Nonergodicity and Simpson's paradox in neurocognitive dynamics of cognitive control

Percy K. Mistry[1]*, Nicholas K. Branigan[1]*, Zhiyao Gao[1], Weidong Cai[1,2], Vinod Menon[1,2,3]

[1] Department of Psychiatry & Behavioral Sciences
[2] Wu Tsai Neurosciences Institute
[3] Department of Neurology & Neurological Sciences

Stanford University, Stanford, CA, United States.

* These authors contributed equally to this work.

Corresponding authors:
Percy K. Mistry, Ph.D. & Nicholas K. Branigan & Vinod Menon, Ph.D.
email: percym@stanford.edu; branigan@stanford.edu; menon@stanford.edu

Abstract

Nonergodicity and Simpson's paradox pose significant and underappreciated challenges for neuroscience. Using stop signal task data from over 4,000 children and a Bayesian computational model of cognitive dynamics, we investigated brain-behavior relationships underlying inhibitory control at both between-subjects and within-subjects levels. Strikingly, between-subjects associations of inhibitory control activations with stop signal reaction times, probabilities of proactivity, and proactive delays were reversed within subjects, revealing the nonergodic nature of these processes. Nonergodicity was observed throughout the brain but was most pronounced in the salience network. Furthermore, within-subjects analysis revealed dissociated brain representations of reactive and proactive processing, and distinct brain-behavior associations for subjects who adaptively and who maladaptively regulated inhibitory control. This work advances our knowledge of the dynamic neural mechanisms of inhibitory control during a critical developmental period and has implications for personalized interventions in cognitive disorders. Embracing nonergodicity is crucial for understanding brain-behavior relationships and developing effective interventions.

Introduction

Since the nineteenth century, notions of ergodicity have had an important role in statistical physics[1-3], but only recently have other disciplines such as economics[4] and psychology[5-7] begun to examine how their own theories and findings rest upon assumptions of ergodicity. In the behavioral sciences, a clear example of nonergodicity is Simpson's paradox: an association between variables in a population may disappear or even reverse when the population is divided into subpopulations[8-11]. One form of the paradox occurs when associations are different in a population and in the individuals that make up the population. A classic case of nonergodicity and Simpson's paradox of this sort is the speed-accuracy tradeoff: in some tasks, speed and accuracy are positively correlated between individuals (faster people are more accurate)[12], but are negatively correlated within individuals (when an individual tries to respond faster, their accuracy decreases)[13]. This highlights how between-subjects and within-subjects inferences can diverge. There is growing evidence that many psychological phenomena may be nonergodic, meaning that inferences about them would differ when data are analyzed across time within an individual versus across individuals at a single point in time[5,7,14-17]

Nonergodicity thus implies that associations at the group level (**Figure 1a**) may fail to generalize to associations at the individual level (**Figure 1b**), and thus may fail to capture the dynamic processes occurring within each subject[5,6,16]. This view suggests that the relationship between behavior and the underlying brain activity may differ substantially when examined within subjects over time compared to between subjects at a single time point[7,18], leading to Simpson's paradox when comparing inferences from these two perspectives (**Figure 1c**). However, applications of this principle to human cognitive neuroscience have been limited, and nonergodic dynamics in brain and cognitive functions remain poorly understood. Nonergodic principles have yet to be fully integrated into the study of human cognition and brain function.

While there has been notable progress in understanding nonergodicity within behavioral contexts[5,6,18-20], its application to human brain function remains largely unexplored. This gap in research is significant because insights into nonergodicity in neural dynamics could profoundly enhance our understanding of brain-behavior relationships. Exploring this concept at the level of neural dynamics is crucial for identifying how an individual's differences in brain activity contribute to variability in their cognitive processes over time. By treating between- and within-subjects dynamics as dissociable, researchers can potentially develop more personalized and effective neuroscientific models and interventions, which consider the unique neural pathways and cognitive strategies employed by individuals. This approach not only promises to advance theoretical neuroscience but also holds practical implications for tailoring therapeutic strategies to better address individual neurological and cognitive differences.

To address these limitations and investigate the potential nonergodicity in neurocognitive processes and mechanisms, we leveraged data from the Adolescent Brain Cognitive Development (ABCD) study, a large-scale, longitudinal study of brain development and behavior in children and adolescents[21]. We used behavioral and brain imaging data from over 4,000 9- to 10-year-old participants to examine the neural mechanisms of inhibitory control at both the between-subjects and within-subjects levels.

Inhibitory control is the ability to withhold or cancel maladaptive actions, thoughts, and emotions, and is a fundamental component of goal-directed behavior[22,23]. This critical cognitive function allows individuals to navigate complex environments, adapt to changing circumstances, and maintain focus on long-term goals in the face of immediate temptations or distractions[24]. Given its central role in cognitive control, understanding the neural mechanisms underlying inhibitory control has been a primary focus of cognitive neuroscience research[25]. Inhibitory control engages a distributed network of cortical and subcortical regions. Previous studies have consistently implicated the involvement of the salience network, particularly the anterior insula and dorsal anterior cingulate cortex, in detecting and processing relevant stimuli and coordinating neural resources for inhibitory control[26-30]. Additionally, the inferior frontal gyrus, presupplementary motor area, and basal ganglia have been shown to play crucial roles in implementing response inhibition[25]. The dynamic activation of these cortical and subcortical regions underlies the successful execution of inhibitory control processes.

Conventionally, studies investigating the neural basis of inhibitory control have relied on between-subjects analyses, which compare brain activity across individuals who differ in their ability to inhibit responses. For example, researchers examine how between-individual differences in a measure of inhibitory control, such as the stop signal reaction time (SSRT) derived from the stop signal task (SST), relate to between-individual differences in brain activation from the task[27,31-37]. These analyses involve aggregating data at the individual-level and making inferences about the neural mechanisms of inhibitory control based on between-subjects comparisons[38].

However, this approach rests on the assumption of ergodicity, which posits that the statistical properties of a system are the same whether measured across time within an individual or across individuals at a point in time[18]. In other words, ergodicity assumes that the average behavior of a group reflects the behavior of each individual within that group[6,39]. If this assumption holds, then inferences made from group-level associations can be generalized to explain the cognitive processes occurring within each individual[7,18].

To enable within-subjects analyses and capture the dynamic nature of inhibitory control, we developed the Proactive Reactive and Attentional Dynamics (PRAD) model of SST behavior. This computational model extends beyond conventional approaches to the SST by incorporating both top-down and bottom-up processes, integrating mechanisms of proactive and reactive control[40]. The PRAD model allows us to infer dynamic, trial-level parameters for each subject, including SSRTs and measures of proactive delaying based on stop signal anticipation. These parameters characterize the latent cognitive constructs governing action execution and inhibition in the SST, providing a more comprehensive view of inhibitory control processes than traditional models.

We combined task fMRI data with the PRAD model to investigate the relationship between task-evoked brain responses and dynamic inhibitory control processes at both between-subjects and within-subjects levels. This approach allows us to directly compare inferences made from traditional group-level analyses (group-level associations) with those derived from a fine-grained examination of within-individual variation in neurocognitive dynamics. We leveraged the dynamic representations of subjects' behaviors from the PRAD cognitive model alongside

simultaneous task fMRI data to probe how moment-to-moment variations in brain activity relate to fluctuations in inhibitory control processes. This analysis strategy promotes a richer understanding of the neural basis of inhibitory control, potentially revealing insights that would be obscured in conventional group-level analyses.

We had five interconnected aims (**Figure 1d**). First, we examined how key parameters of reactive and proactive control[41] from the PRAD model related to brain activity at both between-subjects and within-subjects levels, identifying networks showing significant associations. By comparing these levels, we show the limitations of assuming ergodicity in the study of inhibitory control and underscore the importance of considering within-subjects variability. Second, motivated by the replicability crisis in neuroscience research, we assessed the stability and robustness of our findings. We examined the stability of our within-subjects findings as a function of the sample and sample size through bootstrap resampling. Then, we tested whether our finding of nonergodicity was robust to alternative between-subjects analyses using different measures of brain activation and robust to the use of a directly observed behavioral measure in place of latent cognitive model parameters. Third, we probed the brain's implementation of proactive and reactive control processes underlying inhibitory control. Proactive control involves the anticipation and preparation for stopping, while reactive control involves the actual implementation of response inhibition[41]. By comparing the brain representations of these processes, we improve understanding of how their neural underpinnings relate. Fourth, we identified subgroups of within-subjects results associated with individual differences in adaptive regulation of inhibitory control. Finally, we quantified and compared nonergodicity levels across different brain networks, defining nonergodicity as the fraction of subjects whose within-subjects brain-behavior association showed an opposite sign to the between-subjects association. This approach allowed us to map the distribution of nonergodicity across the brain and explore its hierarchical organization.

By comparing traditional between-subjects analyses with within-subjects analyses that account for dynamic cognitive processes, we highlight the limitations of assuming ergodicity in the study of cognitive control and provide a more precise understanding of the neural mechanisms supporting this critical function. Our findings have important implications for both basic research and clinical applications, emphasizing the need to consider within-subjects variability and the dynamic nature of cognitive processes in the study of brain-behavior relationships. Ultimately, this work contributes to a more complete understanding of the neural basis of inhibitory control.

Results

Dynamic cognitive process model of behavior

We used data from the SST (**Figure 2a**) to investigate dynamic cognitive processes underlying inhibitory control in a large sample (N~4000) from the baseline visit of the ABCD study. We developed PRAD, a dynamic cognitive process model[40] that provides a multidimensional perspective of the elemental cognitive processes governing the reactive and proactive dynamics involved in inhibitory control. This model incorporates latent dynamics that respond to internal cognitive states (endogenous variables) and external environmental contingencies (exogenous

variables), with interaction between these dynamics governed by latent trait measures. The latent dynamic (trial-level) and trait (individual-level) measures are simultaneously inferred within a hierarchical Bayesian framework. The model conceptualizes behavior during the task as arising from competing drift diffusion processes (**Figure 2b**). It infers a measure of reactive inhibitory control, stop signal reaction time (SSRT), and 2 measures of proactive control, probability of proactivity and proactive delaying (**Figure 2c-d**). SSRT measures how long it takes for a subject to inhibit a response, probability of proactivity determines whether they use a strategy of proactive delaying, and proactive delaying measures the length of proactive delays when a proactive strategy is utilized. Importantly, all 3 of these parameters are inferred for each trial (**Figure 2e**). **Figure 2e** shows how these 3 parameters vary over time within an individual. **Figure 3** provides a detailed illustration of the model.

Between-subjects and within-subjects analysis of brain-behavior associations

We examined how 3 key parameters of the PRAD model—SSRT, probability of proactivity, and proactive delaying—related to brain activity at both between-subjects and within-subjects levels (**Figure 4**). To investigate brain-behavior relationships in inhibitory control, we conducted both between-subjects and within-subjects analyses using fMRI data from the stop signal task (SSRT and probability of proactivity $N = 4469$; proactive delaying $N = 4176$). For the between-subjects analyses, we Pearson correlated subject-average brain activation during successful stopping (correct stop versus correct go activation) with subject-average SSRT, probability of proactivity, and proactive delaying across participants[27,33-37]. For the within-subjects analyses, we regressed the fMRI signal on the parameters using fMRI general linear models with SSRT, probability of proactivity, and proactive delaying included as regressors. We thus modeled the fMRI signal as a combination of static trial-type effects and dynamic effects proportional to trial-by-trial variations in the cognitive model parameters. This allowed us to examine how fluctuations in cognitive processes covaried with brain activity within each individual, while adjusting for stimulus types. These complementary approaches enabled us to compare group-level and individual-level trends in brain-behavior relationships.

SSRT: Between subjects, SSRT showed widespread negative correlations with brain activity, including in frontal, parietal, and temporal areas, and in regions implicated in cognitive control. This suggests that individuals with faster inhibitory responses (lower SSRT) show greater activation in these regions during successful stopping, replicating previous findings[37]. In contrast, within subjects, SSRT showed positive associations with brain activity, particularly in frontal and parietal regions. This suggests that on trials where individuals have slower inhibitory responses, they show increased activation in these areas.

Probability of proactivity: Between subjects, probability of proactivity showed minimal correlations with brain activity. In contrast, within subjects, probability of proactivity showed negative associations in parietal, temporal, and lateral frontal regions, and positive associations in default mode network regions.

Proactive delaying: Between subjects, proactive delaying showed widespread positive correlations with brain activity, particularly in frontal and parietal regions. This indicates that individuals who engage in more proactive delaying exhibit higher activation in these areas

during successful stopping. In contrast, within-subjects, proactive delaying showed negative associations in frontal and parietal cortex.

These findings reveal a striking divergence of between-subjects and within-subjects brain-behavior relationships in inhibitory control. The reversal of association directions, particularly for SSRT and proactive delaying, suggests that inferences about the neural mechanisms underlying inhibitory control do not generalize between group and individual levels. This nonergodic pattern highlights the importance of considering both levels of analysis to more fully understand the neurocognitive dynamics of inhibitory control.

Network-level visualization of between- and within-subjects brain-behavior associations

To elucidate the patterns of brain-behavior relationships across different functional brain networks, we visualized the between-subjects and within-subjects associations for SSRT, probability of proactivity, and proactive delaying (**Figure 5**). We used the Shirer network atlas[42,43] for our primary analysis as it includes the basal ganglia, a subcortical system important for the implementation of inhibitory control[27,44].

SSRT: Between-subjects analysis showed consistent negative correlations in most networks, with effects in the posterior salience, frontoparietal, and default mode networks (all $P_{\text{FDR}} < 0.01$). In contrast, within-subjects analysis revealed mixed negative and positive associations, with opposite associations in the posterior salience, precuneus, visuospatial, auditory, and sensorimotor networks (all $P_{\text{FDR}} < 0.01$). The reversal of association directions underscores the nonergodic nature of SSRT-related brain activity.

Probability of proactivity: Between-subjects analysis demonstrated no significantly nonzero correlations in the networks (all $P_{\text{FDR}} \geq 0.05$). However, within-subjects analysis unveiled a more complex pattern. Negative associations were observed in the salience and frontoparietal networks, while positive associations were seen in the dorsal and ventral default mode networks (all $P_{\text{FDR}} < 0.01$). This disparity highlights the importance of examining within-subjects dynamics for proactivity.

Proactive delaying: Between-subjects analysis revealed positive correlations in all the networks (all $P_{\text{FDR}} < 0.01$). Yet, within-subjects analysis showed predominantly negative associations, notably in the salience and frontoparietal networks (all $P_{\text{FDR}} < 0.01$). The ventral default mode network showed a positive within-subjects associations ($P_{\text{FDR}} < 0.01$).

These network-level visualizations emphasize the divergent patterns between group-level and individual-level associations across different functional brain networks. The consistent reversals observed, for both reactive and proactive control measures, reinforce that brain-behavior relationships in inhibitory control are nonergodic. These findings indicate that both between-subjects and within-subjects perspectives are needed to comprehensively understand the neural dynamics underlying inhibitory control.

Stability of within-subjects brain-behavior associations

To assess the stability of our within-subjects findings, we performed bootstrap resampling at varying sample sizes (**Figure 6**). Our results demonstrate that within-subjects associations between brain activity and the model parameters (SSRT, probability of proactivity, and proactive delaying) are stable, even in modest sample sizes. Key findings were consistently observed in samples as small as 25 subjects, with some effects requiring larger samples to emerge reliably. This stability suggests the validity and potential generalizability of our within-subjects approach to understanding neurocognitive mechanisms of inhibitory control. A detailed description of these results is in the Supplementary Materials.

Robustness of nonergodicity to analytical choices

To test the robustness of our findings of nonergodicity, we conducted several control analyses comparing between- and within-subjects brain-behavior relationships. We tested alternative between-subjects analyses using different measures of brain activation and explored whether the observed nonergodicity was specific to the latent model parameters or could also be seen using reaction time on go trials, a directly observed behavioral measure. Nonergodicity persisted for all approaches to between-subjects analysis and both model-derived parameters and the observed behavioral measure. Across all analytical choices, we observed divergent patterns of between-subjects and within-subjects associations. A detailed description of these control analyses and their results is in the Supplementary Materials. Our results strongly suggest that the neurocognitive dynamics of inhibitory control in children are fundamentally nonergodic, with implications for how we interpret findings from traditional group-level associations.

Representational similarity analysis reveals dissociated reactive and proactive representations

Building on our findings of nonergodic brain-behavior relationships, we sought to deepen understanding of how the brain implements proactive and reactive control at the individual level. While the preceding results demonstrate the importance of within-subjects analyses, they leave unresolved how representations of reactive and proactive processes relate to each other within brain networks. To investigate this, we used representational similarity analysis to examine the overlap between brain representations of reactivity (SSRT) and proactivity (probability of proactivity and proactive delaying) within individuals. For each subject and each brain network, we computed Pearson correlations between the subject's brain maps of SSRT and probability of proactivity, SSRT and proactive delaying, and probability of proactivity and proactive delaying, over the voxels in each network (**Figure 7a**).

SSRT showed low similarity with both proactive measures (probability of proactivity and proactive delaying) in all networks, with median correlations ranging from -0.07 to 0.06. In contrast, the 2 proactive measures (probability of proactivity and proactive delaying) exhibited high similarity in all networks, with median correlations ranging from 0.61 to 0.67. In each network, the 3 similarity measures (SSRT and probability of proactivity, SSRT and proactive delaying, and probability of proactivity and proactive delaying) were significantly different from each other in their median values (**Figure 7b**). These findings suggest that representations of reactivity and proactivity are largely dissociated. This dissociation persists across multiple brain networks, indicating a fundamental separation in how the brain encodes reactive and proactive control processes.

Adaptive regulation of inhibitory control associated with distinct within-subjects results over subgroups

To understand how our within-subjects associations related to between-subjects variation in cognitive and task strategies, we examined the within-subjects results between subgroups showing adaptive and maladaptive regulation of reactive and proactive behaviors. These subgroups were identified using PRAD model parameters γ_1 and θ_1.

First, the cognitive model infers for each subject γ_1, which determines whether subjects adaptively ($\gamma_1 > 0$) or maladaptively ($\gamma_1 < 0$) regulate reactivity. Adaptive (maladaptive) regulation of reactivity involves increasing (decreasing) expectancy of a stop trial as the number of successive go trials increases. Since expectancy of a stop trial is one of several determinants of SSRT, γ_1 influences SSRT variation through time. Thus, we examined within-subjects associations between SSRT and brain activity separately among subjects with $\gamma_1 < 0$ ($N = 2513$) and $\gamma_1 > 0$ ($N = 1956$) (**Figure 8a**). We found differences in the distributions of within-subjects SSRT associations between the γ_1 subgroups in all networks examined (all $P_{FDR} < 0.01$). The maladaptive regulation group had a larger (more positive) effect size in every network. In fact, across most networks, SSRT exhibited opposite associations with brain activity in the two subgroups. For example, in the frontoparietal and default mode networks, SSRT displayed a positive association among subjects with $\gamma_1 < 0$ and a negative association among subjects with $\gamma_1 > 0$. Moreover, this analysis revealed that some of the effects in the full sample were driven by subjects belonging to one of the subgroups. For example, in the anterior salience network, brain activity's positive association with SSRT in the full sample was driven by $\gamma_1 < 0$ subjects; among $\gamma_1 < 0$, the effect in the anterior salience had a Cohen's d of ~0.3, twice that of the effect in the full sample, while among $\gamma_1 > 0$, there was no significant effect at all.

Second, the cognitive model infers for each subject θ_1, which determines whether subjects adaptively ($\theta_1 < 0$) or maladaptively ($\theta_1 > 0$) regulate proactivity. Adaptive (maladaptive) regulation of proactivity involves increasing (decreasing) the probability of proactivity following a failed stop trial and decreasing (increasing) the probability of proactivity following a no-response go-trial. Therefore, θ_1 influences the probability of proactivity's variation through time. Thus, we examined the probability of proactivity's within-subjects associations with brain activity separately among subjects with $\theta_1 < 0$ ($N = 3054$) and $\theta_1 > 0$ ($N = 1415$) (**Figure 8b**). We found consistent differences in the probability of proactivity's within-subjects associations between these subgroups. Generally, subjects with $\theta_1 < 0$ had more negative associations between probability of proactivity and brain activity. This negative coupling between proactivity and brain activity was pronounced among $\theta_1 < 0$ in the anterior salience network (Cohen's d of ~0.4). This analysis also clarified how adaptive and maladaptive regulation of proactivity contributed to the effects observed in the full sample. The within-subjects positive association between probability of proactivity and ventral default mode activation was observed in both θ_1 subgroups, but it was twice as large among subjects who maladaptively regulated proactivity.

These results demonstrate that population subgroups related to adaptative and maladaptive regulation of inhibitory control demonstrate different, and even opposite, within-subjects brain-behavior associations.

Brain networks exhibit varying degrees of nonergodicity and a hierarchical organization by nonergodicity

To further understand the distribution of nonergodicity across the brain, we quantified and compared nonergodicity levels in different brain networks. We defined a nonergodicity measure as the fraction of subjects whose within-subjects brain-behavior association showed an opposite sign to the between-subjects association. Values above 0.5 indicate higher nonergodicity, while values below 0.5 suggest more ergodic behavior.

Our analysis revealed substantial variation in nonergodicity across brain networks (**Figure 9a**). Notably, the anterior salience network consistently demonstrated the highest level of nonergodicity for all three cognitive model parameters (SSRT, probability of proactivity, and proactive delaying). The measures of network nonergodicity for associations with the probability of proactivity showed wide confidence intervals, reflecting the weak between-subjects associations between this parameter and brain activation.

To understand how networks related to each other in terms of nonergodicity, we performed hierarchical clustering on the joint nonergodicity measures of the networks with respect to all three cognitive model parameters (**Figure 9b-d**). This analysis revealed a hierarchical organization of brain networks based on their nonergodicity profiles. The anterior salience network emerged as the most dissimilar, forming its own cluster separate from all other networks. Additionally, the dorsal and ventral default mode networks clustered together, suggesting similarities in their nonergodic behavior.

These findings demonstrate that brain networks exhibit varying levels of nonergodicity and reveal a hierarchical organization of brain networks based on their nonergodic properties. This organization suggest a new perspective on the functional architecture of the brain and how it relates to cognitive control processes.

Discussion

A fundamental question we examined in this study is whether between- and within-subjects brain-behavior associations yield convergent findings in understanding neurocognitive processes underlying inhibitory control. Divergence between these levels of analysis would provide evidence for nonergodicity, a phenomenon where inferences drawn from population-level data do not accurately represent individual-level processes. To investigate this, we leveraged a large community sample of children from the ABCD study and employed a dynamic computational model to elucidate the potentially nonergodic nature of neurocognitive dynamics underlying inhibitory control. We combined task fMRI data with a Bayesian model of cognitive dynamics to examine brain-behavior relationships at both between-subjects and within-subjects levels. Our

findings revealed striking differences between these two levels of analysis, providing robust evidence of nonergodicity in inhibitory control processes.

At the between-subjects level, inhibitory control activations in key cognitive control networks were negatively correlated with subject-average SSRTs, aligning with previous studies[27,31-33]. This suggests that individuals with better inhibitory control exhibit greater activation in regions associated with cognitive control. However, within-subjects analysis revealed a markedly different pattern: brain activity in some of these same networks was positively associated with trial-level SSRTs and negatively associated with trial-level proactive control. Thus, within-subjects associations were largely dissonant with the between-subjects findings, revealing the nonergodic nature of inhibitory control processes. Specifically, the opposing patterns demonstrated Simpson's paradox, a type of nonergodicity frequently observed in behavioral studies, where relationships at the group level are absent or reversed at the individual level[5,10,15,17]. This divergence underscores the importance of considering both group-level and individual-level associations in cognitive neuroscience research and highlights the complex, dynamic nature of inhibitory control processes.

Furthermore, representational similarity analysis uncovered distinct neural representations for reactive and proactive cognitive processes. Across all examined brain networks, representations of reactive control (SSRT) showed low similarity with both measures of proactive control, while the two proactive measures showed high similarity with each other. This dissociation suggests that reactive and proactive aspects of inhibitory control rely on distinct neural resources, potentially allowing for independent modulation and development of these strategies.

Modeling trial-level responses at the single-subject level

We used a hierarchical Bayesian model of proactive and reactive control which represents a significant advance in the assessment of inhibitory control, surpassing the capabilities of conventional race models[40]. Unlike traditional approaches, which provide subject-aggregate SSRT estimates as an index of inhibitory control, this model estimates SSRT at the level of each individual trial. This feature is critical as it allows for precise, trial-specific inferences rather than broad generalizations across the entire task. Importantly, the trial-level model infers additional trial-level measures, such as the probability of proactive cognitive states and the length of proactive delaying of responses.

This trial-level granularity facilitates a more comprehensive examination of how neural responses are modulated across individual trials. Specifically, this model allowed us to identify brain areas that tracked neural activity in response to trial-specific SSRT estimated at each stop trial. Moreover, it enabled tracking ongoing neural dynamics associated with temporal fluctuations in proactive control, providing insights into the brain systems that support a key component of cognitive control.

By employing trial-level analyses, we could dissect proactive and reactive control processes to examine how they fluctuate over time within each individual. This approach deepens our understanding of the mechanisms underlying inhibitory control and enables us to uncover whether assumptions of brain-behavior ergodicity are justified.

Nonergodic brain-behavior associations reveal mechanisms of cognitive control

Nonergodicity, a concept originally from statistical physics, refers to situations where ensemble averages and time averages do not converge[1]. In the context of neurocognitive dynamics, nonergodicity means that brain-behavior associations over a population (ensemble) of subjects fail to reflect brain-behavior associations over time in the subjects comprising the population. Nonergodic brain-behavior associations were observed on multiple measures of inhibition encompassing proactive and reactive control. Proactive control involves the anticipation and preparation for stopping, while reactive control involves the actual implementation of response inhibition. Our findings suggest that these two forms of control operate differently at the within-subjects level and the between-subjects level. This divergence highlights the variability of individual cognitive processes, which may not be captured fully by traditional group-level analyses.

We found a positive association within subjects between trial-level SSRTs and brain activity in the anterior and posterior salience networks, which suggests that longer SSRTs, indicating poorer reactive control, are associated with greater neural effort or engagement. This may reflect compensatory mechanisms or the increased demand for cognitive resources when individuals struggle to inhibit their responses. In contrast, between subjects, we observed no significant association and a negative association, respectively, in the anterior and posterior salience networks between SSRT and brain activity.

Within-subjects analysis revealed a negative association between trial-level engagement of proactive control and brain activity in the frontoparietal, salience, and subcortical systems. This suggests that greater proactive control is associated with relative suppression of cognitive control networks to implement successful response inhibition. This finding is consistent with the theory that proactive mechanisms suppress reactive control pathways, and aligns with recent theoretical frameworks proposing that proactive control modulates reactive control via preparatory processes[45]. In contrast, we observed a positive association between trial-level proactive control and default mode network activity. This included the posterior medial cortex and the ventromedial prefrontal cortex, the two core cortical nodes that anchor the default mode network[46]. This may reflect internally oriented processing that supports proactive regulation. Between-subjects analysis failed to capture these dynamic relationships, showing no association between one measure of proactive control and brain activations.

Our findings highlight the importance of considering within-subjects variability and dynamics when studying the neural mechanisms of cognitive control. Conventional between-subjects analyses, which assume ergodicity, may not capture the complex and dynamic nature of proactive and reactive control processes as they unfold within individuals over time. Properly characterizing such dynamics, rather than assuming ergodicity, is thus crucial for advancing our understanding of the latent processes underlying cognitive control. The nonergodic brain-behavior associations observed in our study have important implications for understanding the mechanisms of proactive and reactive control.

Robustness of nonergodicity findings and stability of within-subjects brain-behavior associations

Leveraging the large-scale ABCD dataset, we addressed the critical challenge of replicability in human neuroscience[47]. Our analyses revealed that within-subjects associations were stable and reliably detectable, even in sample sizes typical of cognitive neuroscience studies. Bootstrap resampling analyses showed the reliability of key findings across different sample sizes. For instance, the association between proactivity measures and right anterior insula suppression was consistently observed in over 95% of samples, even with sample sizes of $N = 25$. Moreover, our findings of nonergodicity were robust to various analytical strategies. When comparing the results of various between- and within-subjects approaches, there were variations in the details of brain-behavior inferences, but nonergodic dissociations persisted across every approach to brain-behavior association.

By demonstrating that these patterns are robust and detectable even in modest sample sizes, our study provides a foundation for future research into nonergodicity in brain function. It also suggests that meaningful insights into neurocognitive mechanisms can be gained from studies with more typical sample sizes, although larger samples provide greater precision and the ability to detect subtler effects. The stability and robustness of our findings suggests their applicability to diverse research and clinical contexts, including understanding cognitive processes related to inhibitory control and studying psychiatric disorders.[48]

Nonergodicity between brain networks: Implications for understanding cognitive control

Our analysis of the distribution of nonergodicity between brain networks has implications for understanding the neural mechanisms underlying inhibitory control as well as for cognitive neuroscience research broadly. The consistent finding of high nonergodicity in the anterior salience network across all cognitive model parameters is particularly intriguing. The salience network is a key brain network, known for its role in detecting behaviorally relevant stimuli and coordinating brain network dynamics[49,50]. It is noteworthy that this core network, which is of great interest in the study of cognition[26-29,51] and psychopathology[52-54], appears to exhibit the most pronounced disconnect between group-level and individual-level inferences. Moreover, the hierarchical clustering of networks based on nonergodicity profiles provides a fresh perspective on brain organization. The distinct clustering of the anterior salience network and the grouping of default mode network components suggest that nonergodicity may be an important factor in understanding functional brain architecture.

These findings have several significant implications for cognitive neuroscience research and practice. Methodologically, our results underscore the importance of complementing group-level analyses with individual-level investigations. The high degree of nonergodicity observed, particularly in key networks involved in cognitive control, suggests that solely relying on group-level analyses may lead to incomplete or misleading conclusions about brain-behavior relationships. From the perspective of individual differences, the varying levels of nonergodicity across networks highlight the importance of considering individual variability in brain function. This may be particularly relevant for understanding individual differences in inhibitory control abilities and for developing personalized interventions for disorders characterized by impaired

inhibitory control. Theoretically, the observed nonergodicity challenges simplistic models of brain function and calls for quantitatively rigorous theories that can account for the complex, context-dependent nature of brain-behavior relationships. This may require a shift toward more dynamic, process-oriented models of cognition and brain function.

Distinct neural representations for reactive and proactive cognitive processes

To further elucidate the neural architecture underlying inhibitory control, we employed representational similarity analysis, a powerful method for investigating the informational content of brain activity patterns[55]. This approach allows us to compare the similarity of neural representations across different cognitive processes, providing insights into how the brain organizes and processes information[56]. In our study, we used representational similarity analysis to examine the overlap between brain representations of reactivity (SSRT) and proactivity (probability of proactivity and proactive delaying) within individuals. By comparing representational patterns across different cognitive processes, we sought to determine whether reactive and proactive control rely on shared or distinct neural resources.

Our analysis revealed a striking dissociation between the neural representations of reactive and proactive control processes. Across all examined brain networks, we found low similarity between representations of SSRT and representations of each proactive measure (probability of proactivity and proactive delaying). In contrast, the two proactive measures showed high similarity with each other. This pattern suggests that reactive and proactive control processes are represented orthogonally in the brain. Theoretically, our findings challenge simplistic models of inhibitory control and suggest that reactive and proactive processes, while both contributing to inhibitory control, are implemented through distinct neural mechanisms[41,57].

Given that our study focused on children, the clear separation of reactive and proactive representations may reflect a developmental stage in the organization of cognitive control processes. Our findings of a separation may allow for independent development of reactive and proactive strategies, potentially explaining individual differences in inhibitory control abilities[41]. Future studies could investigate whether this orthogonality persists or changes with age[58].

Attentional modulation and performance monitoring associated with distinct brain-behavior associations

A separate behavioral investigation of our hierarchical Bayesian model highlighted the significance of adaptive regulation of reactive and proactive control in shaping within-subjects variability in SSRT and stop failure rates[40]. Two model parameters are decisive in controlling these dynamics: γ_1 represents individual differences in sustained attention and regulates the trial-level expectancy of stopping, and θ_1 represents individual differences in performance monitoring and regulates the trial-level proclivity for proactive control. These findings point to the importance of considering individual differences in attentional modulation and performance monitoring systems when studying inhibitory control.

Building on these results, we investigated brain-behavior associations between subjects who differed on these traits. We used γ_1 and θ_1 (separately) as a basis for creating subgroups within

our sample. By dividing our participants into subgroups, we could examine how individual differences in attentional regulation and performance monitoring correlate with the relationship between neural activity and cognitive processes within subjects. Subjects stratified based on whether they adaptively or maladaptively regulated stopping expectancy showed distinct within-subjects associations between trial-level SSRTs and brain activity, with different distributions of associations between the subgroups in every network and opposite associations in most networks. Among all subjects, we observed that anterior salience network activation accompanied poorer reactive control at the trial level, but examining these results by γ_1 subgroup revealed that this association only held for subjects who maladaptively regulated reactivity. Similarly, subjects divided by whether they adaptively or maladaptively regulated proactivity showed different within-subjects associations for a measure of proactivity in most networks. The maladaptive regulation group showed weaker suppression of the anterior salience network and stronger activation of the ventral default mode network with greater trial-level proactivity.

Collectively, the findings reveal that groups characterized by adaptive and maladaptive regulation of reactivity and proactivity display notably different patterns of within-subjects associations between brain activity and model parameters. These distinctions hint that individuals' distinct cognitive strategies or profiles relate to the implementation of proactive and reactive control processes in the brain. This variation highlights the personalized nature of cognitive function and stresses the importance of considering individual differences in the neural mechanisms of inhibitory control. Identifying heterogeneity based on cognitive model parameters provides an interpretable approach for studying individual differences in inhibitory control and their neural correlates. This approach moves beyond simple between-subjects comparisons and allows for a theory- and mechanism-driven investigation of the heterogeneity in brain-behavior relationships.

Conclusions

Our study provides evidence for nonergodicity in the neurocognitive processes underlying inhibitory control using a large, community-representative sample of children from the ABCD study. By combining task fMRI data with a dynamic cognitive model, we found that within-subjects associations between brain activity and model parameters differed from between-subjects associations, challenging the assumption of ergodicity in cognitive neuroscience research. The findings demonstrate divergent group-level and individual-level brain-behavior associations, reveal dissociated proactive and reactive control systems, and identify meaningful individual differences in these control processes. Crucially, the study establishes the stability and robustness of within-subjects measures, laying a foundation for characterizing nonergodic processes. The work highlights the value of large, heterogeneous samples, dynamic computational models, and analysis of within-subjects variability for studying nonergodic phenomena. More broadly, it suggests a paradigm shift away from ergodic assumptions and exclusive reliance on group-level analyses of subject-average measures in cognitive neuroscience. Appreciating the nonergodic nature of neurocognitive processes may be essential for advancing our understanding of cognition in both health and disease.

Methods

Inclusion criteria

Data were from the baseline visit of the ABCD study[21] (Collection #2573), $N = 11817$. Subjects were excluded if they did not meet each of the following criteria: meet the ABCD study's SST task-fMRI inclusion recommendations (in abcd_imgincl01.txt, imgincl_sst_include==1; $N = 3546$ excluded); have 2 SST fMRI runs of good quality (in mriqcrp20301.txt, iqc_sst_total_ser==iqc_sst_good_ser==2; $N = 677$ excluded); are successfully fit with the cognitive model of the SST ($N = 562$ excluded); have 2 SST fMRI runs in the release 4.0 minimally processed data ($N = 16$ excluded); have enough volumes acquired to cover the SST experiment (the last SST trial must have happened no more than 2 seconds after the final volume was acquired; $N = 16$ excluded); have mean framewise displacement of less than 0.5 mm for both runs (calculated using the method of [59]; $N = 1986$ excluded); have release 4.0 minimally processed events.tsv files of shape (181,3) for both runs ($N = 7$ excluded); and have consistent release 4.0 behavioral data (in release 4.0, for some subjects, the "sst.csv" files from ABCD Task fMRI SST Trial Level Behavior, abcd_sst_tlb01, disagreed with the minimally processed "events.tsv" files; for example, one trial might be labeled a go trial by one file and a stop trial by the other; $N = 102$ excluded). Then, we excluded siblings by randomly keeping one member from each family (using the genetic_paired_subjectid variables from gen_y_pihat; $N = 436$ excluded) and excluded subjects without scanner serial number recorded (in mri_y_adm_info, missing mri_info_deviceserialnumber; $N = 10$ excluded). Applying these inclusion criteria left us with a sample of $N = 4469$. For analyses involving the proactive delaying, a further 293 subjects were excluded who had no trials with probability of proactivity greater than 0.5 during at least one run, and therefore, by definition, a proactive delaying of 0 for all trials of at least one run. For these subjects, we were unable to examine within-subjects relationships between proactive delaying and brain activity. To maintain comparability of the between- and within-subjects analyses, we also excluded these subjects from the between-subjects analyses involving proactive delaying. Thus, analyses involving the proactive delaying used a sample of $N = 4176$.

Brain imaging

Imaging acquisition for the ABCD SST is detailed in other work.[60] We used the minimally processed data from ABCD release 4.0 (Collection #2573), which included distortion correction and motion correction[61]. We then further processed the images using Nilearn and FSL FLIRT: (1) initial volumes were removed (Siemens: 8, Philips: 8, GE DV25: 5, GE DV26 and other GE versions: 16); (2) the mean image in the time dimension was computed using mean_img from the Nilearn Image module; (3) the mean image was registered to an echo-planar imaging template in MNI152 space (SPM12's toolbox/OldNorm/EPI.nii) using FSL FLIRT, and an affine of this transformation was obtained; (4) the time-series of images was spatially normalized to MNI152 space with the affine from the previous step using FSL FLIRT; and (5) the images were smoothed with a Gaussian filter with a full-width at half maximum of 6 mm using smooth_img from the Nilearn Image module.

Bayesian modeling of cognitive dynamics

The PRAD model[40] incorporates latent dynamics that respond to endogenous and exogenous variables, with trait measures governing the interaction of such endogenous and exogenous variables with latent processes, giving rise to non-stationary dynamics. This allows the PRAD model to account for violations of context and stochastic independence. Overall, the PRAD model incorporates separate evidence accumulation (drift-diffusion) processes for the go and stop processes, similar to a canonical horse-race model[62]. However, in addition to typical drift-diffusion process parameters, PRAD includes (i) an explicit proactive inhibitory control mechanism that governs state switching between proactive and non-proactive cognitive states (governed by parameter θ_0), and hierarchical dynamics affecting: (ii) adaptive or maladaptive modulation of executive processes that respond to endogenous variables (error and performance monitoring, via parameter θ_1) to modulate the proclivity for activating proactive inhibitory control; (iii) adaptive or maladaptive tracking of environmental contingencies that governs how proactive inhibitory control responds to exogenous variables (stop signal delay, via parameter μ) to update beliefs about the stopping contingences and modulate the proactive delays in responding (PDR); (iv) adaptive or maladaptive attentional modulation of stopping expectancy (AMS) over trials that governs how the stopping process and SSRT respond to exogenous variables (number of trials since the last stop signal, n(SSD), via parameters γ_0, γ_1); (v) modulation of the go process threshold based on response to endogenous variables (choice errors, via parameters α_{G1}, α_{G2}); and (vi) trial-level modulation of the drift rate inferred from behavior. More details can be found in ref. [40].

General linear model analysis of fMRI

We fit general linear models to the fMRI BOLD recordings using Nilearn's FirstLevelModel. Condition and parametric regressors were modeled as impulses, with a duration of 0, and convolved with the SPM software's double gamma hemodynamic response function and the function's time derivative. Before fitting, the BOLD signal was scaled to percent signal-change from the mean in the time dimension. An AR1 model was used to whiten the data and design matrices to account for temporal autocorrelation in the BOLD signal.

To investigate the brain activation associated with SST conditions, for each subject and voxel, we fit the model
$$\text{BOLD}(t) = \beta_0 + (\text{HRF} * \text{Conditions})(t) + \text{Nuissance}(t) + \epsilon(t). \qquad \text{(Model 1)}$$
$\text{BOLD}(t)$ is the BOLD signal of the voxel at time t ($t \in \{1, \ldots, T\}$ for T the total number of volumes acquired); $(\text{HRF} * \text{Conditions})(t)$ is the value at t of the convolution with the hemodynamic response function HRF of condition regressor(s) Conditions; $\text{Nuissance}(t)$ is the effect at t of nuissance regressors, which were 6 motion parameters (translational and rotational displacement along each of three axes) and 6 cosine basis functions (corresponding to high-pass filtering at 0.01 Hz); and $\epsilon(t)$ is the model's error at t. We fit models with three sets of condition regressors:
$$\text{Conditions} = \beta_1(I_{\text{Go}} + I_{\text{Stop}}). \qquad \text{(Conditions 1)}$$
$$\text{Conditions} = \beta_1 I_{\text{Go}} + \beta_2 I_{\text{Stop}}. \qquad \text{(Conditions 2)}$$
$$\text{Conditions} = \beta_1 I_{\text{Correct go}} + \beta_2 I_{\text{Incorrect go}} + \beta_3 I_{\text{Correct late go}} + \beta_4 I_{\text{Incorrect late go}}$$
$$+ \beta_5 I_{\text{No response go}} + \beta_6 I_{\text{Correct stop}} + \beta_7 I_{\text{Incorrect stop}} + \beta_8 I_{\text{SSD stop}}. \qquad \text{(Conditions 3)}$$
$I_{\text{Condition}}$ is an indicator function indicating when the subject experiences Condition (for example, I_{Go} is 0 except at the moment when a subject is presented with a go trial). For between-

subjects analyses, we used Model 1 with Conditions 1 to obtain task activation (β_1); used Model 1 with Conditions 2 to obtain go activation (β_1) and stop activation (β_2); and used Model 1 with Conditions 3 to obtain correct stop versus correct go activation ($\beta_6 - \beta_1$), correct stop versus incorrect go activation ($\beta_6 - \beta_2$), and incorrect stop versus correct stop activation ($\beta_7 - \beta_6$).

To determine within-subjects associations between brain activity and trial-level variables, for each subject and voxel, we fit the model

$$\text{BOLD}(t) = \beta_0 + (\text{HRF} * \text{Conditions})(t) + (\text{HRF} * \text{Modulation})(t)$$
$$+ \text{Nuissance}(t) + \epsilon(t). \qquad \text{(Model 2)}$$

$(\text{HRF} * \text{Modulation})(t)$ is the value at t of the convolution with the hemodynamic response function of the parametric regressor Modulation. To investigate within-subjects associations between brain activity and SSRT on stop trials, probability of proactivity on all trials, proactive delaying on all trials, and observed reaction time on go trials, we set Modulation $= \beta_3 \text{SSRT}$, Modulation $= \beta_3 P(\text{Proactive})$, Modulation $= \beta_3 \text{Proactive delaying}$, and Modulation $= \beta_3 \text{Go RT}$, respectively, and used Model 2 with Conditions 2. Each of SSRT, $P(\text{Proactive})$, Proactive delaying, and Go RT was standardized over the conditions during which it assumed values by subtracting its mean and dividing by its standard deviation; that is, SSRT had mean 0 and standard deviation 1 over correct and incorrect stop trials (and was 0 on all other trials), $P(\text{Proactive})$ and Proactive delaying had mean 0 and standard deviation 1 over all trials, and Go RT had mean 0 and standard deviation 1 over go trials with a recorded response (and was 0 on all other trials).

We fit these regression models for each subject and each of their 2 SST runs. For each model, subject, and voxel, we combined the regression results from the 2 runs with a fixed effects model through FirstLevelModel's compute_contrast method. For each model and voxel, we estimated the effect of each scanner as the mean of the regression coefficients of the subjects who were scanned by it minus the grand mean of the regression coefficients of all subjects. Then, we adjusted the regression coefficients by subtracting the estimated scanner effects. All analyses of the regression coefficients used these adjusted values. We used a 1-sample Cohen's d (sample mean divided by sample standard deviation) to measure the effect sizes of regression coefficients.

Networks and regions of interest

We extracted the whole-brain regression coefficients in 2 networks and 3 sets of regions. We used the Shirer networks for our primary analyses and used the Yeo-17 networks, Shirer regions, cognitive control regions, and subcortical regions to test the stability of our within-subjects findings. For each subject and each regression coefficient of interest, we obtained the coefficient's value in each area (network or region) by calculating the mean of the subject's coefficients over the voxels belonging to the area. We used these area-average regression coefficients of each subject: to compute Cohen's d and Pearson r values for the network-level comparison of between- and within-subjects associations (Figure 5) and for the measurement of nonergodicity (Figure 9); and to compute Cohen's d values for the stability (Figure 6) and subgroup (Figure 8) analyses.

The Shirer networks and regions were obtained from[43]. To obtain the voxel-coordinates of the Yeo-17 networks, we used a mapping between the Brainnetome[63] and Yeo atlases[42]. We assembled the cognitive-control regions to include areas activated by the SST, two core default mode areas, and one core salience network and error-processing area. The regions activated by the SST were taken from a metaanalysis of 70 inhibitory control studies[26] (right anterior insula, right caudate, right inferior frontal gyrus, right middle frontal gyrus, right presupplementary motor area, and right supramarginal gyrus) and a study that segmented high-resolution structural MRI[64] (left and right subthalamic nucleus). To obtain the dorsal anterior cingulate cortex, we retrieved a Neurosynth automatic metaanalysis of 464 studies for the term "error" on 2023-09-21 and defined the region to be the 6 mm cube centered on the voxel with the highest metaanlysis Z-score. To obtain the posterior cingulate cortex and ventromedial prefrontal cortex, we retrieved a Neurosynth automatic metaanalysis of 777 studies for the term "default mode" on 2024-02-05; extracted clusters from this map using the connected_regions function in Nilearn's regions module with keyword argument "extract_type" set to "connected_components"; identified by eye the clusters corresponding to the posterior cingulate and ventromedial prefrontal cortex; and for each cluster defined the region to be the 6 mm cube centered on the voxel with the highest metaanalysis Z-score in the cluster. We obtained subcortical regions from a subcortical probabilistic atlas[65]. We resampled the atlas's probabilistic subcortical labels in 1 mm cubed MNI152 2009c nonlinear asymmetric space to the 2 mm cubed MNI152 space of our SPM echo-planar imaging template using resample_to_img from Nilearn's image module, and then thresholded these probabilistic maps at 0.5 to obtain region masks.

Stability analysis

To assess the stability of the within-subjects results, regression coefficients were resampled at varying sample sizes and the correlation was evaluated against the results in the full sample. Specifically, for each of SSRT, probability of proactivity, and proactive delaying, in each set of networks or regions: 10,000 samples of n subjects were drawn with replacement; the Cohen's d's of each sample's regression coefficients were calculated and Pearson correlated with the Cohen's d's of the full sample over the regions or networks; and the mean and 95% and 99% bootstrap confidence intervals were calculated of the correlation (SSRT and probability of proactivity $n = 25, 40, 70, 120, 200, 335, 560, 945, 1585, 2660, 4469$; proactive delaying $n = 25, 40, 70, 115, 195, 325, 540, 900, 1500, 2505, 4176$). The 95% and 99% bootstrap confidence intervals were calculated, respectively, as the intervals covering the 2.5th to 97.5th percentiles and 0.5th to 99.5th percentiles of the 10,000 correlations at each n.

We also directly examined the distributions of the Cohen's d's of the resamples as a function of n in regions of interest. Specifically, for each of SSRT, probability of proactivity, and proactive delaying, in each region of interest: 10,000 samples of n subjects were drawn with replacement; the Cohen's d of each sample's regression coefficients was calculated; and the mean and 95% and 99% bootstrap confidence intervals were calculated of the Cohen's d (SSRT and probability of proactivity $n = 25, 40, 70, 120, 200, 335, 560, 945, 1585, 2660, 4469$; proactive delaying $n = 25, 40, 70, 115, 195, 325, 540, 900, 1500, 2505, 4176$). The 95% and 99% bootstrap confidence intervals were calculated as the intervals covering, respectively, the 2.5th to 97.5th percentiles and 0.5th to 99.5th percentiles of the 10,000 Cohen's d's at each n.

Representational similarity analysis

For each Shirer network, and for each subject, we computed the correlation between the subject's within-subjects brain maps of SSRT and probability of proactivity, SSRT and proactive delaying, and probability of proactivity and proactive delaying over the voxels in the network. Density estimates used Seaborn's kdeplot function with each distribution of correlations over subjects normalized to 1 (common_norm=False) and limited to values between -1 and 1 (clip=[-1,1]); all other parameters, including those determining the kernel smoothing bandwidth, were kept at their defaults. We tested, for each network, whether there was a difference in the median correlation of SSRT and probability of proactivity and the median correlation of SSRT and proactive delaying; the median correlation of SSRT and probability of proactivity and the median correlation of probability of proactivity and proactive delaying; and the median correlation of SSRT and proactive delaying and the median correlation of probability of proactivity and proactive delaying.

Measuring nonergodicity of brain networks

For each of SSRT, probability of proactivity, and proactive delaying, and for each of the Shirer networks, we computed bootstrap distributions of a measure of nonergodicity: we drew 10,000 samples of n subjects with replacement; for each resample, using the resample's between- and within-subjects associations, we computed the fraction of subjects whose (within-subjects) brain association with the parameter had the opposite sign of the between-subjects brain association with the parameter, in the network; then, we computed the mean and 95% confidence interval over the resamples of the fraction of opposite signs (SSRT and probability of proactivity $n = 4469$; proactive delaying $n = 4176$). The between-subjects association was the Pearson correlation between correct stop versus correct go activation and the subject-average parameter, which was recomputed for the subjects in each resample. The 95% confidence interval was calculated as the interval covering the 2.5th to 97.5th percentiles of the 10,000 fractions of opposite signs. The goal of using bootstrapping was to account for the strength of the between-subjects results. Next, the nonergodicity of each network was defined as the three-dimensional vector whose ith coordinate was the mean (over bootstrap resamples) fraction of subjects with opposite signs in the network for the ith cognitive model parameter. Then, the Euclidean distances were computed between the vectors and hierarchical clustering was performed on the distances using the linkage function in scipy's cluster subpackage, hierarchy module with "method" set to "average".

Significance testing

Permutation tests were used for all significance testing. The tests used two-sided alternatives and 10,000 resamples and were performed with Scipy's permutation_test function. FDR correction was performed using the Benjamini-Hochberg procedure with Scipy's false_discovery_control function. P_{FDR} denotes an FDR-corrected P value. For a set of brain areas and an fMRI regression, we tested the null hypothesis for each brain area that the regression coefficients in the area had a mean of 0 (Figure 5) by computing the means of resamples in which the signs of the coefficients were randomly chosen (permutation_test permutation_type='samples'). Then, FDR correction was applied to the P values of all brain areas in the set (e.g., FDR correction was

applied to the P's of SSRT's regression coefficients over the Shirer networks). For a set of brain areas, an fMRI regression, and a behavioral measure, we tested the null hypothesis for each brain area that the Pearson correlation between subject-average regression coefficients in the brain area and subject-average behavioral measures was 0 (Figure 5) by computing the Pearson correlations of resamples in which regression coefficients were randomly paired with behavioral measures (permutation_test permutation_type ='pairings'). Then, FDR correction was applied to the P values of all brain areas in the set (e.g., FDR correction was applied to the P's of correlations between correct stop versus correct go activation and SSRT over the Shirer networks). For a set of brain areas and a param γ_1 or θ_1, we tested the null hypothesis for each brain area that the mutually exclusive subgroups of subjects with param < 0 and with param > 0 had different mean regression coefficients (Figure 8) by computing the differences between the means for resamples in which coefficients were randomly assigned to param < 0 and param > 0 (permutation_test permutation_type='independent'). Then, FDR correction was applied to the P values of all brain areas in the set (e.g., FDR correction was applied to the P's of mean differences between $\gamma_1 < 0$ and $\gamma_1 > 0$ over the Shirer networks). For the Shirer networks, we tested the null hypothesis for each brain area that there was no difference in the area between the median correlation of SSRT and probability of proactivity and the median correlation of SSRT and proactive delaying; the median correlation of SSRT and probability of proactivity and the median correlation of probability of proactivity and proactive delaying; and the median correlation of SSRT and proactive delaying and the median correlation of probability of proactivity and proactive delaying (Figure 7). We computed the differences between the medians of resamples in which correlations were randomly exchanged within subjects (permutation_test permutation_type='samples'). Then, FDR correction was applied to the P values of all comparisons in all areas. (Since there are 14 Shirer networks and 3 tests per network, FDR correction was applied over 14×3 P's.)

Software

Data were processed and analyzed using Python (version 3.9.16), Scipy (version 1.11.4), Seaborn (version 0.13.2), Nilearn (version 0.10.1), and FSL FLIRT (version 6.0).

References

Lebowitz, J. L. & Penrose, O. Modern ergodic theory. *Physics Today* **26**, 23-29 (1973).
Badino, M. The foundational role of ergodic theory. *Foundations of science* **11**, 323-347 (2006).
Moore, C. C. Ergodic theorem, ergodic theory, and statistical mechanics. *Proceedings of the National Academy of Sciences* **112**, 1907-1911 (2015).
Peters, O. The ergodicity problem in economics. *Nature Physics* **15**, 1216-1221 (2019).
Fisher, A. J., Medaglia, J. D. & Jeronimus, B. F. Lack of group-to-individual generalizability is a threat to human subjects research. *Proc Natl Acad Sci U S A* **115**, E6106-E6115, doi:10.1073/pnas.1711978115 (2018).
Molenaar, P. C. On the implications of the classical ergodic theorems: Analysis of developmental processes has to focus on intra-individual variation. *Developmental Psychobiology: The Journal of the International Society for Developmental Psychobiology* **50**, 60-69 (2008).
Medaglia, J. D., Ramanathan, D. M., Venkatesan, U. M. & Hillary, F. G. The challenge of non-ergodicity in network neuroscience. *Network: Computation in Neural Systems* **22**, 148-153 (2011).
Simpson, E. H. The interpretation of interaction in contingency tables. *Journal of the Royal Statistical Society: Series B (Methodological)* **13**, 238-241 (1951).
Blyth, C. R. On Simpson's paradox and the sure-thing principle. *Journal of the American Statistical Association* **67**, 364-366 (1972).
Sprenger, J. & Weinberger, N. Simpson's paradox. *The Stanford Encyclopedia of Philosophy* (2021).
Wagner, C. H. Simpson's paradox in real life. *The American Statistician* **36**, 46-48 (1982).
Dhakal, V., Feit, A. M., Kristensson, P. O. & Oulasvirta, A. in *Proceedings of the 2018 CHI conference on human factors in computing systems.* 1-12.
Heitz, R. P. The speed-accuracy tradeoff: history, physiology, methodology, and behavior. *Frontiers in neuroscience* **8**, 86875 (2014).
Molenaar, P. C. A manifesto on psychology as idiographic science: Bringing the person back into scientific psychology, this time forever. *Measurement* **2**, 201-218 (2004).
Kievit, R. A., Frankenhuis, W. E., Waldorp, L. J. & Borsboom, D. Simpson's paradox in psychological science: a practical guide. *Frontiers in psychology* **4**, 513 (2013).
Molenaar, P. & Newell, K. M. *Individual pathways of change: Statistical models for analyzing learning and development.* (American Psychological Association, 2010).
Hamaker, E. L. Why researchers should think" within-person": A paradigmatic rationale. (2012).
Molenaar, P. C. & Campbell, C. G. The new person-specific paradigm in psychology. *Current directions in psychological science* **18**, 112-117 (2009).
Castro-Schilo, L. & Ferrer, E. Comparison of nomothetic versus idiographic-oriented methods for making predictions about distal outcomes from time series data. *Multivariate Behavioral Research* **48**, 175-207 (2013).
Hamaker, E., Ceulemans, E., Grasman, R. & Tuerlinckx, F. Modeling affect dynamics: State of the art and future challenges. *Emotion Review* **7**, 316-322 (2015).
Volkow, N. D. *et al.* The conception of the ABCD study: From substance use to a broad NIH collaboration. *Dev Cogn Neurosci* **32**, 4-7, doi:10.1016/j.dcn.2017.10.002 (2018).

22 Diamond, A. Executive functions. *Annual review of psychology* **64**, 135-168 (2013).

23 Logan, G. D. & Cowan, W. B. On the ability to inhibit thought and action: A theory of an act of control. *Psychological review* **91**, 295 (1984).

24 Verbruggen, F. & Logan, G. D. Response inhibition in the stop-signal paradigm. *Trends in cognitive sciences* **12**, 418-424 (2008).

25 Aron, A. R., Robbins, T. W. & Poldrack, R. A. Inhibition and the right inferior frontal cortex: one decade on. *Trends in cognitive sciences* **18**, 177-185 (2014).

26 Cai, W., Ryali, S., Chen, T., Li, C. S. & Menon, V. Dissociable roles of right inferior frontal cortex and anterior insula in inhibitory control: evidence from intrinsic and task-related functional parcellation, connectivity, and response profile analyses across multiple datasets. *J Neurosci* **34**, 14652-14667, doi:10.1523/JNEUROSCI.3048-14.2014 (2014).

27 Cai, W. *et al.* Hyperdirect insula-basal-ganglia pathway and adult-like maturity of global brain responses predict inhibitory control in children. *Nature communications* **10**, 4798 (2019).

28 Cai, W., Griffiths, K., Korgaonkar, M. S., Williams, L. M. & Menon, V. Inhibition-related modulation of salience and frontoparietal networks predicts cognitive control ability and inattention symptoms in children with ADHD. *Molecular psychiatry* **26**, 4016-4025 (2021).

29 Cai, W., Chen, T., Ide, J. S., Li, C.-S. R. & Menon, V. Dissociable fronto-operculum-insula control signals for anticipation and detection of inhibitory sensory cue. *Cerebral cortex* **27**, 4073-4082 (2017).

30 Friedman, N. P. & Robbins, T. W. The role of prefrontal cortex in cognitive control and executive function. *Neuropsychopharmacology* **47**, 72-89 (2022).

31 Zhang, S. *et al.* Independent component analysis of functional networks for response inhibition: Inter-subject variation in stop signal reaction time. *Human brain mapping* **36**, 3289-3302 (2015).

32 Boehler, C. N., Appelbaum, L. G., Krebs, R. M., Hopf, J.-M. & Woldorff, M. G. The influence of different Stop-signal response time estimation procedures on behavior–behavior and brain–behavior correlations. *Behavioural brain research* **229**, 123-130 (2012).

33 Aron, A. R. & Poldrack, R. A. Cortical and subcortical contributions to stop signal response inhibition: role of the subthalamic nucleus. *Journal of Neuroscience* **26**, 2424-2433 (2006).

34 Li, C. S., Huang, C., Constable, R. T. & Sinha, R. Imaging response inhibition in a stop-signal task: neural correlates independent of signal monitoring and post-response processing. *J Neurosci* **26**, 186-192, doi:10.1523/JNEUROSCI.3741-05.2006 (2006).

35 Li, C. S. *et al.* Neural correlates of impulse control during stop signal inhibition in cocaine-dependent men. *Neuropsychopharmacology* **33**, 1798-1806, doi:10.1038/sj.npp.1301568 (2008).

36 Chevrier, A. & Schachar, R. J. BOLD differences normally attributed to inhibitory control predict symptoms, not task-directed inhibitory control in ADHD. *J Neurodev Disord* **12**, 8, doi:10.1186/s11689-020-09311-8 (2020).

37 Chaarani, B. *et al.* Baseline brain function in the preadolescents of the ABCD Study. *Nat Neurosci* **24**, 1176-1186, doi:10.1038/s41593-021-00867-9 (2021).

Swick, D., Ashley, V. & Turken, U. Are the neural correlates of stopping and not going identical? Quantitative meta-analysis of two response inhibition tasks. *Neuroimage* **56**, 1655-1665 (2011).

Mangalam, M. & Kelty-Stephen, D. G. Point estimates, Simpson's paradox, and nonergodicity in biological sciences. *Neurosci Biobehav Rev* **125**, 98-107, doi:10.1016/j.neubiorev.2021.02.017 (2021).

Mistry, P. K., Warren, S. L., Cai, W., Branigan, N. K. & Menon, V. Proactive, reactive, and attentional dynamics in inhibitory control. (In Prep.).

Braver, T. S. The variable nature of cognitive control: a dual mechanisms framework. *Trends in cognitive sciences* **16**, 106-113 (2012).

Yeo, B. T. *et al.* The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J Neurophysiol* **106**, 1125-1165, doi:10.1152/jn.00338.2011 (2011).

Shirer, W. R., Ryali, S., Rykhlevskaia, E., Menon, V. & Greicius, M. D. Decoding subject-driven cognitive states with whole-brain connectivity patterns. *Cereb Cortex* **22**, 158-165, doi:10.1093/cercor/bhr099 (2012).

Aron, A. R. *et al.* Converging evidence for a fronto-basal-ganglia network for inhibitory control of action and cognition. *Journal of Neuroscience* **27**, 11860-11864 (2007).

van den Wildenberg, W. P., Ridderinkhof, K. R. & Wylie, S. A. Towards Conceptual Clarification of Proactive Inhibitory Control: A Review. *Brain Sciences* **12**, 1638 (2022).

Menon, V. 20 years of the default mode network: A review and synthesis. *Neuron* (2023).

Ryali, S., Zhang, Y., de Los Angeles, C., Supekar, K. & Menon, V. Deep learning models reveal replicable, generalizable, and behaviorally relevant sex differences in human functional brain organization. *Proceedings of the National Academy of Sciences* **121**, e2310012121 (2024).

Meyer-Lindenberg, A. The non-ergodic nature of mental health and psychiatric disorders: implications for biomarker and diagnostic research. *World Psychiatry* **22**, 272 (2023).

Craig, A. D., Chen, K., Bandy, D. & Reiman, E. M. Thermosensory activation of insular cortex. *Nature neuroscience* **3**, 184-190 (2000).

Menon, V. Insular cortex: A hub for saliency, cognitive control, and interoceptive awareness. (2024).

Molnar-Szakacs, I. & Uddin, L. Q. Anterior insula as a gatekeeper of executive control. *Neuroscience & Biobehavioral Reviews* **139**, 104736 (2022).

Uddin, L. Q. & Menon, V. The anterior insula in autism: under-connected and under-examined. *Neuroscience & Biobehavioral Reviews* **33**, 1198-1203 (2009).

Segal, A. *et al.* Regional, circuit and network heterogeneity of brain abnormalities in psychiatric disorders. *Nature Neuroscience* **26**, 1613-1629 (2023).

Xie, C. *et al.* A shared neural basis underlying psychiatric comorbidity. *Nature medicine* **29**, 1232-1242 (2023).

Kriegeskorte, N., Mur, M. & Bandettini, P. A. Representational similarity analysis-connecting the branches of systems neuroscience. *Frontiers in systems neuroscience* **2**, 249 (2008).

Haxby, J. V., Connolly, A. C. & Guntupalli, J. S. Decoding neural representational spaces using multivariate pattern analysis. *Annual review of neuroscience* **37**, 435-456 (2014).

Aron, A. R. From reactive to proactive and selective control: developing a richer model for stopping inappropriate responses. *Biological psychiatry* **69**, e55-e68 (2011).

Luna, B., Marek, S., Larsen, B., Tervo-Clemmens, B. & Chahal, R. An integrative model of the maturation of cognitive control. *Annual review of neuroscience* **38**, 151-170 (2015).

Power, J. D., Barnes, K. A., Snyder, A. Z., Schlaggar, B. L. & Petersen, S. E. Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *Neuroimage* **59**, 2142-2154, doi:10.1016/j.neuroimage.2011.10.018 (2012).

Casey, B. J. *et al.* The Adolescent Brain Cognitive Development (ABCD) study: Imaging acquisition across 21 sites. *Dev Cogn Neurosci* **32**, 43-54, doi:10.1016/j.dcn.2018.03.001 (2018).

Hagler, D. J., Jr. *et al.* Image processing and analysis methods for the Adolescent Brain Cognitive Development Study. *Neuroimage* **202**, 116091, doi:10.1016/j.neuroimage.2019.116091 (2019).

Band, G. P., Van Der Molen, M. W. & Logan, G. D. Horse-race model simulations of the stop-signal procedure. *Acta psychologica* **112**, 105-142 (2003).

Fan, L. *et al.* The Human Brainnetome Atlas: A New Brain Atlas Based on Connectional Architecture. *Cereb Cortex* **26**, 3508-3526, doi:10.1093/cercor/bhw157 (2016).

Forstmann, B. U. *et al.* Cortico-subthalamic white matter tract strength predicts interindividual efficacy in stopping a motor response. *Neuroimage* **60**, 370-375, doi:10.1016/j.neuroimage.2011.12.044 (2012).

Pauli, W. M., Nili, A. N. & Tyszka, J. M. A high-resolution probabilistic in vivo atlas of human subcortical brain nuclei. *Sci Data* **5**, 180063, doi:10.1038/sdata.2018.63 (2018).

Figures

Figure 1. Conceptual overview of the study and key findings. The figure illustrates the methodology for between-subjects and within-subjects analyses and the concept of nonergodicity, and previews the main results. **a,** Between-subjects analysis. Subject-average brain activation in each voxel is correlated with a subject-average cognitive measure across the population. **b,** Within-subjects analysis. For each individual, the time-series of brain activity in each voxel is associated with the time-series of a cognitive measure. **c,** Simpson's paradox. Simpson's paradox occurs when associations between subjects and within subjects show conflicting directions; it exemplifies nonergodicity in the behavioral sciences. **d,** Study aims. We examined brain-behavior associations for nonergodicity; tested our within-subjects results for stability and robustness; used these results to probe the brain implementations of proactive and reactive control and adaptive regulation of inhibitory control; and investigated how nonergodicity varied by brain network. Nonergodic patterns in brain-behavior associations were consistently observed, revealing that group-level (between-subjects) and individual-level (within-subjects) associations yield divergent results for inhibitory control processes. This challenges the common assumption that findings from such group-level analyses can be directly applied to understand individual-level cognitive processes.

Figure 2. Computational modeling of inhibitory control dynamics in the stop signal task. a, Stop signal task. On go trials subjects should respond by pressing a button to indicate the arrow direction, and on stop trials should inhibit their response when the stop signal appears. **b,** Race model. The computational model's foundation is that go and stop processes compete, with the first to finish determining the behavioral outcome. **c,d,** Key model parameters. Proactive delaying is the delay in initiating the go process when a subject uses a proactive control strategy. SSRT is the time it takes for the stop process to complete. Probability of proactivity is the probability that a subject uses a proactive control strategy. **e,** Trial-by-trial dynamics. SSRT, probability of proactivity, and proactive delaying are inferred for each trial for each subject. The computational model allows for a detailed, dynamic analysis of inhibitory control processes with trial-level temporal resolution, enabling the investigation of within-subjects variability and nonergodic patterns in cognitive control.

Figure 3. The PRAD cognitive model. a, The PRAD model infers latent variables for each subject from their observed go and stop failure rates, response times, and choice accuracy. The latent variables relate to three mechanisms of dynamic inhibitory control: the basic reactive inhibitory process (red 1), proactive delaying of responses (green 2), and modulation of stopping expectancy (yellow 3). **b,** Visualization of the three mechanisms of dynamic inhibitory control. **c,** Mathematical details of how the model parameterizes the go and stop processes.

Figure 4. Divergence of between-subjects and within-subjects brain-behavior associations in inhibitory control. a, Between-subjects analysis. Whole-brain correlation maps showing associations between subject-average brain activation (correct stop versus correct go activation) and subject-average cognitive model parameters (SSRT, probability of proactivity, and proactive delaying). Thresholded at Pearson $r \geq 0.05$. **b,** Within-subjects analysis. Whole-brain Cohen's d maps showing associations between trial-by-trial brain activity and cognitive model parameters (SSRT, probability of proactivity, and proactive delaying). SSRT associations were computed on stop trials; probability of proactivity and proactive delaying associations were computed on all trials. Thresholded at Cohen's $d \geq 0.1$. For both between- and within-subjects analyses: SSRT and probability of proactivity $N = 4469$; proactive delaying $N = 4176$. Striking differences were observed comparing between-subjects and within-subjects associations across multiple brain regions. This divergence provides evidence for nonergodicity in inhibitory control processes, challenging the assumption that group-level findings can be directly applied to understand individual-level cognitive dynamics.

Figure 5. Network-level comparison of between-subjects and within-subjects brain-behavior associations in inhibitory control. The between-subjects analysis correlated subject-average brain activation (correct stop versus correct go activation) and subject-average cognitive model parameters (SSRT, probability of proactivity, and proactive delay). The within-subjects analysis regressed trial-by-trial brain activity on trial-by-trial cognitive model parameters (SSRT, probability of proactivity, and proactive delay). Effect sizes are shown for both analyses (between-subjects: Pearson r; within-subjects: Cohen's d). Statistical significance is indicated by colored asterisks: red for between subjects ($P_{FDR} < 0.01$) and blue for within subjects ($P_{FDR} < 0.01$). Networks are based on the Shirer parcellation.[43] Differences in the existence and direction of associations between between-subjects and within-subjects analyses were observed across multiple brain networks, including the anterior and posterior salience, left and right frontoparietal, and dorsal and ventral default mode networks.

Figure 6. Stability analysis of within-subjects associations. a-c, Stability plots for SSRT **(a)**, probability of proactivity **(b)**, and proactive delaying **(c)**. In each panel, the top plots show the distributions of correlation between the effect in resamples and the effect in the full sample, while the bottom plots show the distributions of the effect in selected regions. For each plot, 10,000 resamples of n subjects were drawn with replacement for each sample size n, and Cohen's d was calculated for each resample. For top panel plots, correlations were then computed for each resample over the areas belonging to the set of brain areas. The dashed lines depict 95% and 99% bootstrap confidence intervals. Within-subjects associations demonstrated stability for all 3 cognitive model parameters across 5 different collections of brain areas and in 5

regions of interest, even at modest sample sizes. This stability supports the reliability of brain-behavior associations in inhibitory control processes.

Figure 7. Dissociated brain representations of reactive and proactive control processes. For each subject, correlations were computed between pairs of within-subjects brain maps for the model parameters (SSRT, probability of proactivity, and proactive delaying) over each brain network. **a,** Distributions of correlations over subjects. Kernel density estimation was performed. **b,** Average correlations over subjects. The height of the bars is the median. Statistical significance is indicated by asterisks (* $P_{FDR} < 0.01$). Reactive (SSRT) and proactive (probability of proactivity, proactive delaying) control processes showed dissociated representations across all brain networks. This dissociation suggests that the brain employs distinct neural resources for reactive and proactive aspects of inhibitory control. In contrast, the two proactive measures showed high similarity, validating their representation of related cognitive processes.

Figure 8. Distinct brain-behavior associations for adaptive and maladaptive regulators of inhibitory control. a, Within-subjects associations of SSRT with brain activity between γ_1 subgroups. We identified two distinct subgroups of subjects with opposite profiles of attentional modulation. Subjects showed either maladaptive regulation ($\gamma_1 < 0$, $N = 2513$) or adaptative regulation ($\gamma_1 > 0$, $N = 1956$) of their expectancy of stopping over time. The two groups showed differences in within-subjects associations between SSRT and brain activity in various networks. This demonstrates that individual differences in attentional dynamics play a role in shaping the relationship between neural activity and inhibitory control processes. A colored asterisk indicates that a network or region's associations were nonzero among subjects with $\gamma_1 < 0$ (blue * $P_{FDR} < 0.01$) or among subjects with $\gamma_1 > 0$ (purple * $P_{FDR} < 0.01$). **b,** Within-subjects associations of probability of proactivity with brain activity between θ_1 subgroups. We identified two distinct subgroups of subjects with opposite profiles of performance monitoring. Subjects showed either adaptive regulation ($\theta_1 < 0$, $N = 3054$) or maladaptive regulation ($\theta_1 > 0$, $N = 1415$) of their proactivity over time. The two groups showed differences in within-subjects associations between probability of proactivity and brain activity in various networks. This demonstrates that the neural correlates of proactive control are influenced by an individual's strategy for adjusting proactivity in response to task outcomes. A colored asterisk indicates that a network or region's associations were nonzero among subjects with $\theta_1 < 0$ (blue * $P_{FDR} < 0.01$) or among subjects with $\theta_1 > 0$ (purple * $P_{FDR} < 0.01$). For both panels, a black asterisk indicates that associations had different distributions between the two subgroups (* $P_{FDR} < 0.01$, n.s. $P_{FDR} \geq 0.01$).

Figure 9. Nonergodicity across brain networks. a, Degree of nonergodicity of brain networks. Bar plots show the extent of nonergodicity for each brain network across three cognitive model parameters (SSRT, probability of proactivity, proactive delaying). The height of the bars is the mean fraction of subjects showing opposite-sign associations compared to between-subjects results. The bars are shaded by the magnitude of the between-subjects correlations. Error bars show 95% bootstrap confidence intervals. **b,** 3-dimensional embedding of each network. Networks are represented as points based on their nonergodicity measures for the three parameters. **c-d,** Hierarchical clustering of networks based on their nonergodicity profiles. **c,** Dendrogram of clustering. **d,** Euclidean distances between network embeddings. Brain networks exhibited varying degrees of nonergodicity. The anterior salience network showed the highest level of nonergodicity and a unique profile of nonergodicity distinct from that of all other networks.