red, green, yellow, blue), which is determined at random. The goal of the participants is to end on identically-coloured quadrants and, if they do, score a point. Players can move within their own grid and see movements in both grids, but can only see the colours of their own quadrants, showing the others' quadrants as grey (Figure 2.1a). Once finished moving, the colours of all quadrants are revealed to both players (Figure 2.1b) as a means of feedback. The colours of the quadrants and starting positions of both players are randomly chosen with the proviso that there is always one overlapping colour between both worlds, such that it is always possible to score a point. Players are informed that their goal is to score as many consecutive points as possible, meaning that players cannot win by playing many games but must instead find a way to communicate reliably and coordinate behaviours with each other (see Scott-Phillips et al. (2009) for a more elaborate explanation).

The setup of this experiment required participants to coordinate their behaviours by agreeing on what behaviours correspond to what meaning, and they had to find a way to signal that these behaviours were of communicative intent. Crucially, this problem can be solved when movements between the quadrants eventually come to be understood as communicative. It turned out to be a non-trivial task since only 7 out of 12 pairs managed to co-opt one's movements for the purpose of communication. Scott-Phillips et al. (2009) conclude that the problem of mapping form onto meaning is solved by finding sufficient common ground and bootstrapping new meanings upon that. As such, the authors suggest that the latter significantly increases the likelihood that a symbolic communication system emerges and that the emergence of dialogue is a crucial step in the development of a system that can be employed to achieve shared goals.

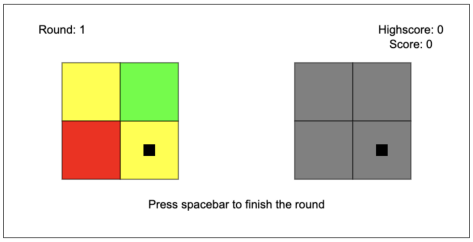

(a) The view while participants are playing.

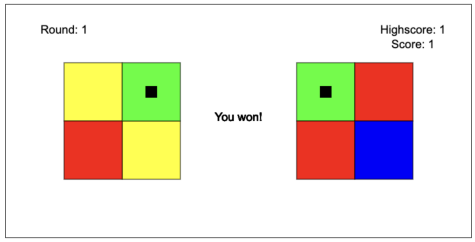

(b) The view after both players ended the round.

Figure 2.1: The game environment, figure (a) shows the view while players are moving, where movements from both, but only the colours from the participants' own world are visible. Figure (b) shows the environment after both players are done with their movements. The colours of all quadrants are revealed to both players as a means of feedback.

2.2.2 Successful interactions and shared expectations

Many studies have involved the experimental emergence of artificial languages, where participants are not permitted to use conventional language systems (e.g. Steels, 2006; Scott-Phillips and Kirby, 2010; Tylén et al., 2013). A task that is somewhat related to the ECG was studied in an experiment by Galantucci (2005). Here participants played a collaborative computer game and were required to develop new semiotic conventions, which map signals and meanings, to communicate information regarding their location using a novel communicative channel. Similar to the findings of Scott-Phillips et al. (2009), not all pairs succeeded in this task. Moreover, pairs who did succeed differed widely in the manner and rate at which they managed to solve the game. Success in such tasks is typically attributed to feedback, alignment, shared biases, and similarities between pairs; however, a specific focus on the underlying mechanisms that allow some pairs to converge on a system while others can not achieve this is lacking. We are interested in precisely these dynamics and investigate how the diversity of preferences and biases in pairs influences collaborative tasks.

2.2.3 Personal need for structure

Individual differences in the desire for structure may influence how people understand and interact with their worlds. This desire can be measured by means of the Personal Need for Structure Scale (Neuberg and Newsom, 1993). It consists of 12 statements (e.g. "I enjoy having a clear and structured mode of life") that are answered on a 6-point Likert scale, which measure the tendency to seek structure in chaotic environments. It is characterised by a representation of simplified information and generalisation of previous experience into fewer complex categories that an individual uses in new and ambiguous situations (Svecova and Pavlovicova, 2016). Two conceptually different sub-factors are identified: the desire for structure in unstructured

environments (*F1*) and an individual's response to the lack of structure (*F2*).

2.3 Current study

As mentioned before, reports on cooperative games and the emergence of communication often emphasise the importance of common ground and the reliance on shared biases and expectations between interacting partners. However, we expect that differences can also play a role as interacting partners that differ might complement each other's shortcomings, which possibly aids cooperation. Arguably, the initial states of the ECG can be considered as an unstructured environment and thus may evoke different responses in humans that differ in PNS. We investigate precisely how PNS might affect the evolution of a communication system that is created de novo.

2.4 Methods

Participants ($N = 40$: 31 females, 9 males; $M_{\text{age}} = 22.12$, $SD_{\text{age}} = 3.56$) were recruited via two methods: the participant recruitment website from the Psychology department of Leiden University, and by the experimenters during lectures or other events. As a result, 20 pairs played the ECG. Upon arrival, they were given instructions about the experimental procedure and seated behind a computer in two separate rooms. The entire experiment took place on two connected computers via a web application. Participants then read instructions explaining the goal of the game, its mechanics, and were allowed to ask clarifying questions solely concerning the mechanics. This setup ensured that no conventional communication was possible and that the problem of signalling signalhood had to be solved by the participants themselves. The pairs then played the game for 40 minutes uninterrupted for, on average, 255 rounds. Both players could move between the centres of each of their own quadrants using the arrow keys and finalised their movements with the spacebar, after which both players received feedback on their performance (Figure 2.1b) and continued to the next round. The game was stopped after 40 minutes. Participants then completed the PNS questionnaire and reported whether they thought that any communication had occurred. If any, they described the communication systems they developed or attempted to develop. Finally, they were debriefed and given the opportunity to discuss their experience. This study was approved by the Psychology Research Ethics Committee of Leiden University.[1]

[1] All code, materials, and data are available on OSF: `https://osf.io/n3uj6/`.

2.4.1 Measures

Game performance was measured using a *score* and *high score*. The *score* was increased by one point when both players ended on a quadrant with identical colours. The *high score* represents the number of consecutive successful rounds. *PNS* and its sub-factors were measured using a survey of 12 statements (see Neuberg and Newsom (1993) to see all statements), where the sum of all answers defines PNS; a higher sum corresponds to a higher *need for structure*. Here, items 3, 4, 6, and 10 correspond to sub-factor *F1* (i.e., the desire for structure in unstructured environments) and items 1, 2, 4, 7, 8, 9, 11, 12 sum to *F2* (i.e., the response to the lack of structure). Finally, participants described the communication system they developed via three open questions. We manually cross-checked the post-game descriptions, in which the participants described their communication systems, with the corresponding game data to validate whether both players reported identical systems, and to identify emerging patterns.

2.5 Results

Statistical analyses were performed using *R* 4.1.0 (R Core Team, 2023) and the *BayesFactor* 0.9.12-4.2 package (Morey et al., 2018). Our results align with those of Scott-Phillips et al. (2009), who found that out of 20 pairs, only 11 pairs managed to create a robust communicative system, confirming that this is not a trivial task. Participants perform on average 6.87 moves ($SD = 5.86$) per round and obtain a mean high score of 29.9 ($SD = 31.4$).

2.5.1 Emergence

The emergence of communicative systems happened in a similar manner to what was reported by Scott-Phillips et al. (2009); hence, we refer the reader to their work for a more elaborate description. Successful pairs typically converged on a default colour, allowing them to score above chance levels. This happened for 12 out of 20 pairs (note that one pair was not able to further develop a communication system beyond a default colour). Logically, this strategy failed when, in some of the following rounds, the default colour was not present in the quadrants. In such cases, a new convention had to be formed. Players typically did so by moving between quadrants in initially random directions. An initial convention was formed if these behaviours were recognised as communicative signals. Specifically, these random movements between quadrants could be recognised by the interlocutor as a communicative signal (e.g. “No, not the standard colour”). If the interlocutor did recognise this and, by mere chance, both players would finish on identically coloured quadrants, the initially random movements could be recognised as communicative signals and, henceforth, mapped to the finishing colour. From here, players could bootstrap their signalling behaviour when there were no colours available for which a signal exists and establish new signal-meaning mappings. These elaborate behaviours quickly

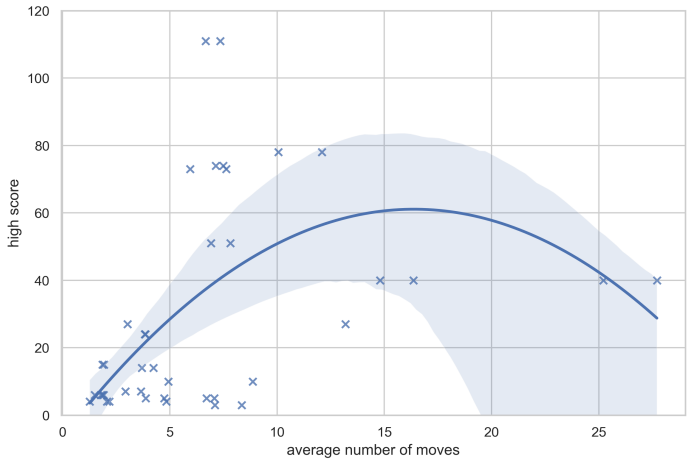

Figure 2.2: The quadratic relationship between the average number of steps over all rounds and high score.

became symbolic signals that participants explicitly recalled in their reports. The timing of convergence on a default colour was crucial towards a high score; pairs that quickly settled on a default colour typically evolved more elaborate and robust systems. Since the action space was rather limited, we observed patterns that are similar to the study by Scott-Phillips et al. (2009), namely oscillating up and down between quadrants, moving in (anti-)clockwise circles, forming U-shapes, or L-shapes. Despite having much overlap between communicative signals across experiments, the mapping to different meanings, i.e., colours, was specific to each experiment. Hence, the evolved systems were idiosyncratic to the pairs that evolved them and consequently would not be useful to immediately communicate successfully with new unseen partners. An example system of a successful pair is as follows: red was the standard colour, move there and wait for other signals. Moving in anti-clockwise circles indicated green, yellow was signalled by clockwise circles, and horizontal oscillations indicated blue.

Successful pairs agreed on a colour through dialogue. In a typical dialogue, one player initiated a signal after which the other copied it to confirm that colour. However, when that colour was not available, the recipient became the signaller and suggested another colour by using its corresponding signal. Such behaviour continued until both players agreed on a certain colour and finished the round. This robust system enabled participants to communicate successfully and gain high scores. We found that this is also reflected in the average number of moves participants made, where dialogue, quantified by the mean number of moves, has a quadratic relationship to higher scores, $F(2, 37) = 7.29, p = .002, R^2 = .28, R^2_{adjusted} = .24$ (see Figure 2.2). We also tested a linear relation between dialogue and *high score*, but found that this resulted in a lower fit ($F(1, 38) = 5.24, p = .02, R^2 = .12, R^2_{adjusted} = .09$). Moreover,

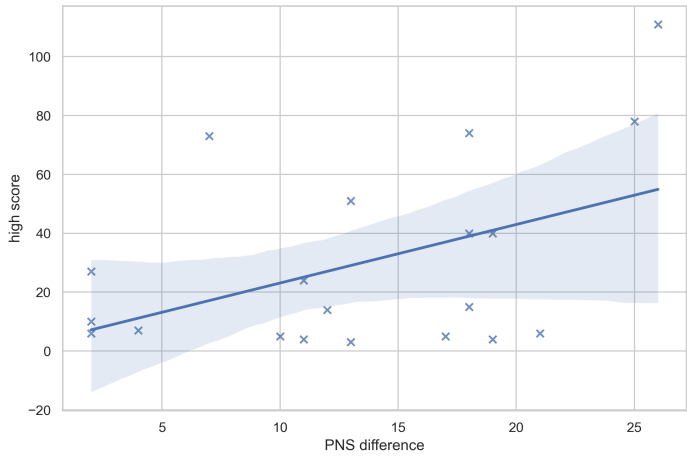

Figure 2.3: Pairs' difference in PNS score positively influences high score.

the quadratic relation remains best when outliers, the two points larger than 25 moves, are removed $R^2_{adjusted} = .28$ for quadratic regression and $R^2_{adjusted} = .19$ for linear regression. Taken together, this suggests that there appears to be an optimum number of moves: too few movements cannot convey communicative content, while too many movements can become confusing.

The reports of non-successful pairs typically describe that at least one participant tries to stick to their own system, not paying attention to the behaviours of the other. In some cases, participants even report having actively tried to communicate, whilst realising that their teammate did not notice and thus decided to unsuccessfully submit to their dominance. This is not trivial and often fails. This again shows that settling on conventions and the emergence of a communicative system requires *all* members to cooperate and interact actively.

2.5.2 Need for structure

Simple linear regression showed no relation between *PNS* ($M = 41.8, SD = 8.78$), *F1* ($M = 15.4, SD = 3.48$), *F2* ($M = 26.4, SD = 6.53$), and *high score* or the average number of moves on an individual level. However, the ECG enforces team cooperation of both players; we therefore combined individual scores to calculate team scores and assess team performance. We computed the difference in PNS between the two participants, and Figure 2.3 reveals that pairs with individuals that have a large difference in PNS score higher, $F(1, 18) = 4.869, p = .041, R^2 = .21, R^2_{adjusted} = .17$. This means that partners that respond differently to chaotic environments perform better in the ECG than those that have both either a high or low personal need for structure.

2.5.3 Comparing teams

As mentioned earlier, not all pairs were able to form a robust communication system and successfully convey their intentions. To further investigate why some are successful and some are not, we labelled games based on self-reports that describe the communication system that was used. After playing the game, participants individually reported on the communication system they thought was present, and the answers to these questions were cross-checked between pairs and used to split the pairs into groups. Teams were labelled as *good* ($n = 11$) when both participants individually reported identical signals for the same colours. They are labelled *medium* ($n = 3$) when there was partial overlap or when there was only a default colour, and *bad* ($n = 6$) otherwise. An analysis of variance (one-way ANOVA) showed that the mean high scores of these groups were significantly different, $F(2, 17) = 7.91, p = .004$. When we combined *medium* and *bad* performing pairs to have roughly equal sample sizes, the mean high scores were again significantly different, $t(18) = 4.07, p < .001$. This is expected because when two players can both recall the same systems, communication was probably successful in many consecutive rounds.

Although Figure 2.2 shows that pairs which use more movements do not necessarily reach higher scores, when comparing the two groups we do see that teams that performed well in the ECG, on average, moved more than those who performed worse ($M_{\text{good}} = 9.50, SD_{\text{good}} = 6.58, M_{\text{bad}} = 3.65, SD_{\text{bad}} = 2.29, t(27) = 3.89, p < .001$). This supports the assumption that well-performing teams have sufficient dialogue, which indicates that a pair can be considered to have a robust system (Scott-Phillips et al., 2009).

Figure 2.4 shows the correlations between team measures for pairs labelled as *good* and *medium* or *bad* performing pairs. A significant relationship between the difference in PNS scores and the high score is present for *good* teams ($r = .693, p = .018$). Since *PNS* is the sum of *F1* and *F2*, it allows us to investigate the main contributor to this effect. Differences in desire for structure (*F1*) do not explain higher scores ($r = -.107, p = .754$), yet differences in the response to the lack of structure (*F2*) do ($r = .78, p = .004$).

A Bayesian test for correlation between *PNS* difference and *high score* on good performing pairs yielded $\text{BF}_{10;\kappa=2} = 3.69$, indicating *PNS* difference positively influences the high score. For *F2* difference and high score, $\text{BF}_{10;\kappa=2} = 13.25$, confirming that a greater difference in the response to a lack of structure predicts higher high scores. We did not find these relationships in the group of medium and bad-performing pairs. This could be expected since high scores, in general, were lower for these pairs. Figure 2.3 shows that, although pairs with the largest differences in *PNS* or *F2* tend to score the highest, a relatively large difference in *PNS* or *F2* does not necessarily lead to a higher *high score*. We also observe pairs with a medium difference in *PNS* or *F2* that do not perform better than the lowest-scoring pairs in general. This indicates that diverse reactions to chaotic environments may be beneficial in establishing communication systems, but it does not guarantee success.

Figure 2.4: Comparison of the relations between all measures for good performing pairs (left) and medium or bad performing pairs (right). For good pairs, there is a positive correlation between the difference in Personal Need for Structure (*PNS_dif*), difference in the response to the lack of structure (*F2_dif*) and the obtained *high score*. These relations are not present for medium or bad pairs. In both groups *F2_dif* correlates with *PNS_dif*, while *F1_dif* does not, indicating that *F2_dif* is the main contributor of the relation between *PNS_dif* and *high score*. Note: the colour represents the correlation coefficient and the annotations correspond to p-values.

2.6 Discussion

In this chapter, we describe an experiment in which participants played the Embodied Communication Game from Scott-Phillips et al. (2009) and replicated their findings, while also introducing a novel way of comparing differences in game success. Paired participants had a shared goal without having access to conventional means of communication. This meant that they had to create a novel communication system that allowed them to coordinate their intentions. This non-trivial cooperation problem was typically solved through the formation of initial conventions (common ground) and a bootstrapping process. We extended the original work by incorporating a measure that allowed us to compare the cognitive traits of cooperating individuals. Results showed that a difference in personal need for structure between partners influenced the emergence of the communication systems in this game.

It is important to note that the current sample size limits the possibility of making far-reaching generalisations. Still, the results reveal intriguing relationships that provide insight into the working mechanisms of the emergence of communication systems and may inspire future work. When examining individual participants, no measure of personal need for structure, *PNS*, *F1*, and *F2* correlated with high scores. However, when comparing partners in a team, we found that team measures—defined as the difference in pairs' individual scores—influenced performance. Greater differences in *PNS* and *F2* positively correlated with a team's *high score*.

Situated in the ECG, this entails that pairs of individuals who respond differently to unstructured situations were more successful in building a communication system together. A split of pairs into *good*, *medium*, and *bad* teams revealed that this relation is only present for well-performing teams. We, therefore, concluded that, while our results indicate that diverse reactions to a lack of structure may be beneficial in creating a communication system together, this difference does not necessarily *guarantee* better performance in the ECG. Many other factors, of course, influence the complex process of social coordination, and here we have identified one of them. Therefore, we suggest also to study other factors and interactions between people. We propose not only to further investigate the relation of *PNS* to the creation of novel communication systems but also to include analyses of other personality traits, such as the Big Five personality inventory (McCrae et al., 2005) or other questionnaires that assess personality traits (e.g. leadership, submissiveness). This would allow us to investigate further how various combinations of traits influence the creation of novel communication systems and create a deeper understanding of what might lead to success in collaborative tasks.

Human language is highly structured. It is suggested that systematic patterns emerged in language because humans are naturally biased towards compressible systems, through a general preference for simplicity (Kemp and Regier, 2012; Kirby et al., 2015). Here, we investigated the influence of such a *bias for structure* in a task where participants had to cooperate and coordinate their signals. These biases also significantly affect the emergence of structure in language as languages are learned and transmitted across generations (Kirby et al., 2008; Theisen-White et al., 2011; Verhoef, 2012; Kirby et al., 2015). Such experiments of iterated transmission often also expose participants to initially unstructured systems, which then gradually become structured over generations of transmission. Yet, diversity in the bias for structure has never been used as a factor in these studies, as such we propose there is an opportunity to further investigate this by assessing how differences in PNS may affect the emergence of patterns in transmission chain experiments like those of Kirby et al. (2008); Theisen-White et al. (2011); Verhoef (2012). This could reveal whether, besides the processes of transmission and interaction (Kirby, 2017), a direct individual need for structure, or differences therein, indeed affect the evolution of signals. If the latter is true, this would provide more evidence for the benefits of diverse members in collaborative tasks. The effect of diverse members in groups on the emergence of signalling systems can also be investigated when the ECG is adapted to accommodate groups instead of pairs. It has been found that communicating with multiple interaction partners introduces pressures that result in more stable shared vocabularies (Raviv et al., 2019a). In combination with our findings (i.e. that the ECG is a non-trivial task for pairs), we speculate that establishing common ground and emerging signals in an adapted ECG will be more difficult for groups, but that once these are in place, they will be more robust. We, moreover, expect that groups consisting of diverse members that score differently on PNS will benefit from this and obtain higher high scores.

It seems obvious why alignment in expectations may aid cooperation; it makes it easier to

coordinate and predict the moves of another player. The reason why diversity in expectations may be beneficial in cooperation tasks may be less intuitive, but we suggest that differences between interacting partners might complement each other's weaknesses, possibly aiding cooperation. In light of the ECG, this happens when one partner actively tries to create structure, while the other is looking for structure. This simultaneously raises questions.

For instance, what is the relationship between an individual's need for structure and their willingness to accommodate during moments of misunderstanding or ambiguity? While a high personal need for structure might imply a preference for clarity and predictability, this could manifest either as rigid insistence on one's own expectations or as a willingness to accommodate the other's framework to re-establish clarity. Conversely, a low need for structure might afford greater flexibility in interpretation, but also less urgency to coordinate or accommodate under pressure. These nuances, moreover, suggest that establishing common ground may benefit from specific combinations of individuals' preferences. A dyad composed of two high-PNS individuals might appear aligned in their desire for structure, but diverge when their preferred structures conflict. Slightly mismatched levels of structure preference—where one individual seeks guidance and the other provides it—may, in such cases, result in smoother coordination. This functional complementarity could enable teams to balance the need for structure with adaptive responsiveness. We suggest targeted team formation based on different cognitive preferences as a fruitful research direction.

2.7 Conclusion

In general, we argue that novel insights can be obtained if we do not only focus on the systems invented by successful pairs in communication game studies but also investigate what might separate those who score high from those who perform worse. Contrary to what is usually assumed, namely that overlap in cognitive biases and similarities in expectations drives the emergence of shared systems (Tylén et al., 2013; Scott-Phillips and Kirby, 2010), we found that differences in personal need for structure also matter in cooperative tasks and that diversity of communication partners might be beneficial for the emergence of new communication systems. While more evidence is needed to support this benefit, we speculate that differences in biases or personalities can aid by complementing the weaknesses of partners in unfamiliar collaborative situations, such as language evolution. We propose that novel insights can be obtained by focusing on targeted differences between interacting pairs that have been unable to communicate successfully. Finally, we suggest including other personality traits and investigating the exact workings of the dynamics between mixed prior expectations, personality traits and the emergence of novel communication systems.