Uncertainty Tolerance Dissociates Sensation Seeking and Impulsivity

Ern Wong[1,2*], Tobias U. Hauser[3,4,5], Aswath Chandrasekaran[2], Pietro Pietrini[1], and Charley M. Wu[2,6,7]

[1]MoMiLab, IMT School for Advanced Studies Lucca, Lucca, Italy
[2]Human and Machine Cognition Lab, University of Tübingen, Tübingen, Germany
[3]Department of Psychiatry and Psychotherapy, Faculty of Medicine, University of Tübingen, Tübingen, Germany
[4]German Centre for Mental Health (DZPG), Tübingen, Germany
[5]Max Planck UCL Centre for Computational Psychiatry and Ageing Research, University College London, London, UK
[6]Centre for Cognitive Science, Technische Universität Darmstadt, Darmstadt, Germany
[7]Hessian AI, Darmstadt, Germany
[*]ern.wong@imtlucca.it

ABSTRACT

Sensation seeking (SS) is the drive to pursue novel and intense experiences, often despite associated risks. Thus, it is commonly assumed to reflect a heightened valuation of stimulation. Yet, intense experiences are also highly informative, raising a key question: Does SS reflect a preference for intensity or for information/uncertainty? We address this in two studies. In Experiment 1, we fit learning-based choice models to a prior dataset and found that an uncertainty-weighted account explained behaviour better than a purely intensity-weighted account, with higher SS linked to reduced uncertainty aversion. In a preregistered Experiment 2, we investigated the joint contribution of intensity and uncertainty. We demonstrate that SS contributions are context-dependent: when stimulation was valued less, higher SS predicted greater uncertainty-tolerance (i.e., higher information-seeking). In contrast, when stimulation was highly valued, higher SS shifted toward uncertainty avoidance. By contrast, impulsivity (IP) showed a distinct profile consistent with uncertainty aversion when stimulation was valued less. These results reveal SS as a joint, context-dependent function of value and uncertainty processing, rather than a single process. This framework reconciles how SS can support adaptive exploration in some settings yet promote maladaptive risk in others, and it provides a mechanistic dissociation from impulsivity with implications for targeted intervention and developmental theory.

Introduction

Sensation seeking (SS) is a personality trait characterised by a preference for "varied, novel, complex, and intense sensations and experiences," often pursued with a disregard for potential aversive consequences across physical, social, legal, and financial domains[1]. This trait peaks during adolescence, driving a range of exploratory and risky behaviours that are both a hallmark of this developmental period[2,3] and a significant risk factor for maladaptive outcomes with high societal costs[4,5]. For instance, SS is a robust predictor of the initiation of alcohol and substance use[6,7] and engagement in antisocial behaviours[8,9].

However, SS is also linked to long-term adaptive outcomes, including enhanced psychological well-being[10,11] and forms of "positive risk-taking", such as social exploration and pursuing challenging goals, which are developmentally normative and beneficial[12–15]. This paradox is mirrored at the neurobiological level, where dopamine (DA) systems, crucial for processing incentive salience and reward learning, are fundamentally re-organised during adolescence, leading to individual differences in SS[16]. High sensation seekers exhibit distinct dopaminergic profiles, including elevated DA in the caudate nucleus and attenuated DA turnover rates[17], yet the precise cognitive mechanisms that channel this drive towards functional or dysfunctional ends remain unclear.

An important difficulty is that SS is often conflated with impulsivity (IP), as both traits are associated with risk-taking. However, a wealth of evidence now supports a dual-process model of adolescent risk-taking, which posits these are distinct constructs subserved by dissociable neurocognitive systems with divergent developmental timelines[3,20–22]. SS is thought to be driven by the early-maturing, hyper-responsive socioemotional system, while IP reflects the relative immaturity of the top-down cognitive control system. Empirically, the distinction is clear. SS and IP show only modest correlations and are differentially associated with real-world behaviours[21,23]. For example, while SS predicts initial experimentation with novel substances (i.e. alcohol, drugs), IP more strongly predicts an escalation to substance use disorders and difficulties with sustained, goal-directed behaviour[24–26].

Thus, a critical gap remains in our understanding of how these traits differentially shape decision-making under uncertainty. In the lab, IP is often associated with poorer response inhibition and with steeper delay discounting, reflecting a stronger preference for smaller, immediate rewards[27–29]. SS, by contrast, shows more robust associations with risk-taking on gambling-like tasks such as the Balloon Analogue Risk Task[30], with its links to classic inhibitory-control deficits be-

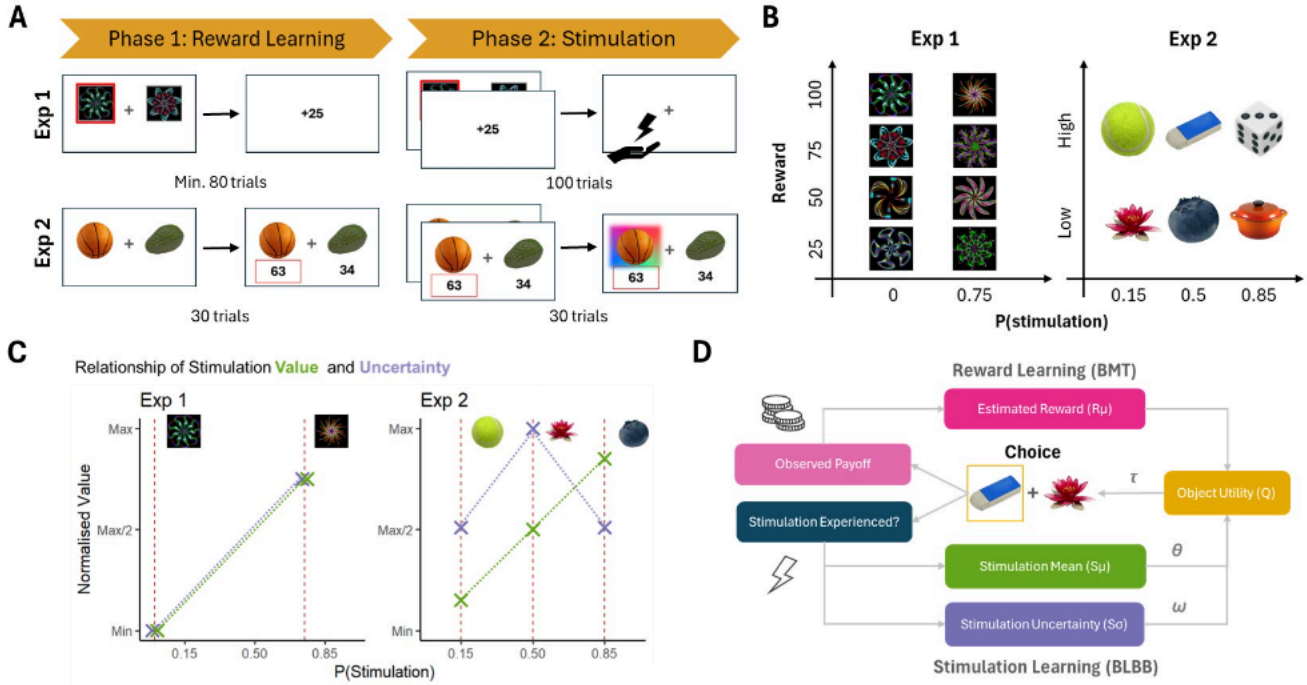

Figure 1. (**A**) Task overview. In both experiments, participants completed a reward-learning phase followed by a stimulation phase. In *Exp. 1*, participants first learned the reward payoffs of eight bandits until they reached a performance criterion. In the test phase, four bandits were associated with stimulation and four with no stimulation; each trial paired a stimulation bandit with a non-stimulation bandit. In *Exp. 2*, participants completed four rounds of the same two-phase structure (60 trials per round) with six objects per round defined by crossing reward level with stimulation probability. Rewards were learned with full feedback, then the same reward contingencies were tested under stimulation, with all object pairs presented repeatedly in random order. (**B**) Stimulus sets. *Exp. 1* used eight fractals parameterised by payoff (25/50/75/100) and stimulation assignment ($p = 0$ or 0.75); *Exp. 2* used six Bank of Standardized Stimuli (BOSS)[18,19] images parameterised by reward level (low/high) and stimulation probability (0.15/0.50/0.85). (**C**) Sampling scheme. In *Exp. 1*, the limited probability set caused the stimulation *mean* (green) and *uncertainty* (purple) to be correlated. *Exp. 2* removed this confound by sampling a wider range of probabilities, thereby orthogonalising intensity and uncertainty. (**D**) Computational model schematic: object utility (U) depends on estimated reward ($R\mu$) and contributions from stimulation mean ($S\mu$) weighted by θ, and uncertainty ($S\sigma$) weighted by ω, given observed payoffs and whether stimulation occurred. Object utilities are then passed through a softmax controlled by temperature τ, which governs the stochasticity of the choice. These quantities are updated using the learning rules described in the main text. BMT = Bayesian Mean Tracker; BLBB = Bayesian Learner Beta–Binomial.

ing weaker and less consistent, consistent with work suggesting that sensation seeking is partly dissociable from executive control[31]. Despite these distinctions, few studies have directly compared the cognitive processes underlying their influence on exploratory choice. While dual-process models offer only coarse-grained accounts of why adolescence is a period of peak risk-taking, they provide less insight into the precise cognitive mechanisms that differentiate how sensation seeking and impulsivity operate during decision-making.

Here, we suggest that SS and IP can be more precisely understood from the lens of exploratory behaviour. The exploration-exploitation trade-off offers a powerful computational framework for understanding these underlying mechanisms[32–35]. In any new environment, an agent must balance exploiting known options for a reliable reward against exploring uncertain alternatives to gather information that could lead to better long-term outcomes. Excessive exploitation leads to inflexibility, while excessive exploration can be costly and dangerous. For adolescents, navigating this trade-off is a core developmental task, as they must build a rich model of the world to guide future decisions[36,37]. The high-risk, high-exploration signature of adolescence suggests this framework is vital for mechanistically understanding traits like SS and IP[38,39].

This trait-exploration connection is supported by key findings in developmental science. Adolescents are not globally risk-averse; rather, they show a specific tolerance for ambiguity. Compared to adults, they are more likely to make risky choices when outcome probabilities are unknown, yet are similarly risk-averse when probabilities are explicit[40,41]. This suggests that adolescent risk-taking is not a simple failure of control, but may reflect a potent, information-seeking drive

to reduce epistemic uncertainty about the environment. This positions adolescence as a "developmental window" for exploration, where risky choices serve the adaptive function of learning[42,43]. Correspondingly, a parallel increase in trait SS during adolescence may reflect this epistemic drive[4].

Thus, whether or not SS represents an epistemic drive has been clouded by a fundamental confound: the experiences that feel most intense also tend to be the most informative. Subjective intensity tracks surprise and contrast, while information gain tracks prediction error and entropy reduction. Yet, both peak when outcomes are uncertain, extreme, or high-contrast, such that intensity and information co-occur. For example, moving to an entirely foreign country is intense and yields significant information gains about norms, language, and networks, whereas moving to a nearby town is of lower intensity and typically less informative.

To move beyond traditional, intensity-centric accounts of sensation seeking, we therefore need to separate the influence of intensity and information. Traditional conceptualisations of SS have largely emphasised heightened sensitivity to stimulus intensity and reward magnitude[25,44]. However, high-variance options are both intensely stimulating and maximally uncertain, thus making it difficult to disentangle a preference for intensity from a preference for information[4]. Recent computational work has also begun to challenge the intensity-centric view. Agent-based simulations demonstrate that a peak in adolescent risk-taking can emerge simply from a drive to learn efficiently in a changing world, without invoking any special motivational changes or a "sensation-seeking" parameter[36]. This suggests that an information-seeking drive may be sufficient to explain SS-related behaviours. However, behavioural tasks that can rigorously and orthogonally disentangle intensity-driven preferences from uncertainty-driven exploration are critically lacking.

In this present study, we address this gap through a two-part approach. In Experiment 1, we demonstrate that previous SS-related behaviours can be better explained by an uncertainty account rather than an intensity account. In Experiment 2, we extended this work with a new, preregistered, web-based experiment to demonstrate that greater uncertainty tolerance in SS is context-dependent, and that SS and IP show distinct relationships with uncertainty. Together, these results propose a novel mechanistic framework for SS that separates it from superficially related risky behaviours, which have important implications for developmental theory, personality psychology, and psychiatry.

Results

We report our results in two parts. Exp. 1 reanalyses prior data[16] to adjudicate whether intensity or uncertainty better accounts for SS-related behaviour. Exp. 2 extends our findings and uses a preregistered[45], orthogonalised design to test their joint and unique roles and to compare SS with IP using behavioural and model-based analyses.

Exp. 1: Uncertainty Not Intensity Accounts For SS

We reanalysed the Norbury et al. (2015)[16] dataset to investigate whether SS is more closely related to responses to stimulation *intensity* (expected probability, $S\mu$) or to stimulation *uncertainty* (variance around that probability, $S\sigma$). Unlike the original analysis, which assumed full and stable knowledge of reward and stimulation outcomes for each choice, we fit hierarchical generative Bayesian learning models (Bayesian Learner Beta Binomial, BLBB; see Methods for details) that updated beliefs about stimulation probability over trials.

Critically, because the original task design made intensity and uncertainty highly correlated (Fig. 2C), we were unable to estimate their interaction effects in a single model. We therefore compared a stimulation intensity learner ($R-\theta$) versus a stimulation uncertainty learner ($R-\omega$), each assigning a bonus on either learned intensity ($S\mu$) or uncertainty ($S\sigma$) values. We compare these Bayesian learning models alongside the original full-knowledge model (Norbury 2015).

Bayesian model comparisons (Fig. 2A) favoured the $R-\omega$ (LOOIC: 6712) over $R-\theta$ (LOOIC: 7673) and original (LOOIC: 7476) alternatives, indicating that allowing uncertainty to guide choice better captured behaviour in this design (see Methods for details on model specifications and fitting). We found a similar association between SS and learned stimulation probability ($r = 0.3$, $p = .01$; Fig. 2B), which replicated the original correlation reported in Norbury 2015 using a model with full knowledge of stimulation probabilities (Fig. 2C). However, in our winning model, we found a link between the uncertainty weight, ω, and SS scores ($r = 0.27$, $p = .03$; Fig. 2D). Although the strength of this correlation was similar to that found in the other two models for stimulation intensity, the uncertainty-weighted learner provided a better prediction fit based on relative LOO-IC. Furthermore, ω estimates were negative on average, indicating that most participants treated uncertainty as something to be avoided – consistent with the fact that uncertainty was not informative about reward in this task. However, the positive association between ω and SS suggests that individuals with higher SS showed less uncertainty avoidance. In other words, while most participants discounted uncertain options, those high in sensation seeking were more willing to sample them, reflecting a greater tolerance for uncertainty even when it offered no strategic benefit.

To further corroborate this result, we constructed a series of logistic mixed-effects models to assess the impact of $S\mu$ and $S\sigma$ on the probability of choosing an option in a given trial. We investigated how the probability of each choice was influenced by the reward payoff ($R\mu$) and either estimated stimulation probability or stimulation uncertainty, including interactions with SS scores. To account for individual variations, we included participant-specific random intercepts. Ranking models by AIC, the best model was again $R-\omega$, which used stimulation uncertainty and interaction with SS scores (AIC: 6394.9; Fig. 2E), compared to the second-best-fitting model, $R-\theta$ (AIC: 7718.2).

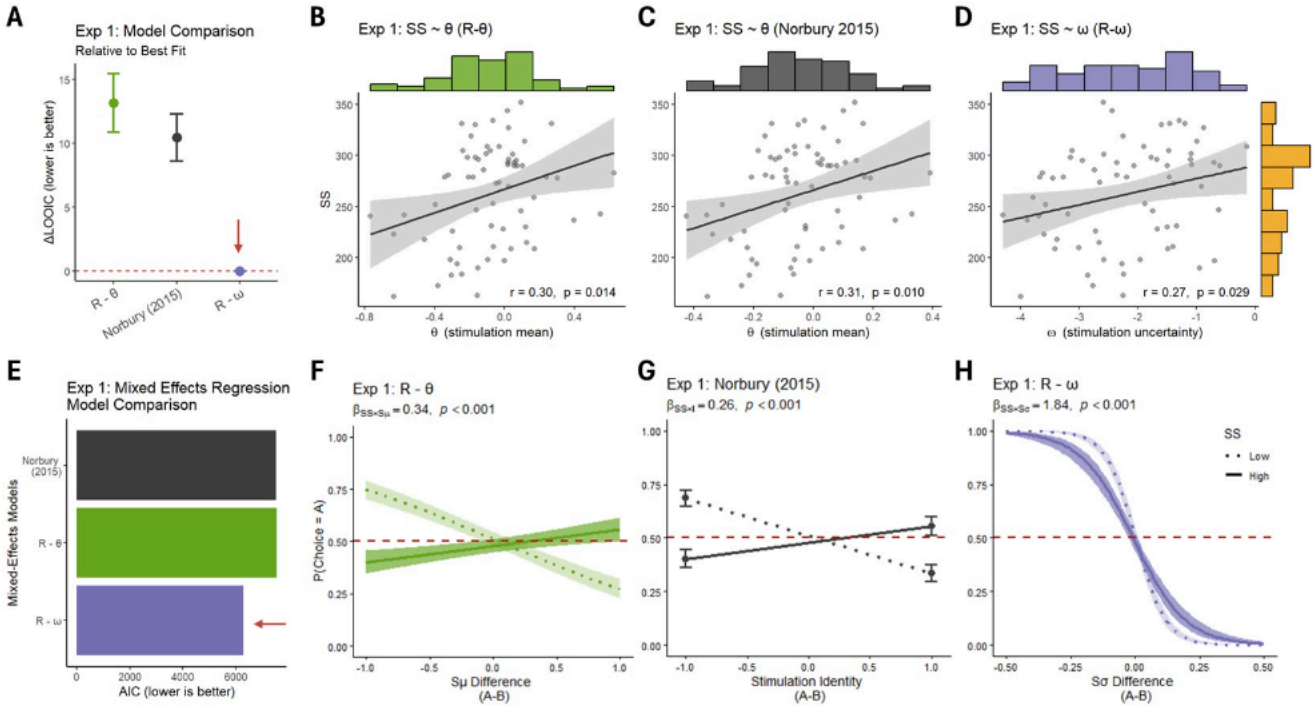

Figure 2. Experiment 1: re-analysis and model comparison. We re-analysed Norbury et al. (2015)[16] to test whether sensation seeking (SS) is better captured by an intensity account (weight on learned stimulation mean, θ) or an information account (weight on learned stimulation uncertainty, ω). A Beta–Bernoulli Bayesian learner provided trial-wise estimates of stimulation probability and uncertainty; choices were fit with single-process utilities—intensity ($R-\theta$) and uncertainty ($R-\omega$)—and compared with the original Norbury model. (**A**) Out-of-sample predictive fit (LOO, relative to best) favoured $R-\omega$. (**B**) Replicating the original intensity effect with our learner: SS correlated with θ ($\rho = 0.30$, $p = .014$), though this was not the best-fitting model. (**C**) Replication using the Norbury (2015) model with hierarchical Bayesian fitting ($\rho = 0.31$, $p = .010$). (**D**) From the best-fitting model, SS also correlated with the uncertainty weight ω ($\rho = 0.27$, $p = .029$); because ω was negative on average, higher SS implies greater tolerance for uncertainty. (**E**) Mixed-effects regression comparison (AIC): models including stimulation uncertainty provided the best fit. We visualise the interaction of either stimulation mean ($S\mu$) or uncertainty ($S\sigma$) with SS at $\pm 2SD$. (**F**) $S\mu \times$ SS interaction under the $R-\theta$ model. (**G**) $S\mu \times$ SS interaction under the Norbury (2015) model. (**H**) $S\sigma \times$ SS interaction under the $R-\omega$ model. Note. Stimulation mean and uncertainty were collinear in this design, so models were evaluated separately; a joint model is tested in Experiment 2, where these factors were orthogonalised.

We replicated the positive interaction between stimulation intensity and SS in our learning model R-θ ($\beta_{S\mu \times SS} = 0.34$, $p < .001$; Fig. 2F), which was similar to that observed in the Norbury mode ($\beta_{S_{Identity} \times SS} = 0.26$, $p < .001$; Fig. 2F). In both cases, high-SS individuals were more likely to select options associated with stimulation, whereas low-SS individuals tended to avoid them. In our *winning* model, R-ω, there was a strong negative relationship between the probability of selecting some choice A (over alternative B) and the difference in uncertainty ($\beta_{S\sigma} = -13.6$, $p < .001$), indicating that participants tended to avoid bandits with greater stimulation uncertainty. However, this trend was less pronounced among individuals with high SS tendencies ($\beta_{S\sigma \times SS} = 1.84$, $p < .001$), who exhibited flatter choice–uncertainty curves (Fig. 2H).

In summary, our reanalysis of Norbury (2015) points to a different interpretation of SS than the one originally proposed. Whereas the original analyses framed SS primarily as a general attraction to intense stimulation, our modelling suggests that SS is better characterised by a reduced avoidance of uncertainty— high SS individuals are more willing to sample uncertain options, rather than simply being drawn to stimulation per se. However, there are still limitations to this dataset, because intensity and uncertainty were intrinsically collinear. Thus, we designed Experiment 2 to orthogonalise these two factors, allowing us to disentangle their unique and interactive effects on SS.

Exp. 2: Context Dependent Uncertainty Tolerance

Experiment 2 (Fig. 1B) was a new, preregistered experiment[45] in which we removed the collinearity between stimulation intensity (mean) and uncertainty (variance) that had constrained interpretation in Experiment 1 (Fig. 1C). This allowed us to ask, at the behavioural level, whether sensation seeking

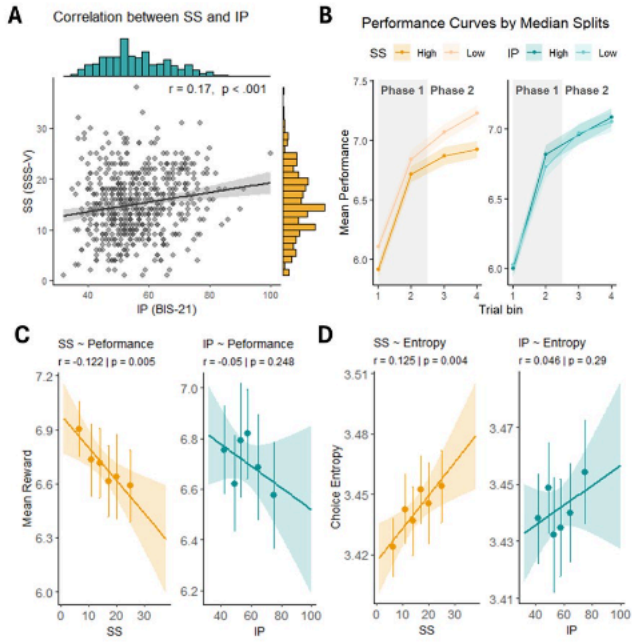

Figure 3. Experiment 2 Behavioural Results (A) Pearson's correlation between SSS-V and BIS-21 scores ($r = 0.17, p < 0.001$); the solid line is the fitted regression with 95% CI. (**B**) Learning/performance curves for high- vs low-trait groups (median split); shaded ribbons denote ± SD across participants. The trial sequence was grouped into bins of 15 trials, corresponding to one full presentation cycle of all object combinations. (**C**) Correlation of trait and performance, indexed as the mean reward earned across rounds. For SS, we see a negative correlation with performance ($r = -0.12, p = 0.005$). For IP, we found no significant association ($r = -0.05, p = 0.25$) (**D**) Correlation of trait and average choice entropy (Shannon entropy of the stimulus-choice distribution) across rounds, with higher values indicating broader sampling. For SS, we observed a positive association with choice entropy ($r = 0.13, p = 0.004$). On the other hand, we observed no relationship with IP ($r = 0.05, p = 0.29$).

(SS) and impulsivity (IP) show distinct signatures before turning to model-based analyses. We collected behavioural and questionnaire data from $n = 619$ Prolific participants. Based on our two preregistered exclusion criteria, we removed 84 participants who failed at least one of the two instructional manipulation checks (IMCs) and an additional 8 participants who performed below chance (i.e., below 50% of the maximum score). Thus our final sample was $n = 527$ participants (264 female; $M_{age} = 39.8 \pm 13.0$ years).

SS Show Greater Choice Variability

SS and IP were modestly positively correlated in our sample ($r = 0.17, p < 0.001$; Fig. 3A), but they showed clearly different relationships with behaviour. We first examined task performance and learning dynamics as the mean reward within 15-trial bins, chosen because each bin contains all pairwise stimulus combinations. Using a regression of performance on trial bin and SS, we found a clear effect of learning ($\beta_{bin} = -0.75$, $p < 0.001$) and lower overall reward in high- vs. low-SS participants ($\beta_{SS} = -0.11$, $p = 0.005$; Fig. 3C). Learning curves (Fig. 3B) indicated that this disadvantage was already present in early bins and widened, particularly from the start of phase 2. However, we did not find a significant interaction between SS × bin ($\beta_{SS \times bin} = -0.04$, $p = 0.19$), consistent with a relatively stable performance gap rather than diverging learning curves. IP groups did not differ reliably in performance in these analyses. Follow-up correlations with overall performance (Fig. 3C) showed a negative relationship with SS ($r = -0.12, p < 0.01$) but again no relationship with IP ($r = -0.05, p = 0.25$).

Next, we tested whether SS was associated with broader, high entropy sampling (Fig. 3D).[42,46,47]. We computed the Shannon entropy over choices, which positively correlated with SS scores (Pearson's $r = 0.13, p = 0.004$). Thus, high-SS participants displayed more diffuse, exploratory choice behaviour. In contrast, IP scores did not significantly correlate with choice entropy ($r = 0.05, p = 0.29$), suggesting that this exploratory profile was specific to SS.

We next asked whether the SS-specific behavioural signatures could be explained by reduced task engagement, boredom, or fatigue. Round-wise performance improved across the session (consistent with practice effects), but high-SS participants performed worse overall, and neither SS nor IP altered the trajectory of this improvement (Fig. S2). This pattern suggests that accumulating fatigue or disengagement is not the primary driver of SS-related performance costs.

To rule out a speed–accuracy account, we examined log reaction times (RTs) within trial bins (Fig. S3). RTs decreased over time, indicating increasing efficiency, but SS was unrelated to both overall response speed and the rate of RT change. IP also showed no reliable association with mean RTs; although there was a small interaction with trial bin, this modulation of RT dynamics was not accompanied by any differences in accuracy or reward. Thus, neither trait effect can be explained by a systematic speed–accuracy trade-off.

Finally, we examined omission rates as an index of attention lapses (Fig. S4). High-SS, but not high-IP, participants missed more trials, and higher omission rates were associated with poorer performance. Crucially, SS remained a negative predictor of performance even after statistically controlling for omission rate, indicating that reduced engagement contributes to, but does not fully account for, the SS-related performance deficit.

Taken together, high-SS is characterised by more variable, exploratory choice behaviour (higher entropy) and a modest but reliable performance cost that cannot be attributed solely to fatigue, global disengagement, or speed–accuracy trade-offs, even though high-SS individuals show slightly elevated omission rates. In contrast, IP does not show the same profile. These behavioural results thus provide a focused, trait-specific

foundation for the subsequent model-based analyses.

Distinct Trait Signatures in Choice Behaviour

We then asked what computational features underpinned the higher choice variability observed in high-SS participants. Using our learning models of reward (Bayesian Mean Tracker; BMT) and stimulation (Bayesian Learner Beta–Binomial; BLBB), we derived trial-by-trial predictors of reward mean (Rμ) and reward uncertainty (Rσ) (Fig. 4A, left), as well as stimulation mean (Sμ) and stimulation uncertainty (Sσ) (Fig. 4A, right).

We first compared two formulations of stickiness using AIC scores: a model with motor perseveration (previous response side) ($AIC = 141135.4$), which outperformed the stimulus perseveration (previous chosen bandit) model ($AIC = 141098$), which we had preregistered (Table S1). Therefore, subsequent analyses used the better motor perseveration model, which revealed a significant effect of motor perseveration ($\beta_{sticky} = 0.015$, $p = 0.03$).

We then entered these trial-level predictors into a mixed-effects logistic regression (Fig. 4B–H). Reward uncertainty Rσ was not included as a predictor, as it was experimentally controlled by providing full feedback of reward payoffs on every trial. Choices were modelled as a function of the difference (denoted Δ) in expected reward ($\Delta R\mu$), stimulation mean ($\Delta S\mu$), stimulation uncertainty ($\Delta S\sigma$), and a motor stickiness term, with participant-specific random intercepts and random slopes for $\Delta R\mu$. As preregistered, participant choices reliably tracked value and avoided uncertainty: larger $\Delta R\mu$ increased the odds of choosing right ($\beta_{R\mu} = 3.38, p < 0.001$), whereas larger $\Delta S\sigma$ decreased it ($\beta_{S\sigma} = -2.32, p < 0.001$). The stimulation mean was trending towards significance ($\beta_{S\mu} = 0.06, p = 0.058$) in the expected direction (Fig. 4H).

We then tested whether SS and IP traits modulate these computational signals by interacting with $\Delta R\mu$ and $\Delta S\sigma$, while also including the motor stickiness term (with all two- and three-way terms). Mirroring behavioural patterns, SS reduced reward sensitivity ($\beta_{SS\times R\mu} = -0.3$, p = 0.003; Fig. 4B), but not for IP ($\beta_{IP\times R\mu} = -0.05, p = 0.60$; Fig. 4C). Both traits also modulated motor perseveration in the same direction ($\beta_{SS\times Sticky} = -0.02, p = 0.02; \beta_{IP\times Sticky} = -0.12, p = 0.01$), suggesting that higher SS and higher IP each dampen value-free side-repetition once reward value and stimulation uncertainty are accounted for (Fig. 4D, E).

Crucially, and consistent with Exp. 1 and our preregistration, SS was associated with less uncertainty aversion: higher SS weakened the negative impact of stimulation uncertainty on choice ($\beta_{SS\times S\sigma} = 0.12$, $p = .03$; Fig. 4F). In contrast, IP showed the opposite pattern, being associated with greater avoidance of stimulation uncertainty ($\beta_{IP\times S\sigma} = -0.13$, $p = .02$; Fig. 4G).

We also observed a three-way interaction between $\Delta S\sigma$ and both SS and IP traits ($\beta_{SS\times IP\times S\sigma} = 0.14, p = 0.01$). To interpret this three-way interaction, we plotted the conditional slopes of $\Delta S\sigma$ across the observed SS and IP space (Fig. 4I) Here, we observe how the $\Delta S\sigma$ slope becomes less negative as SS increases, but more negative as IP increases. The strongest levels of uncertainty aversion are thus observed at low SS and high IP (top-left corner), while for individuals high in both SS and IP (top-right corner), the effect of $\Delta S\sigma$ on choice is close to zero and not reliably different from zero, indicating that uncertainty no longer systematically deters choice in this subgroup.

Trait Dissociation Based on an Interaction of Information and Intensity

We then fit a hierarchical, generative Bayesian choice model that decomposes trial-wise decisions into $R\mu$, $S\mu$, $S\sigma$, and a motor–stickiness. This yielded subject-level parameters θ (bonus on $S\mu$), ω (bonus on $S\sigma$), ρ (stickiness bonus), and τ (softmax temperature; choice stochasticity). Model comparison indicated that the full specification (reward, stimulation mean, stimulation uncertainty, and motor stickiness) outperformed reduced alternatives (Fig. 5A; see Methods).

Using subject-specific posterior estimates (Fig. 5B), we computed partial Pearson correlations between traits and model parameters while controlling for IP, age, IQ, and gender. Contrary to our preregistered prediction, SS was not significantly associated with ω ($r_{SS-\omega} = -0.08, p = 0.07$), but positively correlated with τ ($r_{SS-\log(\tau)} = -0.13, p = 0.003$; Fig. 5C), consistent with higher choice randomness among high-SS participants. No IP-parameter associations reached significance (all $p > 0.21$).

Motivated by the moderation pattern observed (Fig. 4I), we tested if SS/IP may also depend on a *combination* of mechanisms. To achieve this, we regressed our model parameters and their interactions, adjusting for covariates, on our traits of interest (Fig. 5D). We observed that SS remained positively associated with choice stochasticity ($\beta = 0.12$, $p = .01$). More critically, we observed a *negative* $\theta \times \omega$ interaction ($\beta = -0.19$, $p = .003$). A Johnson–Neyman[48,49] (J-N) analysis (Fig. 5E) into the interaction restricted within our observed moderator range ($-3.67 \leq \theta_z \leq 3.46$) showed that the conditional slopes of ω on SS was *positive* when $\theta_z \leq -1.01$ but *negative* when $\theta_z \geq -0.36$. Thus, uncertainty tolerance corresponds to higher SS when stimulation value weighting is low, whereas greater uncertainty aversion relates to higher SS when the weighting of stimulation value is high. These patterns support a mixed account in which SS reflects reduced uncertainty aversion when stimulation is valued less than average, but increased intolerance to uncertainty when stimulation is strongly valued. We then applied the same model to IP. Although no main effects of the HB parameters emerged, there was a *positive* $\theta \times \omega$ interaction ($\beta = 0.15$, $p = .02$); J–N analysis indicated that the conditional slope of ω on IP was significantly negative when $\theta_z \leq -0.24$.

Together, these results suggest a mechanistic dissociation between SS and IP, driven by an interplay between stimulation value and uncertainty weighting. Intuitively, the association between SS and uncertainty weighting is *context-dependent*: when stimulation carries relatively little value (low θ), higher

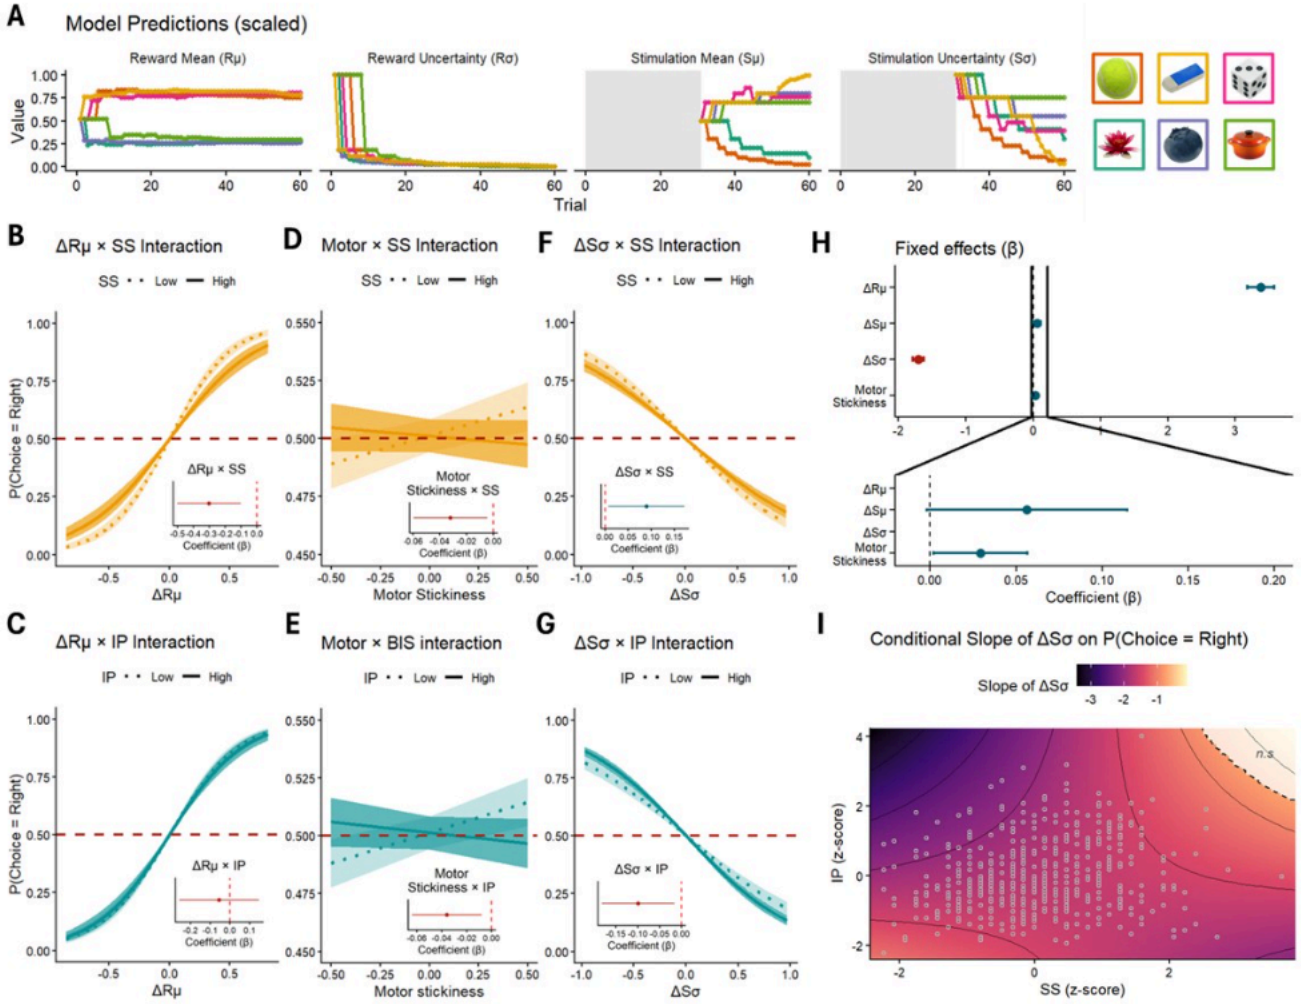

Figure 4. Model predictors, trait interactions, and fixed-effect estimates for the choice model. **(A)** Scaled time courses of model-derived predictors across a block: reward mean ($R\mu$), reward uncertainty ($R\sigma$), stimulation mean ($S\mu$), and stimulation uncertainty ($S\sigma$) for one example participant. **(B, C)** Predicted P(choice=Right) as a function of $\Delta R\mu$ at ± 2SD of Sensation Seeking (SS) and Impulsivity (IP), respectively; ribbons show 95% CIs and the red dashed line marks $P = 0.5$. **(D, E)** Motor stickiness × SS and motor stickiness × IP effects on choice. **(F, G)** $\Delta S\sigma$ (stimulation uncertainty) × SS and $\Delta S\sigma$ × IP effects on choice. Insets in panels (B–G) display the corresponding fixed-effect coefficients (β) with 95% CIs. **(H)** Forest plot of fixed effects for the full GLMM (points: β; bars: 95% CIs). **(I)** Conditional effect (slope) of $\Delta S\sigma$ on P(choice=Right) as a function of SS and IP (Johnson–Neyman–style surface).

SS is linked to greater tolerance of uncertainty; when stimulation is highly valued (high θ), higher SS instead corresponds to greater avoidance of stimulation uncertainty. For IP, in contrast, higher trait levels are primarily associated with increased aversion to stimulation uncertainty when stimulation value weighting is low, without a comparable reversal at higher θ.

Discussion

SS is often framed as a tendency to pursue intense stimulation, yet high-intensity options can also be especially informative. We therefore tested, across two experiments, whether SS reflects a preference for stimulation intensity (the estimated mean probability of receiving stimulation) or for uncertainty (variance around that probability) and how it can be distinguished from IP. In both datasets, participants integrated expected payoff, intensity, and uncertainty when making decisions. However, SS scores did not consistently increase the weight placed on intensity or uncertainty across experiments. In other words, higher SS did not reliably pull people toward more intense or more uncertain options. Instead, SS-related differences appeared when intensity and uncertainty interacted.

In Experiment 1, participants with high SS scores were less averse to stimulation uncertainty. When the estimated probability of stimulation was poorly specified, they sampled

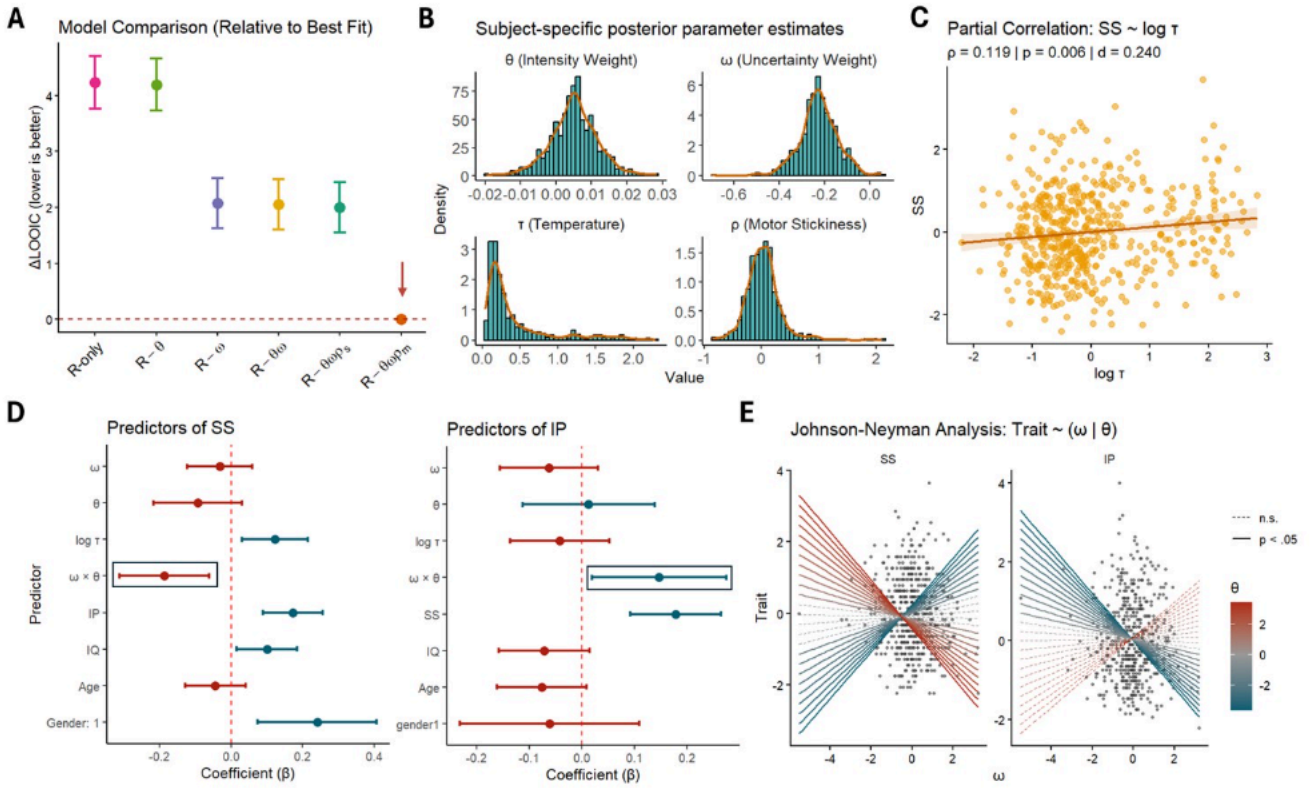

Figure 5. Hierarchical Bayesian Modelling results. (**A**) Model comparison with reduced versions of our full model, as well as an alternative version of stickiness, which is based on stimulus identity. We find that our full model specification, which includes a motor stickiness bonus, best accounts for individual differences in choice behaviour. (**B**) Distribution of model posterior estimates. (**C**) Partial correlation of sensation seeking (SS) with temperature $\log(\tau)$ after controlling for impulsivity (IP), age, IQ and gender. We observed a positive association ($\rho_{SS-\log(\tau)} = -0.13, p = 0.003$), suggesting that higher SS is associated with a global increase in choice stochasticity or random exploration. (**D**) Regression coefficients predicting traits as interactions between computational parameters controlling for various correlates. We observe that SS is significantly predicted by increasing temperature, as well as a *negative* interaction between intensity θ and information ω accounts (boxed). Conversely, IP is only predicted by a *positive* interaction between θ and ω. This highlights a mechanistic dissociation between the two traits. All predictors were z-scored before analysis. (**E**) Johnson-Neyman fanplots of the $\theta \times \omega$ interaction. Conditional slopes of ω against traits within the observed range of moderator θ. Bold solid lines indicate statistically significant relationships ($p < 0.05$), while dotted lines did not reach significance. Dots are raw subject points. We observed that IP is associated with a decrease in ω at below-average values of θ. Conversely, SS is associated with increasing values of ω at below average values of θ, but decreasing ω at above average values of θ.

those options more often. In our preregistered Exp. 2[45], SS was associated with greater choice stochasticity and a negative $\theta - \omega$ interaction. The relation between SS and uncertainty depended on how strongly a person weighted intensity. When stimulation carried low subjective value (smaller θ), higher SS predicted greater tolerance for uncertainty— more sampling of unpredictable options. When stimulation carried high subjective value (larger θ), higher SS predicted lower tolerance for uncertainty—choices shifted toward predictable options. In short, SS is neither purely "intensity-seeking" nor purely "information-seeking". It is context-dependent: when stimulation matters less, SS expresses as uncertainty-tolerant exploration; when stimulation matters more, SS expresses as exploitation. This pattern aligns with computational work demonstrating that humans employ a mix of directed (information-seeking) and random (stochastic) exploration strategies, and that the balance between them shifts in response to goals and context[50,51].

This context dependence helps reconcile mixed findings in the broader literature, with SS linked to both beneficial outcomes (e.g., novelty seeking that supports learning, exploration, and well-being) and harmful outcomes (e.g., substance use and risky behaviour)[4,6,9–11,52,53]. Our results offer a principled account of this duality: when the environment carries lower subjective stimulation value, SS may facilitate adaptive information-gathering; but when stimulation is highly valued, SS can tip behaviour towards exploitation of intense options, increasing the risk of maladaptive pursuits.

Having established the SS profile, we next asked whether IP shows the same pattern. IP instead showed a dissociable profile. Contrary to previous reports that IP generally manifests as noisier, more random choice[39], we observed no robust increases in value-free stochasticity or reliable performance decrements as a function of IP in this task. Instead, participants with high IP scores biased choices away from uncertainty, especially when stimulation had low subjective value. One interpretation is that the exploration signature of IP depends on what is uncertain: prior links between IP and value-free randomness come from reward-only or affectively neutral sampling tasks[39,54,55], whereas here the uncertainty concerned the probability and intensity of *stimulation* itself. In such contexts, uncertainty-averse policies may dominate. Independent evidence shows that aversive or affect-laden states suppress directed exploration and promote uncertainty avoidance[56,57]. We neither manipulated affect nor could test for IP-affect interactions; thus, whether affect moderates the influence of impulsivity on uncertainty weighting remains open. Future work should orthogonally manipulate the outcome domain (reward vs. stimulation) and affect to test this mechanism directly.

Collectively, our findings point to a *mechanistic dissociation* between SS and IP. This pattern is *consistent* with broader developmental and personality frameworks that distinguish SS from IP by emphasising distinct neurocognitive systems for incentive responsiveness (SS) versus control and constraint (IP)[10,58–60].

Limitations and Future Directions

Several limitations should be considered. First, Experiment 1's design confounded stimulation intensity and uncertainty, so those results—while suggestive—required the orthogonalisation we implemented in Experiment 2. Second, although we used a neutral visual stimulation, subjective valence/arousal likely varies across individuals and may be an unmodeled contributor to θ; future studies should manipulate valence and arousal across different modalities of sensory stimulation to concretely test whether SS shifts along the $\theta - \omega$ surface as affective engagement changes[61].

Despite these limitations, our framework yields concrete predictions and translational applications. Experimentally, increasing the *subjective* value of stimulation (e.g., making it more pleasant or goal-relevant) should shift high-SS participants toward stronger exploitation, whereas decreasing its value should encourage exploration. The gambling literature offers a robust parallel: in slot machines, salient audiovisual *cues* increase subjective reward value and predict immersion scores[62], which can bias choices toward riskier, more exploitative policies[63,64]. Structural features such as "Losses Disguised as Wins" and near-misses mimic wins, elevate arousal, and increase persistence despite negative expected value[65–70]. In our terms, such cues and situations raise θ, pushing high-SS decision policies toward exploitation of intense options. This dovetails with our findings and points to practical levers: reducing cue salience (attenuating sound/light intensity and neutralising win animations) to lower θ, and increasing usable information (transparent odds/volatility displays), thereby re-exposing exploration and dampening exploitative momentum among high-SS individuals.

By making the context dependence explicit and separating stimulation value from stimulation uncertainty, our results offer a mechanistic account of how the same trait can underwrite both beneficial and harmful outcomes, and a roadmap for tipping the balance toward the former.

Methods

Participants

In Experiment 1, we reanalysed a dataset previously collected by Norbury et al.[16]. Below, we provide a brief overview of the task design and participant demographics; further details are available in the original publication. Norbury et al. (2015)[16] reported two datasets. Dataset 1 recruited 45 healthy adults (28 female; $M_{\text{age}} = 24.3 \pm 3.55$ years). Following behavioural quality checks, 3 participants were excluded; sensation-seeking scores were not recorded for 6 participants, leaving $n = 39$ with complete trait data for reanalysis. Dataset 2 (pharmacological) recruited 30 healthy males ($M_{\text{age}} = 22.3 \pm 2.74$ years); 2 were excluded in the original report, yielding $n = 28$. To increase sensitivity for behavioural analyses, we pooled data from both datasets, leaving 73 subjects for reanalysis. All participants completed the revised version of the Sensation-Seeking Scale (SSS-V)[58,71].

Experiment 2 was preregistered under https://doi.org/10.17605/OSF.IO/H4P3T. A total of 619 participants were recruited via Prolific. Inclusion criteria were residency in the United Kingdom or the United States and an approval rate of minimum 95% on the platform. From our recruited sample, 84 subjects failed IMC and 5 subjects performed below chance. After exclusion, our final sample consisted of 530 subjects. They were paid a base fee of £5, with an additional bonus of up to £5 calculated based on the performance of one randomly selected round. All participants were above 18 years and from the UK/US. The study was approved by the Ethics Committee of the University of Tübingen, and all procedures were conducted in accordance with relevant guidelines and regulations. Informed consent was obtained from all participants prior to participation.

Study Design

Experiment 1

Participants completed a two-phase, two-armed bandit task designed to measure the value assigned to stimulation. Here, sensory stimulation was a mild, non-aversive electric stimulation delivered to the hand. The task comprised a reward-learning phase followed by a test phase with stimulation. During reward learning, participants learned the point values associated with eight bandits. Two bandits were assigned to each of four fixed payoffs (25, 50, 75, or 100 points). On every trial, two bandits were presented side by side, and one choice was

required; this pairing scheme produced 10 unique pair types (combinatorial pairs of the four payoffs, with two exemplars per payoff). The phase continued for a minimum of 80 trials and progressed until a pre-specified performance criterion was met, ensuring adequate learning of the reward structure prior to the introduction of stimulation.

In the subsequent test phase, half of the bandits were designated stimulation-paired (probability of stimulation = 0.75) and half non-stimulation (probability = 0). Each test trial paired one stimulation-paired bandit with one non-stimulation bandit, yielding three payoff scenarios in equal proportion: the stimulation-paired bandit offered a higher, lower, or equal point payoff relative to its partner. Participants completed 100 test trials (10 per trial type), allowing us to estimate preferences for stimulation-paired options conditional on relative monetary value. Prior to any task exposure, participants provided liking ratings for each bandit on a continuous visual analogue scale. Identical ratings were collected again after reward learning and after the test phase. The change from post-learning to post-test indexed overall stimulation liking at the end of the experiment.

Experiment 2

Participants completed four rounds of a web-based adaptation of the operant sensation task[16,72]. Similarly, each round consists of Phase 1 (reward learning) and Phase 2 (test with stimulation), totalling 60 trials per round (30 trials per phase). In this paradigm, participants could self-administer visual stimulation (a coloured spinning animation; 1.4s) while engaging in a two-armed bandit task. On every trial, participants had up to 6 seconds to respond; non-responses were recorded as missed trials. Before proceeding with the main task, participants provided informed consent, read on-screen instructions, and completed four comprehension questions; all questions had to be answered correctly to proceed.

At the start of each round, we sampled six object images from the Bank of Standardized Stimuli (BOSS)[18,19] and assigned two orthogonal parameters to each object. First, an unscaled reward level that is high ($N \sim (10, 2.5)$) vs. low ($N \sim (-10, 2.5)$) and a stimulation probability (0.15, 0.50, or 0.85). Assignments were fixed within a round and balanced so that across the six objects, the reward factor and stimulation-probability factor were independent by design, enabling separate identification of effects due to stimulation mean (intensity) and stimulation uncertainty (variance). Phase 1 (reward learning) was explicitly structured to minimise epistemic sampling for reward: on each of 30 trials, participants chose between two of the six objects and then received full feedback—the realised reward for the chosen and the unchosen object—expediting belief formation about object values. Phase 2 (test with stimulation here termed "glitch") introduced stimulation while keeping the learned reward contingencies unchanged. On each of 30 trials, the visual stimulation was delivered on the chosen object with its assigned probability (0.15, 0.50, or 0.85). Participants were reminded at the transition that stimulation does not influence rewards, which remained governed solely by the Phase-1 reward structure. Within each round, we pre-generated all 15 unique object pairs (6 choose 2) and presented each pair exactly four times, yielding 60 trials. The order of pair presentation and left/right positions was independently randomised for each participant across trials. This scheme ensured that each object appeared in a comparable number of pairings, preserved an identical pairing structure across participants (with only trial order randomised), and controlled for spatial biases.

We applied a random scaling of rewards across rounds (i.e., different minimum and maximum rewards) to prevent participants from immediately recognising when they had chosen the optimal option. For each round, we sampled a value from a uniform distribution $U \sim (1, 5)$, which scaled the rewards after shifting the rewards by a constant value of 10. We also truncated rewards to ensure they were always non-negative. In order to convey intuitions about the random scaling of rewards, payoffs were presented using a different fictional currency in each round, such that the absolute value was unknown, but higher payoffs were always better.

$$\text{payoff} = (\text{rawScore} + 10)\, s, \qquad s \sim U(1, 5). \tag{1}$$

Following each round, participants provided subjective ratings of stimulation valence (unpleasant-pleasant) and arousal (sleepy-wakeful) on Likert scales [0-100]. The task was delivered online with standard display instructions (full-screen mode encouraged); responses were collected via keyboard, and reaction times were recorded on each trial. Prior to main analyses, we implemented preregistered quality controls (instructional manipulation checks and basic performance criteria) and excluded participants who failed checks or performed at/below chance (further details in the Participants and Data Quality section).

Psychiatric Questionnaires

After completing the task, participants completed several questionnaires assessing psychiatric dimensions and general cognitive ability. These include the Sensation Seeking Scale V[58,71] (SSS-V), Barratt Impulsiveness Scale[73] (BIS-11), Depression Anxiety and Stress Scale 21[74] (DASS-21), Obsessional Compulsive Inventory – Revised[75] (OCI-R), Alcohol Use Disorders Identification Test[76] (AUDIT), Drug Abuse Screening Test[77] (DAST-10), and the International Cognitive Ability Resource[78] (ICAR-16). To ensure data quality, instructional manipulation checks[79] (IMCs) were included as attentional checks. Participants failing any one of these IMCs were excluded from the analysis.

Statistical Analysis

Analyses were conducted in R (version 4.4.1) using RStudio. Hierarchical generative modelling was carried out with `rstan` (version 2.32.6). Mixed-effects regressions were performed using `lme4` (version 1.1-35.5). Regression model comparisons were conducted with the `performance` package (version 0.15.0). Johnson–Neyman interaction analyses were conducted using `emmeans` (version 1.10.4).

Data Validation and Exclusion Criteria

Our preregistration specified excluding participants who achieved less than 50% of the total maximum score. We adopted a more stringent, theory-driven criterion that targets learning directly. Because participants were expected to have fully learned the reward contingencies by the end of Phase 1, we evaluated performance in Phase 1 only and excluded participants who scored < 50% of the Phase-1 maximum (operationalised as the sum, over trials, of the higher of the two available outcomes on each trial). This approach provides a more targeted measure of contingency learning, consistent with our models' assumption that reward contingencies are fully learned and stabilised by the end of Phase 1. Additionally, we excluded participants who failed the instructional manipulation checks (IMCs) in the questionnaires and those with missing questionnaire data due to technical issues.

Computational Modeling

We model trialwise choices as the integration of three components: (i) expected point payoff, (ii) Beliefs about stimulation (its mean and uncertainty) and (iii) a value-free tendency to repeat either the previous button pressed (motor stickiness) or the previous bandit selected (bandit stickiness). Beliefs about expected payoff and stimulation are updated by learning rules elaborated below.

Reward Learning (Bayesian Mean Tracker. In Experiment 1, rewards were fixed after a training phase, so object payoff values $R\mu_j$ are treated as known constants (no trialwise updating). In contrast, the reward payoffs in Experiment 2 were Gaussian but stationary. To capture participants' evolving expectations in this environment, we employed a Bayesian Mean Tracker (BMT)[80–82], a type of Kalman filter specifically designed for stationary reward settings. Briefly, for each bandit j we track a normal posterior:

$$P(R\mu_{j,t}|\mathscr{D}_t) = \mathscr{N}(R\mu_{j,t}, R\sigma_{j,t}) \tag{2}$$

The posterior mean $R\mu_{j,t}$ and variance $R\sigma_{j,t}$ are updated iteratively based on observed rewards y_t. In our experiment setup, the posterior is updated for both chosen and unchosen bandits due to complete feedback.

$$R\mu_{j,t} = R\mu_{j,t-1} + G_{j,t}\left[y_t - R\mu_{j,t-1}\right] \tag{3}$$
$$R\sigma_{j,t} = [1 - G_{j,t}]R\sigma_{j,t-1} \tag{4}$$

Here, the Kalman Gain $G_{j,t}$ modulates the size of each update and is calculated as:

$$G_{j,t} = \frac{R\sigma_{j,t-1}}{R\sigma_{j,t-1} + \zeta_\varepsilon^2} \tag{5}$$

The parameter ζ_ε^2 represents the error variance, controlling the model's sensitivity to new information. Lower values of ζ_ε^2 result in larger updates and faster uncertainty reduction. Thus, the BMT offers a flexible, uncertainty-aware framework for modelling learning in environments with stationary reward distributions.

Stimulation Learning. We implemented a Bayesian model that assumes stimulation probabilities for each object j are represented as beta distributions. This approach is most appropriate for learning probabilities since the distribution is between 0 and 1, and the variance estimates uncertainty. Similar models have previously been used to model value-based learning[83] and aversive learning tasks[84,85].

$$P_{j,t}(stim) \sim Beta(\alpha_{j,t}, \beta_{j,t}) \tag{6}$$

Intuitively, the model describes how the evidence for stimulation depends on the number of stimulations previously observed. These counts can be represented for each bandit by the parameter α_j, which is incremented depending on whether a stimulation is observed ($S_{j,t} = 1$). Similarly, the counts of no stimulation ($S_{j,t} = 0$) are tracked by a complementary parameter β_j. For simplicity, we assume optimal updating, where α and β are updated by a value of 1 after observing the respective outcome:

$$\alpha_{j,t+1} = \alpha_{j,t} + S_{j,t} \tag{7}$$

$$\beta_{j,t+1} = \beta_{j,t} + (1 - S_{j,t}) \tag{8}$$

The expected probability of stimulation $S\mu$ for bandit j at trial t is:

$$S\mu_{j,t} = \frac{\alpha_{j,t}}{\alpha_{j,t} + \beta_{j,t}} \tag{9}$$

and the associated uncertainty $S\sigma$ is:

$$S\sigma_{j,t} = \sqrt{\frac{\alpha_{j,t}\beta_{j,t}}{(\alpha_{j,t} + \beta_{j,t})^2(\alpha_{j,t} + \beta_{j,t} + 1)}} \tag{10}$$

Value-free stickiness. We compare two value-free perseveration effects, added *outside* the temperature scaling: motor (repeat the previous *button*, ρ_m) and stimulus (repeat the previous *object identity*, ρ_s).

Given $a_{t-1} \in \{\mathrm{L}, \mathrm{R}\}$ as the previous action and o_{t-1} the previously chosen object (so $o_{t-1} = o_{t-1}(a_{t-1})$).

$$\rho_t^{\mathrm{m}}(a) = \mathbb{I}[a = a_{t-1}], \tag{11}$$
$$\rho_t^{\mathrm{s}}(a) = \mathbb{I}[o_t(a) = o_{t-1}], \tag{12}$$

with $\rho_1^{\cdot}(a) = 0$ on the first trial. Intuitively, this assigns a bonus to either the button side that was chosen or the bandit that was chosen on the preceding trial.

Decision Rule. Q-values for each bandit j were calculated as a linear sum of their estimated underlying point payoff $R\mu_j$

and either the estimated mean probability of stimulation, $S\mu_j$, scaled by free parameter θ, or the associated uncertainty $S\sigma_j$, scaled by free parameter ω.

$$Q_{j,t} = \frac{R\mu_{j,t} + \theta S\mu_{j,t} + \omega S\sigma_{j,t}}{\tau} + \rho_{j,t} \tag{13}$$

Given bandits A and B, the utility values are then put through a softmax:

$$P(choice_t = A) = \frac{e^{Q_A}}{e^{Q_A} + e^{Q_B}} \tag{14}$$

For brevity, the full model($R - \theta\omega\rho$) is described as above. We also compared reduced versions of the model via model ablation. In Experiment 1, we compared single-process models against each other, i.e., $R - \theta$ and $R - \omega$, respectively (See SI for more details on reduced models).

Mixed Effect Regressions

We quantified how reward and stimulation signals shaped choice using trial-level mixed-effects logistic regression. For each trial, we constructed difference predictors as right minus left model estimates (differences are denoted by Δ): expected reward ($\Delta R\mu$), stimulation mean ($\Delta S\mu$), and stimulation uncertainty ($\Delta S\sigma$). Trials with no response were excluded. All continuous predictors and trait scores were z-scored. Our preregistered primary model predicted the probability of choosing the right option with a logit link, including fixed effects of $\Delta R\mu$, $\Delta S\mu$, and $\Delta S\sigma$, and a random intercept and random $\Delta R\mu$ slope by participant, which allowed value sensitivity to vary across individuals:

$$\begin{aligned}\Pr(\text{choice}_{i,t} = \text{Right}) \sim \Delta R\mu_{i,t} + \Delta S\mu_{i,t} + \Delta S\sigma_{i,t} + \text{Sticky}_{i,t} \\ + (1 + \Delta R\mu_{i,t} \mid \text{Subject}_i).\end{aligned}$$

We also evaluated the two alternative versions of stickiness as mentioned above and compared them using AIC/BIC. The final trait-moderation analyses retained the preregistered value and stimulation predictors, with motor stickiness (Table S1).

To further test whether traits modulate these computational signals, we added cross-level interactions with SS and IP, entered as z-scores. All lower-order terms were included.

Models were fit in `lme4::glmer` (binomial family, logit link); convergence and singularity were checked. We report the β coefficients with 95% CIs, and compared random-effects structures and fixed-effect subsets using AIC and likelihood-ratio tests where appropriate. Model diagnostics included overdispersion checks and DHARMa residual simulations. For statistically significant higher-order interactions, we probed simple slopes (estimated marginal trends) at representative moderator levels and, where relevant, Johnson–Neyman regions restricted to the observed range of the moderator.

Generative Model Fitting and Comparison

All models were estimated hierarchically using custom-written STAN code. Specifically, we used Hamiltonian MCMC with a No-U-Turn sampler to estimate the group-level mean, μ_0, and variance σ_0^2 for all model parameters across participants. Weakly informative $N \sim (0, 1)$ priors were assigned to group-level parameters. Chain convergence was assessed using the $\widehat{R}$ statistic, where $1 \leq \widehat{R} \leq 1.01$ was acceptable. The model was estimated over four chains of 4000 iterations, with a burn-in period of 1000 samples and a proposal acceptance probability set to 0.99. The point payoff for all bandits was scaled to the range $[0, 1]$ before fitting.

Model comparison was conducted using the loo package in R, leveraging a version of the loo estimate optimised through Pareto smoothed importance sampling (PSIS) methodology[86]. The loo approach assesses the out-of-sample predictive accuracy of the model, essentially evaluating how well the entire dataset, excluding one data point, predicts the outcome for the excluded point.

Linking Model Parameters to Traits

We related individual differences in the hierarchical model to questionnaire traits in two steps. First, we quantified partial correlations between trait scores and subject-level parameters from the hierarchical Bayesian choice model, θ, ω, τ, and ρ. For interpretability, we analysed $\log(\tau)$. All variables were z-scored prior to analysis. Partial correlations were computed using Pearson's r, adjusting for age, IQ, gender, and the other trait (i.e., SS partials controlled for IP and vice versa). Distributional assumptions were verified by testing univariate normality with the Shapiro–Wilk test (all $p > 0.05$).

Second, to test whether traits reflect a joint contribution of value and uncertainty mechanisms, we fit subject-level linear regressions with all lower-order terms and the three-way interaction between $\omega, \theta, and \log(\tau)$, while adjusting for covariates (age, IQ, gender) and the other trait. We did not include ρ (motor stickiness) in the interaction space because our a priori question targeted the synergy between intensity (θ) and uncertainty (ω); adding ρ interactions would inflate the model without clear theoretical gain and reduce power. As a robustness check, adding ρ as an additional covariate did not substantially change our results.

The models were:

$$\text{SS} \sim \omega \times \theta \times \log(\tau) + \text{IP} + \text{age} + \text{IQ} + \text{gender} \tag{15}$$

$$\text{IP} \sim \omega \times \theta \times \log(\tau) + \text{SS} + \text{age} + \text{IQ} + \text{gender} \tag{16}$$

Model diagnostics included residual inspection and variance-inflation factors (all VIFs < 5). When higher-order interactions were significant, we probed simple slopes using

standard linear-model contrasts and Johnson–Neyman intervals, expressly restricting inference to the observed moderator range and reporting the percentage of participants falling within each region of significance.

References

1. Zuckerman, M. The sensation seeking motive. *Prog. experimental personality research* **7**, 79–148 (1974).
2. Shulman, E. P., Harden, K. P., Chein, J. M. & Steinberg, L. Sex differences in the developmental trajectories of impulse control and sensation-seeking from early adolescence to early adulthood. *J. youth adolescence* **44**, 1–17 (2015).
3. Steinberg, L. *et al.* Age differences in sensation seeking and impulsivity as indexed by behavior and self-report: evidence for a dual systems model. *Dev. psychology* **44**, 1764 (2008).
4. Chase, H. W. & Ghane, M. Seeking pleasure, finding trouble: Functions and dysfunctions of trait sensation seeking. *Curr. Addict. Reports* 1–9 (2023).
5. Peritogiannis, V. Sensation/novelty seeking in psychotic disorders: A review of the literature. *World journal psychiatry* **5**, 79 (2015).
6. Evans-Polce, R. J., Schuler, M. S., Schulenberg, J. E. & Patrick, M. E. Gender-and age-varying associations of sensation seeking and substance use across young adulthood. *Addict. behaviors* **84**, 271–277 (2018).
7. Hittner, J. B. & Swickert, R. Sensation seeking and alcohol use: A meta-analytic review. *Addict. behaviors* **31**, 1383–1401 (2006).
8. Hammerton, G. *et al.* Low resting heart rate, sensation seeking and the course of antisocial behaviour across adolescence and young adulthood. *Psychol. Medicine* **48**, 2194–2201 (2018).
9. Mann, F. D. *et al.* Sensation seeking and impulsive traits as personality endophenotypes for antisocial behavior: Evidence from two independent samples. *Pers. individual differences* **105**, 30–39 (2017).
10. Ravert, R. D. & Donnellan, M. B. Impulsivity and sensation seeking: Differing associations with psychological well-being. *Appl. Res. Qual. Life* **16**, 1503–1515 (2021).
11. Yoneda, T., Ames, M. E. & Leadbeater, B. J. Is there a positive side to sensation seeking? trajectories of sensation seeking and impulsivity may have unique outcomes in young adulthood. *J. Adolesc.* **73**, 42–52 (2019).
12. Duell, N. & Steinberg, L. Differential correlates of positive and negative risk taking in adolescence. *J. youth adolescence* **49**, 1162–1178 (2020).
13. Duell, N. & Steinberg, L. Adolescents take positive risks, too. *Dev. Rev.* **62**, 100984 (2021).
14. Hansen, E. B. & Breivik, G. Sensation seeking as a predictor of positive and negative risk behaviour among adolescents. *Pers. individual differences* **30**, 627–640 (2001).
15. Fischer, S. & Smith, G. T. Deliberation affects risk taking beyond sensation seeking. *Pers. Individ. Differ.* **36**, 527–537 (2004).
16. Norbury, A., Kurth-Nelson, Z., Winston, J. S., Roiser, J. P. & Husain, M. Dopamine regulates approach-avoidance in human sensation-seeking. *Int. J. Neuropsychopharmacol.* **18**, pyv041 (2015).
17. Chang, N. H. S. *et al.* On the learning of addictive behavior: Sensation-seeking propensity predicts dopamine turnover in dorsal striatum. *Brain Imaging Behav.* **16**, 355–365 (2022).
18. Brodeur, M. B., Dionne-Dostie, E., Montreuil, T. & Lepage, M. The bank of standardized stimuli (boss), a new set of 480 normative photos of objects to be used as visual stimuli in cognitive research. *PloS one* **5**, e10773 (2010).
19. Brodeur, M. B., Guérard, K. & Bouras, M. Bank of standardized stimuli (boss) phase ii: 930 new normative photos. *PloS one* **9**, e106953 (2014).
20. Casey, B. J., Getz, S. & Galvan, A. The adolescent brain. *Dev. review* **28**, 62–77 (2008).
21. Shulman, E. P. *et al.* The dual systems model: Review, reappraisal, and reaffirmation. *Dev. cognitive neuroscience* **17**, 103–117 (2016).
22. Luna, B. & Wright, C. Adolescent brain development: Implications for the juvenile criminal justice system. (2016).
23. Holmes, A. J., Hollinshead, M. O., Roffman, J. L., Smoller, J. W. & Buckner, R. L. Individual differences in cognitive control circuit anatomy link sensation seeking, impulsivity, and substance use. *J. neuroscience* **36**, 4038–4049 (2016).
24. Khurana, A. *et al.* Working memory ability predicts trajectories of early alcohol use in adolescents: the mediational role of impulsivity. *Addiction* **108**, 506–515 (2013).
25. Norbury, A. & Husain, M. Sensation-seeking: Dopaminergic modulation and risk for psychopathology. *Behav. brain research* **288**, 79–93 (2015).
26. Mitchell, M. R. & Potenza, M. N. Addictions and personality traits: impulsivity and related constructs. *Curr. behavioral neuroscience reports* **1**, 1–12 (2014).
27. Reynolds, B., Ortengren, A., Richards, J. B. & De Wit, H. Dimensions of impulsive behavior: Personality and behavioral measures. *Pers. individual differences* **40**, 305–315 (2006).
28. Jauregi, A., Kessler, K. & Hassel, S. Linking cognitive measures of response inhibition and reward sensitivity to trait impulsivity. *Front. psychology* **9**, 2306 (2018).

29. Yeh, Y.-H., Myerson, J. & Green, L. Delay discounting, cognitive ability, and personality: What matters? *Psychon. Bull. & Rev.* **28**, 686–694 (2021).

30. Lauriola, M., Panno, A., Levin, I. P. & Lejuez, C. W. Individual differences in risky decision making: A meta-analysis of sensation seeking and impulsivity with the balloon analogue risk task. *J. Behav. Decis. Mak.* **27**, 20–36 (2014).

31. Blaskey, L. G., Harris, L. J. & Nigg, J. T. Are sensation seeking and emotion processing related to or distinct from cognitive control in children with adhd? *Child Neuropsychol.* **14**, 353–371 (2008).

32. Addicott, M. A., Pearson, J. M., Sweitzer, M. M., Barack, D. L. & Platt, M. L. A primer on foraging and the explore/exploit trade-off for psychiatry research. *Neuropsychopharmacology* **42**, 1931–1939 (2017).

33. Mehlhorn, K. *et al.* Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision* **2**, 191 (2015).

34. Sercombe, H. Risk, adaptation and the functional teenage brain. *Brain cognition* **89**, 61–69 (2014).

35. Willoughby, T., Good, M., Adachi, P. J., Hamza, C. & Tavernier, R. Examining the link between adolescent brain development and risk taking from a social–developmental perspective (reprinted). *Brain cognition* **89**, 70–78 (2014).

36. Ciranka, S. & Van Den Bos, W. Adolescent risk-taking in the context of exploration and social influence. *Dev. Rev.* **61**, 100979 (2021).

37. Ciranka, S. & Hertwig, R. Environmental statistics and experience shape risk-taking across adolescence. *Trends cognitive sciences* **27**, 1123–1134 (2023).

38. Dubois, M. *et al.* Human complex exploration strategies are enriched by 17 noradrenaline-modulated heuristics. elife 10. *ARTN e59907* **18** (2021).

39. Dubois, M. & Hauser, T. U. Value-free random exploration is linked to impulsivity. *Nat. Commun.* **13**, 4542 (2022).

40. Tymula, A. *et al.* Adolescents' risk-taking behavior is driven by tolerance to ambiguity. *Proc. Natl. Acad. Sci.* **109**, 17135–17140 (2012).

41. Tymula, A., Rosenberg Belmaker, L. A., Ruderman, L., Glimcher, P. W. & Levy, I. Like cognitive function, decision making across the life span shows profound age-related changes. *Proc. Natl. Acad. Sci.* **110**, 17143–17148 (2013).

42. Giron, A. P. *et al.* Developmental changes in exploration resemble stochastic optimization. *Nat. Hum. Behav.* DOI: 10.1038/s41562-023-01662-1 (2023).

43. Cogliati Dezza, I., Schulz, E. & Wu, C. M. (eds.) *The Drive for Knowledge: The Science of Human Information-Seeking* (Cambridge University Press, 2022).

44. Zuckerman, M., Eysenck, S. B. & Eysenck, H. J. Sensation seeking in england and america: cross-cultural, age, and sex comparisons. *J. consulting clinical psychology* **46**, 139 (1978).

45. Wong, E., Chandrasekaran, A., Hauser, T. U., Pietrini, P. & Wu, C. M. Preference for exploring information-rich contexts in sensation-seeking. OSF preregistration, DOI: 10.17605/OSF.IO/H4P3T (2025). Preregistration.

46. Gopnik, A. *et al.* Changes in cognitive flexibility and hypothesis search across human life history from childhood to adolescence to adulthood. *Proc. Natl. Acad. Sci.* **114**, 7892–7899 (2017).

47. Gopnik, A. Childhood as a solution to explore–exploit tensions. *Philos. Transactions Royal Soc. B* **375**, 20190502 (2020).

48. Johnson, P. O. & Neyman, J. Tests of certain linear hypotheses and their application to some educational problems. *Stat. research memoirs* (1936).

49. Preacher, K. J., Curran, P. J. & Bauer, D. J. Computational tools for probing interactions in multiple linear regression, multilevel modeling, and latent curve analysis. *J. educational behavioral statistics* **31**, 437–448 (2006).

50. Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A. & Cohen, J. D. Humans use directed and random exploration to solve the explore–exploit dilemma. *J. Exp. Psychol. Gen.* **143**, 2074 (2014).

51. Wilson, R. C., Bonawitz, E., Costa, V. D. & Ebitz, R. B. Balancing exploration and exploitation with information and randomization. *Curr. opinion behavioral sciences* **38**, 49–56 (2021).

52. Duell, N. & Steinberg, L. Positive risk taking in adolescence. *Child development perspectives* **13**, 48–52 (2019).

53. Zhang, X., Qu, X., Tao, D. & Xue, H. The association between sensation seeking and driving outcomes: A systematic review and meta-analysis. *Accid. Analysis & Prev.* **123**, 222–234 (2019).

54. Hauser, T. U. *et al.* Role of the medial prefrontal cortex in impaired decision making in juvenile attention-deficit/hyperactivity disorder. *JAMA psychiatry* **71**, 1165–1173 (2014).

55. Williams, J. & Taylor, E. The evolution of hyperactivity, impulsivity and cognitive diversity. *J. Royal Soc. Interface* **3**, 399–413 (2006).

56. Li, N. *et al.* Directed exploration is reduced by an aversive interoceptive state induction in healthy individuals but not in those with affective disorders. *Mol. Psychiatry* 1–10 (2025).

57. Smith, R. *et al.* Lower levels of directed exploration and reflective thinking are associated with greater anxiety and depression. *Front. Psychiatry* **12**, 782136 (2022).

58. Zuckerman, M. *Behavioral expressions and biosocial bases of sensation seeking* (Cambridge university press, 1994).

59. Zuckerman, M. The psychophysiology of sensation seeking. *J. personality* **58**, 313–345 (1990).

60. Wasserman, A. M. *et al.* The age-varying effects of adolescent stress on impulsivity and sensation seeking. *J. Res. on Adolesc.* **33**, 1011–1022 (2023).

61. Seow, T. X. & Hauser, T. U. Reliability of web-based affective auditory stimulus presentation. *Behav. research methods* **54**, 378–392 (2022).

62. Arshad, F. *et al.* Effects of audiovisual cues on game immersion during simulated slot machine gambling. *J. Gambl. Stud.* 1–18 (2025).

63. Dixon, M. J. *et al.* The impact of sound in modern multiline video slot machine play. *J. Gambl. Stud.* **30**, 913–929 (2014).

64. Cherkasova, M. V. *et al.* Win-concurrent sensory cues can promote riskier choice. *J. Neurosci.* **38**, 10362–10370 (2018).

65. Dixon, M. J., Harrigan, K. A., Sandhu, R., Collins, K. & Fugelsang, J. A. Losses disguised as wins in modern multi-line video slot machines. *Addiction* **105**, 1819–1824 (2010).

66. Dixon, M. J., Collins, K., Harrigan, K. A., Graydon, C. & Fugelsang, J. A. Using sound to unmask losses disguised as wins in multiline slot machines. *J. Gambl. Stud.* **31**, 183–196 (2015).

67. Myles, D. *et al.* "losses disguised as wins" in electronic gambling machines contribute to win overestimation in a large online sample. *Addict. Behav. Reports* **18**, 100500 (2023).

68. Clark, L., Lawrence, A. J., Astley-Jones, F. & Gray, N. Gambling near-misses enhance motivation to gamble and recruit win-related brain circuitry. *Neuron* **61**, 481–490 (2009).

69. Sescousse, G. *et al.* Amplified striatal responses to near-miss outcomes in pathological gamblers. *Neuropsychopharmacology* **41**, 2614–2623 (2016).

70. Palmer, L., Ferrari, M. A. & Clark, L. The near-miss effect in online slot machine gambling: A series of conceptual replications. *Psychol. Addict. Behav.* **38**, 716 (2024).

71. Gray, J. M. & Wilson, M. A. A detailed analysis of the reliability and validity of the sensation seeking scale in a uk sample. *Pers. Individ. Differ.* **42**, 641–651 (2007).

72. Norbury, A., Valton, V., Rees, G., Roiser, J. P. & Husain, M. Shared neural mechanisms for the evaluation of intense sensory stimulation and economic reward, dependent on stimulation-seeking behavior. *J. Neurosci.* **36**, 10026–10038 (2016).

73. Patton, J. H., Stanford, M. S. & Barratt, E. S. Factor structure of the barratt impulsiveness scale. *J. clinical psychology* **51**, 768–774 (1995).

74. Lovibond, S. H. Manual for the depression anxiety stress scales. *Syd. psychology foundation* (1995).

75. Foa, E. B. *et al.* The obsessive-compulsive inventory: development and validation of a short version. *Psychol. assessment* **14**, 485 (2002).

76. Saunders, J. B., Aasland, O. G., Babor, T. F., De la Fuente, J. R. & Grant, M. Development of the alcohol use disorders identification test (audit): Who collaborative project on early detection of persons with harmful alcohol consumption-ii. *Addiction* **88**, 791–804 (1993).

77. Skinner, H. A. The drug abuse screening test. *Addict. behaviors* **7**, 363–371 (1982).

78. Condon, D. M. & Revelle, W. The international cognitive ability resource: Development and initial validation of a public-domain measure. *Intelligence* **43**, 52–64 (2014).

79. Oppenheimer, D. M., Meyvis, T. & Davidenko, N. Instructional manipulation checks: Detecting satisficing to increase statistical power. *J. experimental social psychology* **45**, 867–872 (2009).

80. Wu, C. M., Schulz, E., Pleskac, T. J. & Speekenbrink, M. Time pressure changes how people explore and respond to uncertainty. *Scientific Reports* **12**, 1–14, DOI: 10.1038/s41598-022-07901-1 (2022).

81. Gershman, S. J. Deconstructing the human algorithms for exploration. *Cognition* **173**, 34–42 (2018).

82. Speekenbrink, M. & Konstantinidis, E. Uncertainty and exploration in a restless bandit problem. *Top. cognitive science* **7**, 351–367 (2015).

83. de Boer, L. *et al.* Attenuation of dopamine-modulated prefrontal value signals underlies probabilistic reward learning deficits in old age. *elife* **6**, e26424 (2017).

84. Wise, T. & Dolan, R. J. Associations between aversive learning processes and transdiagnostic psychiatric symptoms in a general population sample. *Nat. communications* **11**, 4179 (2020).

85. Wise, T., Michely, J., Dayan, P. & Dolan, R. J. A computational account of threat-related attentional bias. *PLoS computational biology* **15**, e1007341 (2019).

86. Vehtari, A., Gelman, A. & Gabry, J. Practical bayesian model evaluation using leave-one-out cross-validation and waic. *Stat. computing* **27**, 1413–1432 (2017).

Acknowledgements

We thank Agnes Norbury for providing the data that made Experiment 1 possible. TUH has received funding from the Wellcome Trust (316955/Z/24/Z) and the Carl-Zeiss-Stiftung. TUH consults for limbic ltd and holds shares in the company, which is unrelated to the current project.

Author contributions statement

E.W., T.H. and C.M.W. conceived the experiments. E.W. and A.C. programmed and conducted the experiments. E.W. analysed the results. C.M.W. and T.H. supervised the project. E.W. wrote the first draft with contributions from C.M.W. All authors reviewed the final manuscript.

Additional information

Experiment 1 data can be requested from the original authors[16]. The data and code are not publicly available due to ethical guidelines. Experiment 2 data and code are available at `github.com/ErnWg/UncertaintyToleranceSS`

Supporting information

Additional Behavioural/ Questionnaire Analysis

Distribution of Performance and Entropy

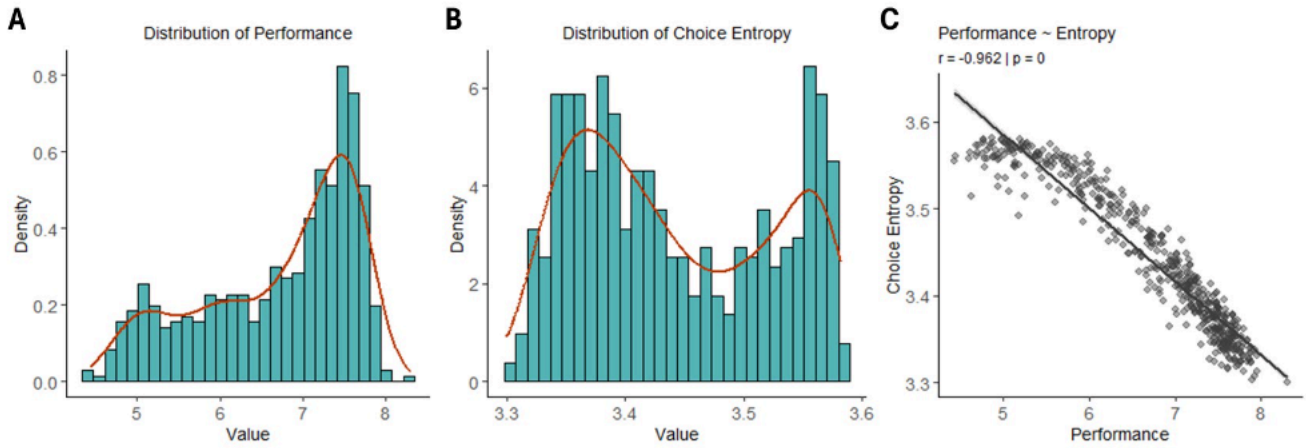

Figure S1. Distributions and association between performance and choice entropy. **(A)** Histogram and smoothed density of *Performance*. **(B)** Histogram and smoothed density of *Choice Entropy*. **(C)** Scatterplot of *Choice Entropy* versus *Performance* with linear fit (95% CI). Pearson correlation: $r = -0.962$, $p < 0.001$.

Effect of Traits on Round Performance

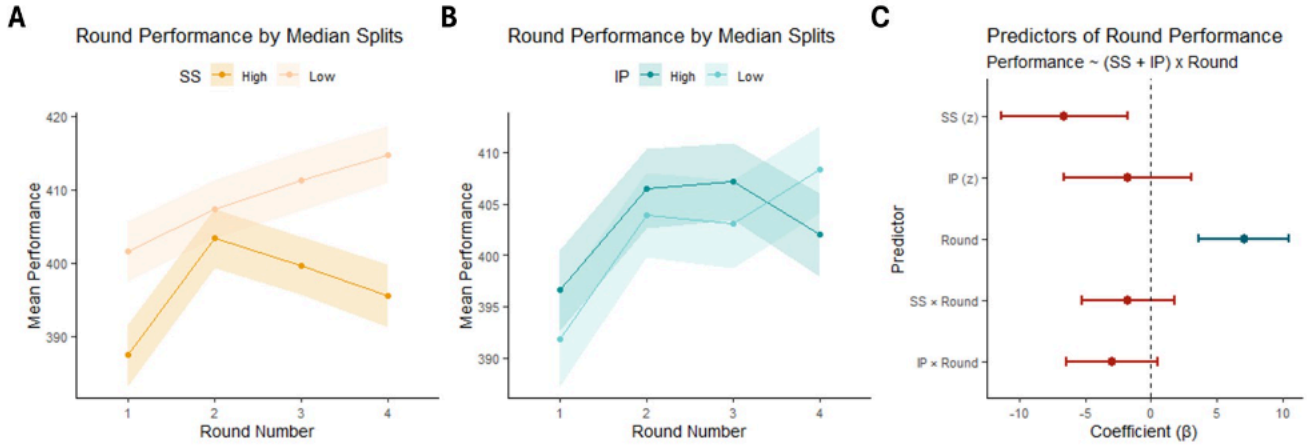

Figure S2. Round Performance × Trait. **(A)** Mean total reward by round for median splits on Sensation Seeking (SS); lines show group means and shaded bands show 95% CIs. **(B)** Mean total reward by round for median splits on Impulsivity (IP) with 95% CIs. **(C)** Fixed–effect coefficients (points) and 95% CIs (bars) from the mixed model $roundScore \sim (SS + IP) \times round + (1 \mid ID)$. We observed a negative association with SS ($\beta_{SS} = -6.55$, $p = 0.008$) and a positive linear effect of round ($\beta_{round} = 7.06$, $p < 0.001$), consistent with practice-related improvement, but no main effect of IP ($\beta_{IP} = -1.73$, $p = 0.48$). Neither SS × round nor IP × round interactions were significant ($\beta_{SS \times round} = -1.71$, $p = 0.34$; $\beta_{IP \times round} = -2.94$, $p = 0.10$), providing no evidence that the SS effect reflects accumulating fatigue or disengagement over time.

Effect of Traits on Reaction Times

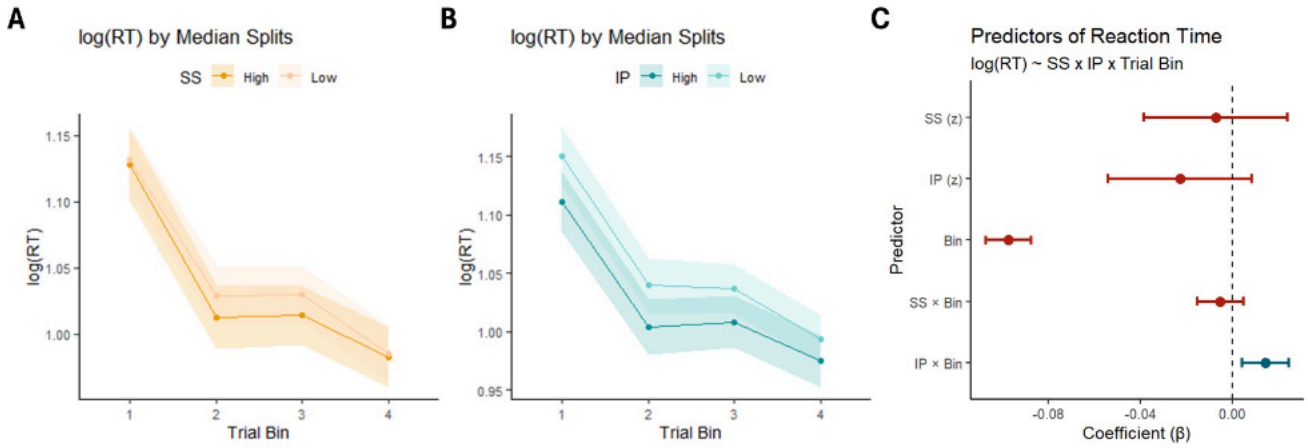

Figure S3. Reaction time by trait and trial bin.**(A)** Mean log(RT) by trial bin for median splits on Sensation Seeking (SS); lines show group means, shaded bands show 95% CIs.**(B)** Mean log(RT) by trial bin for median splits on Impulsivity (IP) with 95% CIs.**(C)** Fixed–effect coefficients (points) and 95% CIs (bars) from the mixed model $\log(\mathrm{RT}) \sim (\mathrm{SS} + \mathrm{IP}) \times \mathrm{Bin} + (1 \mid \mathrm{ID})$. RTs decreased across bins ($\beta_{bin} = 0.10$, $p < 0.001$), reflecting increasing efficiency, but there was no main effect of SS ($\beta_{SS} = -0.01$, $p = 0.65$) and no SS × bin interaction ($\beta_{SS \times bin} = -0.01$, $p = 0.31$). Thus, high-SS participants were generally neither lower nor faster, and they did not exhibit an atypical RT learning trajectory. IP scores showed no main effect on RTs ($\beta_{IP} = -0.02$, $p = 0.15$), although a significant IP × bin interaction ($\beta_{IP \times bin} = 0.01$, $p = 0.005$) indicated a modest modulation of RT dynamics without corresponding performance differences.

Effect of Traits on Performance with Omission Rate as Covariate

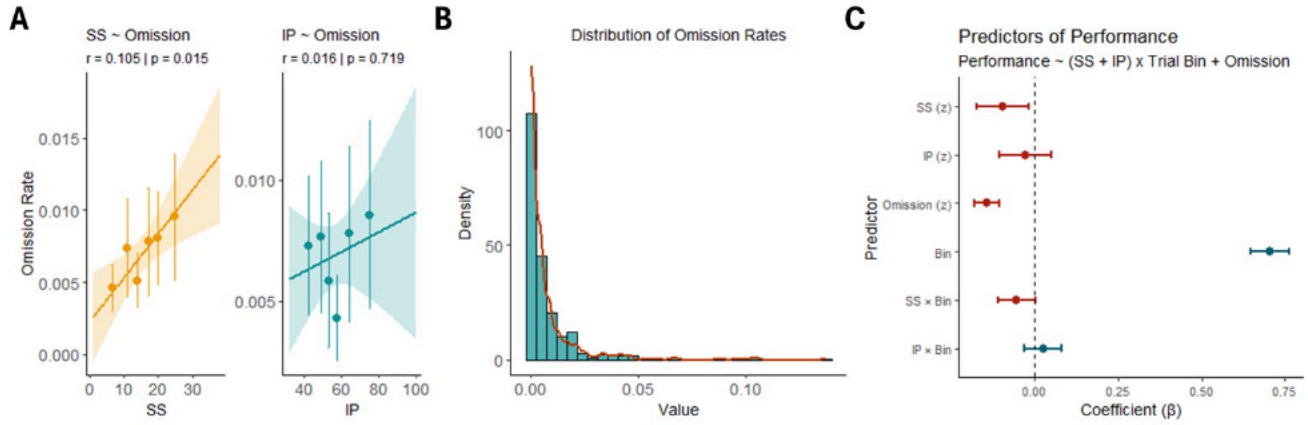

Figure S4. Omission rates and performance. **(A)** Relationship between omission rate and traits: Sensation Seeking (SS) and Impulsivity (IP); points show participant means with 95% CIs and fitted linear trends. **(B)** Distribution of omission rates across trials/participants (histogram with smoothed density). **(C)** Fixed–effect coefficients (points) and 95% CIs (bars) from the mixed model Performance $\sim (\mathrm{SS} + \mathrm{IP}) \times \mathrm{Bin} + \mathrm{Omission} + (1 \mid \mathrm{id})$, with standardised predictors. Omission rate robustly predicted poorer performance ($\beta_{missed} = -0.21$, $p < 0.001$), but SS crucially remained a significant negative predictor even after controlling for omissions ($\beta_{SS} = -0.09$, $p = 0.02$). Thus, increased misses contribute to, but do not fully explain, the SS-related performance cost.

Correlation Matrix of Questionnaires

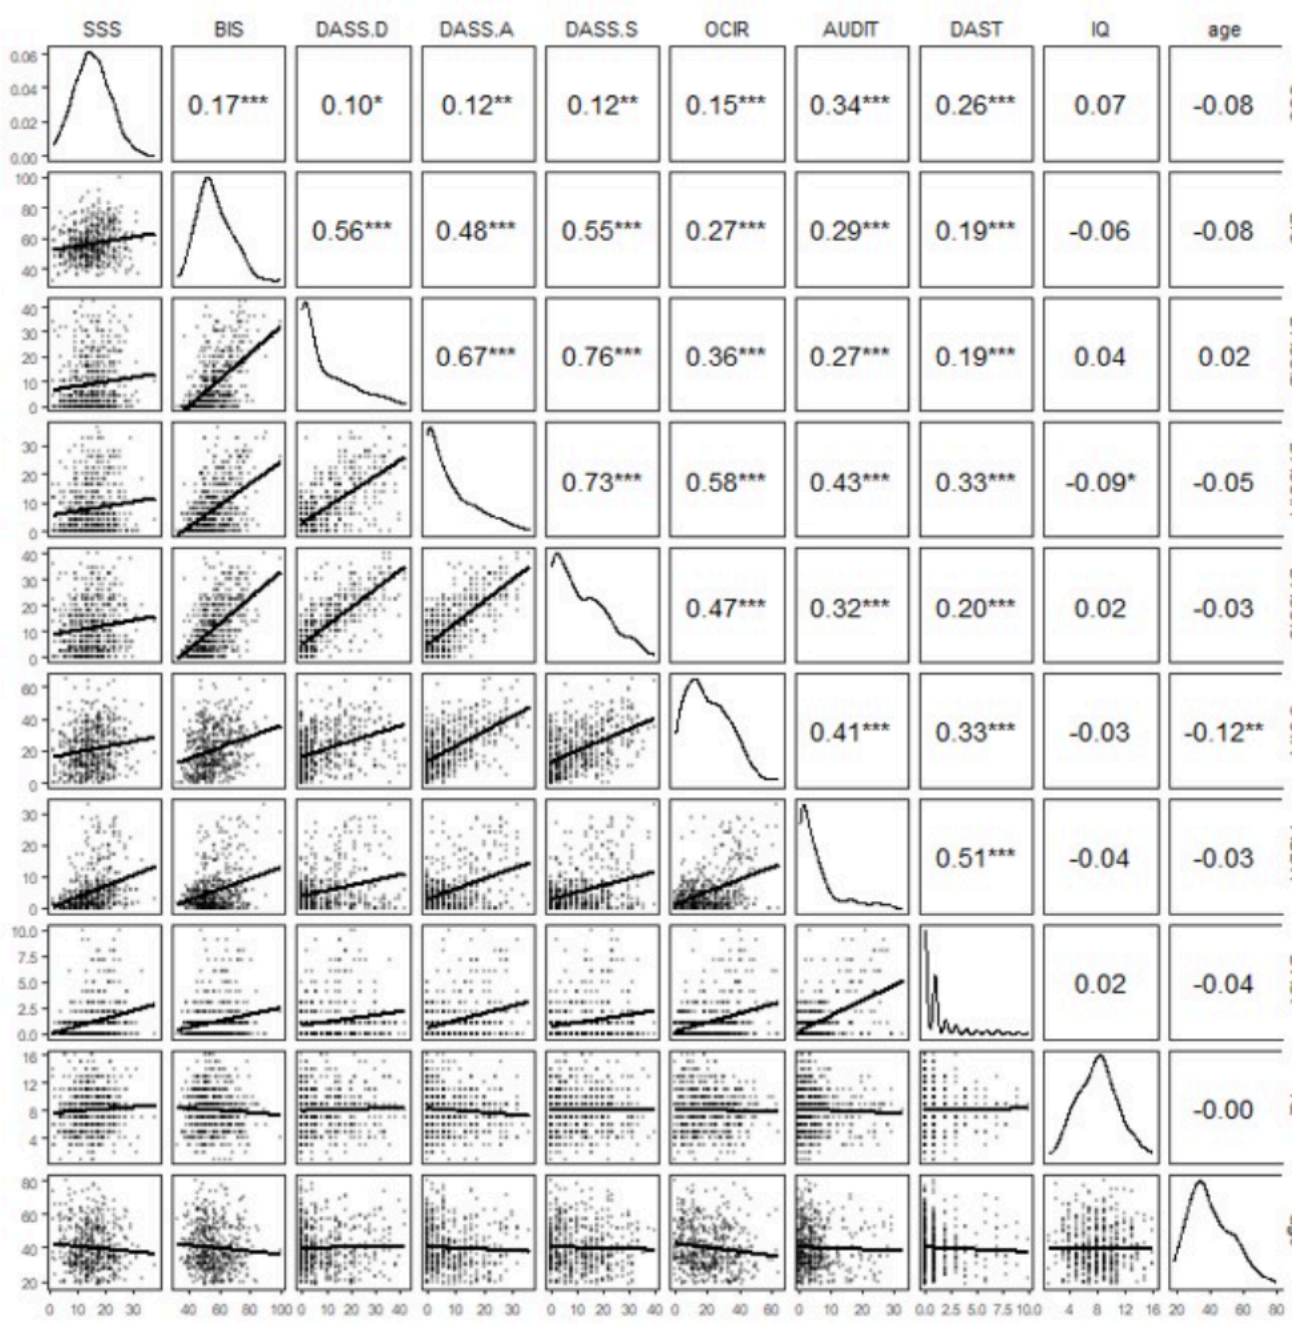

Figure S5. Correlation matrix for questionnaire measures. Upper triangle displays Pearson's r with significance stars ($^{***}p < .001$, $^{**}p < .01$, $^{*}p < .05$); diagonal panels show kernel density estimates; lower triangle shows scatterplot with a best fit line. Abbreviations: SSS = Sensation Seeking Scale; BIS = Barratt Impulsiveness Scale; DASS.D/A/S = Depression, Anxiety, and Stress subscales of the Depression Anxiety Stress Scales (DASS); OCIR = Obsessive-Compulsive Inventory-Revised (OCI-R); AUDIT = Alcohol Use Disorders Identification Test; DAST = Drug Abuse Screening Test; IQ = intelligence quotient.

General Linear Mixed Effect Modelling

Motor vs Stimulus Stickiness

We compared two forms of choice perseveration/ stickiness: a preregistered stimulus stickiness bonus based on the identity of the previous choice ($Sticky_s$, model 2) and a motor stickiness bonus based on the previous response side ($Sticky_m$, model 3). Motor stickiness provided a better overall fit, so we focus on $Sticky_m$ in subsequent analysis. We also evaluated models with random intercepts and random slopes for $R\mu$. Attempts to add random slopes for $S\mu$ and $S\sigma$ produced singular fits. Our best-fitting specification was Model 5, which we therefore used to test trait interactions in the next subsection.

model	AIC	BIC	RSME	Formula
1	141153.9	141192.8	0.4369886	$choice \sim \Delta R\mu + \Delta S\mu + \Delta S\sigma$
2	141135.4	141184.1	0.4369726	$choice \sim \Delta R\mu + \Delta S\mu + \Delta S\sigma + Sticky_s$
3	141098.0	141146.7	0.4368896	$choice \sim \Delta R\mu + \Delta S\mu + \Delta S\sigma + Sticky_m$
4	139950.1	140008.6	0.4319289	$choice \sim \Delta R\mu + \Delta S\mu + \Delta S\sigma + Sticky_m + (1\|ID)$
5	131801.5	131879.4	0.4172918	$choice \sim \Delta R\mu + \Delta S\mu + \Delta S\sigma + Sticky_m + (1 + \Delta R\mu\|ID)$

Table S1. Model comparison for Experiment 2 predicting rightward choice (choice = 1 for "right"). All models include the core decision variables $\Delta R\mu$, $\Delta S\mu$, and $\Delta S\sigma$; optional history terms are Sticky_s (stimulus perseveration) and Sticky_m (motor perseveration). Random-effects structures are denoted $(1 \mid \text{ID})$ and $(1 + \Delta R\mu \mid \text{ID})$. Fit indices reported are AIC, BIC, and RMSE (all lower is better). Adding Sticky_m improves fit over Sticky_s (Model 3 vs. 2); allowing random intercepts improves fit (Model 4); the model with random intercepts and a random slope for $\Delta R\mu$ (Model 5) provides the best overall fit and is used for subsequent inference.

Interactions with Traits

We tested all candidate two- and three-way interactions between core decision variables and our traits of interest—sensation seeking (SS) and impulsivity (IP). Significant interactions are reported in the main text; complete model outputs (all coefficients β weights) are provided in Fig. S6. To assess model specification, we used DHARMa simulation-based residual diagnostics in R, including checks for dispersion, outliers, and normality (Fig. S7).

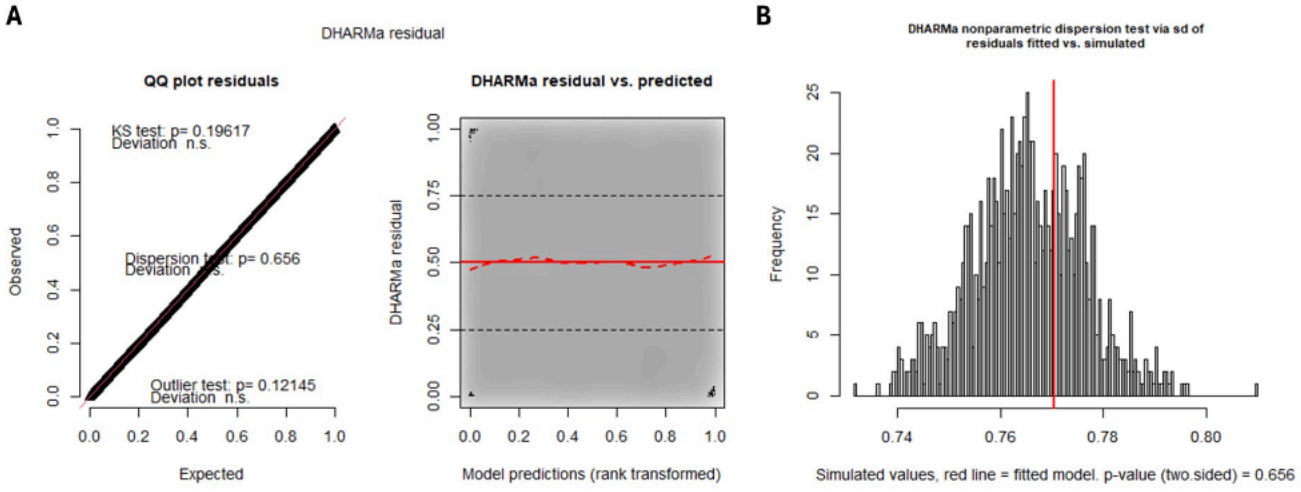

Figure S6. DHARMa diagnostics for the Predictors × Trait GLMM (1,000 simulations). (A) Left: QQ plot of scaled residuals indicates approximate uniformity (KS test, $p = 0.20$). Right: residuals vs. fitted values with DHARMa quantile curves show no systematic pattern; dispersion test n.s. ($p = 0.66$). (B) Nonparametric dispersion check: distribution of simulated residual SDs with the fitted-model value (red line); $p = 0.66$. Outlier test n.s. ($p = 0.12$)

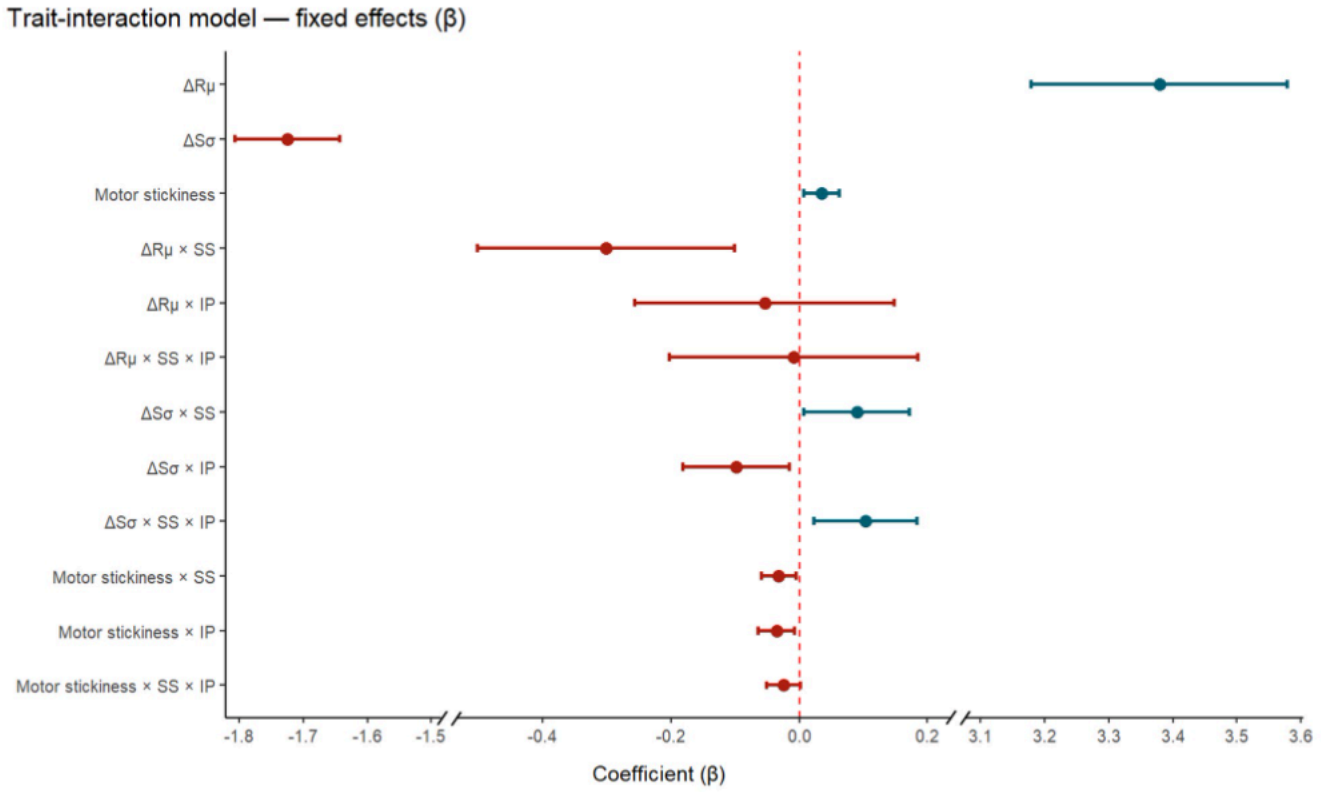

Figure S7. Fixed-effect coefficients (β) from the Trait–Interaction GLMM (binomial link). Points show estimates and horizontal bars show 95% CIs for the main effects of the decision variables ($\Delta R\mu$, $\Delta S\sigma$, Motor stickiness) and their interactions with traits—Sensation Seeking (SS) and Impulsivity (IP)—including two-way and three-way terms (e.g., $\Delta R\mu \times SS$, $\Delta R\mu \times IP$, $\Delta R\mu \times SS \times IP$). The vertical dashed line marks $\beta = 0$; effects whose CIs do not cross zero are considered statistically reliable. Random effects: intercepts and by-participant slopes for $\Delta R\mu$.

Description of Reduced Models

We outline here the reduced models introduced in the main text. These models differ only in their choice rule. For each bandit j at time t, the Q-value $Q_{j,t}$ was defined as a linear combination of the estimated underlying point payoff $R\mu_{j,t}$ and, depending on the model, the stimulation-related terms: the estimated mean probability of stimulation $S\mu_{j,t}$, weighted by free parameter θ, and the associated uncertainty $S\sigma_{j,t}$, weighted by free parameter ω. All Q-values were scaled by the temperature parameter τ:

R-only

$$Q_{j,t} = \frac{R\mu_{j,t}}{\tau} \tag{17}$$

R-θ

$$Q_{j,t} = \frac{R\mu_{j,t} + \theta S\mu_{j,t}}{\tau} \tag{18}$$

R-ω

$$Q_{j,t} = \frac{R\mu_{j,t} + \omega S\sigma_{j,t}}{\tau} \tag{19}$$

R-$\theta\omega$

$$Q_{j,t} = \frac{R\mu_{j,t} + \theta S\mu_{j,t} + \omega S\sigma_{j,t}}{\tau} \tag{20}$$

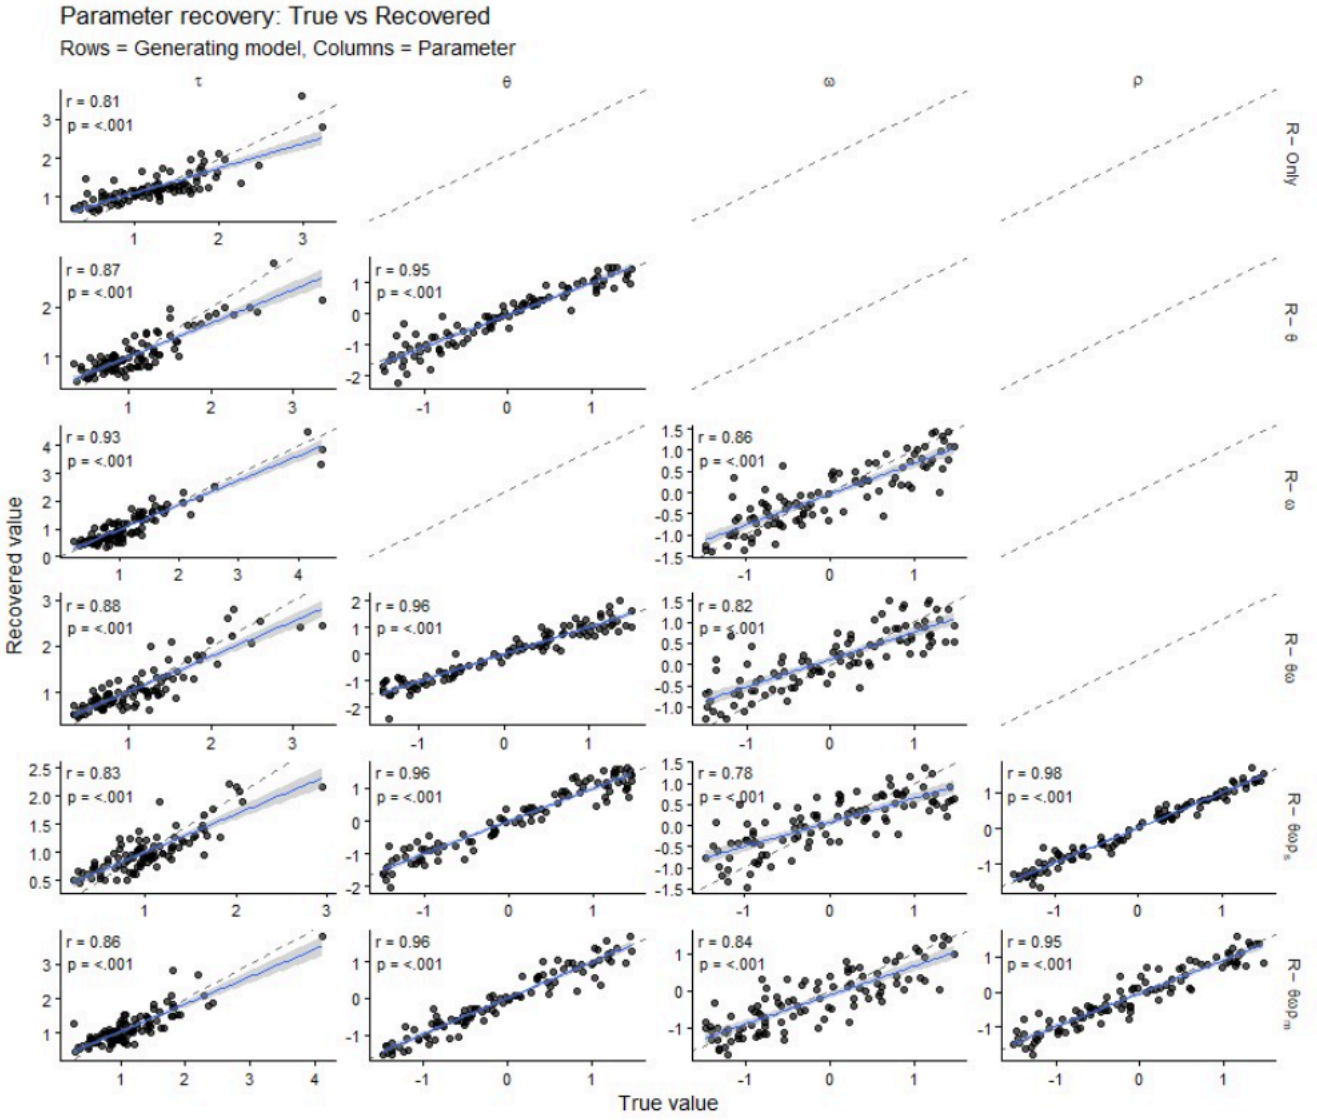

Figure S8. Scatter plots compare true (x-axis) vs recovered posterior mean (y-axis) parameters for a synthetic dataset ($N = 100$). Rows are generating models (R–Only, R–θ, R–ω, R–$\theta\,\omega$, R–$\theta\,\omega\rho_s$, R–$\theta\,\omega\rho_m$); columns are parameters (τ, θ, ω, ρ). The dashed line shows the identity (y = x); the solid line is a regression line fit with 95% CI. Panel insets report Pearson and p-values. Blank panels indicate parameters not present in that generating model.

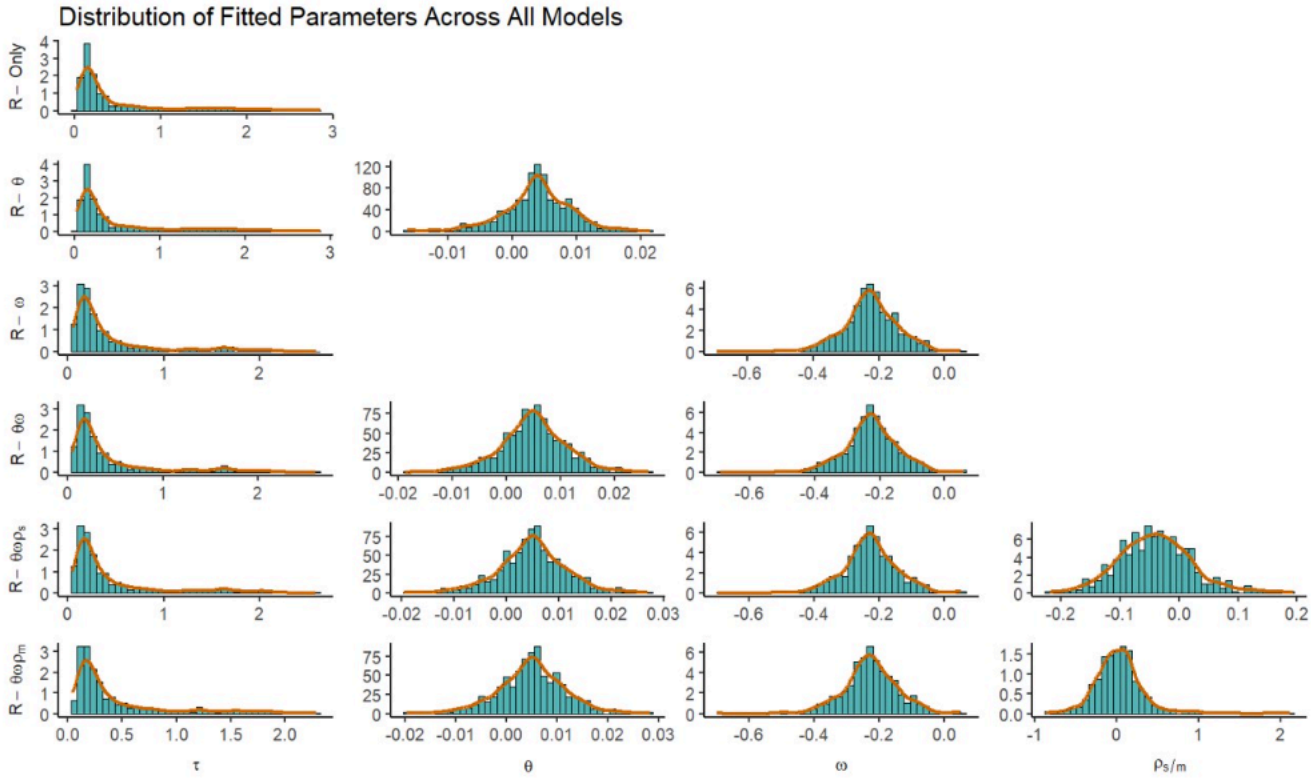

Figure S9. Distributions of fitted parameters across computational models. Columns show parameter estimates for θ, ω, ρ, and τ; rows correspond to model variants. Histograms depict participant-level fits with overlaid kernel density curves (orange). *Note:* In models with stimulus stickiness (ρ_s), κ denotes the bonus assigned to stimulus identity (stimulus stickiness); in models with motor stickiness (ρ_m), κ denotes the action-repetition bonus (motor stickiness).

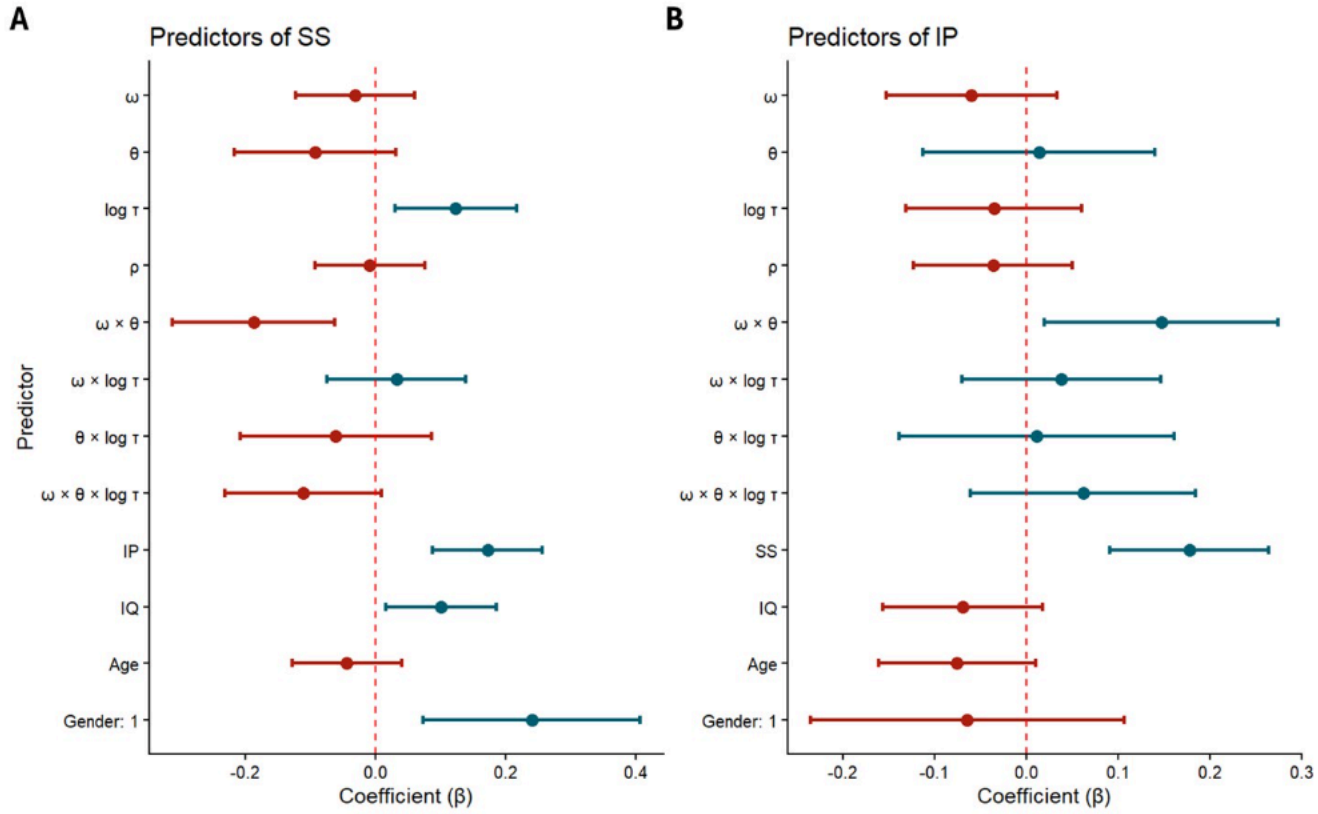

Figure S10. Regression of traits on computational parameters and covariates. Points show fixed–effect coefficients (β) with 95% CIs; the vertical dashed line marks $\beta = 0$. **(A)** Predictors of Sensation Seeking (SS). **(B)** Predictors of Impulsivity (IP). Predictors include model parameters ω, θ, and $\log \tau$, the stickiness parameter ρ (additive main effect), and covariates (IP or SS, IQ, Age, Gender). Interaction terms were limited to $\omega \times \theta$, $\omega \times \log \tau$, $\theta \times \log \tau$, and $\omega \times \theta \times \log \tau$. *Note:* We did not include ρ in the interaction terms, as this was not theoretically motivated and we sought to avoid unnecessary expansion of the interaction space.