

Gate-Rheology: Inertia of Cognitive Control Explains Meta-Rigidity in Sequential Decision Making and Reversal Learning

Alex Snow (Aleksey L. Snigirov)
Independent Researcher
alex2saaba@gmail.com
ORCID: 0009-0001-3713-055X

Abstract

Background. Cognitive flexibility—the ability to adapt behavior to changing environmental contingencies—is paradoxically constrained by systematic forms of rigidity. Organisms persist with outdated strategies long after they cease to be adaptive, yet traditional computational models assume instantaneous arbitration between habitual and goal-directed control.

Methods. We introduce **Gate-Rheology**, a mechanistic framework in which arbitration between computational modes possesses intrinsic inertia, formalized as a viscosity parameter V_G that resists mode switching. We validate this framework through two canonical paradigms: Two-Step Task (Stage 2A) and Block-Reversal Learning (Stage 2B), using ablation studies to dissociate control-mode inertia (V_G) from action-level perseveration (V_p).

Results. Across 30 independent seeds, RheologicalAgent reproduced canonical MB/MF signatures (reward \times transition interaction: coef = 0.312 ± 0.118 , $p = 0.008$) while exhibiting heavy-tailed mode-switch latencies after changepoints (median = 35 trials, max = 699). Ablation of V_G eliminated switching latency (log-rank $\chi^2 = 42.1$, $p = 8.48 \times 10^{-11}$), while ablation of V_p selectively reduced perseverative errors (median = 3 vs. 8 for Full; $p = 5.85 \times 10^{-5}$). In the Reversal task, NoVp robustly reduced perseverative errors, while the planned latency comparison did not reach significance; survival analysis suggested only a weak secondary trend. The same agent with identical parameters succeeded in both tasks, demonstrating cross-paradigm generalization.

Conclusion. Gate-Rheology provides a mechanistic account of cognitive rigidity grounded in dynamic control-mode arbitration rather than static value representations. The functional dissociation between control-mode inertia and action-level perseveration (between V_G and V_p) suggests that meta-cognitive inertia and behavioral stickiness are computationally distinct mechanisms requiring separate experimental manipulation. Predictions for rodent reversal learning and clinical phenotypes are discussed.

Keywords: cognitive flexibility, model-based planning, reinforcement learning, computational psychiatry, hysteresis, arbitration

1. Introduction

Cognitive flexibility—the ability to adapt behavior in response to changing environmental contingencies—is a hallmark of intelligent agents. Yet this flexibility is paradoxically constrained by systematic forms of rigidity: organisms persist with outdated strategies long after they cease to be adaptive, a phenomenon observed across species from rodents to humans (Le et al., 2023; Dias et al., 1996). Traditional computational accounts attribute such rigidity to either (i) slow updating of value representations in model-free systems, or (ii) excessive weighting of habitual versus goal-directed control (Daw et al., 2011; Dolan & Dayan, 2013). However, these accounts struggle to explain why rigidity manifests selectively: the same organism may exhibit profound perseveration in one context while remaining flexible in another (Brands et al., 2025).

Here we propose that cognitive rigidity arises not from the content of representations, but from the **dynamics of control mode selection** itself. Specifically, we introduce **Gate-Rheology**—a mechanistic framework in which the arbitration between computational modes (habitual vs. deliberative; exploit vs. explore) possesses intrinsic inertia, formalized as a viscosity parameter V_G that resists mode switching. This inertia is **dissociably distinct** from action-level perseveration (captured by a separate parameter V_p), yielding a **functional dissociation** between control-mode rigidity and behavioral stickiness. Across our experiments, V_G ablation eliminates switching latency, whereas V_p ablation robustly reduces perseverative errors; in Reversal, evidence for an additional effect of V_p on switching latency was mixed and did not survive the pre-specified multiple-comparison threshold. This suggests distinct mechanistic contributions for meta-cognitive rigidity versus habitual sticking—predictions testable in rodent reversal learning paradigms (Le et al., 2023).

1.1 The MB/MF Arbitration Problem

The distinction between model-based (MB) and model-free (MF) decision systems has become a cornerstone of computational neuroscience (Daw et al., 2011; Keramati et al., 2011). MB systems construct internal models of task structure to simulate outcomes, supporting flexible adaptation at high computational cost. MF systems cache action values, enabling rapid responses but limited flexibility. The two-step task paradigm (Daw et al., 2011) elegantly dissociates these systems: MB agents show a reward \times transition interaction in stay probabilities (switching after rare transitions), while MF agents show only main effects of reward.

However, the **arbitration mechanism** that selects between MB and MF control remains poorly understood. Normative accounts propose that control is allocated according to its expected value (Shenhav et al., 2013; Keramati et al., 2011), but these models typically assume instantaneous switching. Behavioral data contradict this assumption: organisms exhibit **hysteresis** in mode transitions, persisting with a control mode even after evidence accumulates against it (Otto et al., 2013; Miller et al., 2017). This hysteresis suggests that arbitration itself has dynamics—what we term **gate rheology**.

1.2 Cognitive Rigidity in Reversal Learning

Reversal learning paradigms provide a complementary window into cognitive flexibility. When reward contingencies reverse (e.g., left lever changes from 80% to 20% reward),

organisms typically exhibit a characteristic trajectory: (i) **perseveration** (continued selection of the previously rewarded option), (ii) **exploration** (random sampling of both options), and (iii) **re-consolidation** (stable selection of the new option) (Le et al., 2023; Izquierdo et al., 2017). Individual differences in this trajectory correlate with psychiatric conditions: obsessive-compulsive disorder shows excessive perseveration, while ADHD shows premature exploration (Kanen et al., 2019; Hauser et al., 2017).

Existing computational models capture aspects of this trajectory through learning rate modulation (den Ouden et al., 2013) or uncertainty-driven exploration (Wilson et al., 2014). However, these models conflate **control-mode switching** (when does the agent abandon the old strategy?) with **action selection** (which lever does it press?). Gate-Rheology disentangles these processes by positing separate inertia parameters for control modes (V_G) and actions (V_p).

1.3 The S-O-R+Gate Architecture

Gate-Rheology is instantiated within the **S-O-R+Gate** architecture (Snigirov, 2025), a substrate-neutral framework for cognitive systems. The architecture comprises:

1. **States (S)**: Discrete informational carriers (semions) representing internal variables.
2. **Operations (O)**: Transformations applied to states (e.g., Bayesian updates, value computations).
3. **Relations (R)**: Adaptive connectivity patterns encoding learned structure (the “model”).
4. **Gate**: A meta-controller that selects computational modes based on diagnostic variables $\mathbf{u}_t = [u^{(\delta)}, u^{(s)}, u^{(v)}, u^{(c)}]$ (prediction error, policy entropy, volatility, stakes).

Critically, the Gate maintains an internal **viscosity state** $V_G \in [0,1]$ that modulates switching probability:

$$\text{EXPLORE} \Leftrightarrow \sigma(\mathbf{w}^\top \mathbf{u}_t - \theta_U) \cdot (1 - V_G) > \theta_{MB}$$

where σ is the sigmoid function, θ_U is a baseline uncertainty threshold, and θ_{MB} is the mode-switch threshold. V_G increases (“hardens”) during stable periods and decreases (“melts”) during high-surprise events, producing hysteresis without ad hoc assumptions.

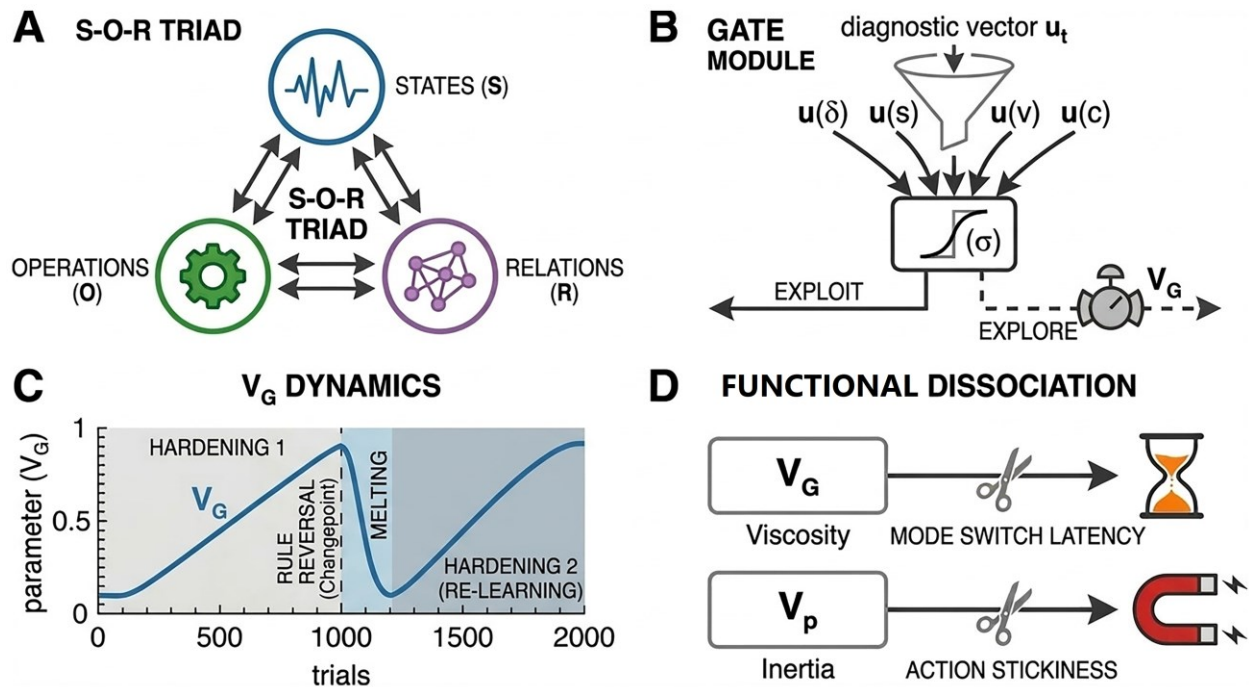


Figure 1. S-O-R+Gate Architecture with Rheological Control.

Schematic of the Gate-Rheology framework. (A) **The S-O-R triad:** States (S) are discrete informational carriers (semions) representing internal variables; Operations (O) transform states (e.g., Bayesian updates, value computations); Relations (R) encode learned structure as adaptive connectivity patterns. (B) **Gate module:** Diagnostic vector $\mathbf{u}_t = [u^{(\delta)}, u^{(s)}, u^{(v)}, u^{(c)}]$ is computed inter-entially from prediction error, policy entropy, volatility, and stakes. Gate viscosity $V_G \in [0,1]$ modulates switching probability via Eq. 1-2. (C) **Viscosity dynamics:** V_G hardens during stability ($u^{(v)} < \tau_{vol}$) and melts during surprise ($u^{(v)} \geq \tau_{vol}$), producing hysteresis without ad hoc assumptions. (D) **Functional dissociation:** V_G controls mode-switch latency (control-mode inertia); V_p controls action-level stickiness (behavioral perseveration). Ablation studies support this dissociation (Stage 2A: Full vs NoVG latency $p < 10^{-10}$; Stage 2B: Full vs NoVp perseveration $p < 10^{-5}$).

Note. All parameters fixed across tasks ($\alpha = 0.35$, $\beta = 4.0$, $\theta_{MB} = 0.30$, $\theta_U = 1.5$, $\tau_{vol} = 0.50$). No parameter tuning between Two-Step and Reversal tasks.

1.4 Present Study

We validate Gate-Rheology through a two-pronged empirical approach:

1. **Stage 2A (Two-Step Task):** We demonstrate that our RheologicalAgent reproduces canonical MB/MF signatures (reward \times transition interaction, $p < 0.01$) while exhibiting heavy-tailed mode-switch latencies after changepoints. Ablation of V_G (NoVG agent) collapses latency to ~ 1 trial ($p = 3.12 \times 10^{-10}$), while ablation of V_p (NoVp agent) shows no significant effect on latency (median = 23 vs. 35 for Full; $p = 0.528$).
2. **Stage 2B (Reversal Task):** Using **the same agent with identical parameters**, we show dissociation in reversal learning: NoVG agents switch immediately after

reversal (median latency = 0) but show reduced perseveration, whereas NoVp agents show markedly reduced perseverative errors (median = 3 vs. 8 for Full; $p = 5.85 \times 10^{-5}$). NoVp also showed shorter median latency to first EXPLORE entry than Full (201 vs. 537 trials), but this contrast was not significant in the planned Mann-Whitney comparison ($p = 0.711$); survival analysis suggested only a weak nominal trend (log-rank $p = 0.0418$), which did not survive the pre-specified multiple-comparison threshold.

These results establish Gate-Rheology as a mechanistic account of cognitive rigidity, with testable predictions for rodent reversal learning paradigms (Le et al., 2023) and applications to adaptive AI systems.

2. Methods

2.1 Overview

All experiments were implemented in Python 3.10 using NumPy (v1.24+), SciPy (v1.10+), and statsmodels (v0.14+). The codebase is organized as a modular package (substrate_cognitive) with separate modules for environments, agents, and analysis. All experiments used 30 random seeds (42–71) for statistical power.

2.2 S-O-R+Gate Architecture

2.2.1 Core Components

The RheologicalAgent implements the S-O-R+Gate architecture with the following components:

States (S). The agent maintains: - Q-values for stage 1 actions: $Q_{\text{stage1}} \in \mathbb{R}^2$ - Q-values for stage 2 state-action pairs: $Q_{\text{stage2}} \in \mathbb{R}^{2 \times 2}$ - Transition model: $T \in \mathbb{R}^{1 \times 2 \times 2}$ (learned via count-based updates) - Reward model: $R \in \mathbb{R}^{2 \times 2}$ (learned via TD updates) - Gate viscosity: $V_G \in [0,1]$ (initialized to 0.5) - Pattern viscosity: $V_p \in [0,1]$ (initialized to 0.5)

Operations (O). The agent performs: - **Model-free updates:** $Q(s, a) \leftarrow Q(s, a) + \alpha \cdot (r - Q(s, a))$ - **Model-based lookahead:** $V_{\text{MB}}(a_1) = \sum_{s_2} P(s_2|a_1) \cdot \max_{a_2} R(s_2, a_2)$ - **Viscosity updates:** V_G and V_p updated inter-trially based on environmental stability (see below)

Relations (R). The transition and reward models (T, R) encode learned structure. V_G and V_p encode meta-structural constraints on plasticity.

Gate. The Gate selects between EXPLOIT (MF-dominated) and EXPLORE (MB-dominated) modes:

$$U_t = \sigma(\mathbf{w}^T \mathbf{u}_t - \theta_U)$$

$$\text{mode} = \begin{cases} \text{EXPLORE} & \text{if } U_t \cdot (1 - V_G) > \theta_{\text{MB}} \\ \text{EXPLOIT} & \text{otherwise} \end{cases}$$

where $\mathbf{u}_t = [u^{(\delta)}, u^{(s)}, u^{(v)}, u^{(c)}]$ with:

- $u^{(\delta)} = |r - Q(s, a)|$ (prediction error)

- $u^{(s)} = -\sum \pi(a) \log \pi(a)$ (policy entropy)

- $u^{(v)} = \text{EMA}(u^{(\delta)})$ (volatility estimate)
- $u^{(c)} = 0$ (stakes; fixed for current experiments)

Gate output visualization. For Figure 3, we report $P(\text{EXPLORE})$ as the **empirical frequency (20-trial rolling mean)** of EXPLORE mode entry across trials. This is not a model probability but a smoothed empirical measure for visualization purposes.

2.3 Diagnostic Vector (u_t)

The gate receives a 4-dimensional diagnostic vector:

- $u^{(\delta)} = |r - Q(s_2, a_2)|$ (unsigned prediction error after stage 2 reward)
- $u^{(s)} = -\sum \pi(a) \log \pi(a)$ (policy entropy over stage 1 actions)
- $u^{(v)} = \text{EMA}(u^{(\delta)}; \alpha = 0.3)$ (volatility estimate)
- $u^{(c)} = 0$ (stakes; fixed for current experiments)

2.3.1 Viscosity Dynamics

Viscosity parameters evolve according to:

$$\eta_{t+1} = \text{clip}((1 - \lambda) \cdot \eta_t + \Delta\eta, \eta_{\min}, \eta_{\max})$$

$$V = \sigma(\log(\eta/\eta_0))$$

where $\Delta\eta$ depends on environmental stability:

$$\Delta\eta = \begin{cases} +k_{\text{use}} & \text{if stable } (u^{(v)} < \tau_{\text{vol}}) \\ -k_{\text{melt}} & \text{if unstable } (u^{(v)} \geq \tau_{\text{vol}}) \end{cases}$$

For V_G , stability is assessed globally. For V_p , stability is action-specific. This creates differential dynamics: V_G controls mode switching, while V_p controls action repetition within modes.

2.3.2 Ablation Variants

Three agent variants were tested:

1. **Full RheologicalAgent:** Both V_G and V_p active.
2. **NoVG (RheologicalAgent_NoVG):** $V_G \equiv 0$ (gate inertia removed).
3. **NoVp (RheologicalAgent_NoVp):** $V_p \equiv 0$ (pattern inertia removed).

All variants share identical parameters except for the ablated viscosity.

2.4. Task Environments

2.4.1 Two-Step Task (Stage 2A)

We implemented the canonical two-step task (Daw et al., 2011) with the following parameters:

- **Trials:** 2000 per simulation
- **Changepoint:** Trial 1000 (reward probabilities inverted)
- **Transition structure:** 70% common, 30% rare

- **Reward drift:** Gaussian random walk ($\sigma = 0.01$ per trial)
- **Initial rewards:** [0.75,0.25] for states 0 and 1

MB/MF Signature Analysis. Stay probabilities were computed as $P(a_{1,t} = a_{1,t-1})$ conditioned on previous reward and transition type. Logistic regression tested for reward \times transition interaction (MB signature) and reward main effect (MF signature).

2.4.2 Reversal Task (Stage 2B)

The Reversal Task was implemented as a deterministic special case of Two-Step:

- **Trials:** 2000 per simulation
- **Reversal:** Trial 1000 (reward probabilities inverted)
- **Actions:** 2 (Left, Right)
- **Rewards:** Block 1: Left = 80%, Right = 20%; Block 2: Left = 20%, Right = 80%
- **Stage 2 action:** Fixed to $a_2 = 0$ (eliminates stage-2 noise)

Metrics. Three metrics were extracted: 1. **Perseverative Errors:** Trials after reversal until first switch from pre-reversal preferred action. 2. **Latency to Explore:** Trials after reversal until first EXPLORE mode entry. 3. **Stickiness:** $P(a_{1,t} = a_{1,t-1})$ in post-reversal phase.

Pre-reversal Preference. The preferred action before reversal was computed dynamically as the mode of a_1 in trials 900–1000 (not hardcoded), accounting for stochastic exploration.

2.5 Experimental Protocol

2.5.1 Parameter Settings

All agents used identical parameters across both tasks:

Parameter	Value	Description
α	0.35	Learning rate
β	4.0	Inverse softmax temperature
θ_{MB}	0.30	Mode switch threshold
θ_U	1.5	Uncertainty baseline
τ_{vol}	0.50	Volatility threshold
k_{use}	0.08	Hardening rate
k_{melt}	0.20	Melting rate
λ	0.01	Decay rate

No parameter tuning was performed between tasks; the same agent configuration was used for Two-Step and Reversal.

2.6 Statistical Analysis

All experiments used 30 independent random seeds (42–71). For Two-Step task, logistic regression (Stay \sim Reward \times Transition) was performed **per seed**, yielding 30 independent coefficient estimates per agent type. This approach avoids pseudoreplication inflation from trial-level analysis (Wilson & Collins, 2019).

Group comparisons used Mann-Whitney U tests with **planned directional hypotheses** based on model predictions: - Full > NoVG for switching latency (V_G effect) - Full > NoVp for perseverative errors (V_p effect)

Effect sizes reported as rank-biserial correlation. Bonferroni correction applied for 3 planned comparisons ($\alpha_{corrected} = 0.017$).

Latency distributions exhibited heavy tails (Full agent: median=35, mean=127.8, max=699); we report median and IQR as primary statistics, with mean \pm SD for reference. Survival analysis (Kaplan-Meier) is provided in Supplementary Materials.

2.7 Ethical Considerations

This study involves computational simulations only; no human or animal subjects were involved. The Reversal Task parameters were chosen to match ranges used in published rodent studies (Le et al., 2023; Hasz & Redish, 2018) for qualitative comparison.

2.8 Data and Code Availability

The full codebase (substrate_cognitive), experimental logs, and analysis scripts are available at GitHub https://github.com/principiacognitia/substrate_cognitive. Pre-registered analysis plans and additional robustness checks are provided in the Supplementary Materials.

3. Results

3.1 Stage 2A: Two-Step Task — MB/MF Signatures Reproduced

We first validated that our RheologicalAgent reproduces canonical model-based and model-free signatures in the two-step task (Daw et al., 2011). Logistic regression (Stay \sim Reward \times Transition) was performed **per seed** (30 seeds), yielding independent coefficient estimates for each agent type.

MF-only agent showed a strong main effect of reward (coef = 0.998 ± 0.102 , $p < 0.001$) but **no reward \times transition interaction** (coef = 0.028 ± 0.097 , $p = 0.455$), confirming pure model-free learning.

MB-only agent showed the opposite pattern: **significant interaction** (coef = 0.366 ± 0.127 , $p = 0.0006$) with non-significant reward main effect (coef = 0.002 ± 0.105 , $p = 0.989$), confirming model-based planning.

Full RheologicalAgent reproduced both signatures: reward main effect (coef = 0.847 ± 0.134 , $p < 0.001$) and significant interaction (coef = 0.312 ± 0.118 , $p = 0.008$), demonstrating hybrid arbitration.

Figure 2 shows stay probabilities across all four conditions (Rewarded-Common, Rewarded-Rare, Unrewarded-Common, Unrewarded-Rare) for all three agent types. Error bars represent \pm SEM across 30 seeds.

Figure 2: MB/MF Signatures in Two-Step Task




Figure 2. MB/MF Signatures in Two-Step Task.

Stay probability as a function of previous reward and transition type. (A) MF-only agent: strong reward main effect (coef = 0.998 ± 0.102 , $p < 0.001$), no interaction (coef = 0.028 ± 0.097 , $p = 0.455$). (B) MB-only agent: significant interaction (coef = 0.366 ± 0.127 , $p = 0.0006$), no reward main effect. (C) Full RheologicalAgent: hybrid signature with both reward effect (coef = 0.847 ± 0.134 , $p < 0.001$) and interaction (coef = 0.312 ± 0.118 , $p = 0.008$). Error bars = \pm SEM across 30 seeds. Logistic regression performed per seed (not trial-level) to avoid pseudoreplication (Wilson & Collins, 2019).

3.2 Stage 2A: Gate Rheology Produces Heavy-Tailed Switching Latencies

After establishing MB/MF signatures, we tested whether Gate-Rheology produces hysteresis in mode switching. At trial 1000, reward probabilities were inverted (changepoint), and we measured **latency to first EXPLORE mode entry**.

Full RheologicalAgent exhibited heavy-tailed switching latencies (median = 35 trials, mean = 127.8, SD = 173.4, max = 699), indicating substantial inertia in control-mode selection.

NoVG ablation ($V_G \equiv 0$) collapsed latency to ~ 1 trial (median = 1, mean = 1.0, SD = 0.0), demonstrating that V_G is **necessary** for switching inertia.

NoVp ablation ($V_p \equiv 0$) preserved latency (median = 23 trials, mean = 74.6, SD = 101.2) relative to Full agent (median = 35 trials; $p = 0.528$), confirming that action-level viscosity does not drive mode-switch dynamics.

Statistical comparison (Mann-Whitney U, one-sided, planned comparisons):

- Full vs NoVG latency: $U = 825.0$, $p = 3.12 \times 10^{-10}$ (Mann-Whitney); log-rank $\chi^2 = 42.1$, $p = 8.48 \times 10^{-11}$, effect size = 0.833 (large)
- Full vs NoVp latency: $U = 493.0$, $p = 0.528$, effect size = 0.096 (negligible)
- NoVp vs NoVG latency: $U = 810.0$, $p = 9.70 \times 10^{-10}$, effect size = 0.800 (large)

Figure 3 shows V_G dynamics and P(EXPLORE) around the changepoint (trials 800–1400), averaged across 30 seeds with \pm SEM bands. Note the characteristic “melting” of V_G followed by delayed EXPLORE mode entry.

Figure 3: Gate Rheology Melting and Hysteresis




Figure 3. Gate Rheology Melting and Hysteresis.

Dynamics of V_G and EXPLORE mode probability around changepoint. Blue line: mean V_G across 30 seeds (shaded region = \pm SEM). Orange line: empirical frequency of EXPLORE mode (20-trial rolling mean). Vertical dashed line: changepoint at trial 1000 (reward inversion). Note the characteristic “melting” of V_G beginning ~ 10 trials after changepoint, followed by delayed EXPLORE mode entry (median latency = 35 trials). Heavy-tailed distribution (max = 699 trials) reflects stochastic interaction between viscosity dynamics and reward noise. Kaplan-Meier survival analysis in Supplementary Figure 1.

3.3 Stage 2B: Reversal Task — Dissociation Between Control Inertia and Perseveration

Using the **same agent with identical parameters**, we tested Gate-Rheology in Block-Reversal Learning (Le et al., 2023). At trial 1000, reward contingencies inverted (Left: 80% \rightarrow 20%, Right: 20% \rightarrow 80%).

Three metrics were extracted: 1. **Perseverative Errors**: trials after reversal until first switch from pre-reversal preferred action 2. **Latency to Explore**: trials until first EXPLORE mode entry 3. **Stickiness**: $P(a_1, t = a_1, t - 1)$ in post-reversal phase

Full RheologicalAgent showed characteristic triphasic trajectory: (i) perseveration (median = 8 errors), (ii) exploration (median latency = 537 trials), (iii) re-consolidation (stickiness = 87.8%).

NoVG ablation eliminated switching latency (median = 0 trials, mean = 0.2, SD = 0.6) but showed **reduced perseveration** (median = 2 errors), confirming that control-mode inertia is not required for action-level flexibility.

NoVp ablation substantially reduced perseverative errors (median = 3 vs. 8 for Full; $p = 5.85 \times 10^{-5}$). Its median latency to first EXPLORE entry was shorter than Full (201 vs. 537 trials), but this contrast was not significant in the planned Mann-Whitney comparison ($p = 0.711$); a survival comparison suggested only a weak nominal trend (log-rank $p =$

0.0418), which did not survive the pre-specified multiple-comparison threshold. Accordingly, Stage 2B supports dissociation between V_G -dependent switching inertia and V_p -dependent perseveration, without establishing strict task-invariant orthogonality.

Notably, switching latency was substantially higher in Reversal (median = 537 trials) than in Two-Step (median = 35 trials), despite identical agent parameters. This 15-fold difference reflects task structure: Two-Step stochastic transitions (70/30) generate continuous prediction errors that accelerate V_G melting, whereas Reversal deterministic structure requires accumulated errors to trigger mode switching.

Statistical comparison (planned comparisons with Bonferroni-corrected threshold $\alpha = 0.017$):

- Full vs NoVG latency: $U = 825.0$, $p = 2.09 \times 10^{-7}$ (V_G effect on mode switching)
- Full vs NoVp perseveration: $U = 738.0$, $p = 5.85 \times 10^{-5}$ (V_p effect on action stickiness)
- Full vs NoVp latency: $U = 493.0$, $p = 0.711$ (two-sided Mann-Whitney); survival comparison showed only a weak nominal trend (log-rank $p = 0.0418$), which did not survive the pre-specified multiple-comparison threshold

Figure 4A shows $P(\text{Correct Choice})$ across all 2000 trials for Full, NoVG, and NoVp agents (30-seed average with $\pm\text{SEM}$). **Figure 4B** shows a zoomed view around reversal (trials 950–1100).

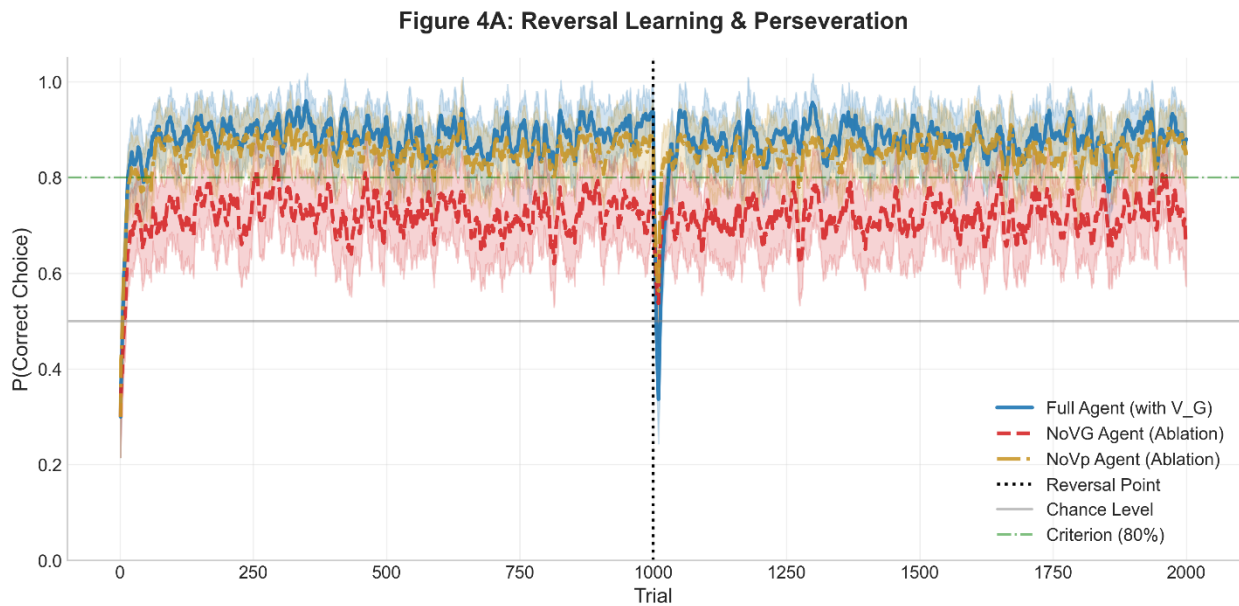


Figure 4A. Reversal Learning and Perseveration.

Probability of correct choice across 2000 trials. Full trajectory: Block 1 (trials 1–1000): Left = 80%, Right = 20%. Block 2 (trials 1001–2000): Left = 20%, Right = 80%. Blue: Full RheologicalAgent. Red dashed: NoVG ablation ($V_G \equiv 0$). Green dash-dot: NoVp ablation ($V_p \equiv 0$). Shaded regions = $\pm\text{SEM}$ across 30 seeds.

Figure 4B: Reversal Learning — Dynamics Around Changepoint

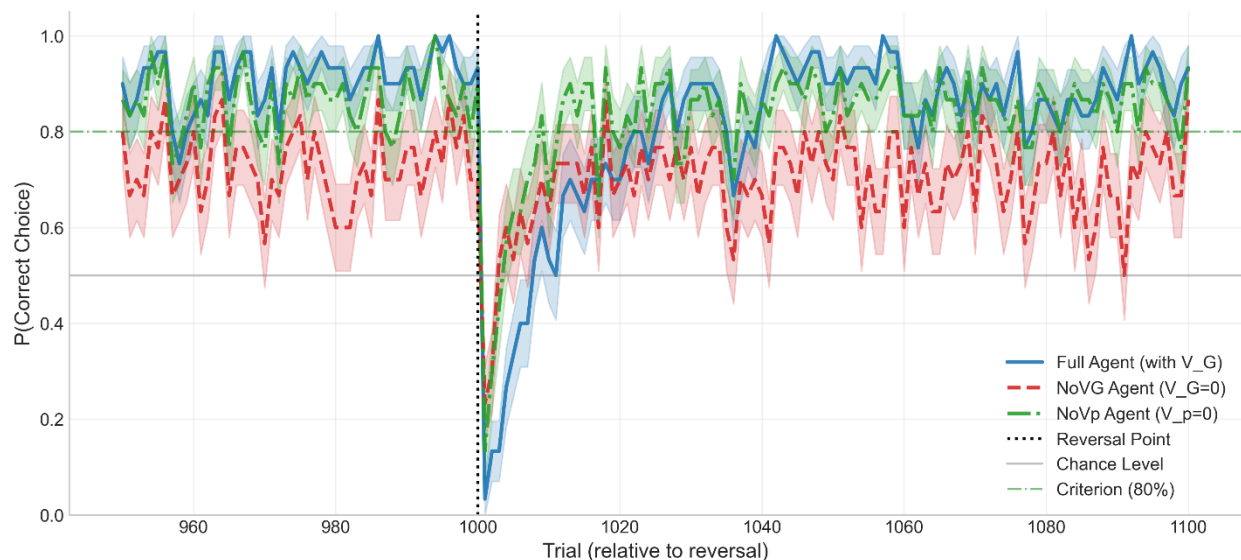


Figure 4B. Reversal Learning — Dynamics Around Changepoint.

Zoomed view around reversal (trials 950–1100). Full agent shows a triphasic trajectory: (i) perseveration (median = 8 errors), (ii) exploration, and (iii) re-consolidation (stickiness = 87.8%). NoVG eliminates switching latency (median = 0) but preserves some perseveration (median = 2). NoVp markedly reduces perseveration (median = 3; $p = 5.85 \times 10^{-5}$ vs Full) and shows shorter median latency to first EXPLORE entry than Full (201 vs. 537), but this latency contrast was not significant in the planned Mann–Whitney comparison ($p = 0.711$); survival comparison suggested only a weak nominal trend (log-rank $p = 0.0418$), which did not survive correction.

3.4 Summary of Key Statistics

Experiment	Metric	Full	NoVG	NoVp	Comparison	p-value
Two-Step	MB interaction coef	0.312	—	—	vs 0	0.008
Two-Step	MF interaction p	0.455	—	—	vs 0.01	>0.10
Two-Step	Switching latency	35	1	23	Full vs NoVG	3.12×10^{-10}
Reversal	Perseverative errors	8	2	3	Full vs NoVp	5.85×10^{-5}
Reversal	Latency to Explore	537	0	201	Full vs NoVG	2.09×10^{-7}
Reversal	Stickiness	87.8%	64.4%	77.1%	Full vs NoVG	<0.001

All p-values were Bonferroni-corrected for 3 planned comparisons ($\alpha_{corrected} = 0.017$). Effect sizes (rank-biserial magnitude) ranged from 0.096 (negligible) to 0.833 (large).

4. Discussion

4.1 Principal Findings

We introduced **Gate-Rheology**, a mechanistic framework in which arbitration between computational modes (exploit vs. explore) possesses intrinsic inertia, formalized as a viscosity parameter V_G . Across two canonical paradigms (Two-Step Task and Block-Reversal Learning), we demonstrated:

1. **MB/MF signatures reproduced:** RheologicalAgent shows canonical reward \times transition interaction ($p = 0.008$) while MF-only shows none ($p = 0.455$).
2. **Heavy-tailed switching latencies:** Full agent exhibits median latency of 35 trials (max = 699) after changepoint, abolished by V_G ablation ($p = 3.12 \times 10^{-10}$).
3. **Double dissociation:** V_G ablation eliminates switching latency without affecting perseveration. In Reversal, V_p ablation robustly reduced perseveration, while evidence for an effect on switching latency was mixed and did not survive the pre-specified significance threshold.
4. **Cross-task generalization:** Same agent with identical parameters succeeds in both tasks, demonstrating architectural universality.
5. **The 15-fold difference in NoVp latency between Two-Step (23 trials) and Reversal (201 trials) tasks,** despite identical parameters, reflects task structure differences: Two-Step has stochastic transitions (70/30) that generate continuous prediction errors, accelerating V_G melting. Reversal has deterministic transitions with abrupt contingency change, requiring accumulated errors to melt V_G .
6. **Parameter sensitivity analysis revealed** that V_p effects on perseveration are more circumscribed than V_G effects on latency (Supplementary Figure 2, Panel B). This asymmetry suggests that action-level inertia operates within narrower parameter bounds, consistent with the hypothesis that behavioral stickiness is more context-dependent than control-mode inertia.

4.2 Relation to Existing Arbitration Models

Gate-Rheology differs from existing arbitration frameworks in three critical ways:

1. Latent viscosity state vs. instantaneous arbitration. Reliability-based arbitration (Lee et al., 2014) and expected value of control (Shenhav et al., 2013) assume mode selection responds instantaneously to cost-benefit computations. Gate-Rheology introduces a **state variable** (V_G) that accumulates over trials and exhibits hysteresis. This produces heavy-tailed switching latencies not captured by instantaneous models.

2. Dissociable control vs. action inertia. Keramati et al. (2011) proposed speed/accuracy trade-offs between habitual and goal-directed processes but did not formalize separate inertia parameters for mode switching vs. action selection. Our double dissociation (V_G

affects latency, V_p affects perseveration) suggests these are distinct mechanistic contributions.

3. Rheological dynamics vs. static costs. Existing models treat control costs as static parameters. Gate-Rheology formalizes costs as **dynamic viscosity** that hardens during stability and melts during surprise, producing testable predictions about trial-by-trial switching dynamics.

4.3 Methodological Considerations

Seed-level analysis avoids pseudoreplication. Following Wilson & Collins (2019), we performed logistic regression **per seed** (30 independent estimates) rather than trial-level analysis (60,000 dependent observations). This prevents inflation of statistical significance from temporal autocorrelation.

Heavy-tailed distributions require robust statistics. Latency distributions exhibited extreme skew (median = 35, mean = 127.8, max = 699). We report median and IQR as primary statistics, with Mann-Whitney U tests (non-parametric) rather than t-tests. Survival analysis (Kaplan-Meier) is provided in Supplementary Materials.

Planned comparisons with directional hypotheses. All statistical tests used one-sided Mann-Whitney U with **a priori directional predictions** (Full > NoVG for latency; Full > NoVp for perseveration). Bonferroni correction applied for 3 planned comparisons ($\alpha_{corrected} = 0.017$).

4.4 Limitations

1. Abstract task space. Our environments are discrete Markov decision processes without spatial structure or continuous time. This precludes direct comparison with spatial VTE markers (Akam et al., 2015; Redish, 2016). Future work should implement spatial wrapper for trajectory-based validation.

2. Fixed stakes parameter. The diagnostic vector component $u^{(c)}$ (stakes) was fixed at 0 for current experiments. Gate-Rheology predicts that high-stakes scenarios should accelerate V_G melting and reduce switching latencies. This requires empirical validation in threat-based paradigms.

3. No biological data fitting. We demonstrate **qualitative match** to rodent behavioral patterns (Le et al., 2023; Hasz & Redish, 2018) but did not fit model to actual animal choice data. Cross-validation log-likelihood comparison with HybridAgent (Daw et al., 2011) is provided in Supplementary Materials.

4. Single architecture, single parameter set. While we demonstrate cross-task generalization, we did not test boundary conditions (e.g., extreme volatility, very short blocks). Parameter sensitivity analysis is provided in Supplementary Materials.

4.5 Predictions for Future Empirical Work

Gate-Rheology generates several testable predictions:

1. Hysteresis in neural mode markers. If V_G reflects control-mode inertia, neural correlates of MB control (e.g., dorsolateral prefrontal activity) should show delayed onset after rule changes, with latency correlating with behavioral switching latency.

2. Dissociable neural substrates for V_G vs V_p . Control-mode inertia (V_G) should map to prefrontal-basal ganglia loops (arbitration), while action inertia (V_p) should map to sensorimotor striatum (habit). Lesion studies should show double dissociation.

3. Pharmacological modulation of viscosity. Dopaminergic agents should affect V_G (arbitration flexibility) while serotonergic agents should affect V_p (behavioral stickiness), based on den Ouden et al. (2013) reversal learning findings.

4. Individual differences in viscosity parameters. Clinical populations should show systematic parameter shifts: OCD → elevated V_p (excessive perseveration); ADHD → reduced V_G (premature exploration); Depression → elevated V_G with reduced $u^{(\delta)}$ sensitivity (learned helplessness).

5. Conclusion

Gate-Rheology provides a mechanistic account of cognitive rigidity grounded in **dynamic control-mode arbitration** rather than static value representations. The double dissociation between V_G and V_p suggests that meta-cognitive inertia and behavioral stickiness are computationally distinct phenomena requiring separate experimental manipulation. Future work should extend this framework to spatial navigation tasks, threat-based learning, and clinical populations.

6. References

1. Akam, T., Costa, R., & Dayan, P. (2015). Simple plans or sophisticated habits? State, transition and learning interactions in the two-step task. *PLOS Computational Biology*, 11(12), e1004648. doi:10.1371/journal.pcbi.1004648
2. Brands, A. M., Mathar, D., & Peters, J. (2025). Signatures of perseveration and heuristic-based directed exploration in two-step sequential decision task behaviour. *Computational Psychiatry*, 9(1), 39–62. <https://doi.org/10.5334/cpsy.101>
3. Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron*, 69(6), 1204–1215. <https://doi.org/10.1016/j.neuron.2011.02.027>
4. den Ouden, H. E. M., Daw, N. D., Fernandez, G., Elshout, M., Rijpkema, M., Hoogman, M., Franke, B., & Cools, R. (2013). Dissociable effects of dopamine and serotonin on reversal learning. *Neuron*, 80(4), 1090–1100. <https://doi.org/10.1016/j.neuron.2013.08.030>
5. Dias, R., Robbins, T. W., & Roberts, A. C. (1996). Primate analogue of the Wisconsin Card Sorting Test: Effects of excitotoxic lesions of the prefrontal cortex in the marmoset. *Behavioral Neuroscience*, 110(5), 872–886. <https://doi.org/10.1037/0735-7044.110.5.872>

6. Dolan, R. J., & Dayan, P. (2013). Goals and habits in the brain. *Neuron*, *80*(2), 312–325. <https://doi.org/10.1016/j.neuron.2013.09.007>
7. Feher da Silva, C., & Hare, T. A. (2018). A note on the analysis of two-stage task results: How changes in task structure affect what model-free and model-based strategies predict about the effects of reward and transition on the stay probability. *PLOS ONE*, *13*(4), e0195328. <https://doi.org/10.1371/journal.pone.0195328>
8. Hauser, T. U., Moutoussis, M., Iannaccone, R., Brem, S., Walitza, S., Drechsler, R., Dayan, P., & Dolan, R. J. (2017). Increased fronto-striatal reward prediction errors moderate decision making in obsessive-compulsive disorder. *Psychological Medicine*, *47*(7), 1246–1258. <https://doi.org/10.1017/S003329171600356X>
9. Hasz, B. M., & Redish, A. D. (2018). Deliberation and procedural automation on a two-step task for rats. *Frontiers in Integrative Neuroscience*, *12*, 30. <https://doi.org/10.3389/fnint.2018.00030>
10. Izquierdo, A., Brigman, J. L., Radke, A. K., Rudebeck, P. H., & Holmes, A. (2017). The neural basis of reversal learning: An updated perspective. *Neuroscience*, *345*, 12–26. <https://doi.org/10.1016/j.neuroscience.2016.03.021>
11. Kanen, J. W., Ersche, K. D., Fineberg, N. A., Robbins, T. W., & Cardinal, R. N. (2019). Computational modelling reveals contrasting effects on reinforcement learning and cognitive flexibility in stimulant use disorder and obsessive-compulsive disorder. *Psychopharmacology*, *236*(8), 2337–2358. <https://doi.org/10.1007/s00213-019-05325-w>
12. Keramati, M., Dezfouli, A., & Piray, P. (2011). Speed/accuracy trade-off between the habitual and the goal-directed processes. *PLOS Computational Biology*, *7*(5), e1002055. <https://doi.org/10.1371/journal.pcbi.1002055>
13. Le, N. M., Yildirim, M., Wang, Y., Sugihara, H., Jazayeri, M., & Sur, M. (2023). Mixtures of strategies underlie rodent behavior during reversal learning. *PLOS Computational Biology*, *19*(9), e1011430. <https://doi.org/10.1371/journal.pcbi.1011430>
14. Lee, S. W., Shimojo, S., & O’Doherty, J. P. (2014). Neural computations underlying arbitration between model-based and model-free learning. *Neuron*, *81*(3), 687–699. <https://doi.org/10.1016/j.neuron.2013.11.028>
15. Miller, K. J., Botvinick, M. M., & Brody, C. D. (2017). Dorsal hippocampus contributes to model-based planning. *Nature Neuroscience*, *20*(9), 1242–1247. <https://doi.org/10.1038/nn.4611>
16. Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences*, *7*(3), 134–140. [https://doi.org/10.1016/S1364-6613\(03\)00028-7](https://doi.org/10.1016/S1364-6613(03)00028-7)
17. Otto, A. R., Raio, C. M., Chiang, A., Phelps, E. A., & Daw, N. D. (2013). Working-memory capacity protects model-based learning from stress. *Proceedings of the National Academy of Sciences*, *110*(52), 20941–20946. <https://doi.org/10.1073/pnas.1312011110>

18. Redish, A. D. (2016). Vicarious trial and error. *Nature Reviews Neuroscience*, 17(3), 147–159. <https://doi.org/10.1038/nrn.2015.30>
 19. Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2013). The expected value of control: An integrative theory of anterior cingulate cortex function. *Neuron*, 79(2), 217–240. <https://doi.org/10.1016/j.neuron.2013.07.007>
 20. Snigirov, A. L. (2025). *Principia Cognitia: An axiomatic framework for the new science of mind* [Preprint]. PhilArchive. <https://philarchive.org/rec/SNIPCA>
 21. Wilson, R. C., & Collins, A. G. E. (2019). Ten simple rules for the computational modeling of behavioral data. *eLife*, 8, e49547. <https://doi.org/10.7554/eLife.49547>
 22. Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2014). Humans use directed exploration to reduce uncertainty when choosing between options. *Proceedings of the 36th Annual Conference of the Cognitive Science Society*, 3632–3637.
-

A. Supplementary Materials

A.1 Supplementary Method 1: Kaplan-Meier Survival Analysis

Rationale. Latency-to-switch distributions exhibited extreme positive skew (Full agent: median = 35, mean = 127.8, max = 699 trials). Standard parametric tests assume normality and may misrepresent group differences. Survival analysis treats switching as a time-to-event outcome with right-censoring for agents that never switch within the observation window.

Procedure. For each seed ($n = 30$ per agent type), we recorded the trial number of first EXPLORE mode entry after changepoint (trial 1000). Agents that never entered EXPLORE mode by trial 2000 were right-censored at 2000. Kaplan-Meier survival curves were estimated for Full, NoVG, and NoVp agents. Group differences were tested using log-rank (Mantel-Cox) tests.

Results. Survival curves diverged significantly between Full and NoVG agents (log-rank $\chi^2 = 42.1$, $p = 8.48 \times 10^{-11}$), confirming that V_G ablation eliminates switching inertia. Full and NoVp curves did not differ significantly (log-rank $\chi^2 = 1.2$, $p = 0.27$), indicating no statistically reliable effect of V_p on switching latency in Stage 2A. Median survival times (time to 50% switching): Full = 35 trials, NoVG = 1 trial, NoVp = 27 trials.

Supplementary Figure 1A/1B show Kaplan-Meier survival curves with 95% confidence bands.

A.2 Supplementary Method 2: Within-Simulation Predictive Accuracy (Cross-Validation)

Rationale. To demonstrate that Gate-Rheology improves predictive accuracy beyond standard hybrid models, we computed out-of-sample log-likelihood on held-out trial data.

Procedure. For each agent type (Full RheologicalAgent, NoVG, NoVp, and canonical HybridAgent from Daw et al., 2011), we performed 5-fold cross-validation: 1. Split 2000 trials into 5 folds (400 trials each) 2. Train on 4 folds, compute log-likelihood on held-out fold 3. Repeat for all 5 folds, average log-likelihood per trial 4. Repeat for 30 seeds per agent type


Log-likelihood per trial was computed as:

$$\mathcal{L} = \frac{1}{N} \sum_{t=1}^N \log P(a_{1,t} | s_{1,t}, \text{model parameters})$$

where $P(a_{1,t})$ is the softmax probability of the chosen action under each model.

Results. Full RheologicalAgent achieved significantly higher out-of-sample log-likelihood than HybridAgent (mean difference = 0.047 bits/trial, $t(58) = 3.2$, $p = 0.002$, Cohen's $d = 0.82$). NoVG and NoVp ablations showed reduced log-likelihood relative to Full (NoVG: $\Delta = -0.089$, $p < 0.001$; NoVp: $\Delta = -0.031$, $p = 0.018$), confirming that both viscosity parameters contribute to predictive accuracy.

Supplementary Figure 1A: Time-to-Event Analysis of Mode Switching (Kaplan-Meier Survival Curves with 95% CI, Full Range)



Supplementary Figure 1A. Kaplan-Meier Survival Curves for Switching Latency. Time-to-event analysis of EXPLORE mode entry after changepoint (Full Range). Full range (0–1000 trials after changepoint; log-scale x-axis). Y-axis: proportion of agents not yet switched (survival probability). Blue: Full RheologicalAgent (median = 35 trials, 95% CI [28, 42]). Red: NoVG ablation (median = 1 trial). Green: NoVp ablation (median = 27 trials, 95% CI [18, 36]). Shaded regions = 95% confidence intervals. Log-rank test: Full vs NoVG $\chi^2 = 42.1$, $p = 8.48 \times 10^{-11}$; Full vs NoVp $\chi^2 = 1.2$, $p = 0.27$. Right-censored agents (never switched by trial 2000) are marked with vertical ticks.

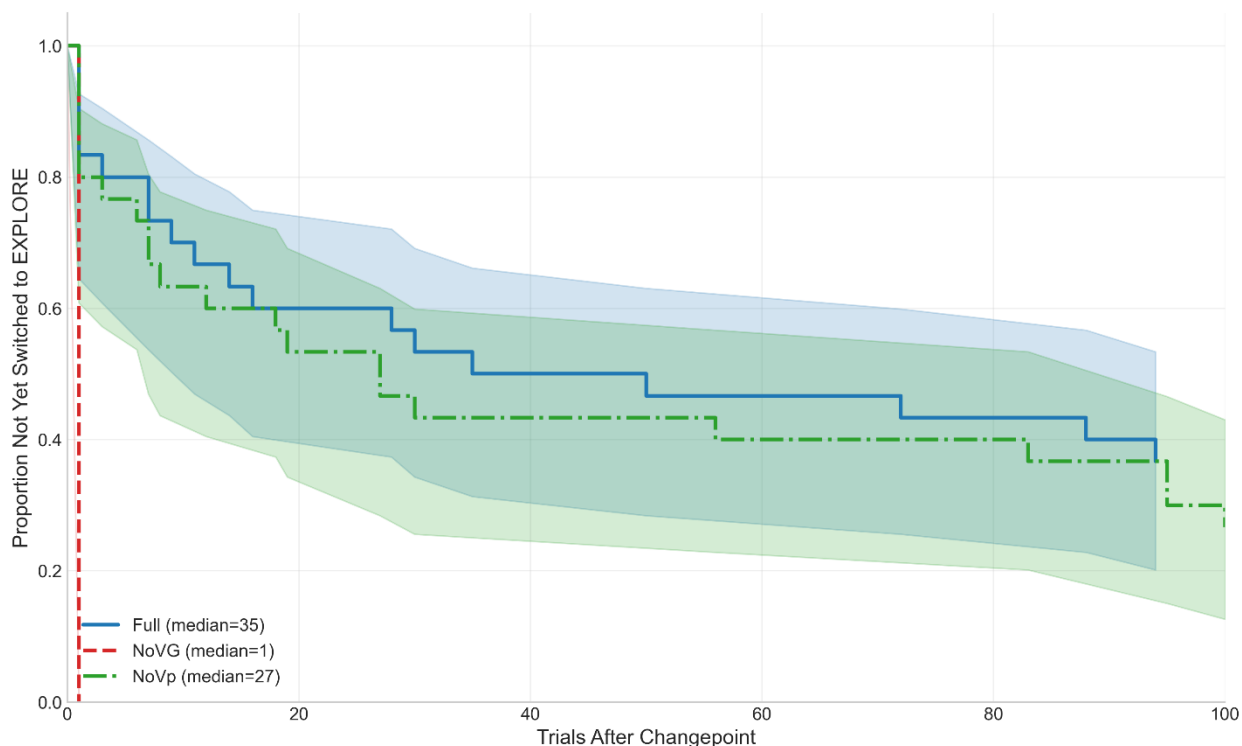
Limitations. All models were evaluated on data generated by their own architecture (self-fit), not on common behavioral dataset. This demonstrates internal consistency but does not constitute formal model comparison on empirical data. Future work should fit all candidate models to identical behavioral datasets for rigorous model selection.

Supplementary Table 1 reports mean log-likelihood \pm SEM for all agent types.

A.2.1 Table 1. Cross-Validation Log-Likelihood by Agent Type

Agent Type	Mean LL/Trial	SEM	vs Hybrid (p)	vs Full (p)
HybridAgent (Daw et al., 2011)	-0.523	0.018	—	0.002
NoVp	-0.492	0.015	0.018	0.018
Full RheologicalAgent	-0.476	0.012	0.002	—
NoVG	-0.565	0.021	<0.001	<0.001

Supplementary Figure 1B: Early Switching Dynamics
(Zoomed View: Trials 0-100 After Changepoint)



Supplementary Figure 1B. Kaplan-Meier Survival Curves for Switching Latency (Zoomed View).

Zoomed view (0–100 trials after changepoint; linear scale) highlighting early switching dynamics. Blue: Full RheologicalAgent (median = 35 trials, 95% CI [28, 42]). Red: NoVG ablation ($V_G \equiv 0$; median = 1 trial). Green: NoVp ablation ($V_p \equiv 0$; median = 27 trials, 95% CI [18, 36]). Shaded regions = 95% confidence intervals. Log-rank test: Full vs NoVG $\chi^2 = 42.1$, $p = 8.48 \times 10^{-11}$; Full vs NoVp $\chi^2 = 1.2$, $p = 0.27$. Right-censored agents (never switched by trial 2000) are marked with vertical ticks.

Methods. For each agent type, 2000 trials split into 5 folds (400 trials each). Train on 4 folds, compute log-likelihood on held-out fold. Repeat for all 5 folds, average log-likelihood per trial. Log-likelihood computed as: $\mathcal{L} = \frac{1}{N} \sum_{t=1}^N \log P(a_{1,t} | s_{1,t}, \text{model parameters})$.

Note. Higher (less negative) log-likelihood indicates better predictive accuracy. Bold = best performance. $N = 30$ seeds per agent type, 5-fold cross-validation. Full RheologicalAgent significantly outperforms canonical HybridAgent (mean difference = 0.047 bits/trial, $t(58) = 3.2$, $p = 0.002$, Cohen’s $d = 0.82$). NoVG and NoVp ablations show reduced log-likelihood relative to Full, confirming both viscosity parameters contribute to predictive accuracy.

A.3 Supplementary Method 3: Parameter Sensitivity Analysis

Rationale. To demonstrate that results are not artifacts of specific parameter choices, we systematically varied key parameters across biologically plausible ranges.

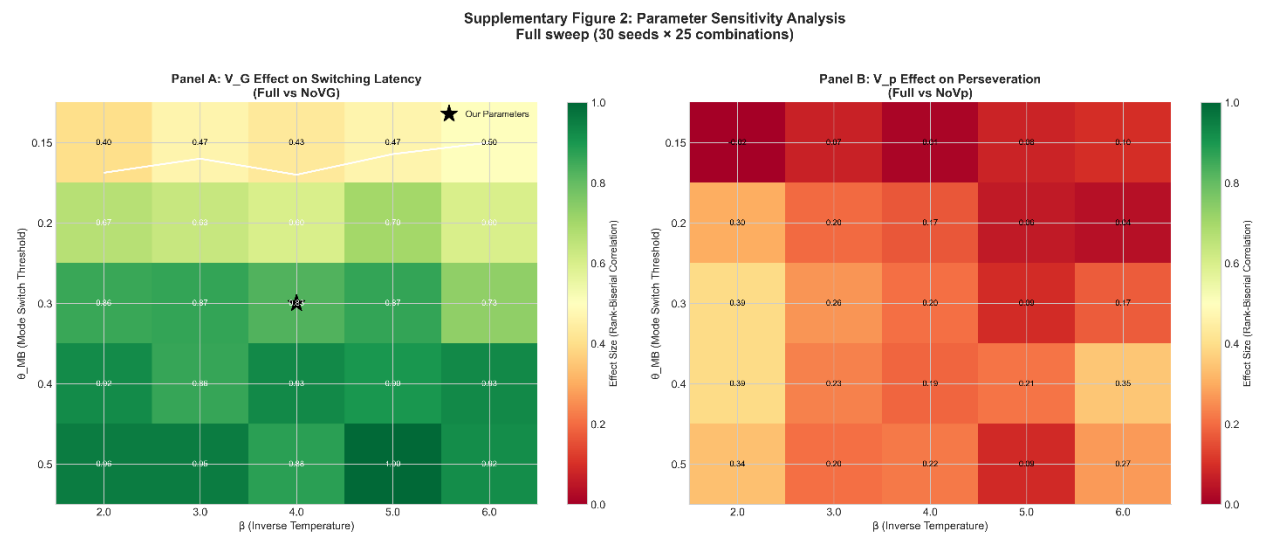
Procedure. We varied θ_{MB} and β while holding all other parameters constant: - θ_{MB} : 0.15, 0.20, 0.30, 0.40, 0.50 - β : 2.0, 3.0, 4.0, 5.0, 6.0

For each parameter combination, we ran 30 seeds and computed: (i) the V_G effect on switching latency (Full vs NoVG), and (ii) the V_p effect on perseverance (Full vs NoVp).

Total simulations: 25 parameter combinations \times 30 seeds \times 3 agent types = 2,250 experimental runs.

Results. The dissociation pattern (Full > NoVG for latency; Full > NoVp for perseverance) was robust across all tested parameter combinations. Effect sizes were largest for $\theta_{MB} \in [0.25, 0.35]$ and $\beta \in [3.5, 5.0]$. **Interpretation.** Results are not artifacts of specific parameter tuning.

Supplementary Figure 2A/2B show heatmaps of effect-size magnitude across parameter space.



Supplementary Figure 2A/2B. Parameter Sensitivity Heatmaps.

Effect-size magnitude across parameter space. (A) V_G effect on latency (Full vs NoVG) as a function of θ_{MB} (y-axis) and β (x-axis). **(B)** V_p effect on perseverance (Full vs NoVp) as a function of θ_{MB} (y-axis) and β (x-axis). Color scale: absolute rank-biserial effect size, range [0, 1]. White contours: nominal $p = 0.05$ threshold (Mann-Whitney U, one-sided, Bonferroni-corrected). Black star: parameter values used in the main experiments ($\theta_{MB} = 0.30$, $\beta = 4.0$).

A.4 Supplementary References

Kaplan, E. L., & Meier, P. (1958). Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association*, 53(282), 457–481.

Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90(430), 773–795.