

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

On the Benefits of Heterogeneity in Cognitive Stability and Flexibility for Collaborative Task Switching

Permalink

<https://escholarship.org/uc/item/6b47b61g>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 46(0)

Authors

Brondetta, Alessandra

Bizyaeva, Anastasia

Lucas, Maxime

et al.

Publication Date

2024

Peer reviewed

On the Benefits of Heterogeneity in Cognitive Stability and Flexibility for Collaborative Task Switching

Alessandra Brondetta^{1,5}, Anastasia Bizyaeva², Maxime Lucas³,
Giovanni Petri^{3,4}, Sebastian Musslick^{5,6}

¹Dept. of Physics, University of Turin, TO, Italy, ²Dept. of Mechanical and Engineering, University of Washington, Seattle, WA, USA,
³CENTAI Institute, Turin, Italy, ⁴Network Science Institute, Northeastern University London, London, UK,
⁵Dept. of Cognitive, Linguistic, and Psychological Sciences, Brown University, Providence, RI, USA,
⁶Institute of Cognitive Science, Osnabrück University, Osnabrück, Germany
Correspondence: alessandra.bronde@edu.unito.it

Abstract

Environments pose antagonistic demands on individual and collective cognition, such as trading off cognitive stability against cognitive flexibility. Manifestations of this tradeoff have been shown to vary across individuals, leading to differences in individual task switching performance. In this simulation study, we examine how individual differences in cognitive stability and flexibility contribute to collective task switching performance. Specifically, we study whether diversity in cognitive stability and flexibility among members of a group can facilitate collaborative task switching. We test this hypothesis by probing task switching performance of a multi-agent dynamical system, and by varying the heterogeneity of cognitive stability and flexibility among agents. We find that heterogeneous (compared to homogeneous) groups perform better in environments with high switch rates, especially if the most flexible agents receive task switch instructions. We discuss the implications of these findings for normative accounts of cognitive heterogeneity, as well as clinical and educational settings.

Keywords: cognitive control; task switching; stability-flexibility tradeoff; opinion dynamics; multi-agent systems.

Introduction

Humans are exceptionally good at focusing on goal-relevant tasks in the face of distraction—an ability attributed to our capacity for cognitive control. Yet, the requirement to focus on a task is often superseded by the demand to quickly switch from one task to another. These antagonistic demands can require humans to balance cognitive stability against flexibility (Goschke, 2000; Musslick & Cohen, 2021).

Individual differences in the balance between cognitive stability and flexibility have been reported to result in differences in task switching performance (Crofts et al., 2001; Ueltzhöffer et al., 2015; Musslick et al., 2019), leading individuals with greater cognitive flexibility to switch faster and individuals with greater cognitive stability to be less distracted. Yet, it is unclear how individual differences in cognitive stability and flexibility contribute to a group’s collective task switching performance, and whether heterogeneity in cognitive stability and flexibility benefits group performance.

Computational analyses suggest that the optimal balance between cognitive stability and flexibility varies with the demand for task switches. Previous modeling efforts formalized this problem in terms of a nonlinear dynamical system, in which an individual’s focus on a particular task is represented by an attractor (Ueltzhöffer et al., 2015; Musslick et al., 2018). Switching between tasks requires a re-configuration of the dynamical system, moving its state from

one attractor to another. The balance between cognitive stability and flexibility can be regulated by means of the attractor depth: deeper attractors promote robustness against distraction (cognitive stability), whereas shallower attractors promote fast switches between tasks (cognitive flexibility). According to these models, limited focus on a task, implemented by shallower attractors, can be beneficial if frequent task switching is required (Musslick & Bizyaeva, 2024).

Computational analyses suggests that, for a given rate of task switches, there is a single optimal balance between cognitive stability and flexibility when attempting to maximize *individual* task switching performance (Musslick et al., 2019). However, it is unclear whether such a balance would also optimize the *collective* performance of a group of individuals in the same task switching environment. Given that individuals can differ in their cognitive stability and flexibility (Crofts et al., 2001; Moustafa et al., 2008; Ueltzhöffer et al., 2015; Musslick et al., 2019), one might wonder whether such differences can improve collective performance.

In this computational study, we test the hypothesis that heterogeneity in cognitive stability and flexibility can benefit the collective performance of a group of individuals. To illustrate this hypothesis, consider a scenario in which a group of students collaboratively study for an exam. Students with high cognitive stability may be able to focus on their exam preparation despite distracting background noise. Students with low cognitive stability may study less effectively, but may notice and communicate environmental cues indicating the demand for a task switch (e.g., the last bus approaching the nearby station). In this scenario, students with low cognitive stability may achieve a low individual outcome on the primary task (studying for the exam), but may contribute to the collective outcome of the group (all students can catch the last bus). Here, we explore the conditions under which heterogeneity in terms of cognitive stability and flexibility can benefit group performance, through the lens of a multi-agent dynamical system. Using this system, we conduct numerical simulations in a task switching environment, comparing the performance of homogeneous groups with shared cognitive traits to heterogeneous groups. We find that heterogeneous groups can outperform homogeneous ones, particularly in scenarios characterized by a high task switching rate and limited external instructions for task switching. Critically, we find that the benefits of heterogeneous groups arise from par-

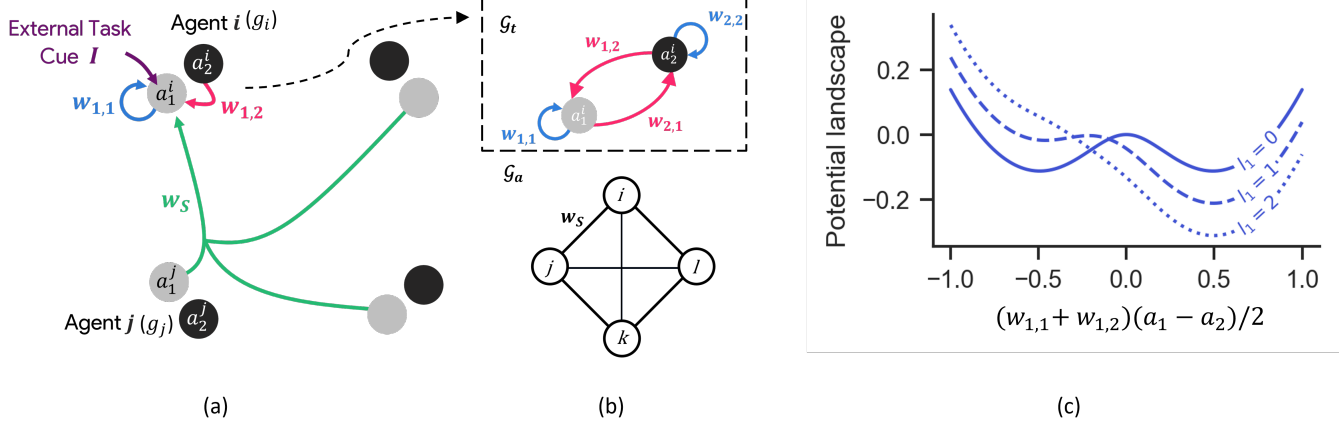


Figure 1: **Multi-agent dynamical system model.** (a) Graphic representation of inputs to a single processing unit for a single agent; a_1^i is the activity (focus) of agent i on Task 1, $w_{1,1}$ represents self-excitation of the respective tasks (blue), while $w_{1,2}$ represents inhibition from the second task a_2^i (pink). The weight w_s modulates the focus on the task derived from social interactions (green). Finally, I is the external relevant task instruction (purple). Every agent has a gain parameter g that encodes their ability to focus. (b) Task G_t graph; the two processing units (one per task) have self-recurrent positive weights and the inhibitory (negative) weight from the other unit (top). Note that $w_{2,1}$ and $w_{2,2}$ are not shown in panel (a) to avoid visual cluttering. Communication graph G_a structure for a group of four agents, all connected to all (bottom). Every connection has a positive weight, reflecting that agents tell each other to perform the task they are currently performing. (c) Effect of varying the external input strength on shape of the normalized gradient potential landscape. The landscape is obtained by restricting Equation (1) to the line $a_1 + a_2 = 0$, with $g = 5$. Local minimum at a positive (negative) value on the horizontal axis corresponds to an attractor for Task 1 (Task 2). Increasing strength of input for Task 1 breaks the symmetry in the potential landscape and removes the attractor for Task 2.

ticular configurations of groups in which flexible agents receive the task switch cues. We conclude by discussing the implications of our findings with respect to interpreting extreme forms of cognitive stability and flexibility in clinical, educational and corporate team contexts.

Model

Here, we formalize collaborative task switching within a multi-agent dynamical system. We begin by introducing a version of this system used to study the stability-flexibility tradeoff within a single individual (Musslick et al., 2019), and then extend the model to study the trade-off within a group of interacting agents. The collaborative task switching model we arrive at is a specialization of a multi-dimensional social belief formation model recently introduced and analyzed in (Bizyaeva et al., 2023b,a).

For the purposes of this study, we consider environments with two tasks and groups of N total agents. Each agent's focus on a task is represented by the activity of two associated processing units, indexed by 1,2. The activity $a_1^i(t)$ of the agent i on Task 1 (for Task 2 we have symmetric equations) evolves over time t according to:

$$\frac{da_1^i(t)}{dt} = \tau^{-1} \left[\underbrace{-d a_1^i(t)}_{\text{decay term}} + \underbrace{\sigma(g_i f_i(a_1^i(t)))}_{\text{internal dynamics}} + \underbrace{I_{ex}^{i,1}(t)}_{\text{external inputs}} \right] \quad (1)$$

where the focus is influenced by a general decay (towards no focus on the corresponding task), internal dynamics of focusing on one task versus another, and external inputs indicating the relevant task to be performed. In this model, τ is a temporal constant, d is the coefficient of the decay term of the

activity, and $\sigma(\cdot) = \tanh(\cdot)$ is a non-linear activation function that bounds the total activity. The stability-flexibility tradeoff of a given individual is regulated by the parameter g_i (gain) which influences the slope of the saturation function. The gain effectively regulates the depth of the task attractors and, thereby, an agent's balance between cognitive stability and flexibility: a higher (lower) value of g indicates greater cognitive stability (flexibility). Specifically, higher g values lead to greater cognitive control allocated to a given task, resulting in increased neural activation and greater cognitive stability. However, this also leads to greater persistence during task switches, causing switch costs, and reducing flexibility. Here, we simulate individual differences in cognitive stability vs. flexibility, by assuming different values of g for different agents.

Inside the saturation function σ , $f_i(a_1^i(t))$ represents the internal dynamics of agent i (illustrated in Figure 1b):

$$f_i(a_1^i(t)) = \frac{1}{N_i+1} (w_{1,1} a_1^i(t) + w_{1,2} a_2^i(t)), \quad (2)$$

which is a linear combination of the unit's own activity $a_1^i(t)$ multiplied by the self-excitatory weight $w_{1,1}$ and the activity $a_2^i(t)$ of the other unit multiplied by an inhibitory weight $w_{1,2}$. The self-excitatory weight implements task-set inertia, making the agent maintain the focus on the currently performed task, whereas the inhibitory weight implements the assumption that the two tasks are mutually inhibitory (McClelland & Rumelhart, 1981).

When considering a single agent alone, the input may be exclusively driven by an external task cue. However, when considering the agent in a group, the input will also be driven by social input derived from the communication between the

agents (illustrated in Figure 1b). Thus, the total external input $I_{\text{ext}}^{i,1}(t)$ received by the agent i for Task 1 is composed of two terms:

$$I_{\text{ext}}^{i,1}(t) = I_{\text{cue}}^{i,1}(t) + I_{\text{social}}^{i,1}(t). \quad (3)$$

The first term, $I_{\text{cue}}^{i,1}(t)$, is the external task instruction specifically for Task 1. Its value is positive when the task cue directs the agent to focus on Task 1 and negative when the task cue instructs the agent to focus on Task 2. It is 0 when the agent does not receive the external task cue. In this study, we will manipulate the number of agents N_{cue} that have access to the same task cue.

The second term, $I_{\text{social}}^{i,1}(t)$, is the external input received from other agents. Since the model is idealized, we assume that sending agents communicate their own activity state directly to the receiving agents, thus biasing the receiving agents towards performing the task that the sending agents are focusing on:

$$I_{\text{social}}^{i,1}(t) = g_i \frac{w_s}{N_i+1} \sum_{j=1, j \neq i}^N A_{ij}^a a_1^j(t), \quad (4)$$

where A_{ij}^a are elements of the adjacency matrices of the communication graph between the agents \mathcal{G}_a ($A_{ij}^a = 1 \quad \forall i, j$ in our case). The social component of the external input is modulated by the parameter w_s . Within the saturation function, there is also a normalization term that considers the number of agents with whom agent i interacts in the communication graph: N_i .

Simulation Experiments

This simulation study aims to examine the performance of heterogeneous and homogeneous groups of agents, with respect to their individual differences in the stability-flexibility balance, in a collaborative task switching environment. As described above, we represent individual cognitive differences in this balance using the gain parameter g . Accordingly, a homogeneous group is formed by agents with identical g values, whereas a heterogeneous group is composed of agents with diverse g values. Individual g values are drawn from a Gaussian distribution with mean μ and coefficient of variation cv . Thus, cv regulates the heterogeneity of the group, with greater values resulting in starker individual differences between agents. We use the coefficient of variation since it is a relative measure of dispersion and gives a standardized measure that is independent of the scale of the gain.

We investigate the effects of heterogeneity across three simulation experiments. The first experiment demonstrates that the model described above can reproduce the stability-flexibility trade-off for an individual agent. Furthermore, we determine the optimal stability-flexibility balance for (a) individual agents, (b) homogeneous groups, and (c) heterogeneous groups as a function of task switch rate. The second experiment analyzes collaborative task switching performance as a function of group heterogeneity. Here, we identify environmental conditions under which heterogeneous groups can outperform homogeneous ones. The third experiment

then examines the characteristics that a heterogeneous group should possess to surpass the performance of a homogeneous group.




Figure 2: **Evolution of task focus and performance.** (a) Evolution of the task focus a_i for a group of four agents with different values of gain g , all receiving the external task cue (black line). a_i is obtained as $a_1^i - a_2^i$; it takes positive values if the agent is more focused on the first task and negative values if more focused on the second task. Vertical dashed lines denote task switches indicated by I_{cue}^1 . (b) The performance metric quantifies how much the agent is focusing on the currently relevant task (green) versus the wrong task (red).

General Simulation Methods

For simplicity, we restrict the scope of our study to an exemplary case with groups involving $N = 4$ agents, each connected to all others with equal weight w_s in the communication network (Figure 1b). As outlined above, the agents are tasked to switch between two tasks. Each task is associated with one of two processing units whose activity— a_1^i or a_2^i —represents the focus on the respective task.

We study the evolution of the task focus for each agent. Specifically, we consider the relative focus on Task 1 versus Task 2, $a_i = a_1^i - a_2^i$ in this task-switching scenario; a_i takes positive values if the agent is more focused on the first task and negative values if it is more focused on the second task. All agents start with initial condition $a_i = 0$ (balanced focus between Tasks 1 and 2). At certain time steps, a subset of agents in the group receives a cue indicating that the relevant task has switched, $I_{\text{cue}}^{i,1}$. Crucially, even agents who do not receive task instructions directly are capable of switching the task on which they are focusing by receiving information from the other agents via $I_{\text{social}}^{i,1}$ about the task on which they are currently focusing (Equation 4).

Throughout the simulations, we keep all model parameters fixed except for four parameters that are varied systematically: (a) the mean and (b) the variance of the distribution of gains g , (c) the rate of task switching ω , and (d) the

number of agents receiving external task cue N_{cue} . The total time of the simulation is kept constant in all simulations ($T = 2200$). We identified the fixed parameters of the model through grid search (Table 1), so to maximize the tradeoff between cognitive stability and flexibility postulated in earlier work (Bizyaeva et al., 2023a; Goschke, 2000; Ueltzhöffer et al., 2015; Musslick & Bizyaeva, 2024; Musslick et al., 2019).

Table 1: Parameters of the Model.

Parameter	Value	Description
g	0 - 20	gain
d	0.2	task decay
$w_{1,1}, w_{2,2}$	0.02	task self-excitation
$w_{1,2}, w_{2,1}$	- 0.01	between-task inhibition
w_s	0.02	social influence
τ	10	integration constant
I_{cue}	-0.1,0.1	external task cue

We quantify (collaborative) task switching performance based on the degree to which the agent(s) are focusing on the currently relevant task. For instance, if $I_{\text{cue}}^1 > 0$ then all agents should be focusing on Task 1, i.e., $a_i > 0$. We quantify performance by computing the area between the a_i curve and the x-axis around zero. We assign positive values to the area where agents focus on the currently relevant task (green in Figure 2) and negative values where they focus on the wrong task (red in Figure 2), considering this area as a penalty. Thus, the performance of an agent i can be quantified as follows:

$$P_i = \int_t a_i(t) \cdot \text{sgn}(I_{\text{cue}}^1) dt$$

where I_{cue}^1 is the function representing the black curve in Figure 2a and $\text{sgn}(\cdot)$ is the sign function.

By summing these areas, we obtain the total performance for all agents in the group, representing the collective task switching performance $P = \sum_i P_i$.

Simulation Study 1: Stability-Flexibility Trade-off and Optimal Gain

The aim of this simulation study is to determine the optimal value of the gain parameter g —reflecting the ideal balance between cognitive stability and flexibility—as a function of the task switch rate ω , the number of agents receiving the task cue N_{cue} , and the group (individual agent, homogeneous group, heterogeneous group).

Simulation Procedure We determined the optimal g for different combinations of ω , ranging from 0.002 to 0.007 (to have 4 to 14 successive switches), and N_{cue} , ranging from 1 to 4 agents receiving the task cue. For each parameter configuration (ω, N_{cue}), we simulated the system for a range of gains g (from 0.1 to 10) for (a) a single agent, (b) a homogeneous group of 4 agents, and (c) a heterogeneous group of 4 agents. For heterogeneous groups, we computed that performance as the average performance over 100 simulations,

each with a newly drawn distribution of g , and reported the mean of the distribution of gains (with $cv = 0.5$). For each of the three scenarios, we determined the g that maximizes the (collective) task switching performance P .

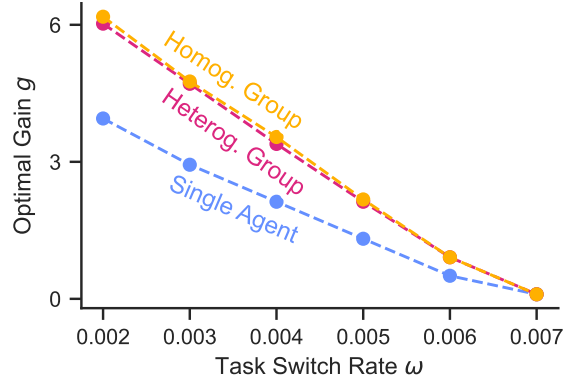


Figure 3: **Optimal gain g as a function of task switching rate ω for a single agent, homogeneous group, and heterogeneous group.** For the groups, each data point represents the average g value obtained across simulations with varying N_{cue} . The optimal g decreases with increasing task switching rate.

Results The simulations of the agents’ activity evolution over time directly indicate a tradeoff between cognitive stability and flexibility: agents with higher stability can focus more on a single task but take longer to switch to the other task (Figure 1a). Furthermore, the simulation results indicate that the optimal gain parameter g decreases with an increase in the task-switching rate ω (Figure 3). This suggests that when the task environment requires more frequent switching between tasks, it is beneficial to prioritize flexibility at the expense of stability. This trend holds true at the individual level, consistent with previous results (Musslick et al., 2019), and extends to both homogeneous and heterogeneous groups of agents. Interestingly, both groups have a higher optimal g compared to single agents across different task switch rates ω , suggesting that the collective task switching performance of groups benefits from greater cognitive stability than individual task switching performance (Figure 3).

Simulation Study 2: Homogeneous Versus Heterogeneous Group Performance

Next, we seek to compare the collaborative task switching performance of homogeneous groups with heterogeneous groups across a wide range of simulation parameters. Specifically, we examine conditions under which heterogeneous groups may outperform homogeneous ones.

Simulation procedure As in the first simulation experiment, we simulated group performance for different combinations of ω and N_{cue} , varying ω from 0.002 to 0.012 in steps of 0.002 and N_{cue} from 1 to 4. Furthermore, we considered different values of g ranging from 1 to 6 in steps of 1. We evaluated the performance for every value of g for




Figure 4: **Homogeneous versus heterogeneous group performance.** (a) Mean group performance value of the heterogeneous group in different settings of N_{cue} and ω . Every point is the mean performance value corresponding to a single set of simulations, with defined N_{cue} and ω . Error bars represent the standard deviation of the simulated average performance across different values of μ and c_V of the Gaussian distribution from which the g vectors of the group are extracted. The black dashed line represents the average trend of the performance value and it is obtained by interpolating a quadratic curve through the original data points. Performance decreases with increasing task switching rate and with fewer agents receiving task cues. (b) Percentage of the heterogeneous groups outperforming homogeneous groups in terms of mean group performance. Points correspond to the mean percentage value of success (heterogeneous group performance $>$ homogeneous group performance), averaged across different sets of μ and c_V . Error bars and black dashed line same as in (a).

the homogeneous group, and then we compared this value with 300 simulations of the heterogeneous group. We simulated different levels of heterogeneity by varying the coefficient of variation c_V with which g is sampled, from 0.25 to 1 in steps of 0.25. We ensured that the mean of the vector elements of g for the heterogeneous group matched the g value of the homogeneous group, isolating the impact of heterogeneity (spread) of g . For each combination of $\mu = g$ and c_V , we assessed the mean group performance \bar{P} over the 300 simulations, the distribution of group performances P for the heterogeneous group, and the percentage of times (out of 300) that the heterogeneous group outperformed the homogeneous one in terms of P .

Results: Comparative Performance Simulation results indicate that heterogeneity in the balance between cognitive stability and flexibility does not universally enhance group performance (Figure 4). In general, we observe that the mean performance of the heterogeneous group decreases as switch rate increases, and increases with the number of agents receiving the task cue (Figure 4a). Interestingly, on average, the heterogeneous group tends to outperform the homogeneous one ($> 50\%$) only in scenarios characterized by high switching rates and a small number of agents receiving the

task cue (Figure 4b). In contrast, if the task switch rate is low or the task cue is provided to all agents, then heterogeneous groups do not outperform homogeneous groups on average. We observe an exception to this pattern for cases with the lowest switching rate, yet even in this instance, there is considerable variability in the performance values within the heterogeneous group. When assessing the best heterogeneous group performance over the 300 simulations in comparison to the homogeneous group, we observed that there is almost always at least one heterogeneous group capable of outperforming the homogeneous one, with the exception being the scenario in which all four agents receive task instructions.

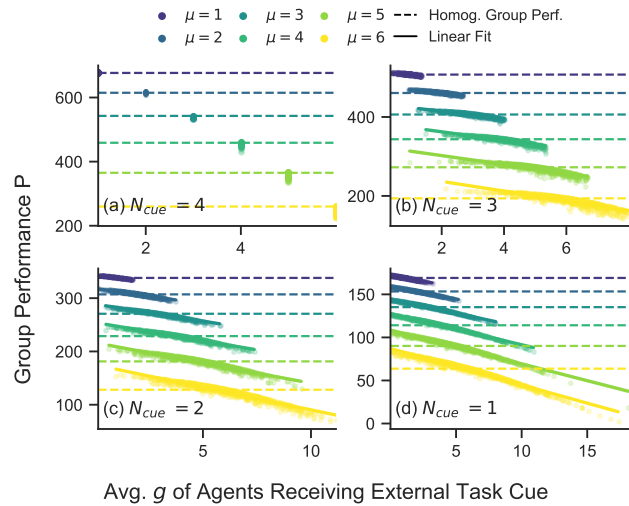


Figure 5: **Group performance as a function of average g for agents receiving I_{cue} .** Simulation data is shown for $\omega = 0.012$ and varying N_{cue} . The points represent the heterogeneous group performance for each simulation. Different colors refer to different values of μ used to sample the g vector for a group. Dashed lines indicate the homogeneous group performances (one single value per $\mu = g_{hom}$). The colored straight lines are obtained doing a linear regression on the performance points. The impact of the average gain of agents receiving the task cue increases with fewer agents receiving the cue.

Results: Characteristics of Well-Performing Heterogeneous Groups After establishing that heterogeneous groups can outperform homogeneous ones, especially in certain scenarios, we examine the characteristics of heterogeneous groups that can result in performance benefits compared to homogeneous groups. As shown in Figure 4b, higher variance in the g vector can either enhance or diminish the performance of the heterogeneous group compared to the homogeneous group. Thus, variance alone appears insufficient to discriminate group performance consistently.

Our follow-up analyses indicated a critical role of the g values for the subset of agents receiving the external task cue. In Figure 5, we examined simulations with $\omega = 0.012$ while decreasing N_{cue} , focusing on the average value of the gain parameter for agents receiving the task cue. In scenario (a) where all the agent receive the I_{cue} ($N_{cue} = 4$), we find that

the mean of the vector of g coincides with the $\mu = g_{\text{hom}}$ values, consistent with the construction of the g vectors. Here, introducing heterogeneity in most cases does not lead to an enhancement in group performance. However, as we reduced the number of agents receiving external instructions, the average gain parameter for the subset of agents receiving the task cue became increasingly critical for the performance of heterogeneous groups. The best-performing heterogeneous groups are the ones where the average g of the agents receiving I_{cue} is lower, meaning that they are, on average, more flexible. Figure 6 focuses on scenarios on which only one agent receives the task cue ($N_{\text{cue}} = 1$), across different task switch rates ω . In a low task switch rate scenario (Figure 6a), heterogeneous performance is improved the more cognitively stable the only agent receiving I_{cue} is (higher g value). Conversely, in scenarios with higher switch rates (Figure 6c-d; cf. Figure 5), the situation is reversed: heterogeneous groups outperform homogeneous groups if the only agent receiving external task instructions is very flexible (lower g value). The situation is more mixed for $\omega = 0.006$ (Figure 6b), and the trend appeared nonlinear. Our findings suggest that, in general, better performance for heterogeneous groups is achieved when, in scenarios with $N_{\text{cue}} < 4$ and high ω , the agents with external instruction are the most flexible in the group. Conversely, when the task-switching rate is low, it is more advantageous for the agents with external instruction to be the most stable ones in the group. However, here in most cases, we did not observe heterogeneous groups outperforming the homogeneous ones.

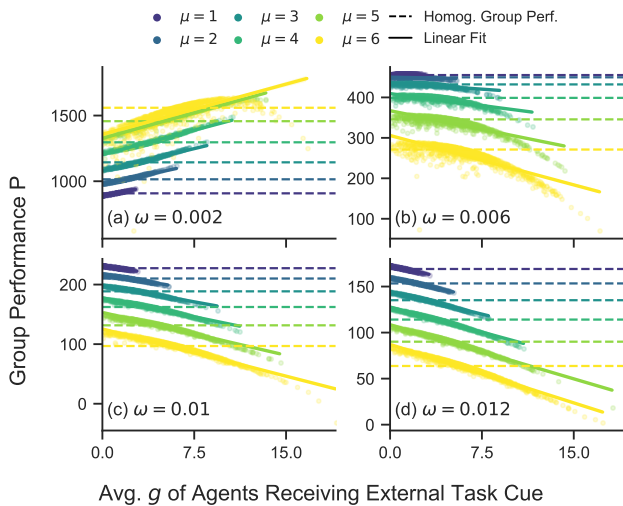


Figure 6: Group performance as a function of g of the agent receiving I_{cue} . Simulation data is shown for $N_{\text{cue}} = 1$ and varying ω . The points represent the heterogeneous group performance for each simulation. Different colors refer to different values of μ used to sample the g vector for a group. Dashed lines indicate the homogeneous group performances (one single value per $\mu = g_{\text{hom}}$). In low task switch rates, better heterogeneous performance is linked to a more cognitively stable agent receiving I_{cue} . Conversely, in high switch rates, better performance is associated with a more flexible agent receiving instructions.

General Discussion and Conclusion

Humans have been found to adapt to varying demands for flexibility, presumably by balancing cognitive stability against flexibility (Braem, 2017; Dreisbach & Fröber, 2019; Goschke, 2000; Mayr & Kliegl, 2000; Monsell & Mizon, 2006; Musslick & Cohen, 2021). Perhaps unsurprisingly, not every participant has been found to exhibit the same balance, with some participants exhibiting greater flexibility at the expense of cognitive stability while others show the reverse pattern (Crofts et al., 2001; Ueltzhöffer et al., 2015; Musslick et al., 2019), although this tradeoff may not always manifest (Egner, 2023; Mayr & Grätz, 2024). In this study, we leverage a multi-agent dynamical systems model to examine whether groups of individuals can benefit from such individual differences when exposed to task switching environments.

Our simulation results indicate that group heterogeneity in the balance between cognitive stability and flexibility does not consistently enhance collective task switching performance compared to group homogeneity. Yet, we found that heterogeneous groups can outperform homogeneous groups if (a) task switches are more frequent and (b) only a few individuals have access to information about the relevant task to perform. In the latter cases, heterogeneous groups perform better if the agents receiving the task cue are the most flexible ones in the group. Conversely, in scenarios with infrequent task switches, better performance is observed when external task cues are provided to the most stable agents.

The simulation results support a normative perspective on diversity in cognitive traits, such as individual differences in the balance between cognitive stability and flexibility. Specifically, our results suggest that individuals with extreme forms of stability or flexibility can benefit group performance despite below-average individual performance. This aligns with a utilitarian perspective on psychiatric dysfunctions associated with extremes on both ends of the balance spectrum. This perspective is also salient in educational settings as contemporary pedagogical frameworks, such as Disability Studies in Education, emphasize strengths-based rather than deficit models of ability/disability (Connor et al., 2008). The same principle may apply in corporate team work environments, where group heterogeneity is an active domain of study (Laureiro-Martínez & Brusoni, 2018; Hmieleski & Ennsley, 2007). Finally, the simulation study complements existing frameworks of rational boundedness Musslick & Cohen (2021); Musslick & Masís (2023), in so far as that it suggests a rational role for individual limitations in cognitive stability or flexibility at the level of group performance.

While this study of group heterogeneity provided initial insights, one must acknowledge that it relied on multiple simplifying assumptions, such as that all agents are informed about the task state of other agents at any point in time. Future analyses are needed to probe the generality of these findings, and to put them to empirical test in collaborative task switching environments.

Acknowledgements

S. M. was supported by Schmidt Science Fellows, in partnership with the Rhodes Trust, as well as the Carney BRAIN-STORM program at Brown University. A. B. gratefully acknowledges the support of the Erasmus+ Traineeship International Mobility program in collaboration with the University of Turin.

References

- Bizyaeva, A., Franci, A., & Leonard, N. E. (2023a). *Multi-topic belief formation through bifurcations over signed social networks*.
- Bizyaeva, A., Franci, A., & Leonard, N. E. (2023b). Nonlinear opinion dynamics with tunable sensitivity. *IEEE Transactions on Automatic Control*, *68*(3), 1415–1430.
- Braem, S. (2017). Conditioning task switching behavior. *Cognition*, *166*, 272–276.
- Connor, D. J., Gabel, S. L., Gallagher, D. J., & Morton, M. (2008). Disability studies and inclusive education—implications for theory, research, and practice. *International Journal of Inclusive Education*, *12*(5-6), 441–457.
- Crofts, H., Dalley, J., Collins, P., Van Denderen, J., Everitt, B., Robbins, T., & Roberts, A. (2001). Differential effects of 6-ohda lesions of the frontal cortex and caudate nucleus on the ability to acquire an attentional set. *Cerebral cortex*, *11*(11), 1015–1026.
- Dreisbach, G., & Fröber, K. (2019). On how to be flexible (or not): Modulation of the stability-flexibility balance. *Current Directions in Psychological Science*, *28*(1), 3–9.
- Egner, T. (2023). Principles of cognitive control over task focus and task switching. *Nature Reviews Psychology*, *2*(11), 702–714.
- Goschke, T. (2000, 01). Intentional reconfiguration and involuntary persistence in task-set switching. *Control of Cognitive Processes: Attention and Performance XVIII*, *331-355* (2000), 18.
- Hmieleski, K. M., & Ensley, M. D. (2007). A contextual examination of new venture performance: Entrepreneur leadership behavior, top management team heterogeneity, and environmental dynamism. *Journal of Organizational Behavior: The International Journal of Industrial, Occupational and Organizational Psychology and Behavior*, *28*(7), 865–889.
- Laureiro-Martínez, D., & Brusoni, S. (2018). Cognitive flexibility and adaptive decision-making: Evidence from a laboratory study of expert decision makers. *Strategic Management Journal*, *39*(4), 1031–1058.
- Mayr, U., & Grätz, D. (2024). Does cognitive control have a general stability/flexibility tradeoff problem? *Current Opinion in Behavioral Sciences*, *57*, 101389.
- Mayr, U., & Kliegl, R. (2000). Task-set switching and long-term memory retrieval.
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: I. an account of basic findings. *Psychological review*, *88*(5), 375–407.
- Monsell, S., & Mizon, G. A. (2006). Can the task-cuing paradigm measure an endogenous task-set reconfiguration process? *Journal of Experimental Psychology: Human Perception and Performance*, *32*(3), 493–516.
- Moustafa, A. A., Sherman, S. J., & Frank, M. J. (2008). A dopaminergic basis for working memory, learning and attentional shifting in parkinsonism. *Neuropsychologia*, *46*(13), 3144–3156.
- Musslick, S., & Bizyaeva, A. (2024). Examining cognitive flexibility and stability through the lens of dynamical systems. *Current Opinion in Behavioral Sciences*, *57*, 101375.
- Musslick, S., Bizyaeva, A., Agaron, S., Leonard, N., & Cohen, J. D. (2019). Stability-flexibility dilemma in cognitive control: A dynamical system perspective. In *Proceedings of the 41st annual meeting of the cognitive science society*.
- Musslick, S., & Cohen, J. D. (2021). Rationalizing constraints on the capacity for cognitive control. *Trends in Cognitive Sciences*, *25*(9), 757–775.
- Musslick, S., Jang, S. J., Shvartsman, M., Shenhav, A., & Cohen, J. (2018, 08). Constraints associated with cognitive control and the stability-flexibility dilemma..
- Musslick, S., & Masís, J. (2023). Pushing the bounds of bounded optimality and rationality. *Cognitive Science*, *47*(4), e13259.
- Ueltzhöffer, K., Armbruster-Genç, D. J., & Fiebach, C. J. (2015). Stochastic dynamics underlying cognitive stability and flexibility. *PLoS computational biology*, *11*(6), e1004331.