

A Unified Model of Learning and Habituation under Mechanical Stimulation in the protozoan *Stentor coeruleus*

Hamza Mohamed

hamzah.mohamed2009@gmail.com

STEM High School for Boys - 6th of October <https://orcid.org/0009-0000-5910-1736>

Research Article

Keywords: *Stentor coeruleus*, associative learning, habituation, mechanical stimulation, computational modeling

Posted Date: October 16th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10346520/v3>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

A Unified Model of Learning and Habituation under Mechanical Stimulation in the protozoan *Stentor coeruleus*

Hamza Mohamed

Correspondence: hamzah.mohamed2009@gmail.com

Abstract

Repeated mechanical stimulation of *Stentor coeruleus* can produce an early enhancement of contraction probability, habituation, recovery, and altered responses after rest. These observations have usually been studied in different protocols and laboratories. In this study, I developed a single model combining response enhancement, habituation, and stimulus history. Each simulated cell updates its states from its own binary contraction events; the population fraction is predicted by averaging across their histories. The model was fitted to a limited subset of the weak–strong pairing experiments of Doan and colleagues, the repeated-bout experiments of Ramdas and colleagues, and recovery curves digitized from the experiments of Rajan and Marshall. Other complete protocols were reserved for evaluation. Mean held-out protocol RMSE in contraction probability was 0.041 for the pairing experiments, 0.057 for the repeated-bout experiments, and 0.240 for the digitized recovery experiments. Mechanistic parameter ablations increased errors in the domains most affected by their corresponding states. The large error in the digitized Rajan and Marshall recovery data, especially after long training at the lower stimulation frequency, shows that a shared stimulus scale does not account for between-cohort and between-laboratory differences. The model therefore connects several forms of learning-like behavior in one quantitative framework while placing a clear limit on its present generality. I propose a reporting framework for mechanical stimulation experiments that can separate stimulus calibration, initial state, and recovery dynamics.

Keywords: *Stentor coeruleus*; associative learning; habituation; mechanical stimulation; recovery; computational modeling


1 Introduction

Associative learning in single cells and habituation are often discussed separately. In one case, the response to a weak stimulus increases after it is paired with a stronger stimulus; in the other, repeated stimulation reduces the response. Yet both changes occur in the same cell, and the response observed on a particular trial depends on the stimulation that came before it. The weak–strong pairing experiments in *Stentor coeruleus* provide one example: contraction to a weak mechanical tap can rise early in training and decline later [1]. A model that only enhances responsiveness misses the later decline, while a model that only habituates cannot explain the early rise.

In the first version of this work, I proposed that a trace left by a weak stimulus can interact with a later strong stimulus to produce a transient enhancement in response probability [2]. That model explained several features of the pairing trajectories but was evaluated principally against one experimental series. Other experiments now make a broader test possible. Single-cell measurements show that apparently gradual population habituation can hide abrupt and variable changes in individual cells [3]. Experiments with different mechanical force levels and training histories describe cross-habituation, recovery, and the effects of stimulation frequency [4]. Repeated training bouts separated by one to five hours provide a systematic view of response recovery and faster re-habituation [6]. These observations should constrain a model that is intended to describe learning under mechanical stimulation rather than one protocol alone.

The studies also present a calibration problem. A weak stimulus in one laboratory is not automatically equivalent to a numbered force level in another. The initial contraction probability can help place stimuli on a common response scale, but it also depends on a cell population’s prior history and response threshold. Accordingly, the model below shares dynamical parameters across datasets while allowing specified stimulus-scale and starting-state calibrations. Its latent states describe effective behavior; they should not be identified with a particular channel, receptor, or measured membrane potential without an independent experiment. Earlier electrophysiology and recent molecular work motivate such tests but do not establish that these specific state equations are cellular mechanisms [5, 8, 9].

I fitted a limited set of complete protocols from each data source and evaluated the remaining protocols without changing the shared parameters. Figure 1 summarizes the experiments and the model.



Measured: contraction at each stimulus. Inferred model states: T , B , L_w , L_s , and Q .

Figure 1: Experimental protocols and shared model. Arrows show the order of delivered stimuli within each study; contraction responses are measured at those stimuli. The model states T , B , L_w , L_s , and Q are inferred rather than directly measured.

2 Model

2.1 Data sources and protocol separation

The Doan data comprise contraction response counts across weak–strong (WS), weak–weak (WW), weak-only, strong-only, and timing conditions [1]. Six complete protocols were fitted: WS and WW at a 1-s within-pair interval and 45-s between-pair interval; weak-only at 1/59; strong-only at 1/45; WS at 10/45; and WS at 1/59. The other eight complete protocols were held out. This split includes additional acquisition and control information in fitting but retains timing changes for evaluation. It is the split used by the V3 export analyzed here, rather than the narrower four-condition V2 panel.

The Ramdas data contain individual binary contraction records from two series of 60 stimuli, separated by a rest interval [6]. The publicly available collated data cover training intervals of 1, 2, and 3 minutes and rest intervals of 1, 2, 3, and 5 hours. Four protocols, crossing 1- and 3-minute training intervals with 1- and 5-hour rests, were fitted; the remaining eight were held out. The model was compared with aggregate contraction fractions at each stimulus, retaining the per-stimulus denominators for uncertainty weighting. Each complete second bout remained in the same split as its first bout.

The Rajan and Marshall behavioral data were taken from published figures because individual-cell records were unavailable for this analysis [4]. Forty-nine approximate recovery-probe fractions were digitized from seven protocols (seven times each): the 20-minute and 12-hour rapid-training comparison in Fig. 5C, the force-level comparison in Fig. 5E, and the frequency comparisons in Fig. 7E/F. The 20-minute rapid-training protocol and both 12-hour force-level protocols were fitted; four other complete recovery protocols were held out. Six early, approximate Fig. 4F alternating-force points provided an additional relative force calibration. Three near-ceiling high-force points were treated as lower bounds on contraction probability, rather than precise observations of 1.00. These figure readings do not provide exact cell-level response counts, replicate structure, or pretraining trajectories and cannot support an exact binomial likelihood. Table 1 records the split and the type of measurement available.

Table 1: Data used for fitting. Entire other protocols were held out: eight Doan, eight Ramdas, and four Rajan recovery conditions. Rajan fitting also used six approximate Fig. 4F force readings.

Source	Fitted	Observations
Doan	6	Response counts; pairing and controls
Ramdas	4	Cell responses; training and second bout
Rajan	3 + 4F	Digitized recovery fractions and force readings

2.2 Model of stimulus response and history

The V3 model advances a separate history state for each simulated cell at every physical stimulus and averages their pre-event contraction probabilities to predict the population response. Intrinsic dynamical parameters are shared across simulated cells; differences in their realized response histories produce state heterogeneity. This does not estimate a distribution of intrinsic cellular parameters. Its parameters are shared across the three studies, except for explicit study-level stimulus scaling and initial-response calibrations. For a weak or strong event, the effective response drives are

$$R_w(t) = R_{w0} \left[1 - A_{\max,w} (1 - e^{-L_w(t)}) \right], \quad (1)$$

$$R_s(t) = R_{s0} \left[1 - A_{\max,s} (1 - e^{-L_s(t)}) \right]. \quad (2)$$

$$P_{\text{contraction}}(t) = \frac{1}{1 + \exp\{-\gamma[V_0 + B(t) + R(t) - 1]\}}. \quad (3)$$

Here $R(t)$ selects R_w or R_s for the actual stimulus. R_{w0} and R_{s0} represent the initial mechanosensory contributions to a common response drive, before attenuation by training. The quantities L_w and L_s accumulate effective weak- and strong-associated habituation, while $A_{\max,w}$ and $A_{\max,s}$ set the maximum fraction of each drive that can be attenuated. These are functional response components; their existence does not establish separate anatomical receptors. The state B adds to response drive, and V_0 specifies background excitability for the experimental population. The threshold is set to 1 by convention, and γ controls how sharply response probability changes as the net drive crosses it. In a population, that sigmoid may summarize biological variability and stochastic responses as well as intracellular thresholding. None of V_0 , R , or B has been measured as a voltage in these experiments.

Between events separated by Δt , the five states B, L_w, L_s, T, Q each decay exponentially, with time constants $\tau_B, \tau_{Lw}, \tau_{Ls}, \tau_P, \tau_Q$, respectively. In the following event equations, all quantities on the right are evaluated after this decay and before the current event's update. For each cell, write $J = R/R_{s0}$ and $H = 1 + q_{\text{gain}}Q$, and draw its realized response $C \sim \text{Bernoulli}(P_{\text{contraction}})$. Here $C = 1$ denotes a contraction, and $C = 0$ its absence; it is not the delivered stimulus label. Stimulus type selects the habituation rates and trace reset; contraction in either stimulus class can recruit enhancement from preceding priming:

$$B^+ = B + B_{\text{update}} T J C (1 - B) \quad (4)$$

$$L_w^+ = L_w + J C H \begin{cases} k_{ww}, & \text{weak stimulus,} \\ k_{sw}, & \text{strong stimulus,} \end{cases} \quad (5)$$

$$L_s^+ = L_s + J C H \begin{cases} k_{ws}, & \text{weak stimulus,} \\ k_{ss}, & \text{strong stimulus,} \end{cases} \quad (6)$$

$$T^+ = \begin{cases} R_w/R_{w0}, & \text{weak stimulus,} \\ 0, & \text{strong stimulus,} \end{cases} \quad (7)$$

$$Q^+ = \min \left\{ 1, Q + \frac{J C (1 - Q)}{60} \right\}. \quad (8)$$

The contraction probability is computed *before* applying the event updates. Its realized response then changes future states. Population contraction fraction is predicted as the mean of these pre-update probabilities across simulated cells. No cell receives that ensemble mean as an input. Replacing every realized C by the population mean would produce a different closure of the nonlinear dynamics, rather than this event model.

2.3 Biological interpretation of the states and update factors

In V1, I described the weak stimulus as leaving a possible subthreshold depolarization [2]. Here, I propose that the relevant trace can instead be a temporary change in channel availability, gating, or regulation. This allows a weak stimulus to leave a state that affects a subsequent stimulus without requiring the membrane to remain depolarized throughout the pairing interval. The model separates this temporary priming state, T , from the longer-lived enhancement state, B . Their biological substrates have not been measured in the pairing experiments, so the interpretation below is a hypothesis that the behavioral fit can constrain and motivate experiments to test.

The state T represents priming of mechanosensitive channels or their associated regulatory machinery by a weak tap. Possible implementations include a reversible conformational change, altered voltage sensitivity, or a modification of a channel-associated protein that changes how the channel responds to a later stimulus. There is a physiological basis for considering channel state rather than voltage alone. In *Paramecium*, voltage-clamp experiments identified a slow, voltage-dependent component of calcium-channel inactivation and recovery on a timescale of tens of seconds [10]. That observation demonstrates history-dependent channel availability in another ciliate; it does not demonstrate facilitatory priming in *Stentor coeruleus*. In *Stentor coeruleus* itself, Wood’s recordings linked habituation to changes in mechanoreceptor-channel voltage dependence [14]. Recent experiments also implicate calcium signaling and protein phosphorylation in habituation [5]. These findings motivate channel gating and channel regulation as candidate substrates, while leaving the proposed direction and pairing dependence of T to be tested.

The implemented rule sets T to R_w/R_{w0} after a weak tap, even if that tap does not cause contraction. This reflects the proposal that mechanosensory activation can prime channels below the contraction threshold. A habituated weak drive leaves less priming. Between taps, T decays with τ_p , giving the pairing window, and a strong tap resets it to zero after its possible contribution to B has been calculated. Resetting is a modeling assumption of consumption or disruption of the primed state; it is not a measured all-or-none channel transition. In these equations, T gates the update of B and does not separately multiply the current stimulus drive or raise the current strong response probability.

The state B represents a longer-lived response-facilitating configuration recruited when a subsequent stimulus produces contraction while weak-induced priming remains. One candidate is that activation of primed channel complexes changes their gating, deactivation, or association with nearby regulatory proteins, leaving a larger effective inward drive after the stimulus. Another is that the accompanying contraction reorganizes a local membrane cortex structure that stabilizes the altered channel configuration. In that interpretation, deformation can trigger the change while a semi-stable organization maintains it after the contraction ends. B summarizes its contribution to later residual excitability; it is not a measurement of membrane potential or a count of modified channels.

The cortical interpretation is motivated by a separate observation in ciliates: experimentally altered cortical organization can be maintained and inherited in *Paramecium* [11]. This establishes that cellular organization can carry persistent information. It does not establish that the same process stores paired mechanical stimulation in *Stentor coeruleus*, or that cortical inheritance and the much shorter-lived B share a molecular mechanism. Likewise, prolonged depolarization can produce persistent changes in *Paramecium* excitability [12], but that is a different preparation and stimulus history. I use these observations as biological precedents for state-dependent excitability and organizational memory, rather than as direct evidence for the proposed B mechanism.

The update $B_{\text{update}} T J C (1 - B)$ requires preceding weak-induced priming and a realized contraction at the current event, whether the delivered stimulus is weak or strong. J scales the effect by the current mechanosen-

sory drive, B_{update} controls its magnitude, and $(1 - B)$ limits accumulation. The state relaxes with τ_B . An unprimed event produces no B update. WW can generate B while habituation still dominates its net response trajectory. This all-stimulus recruitment is a model hypothesis, not an established physiological result. One test would compare mechanoreceptor-current deactivation or background inward current after WS pairing, strong stimulation alone, and WW stimulation. Measuring changes after both WS and WW, and their relation to contraction histories, would test this interpretation more directly than fitted states alone.

The loads L_w and L_s represent reductions in effective mechanosensory drive associated with previous weak and strong stimuli. They can summarize channel gating changes, receptor inactivation, or downstream regulation that reduces the input reaching the contraction threshold. Wood’s electrophysiology supports channel modification, whereas the receptor-inactivation model of Rajan and Marshall provides a different mechanistic account [4, 14]. These alternatives cannot be distinguished by calling L a receptor state. Nor do two functional components establish anatomically distinct weak and strong receptors. The self-update rates k_{ww} and k_{ss} describe the effects of weak and strong events on their corresponding loads; k_{sw} transfers strong stimulation to the weak-associated load, and k_{ws} would transfer weak stimulation to the strong-associated load. The fitted version sets $k_{ws} = 0$ as an explicit negligible transfer assumption. Strong stimuli remaining effective after weak training do not, by themselves, prove that all underlying weak-to-strong modification is absent.

The attenuation fractions $A_{\text{max},w}$ and $A_{\text{max},s}$ determine how much of each initial response drive is habituable. The exponential transformation of L keeps the remaining drive nonnegative when $0 \leq A_{\text{max}} \leq 1$. The constants τ_{Lw} and τ_{Ls} govern effective recovery between events. They summarize recovery within this model; they are not established universal channel or receptor lifetimes. The multiple timescales in the repeated-bout experiments and the errors on digitized recovery curves make that distinction necessary [4, 6].

The factor $J = R/R_{s0}$ is the current effective mechanosensory drive normalized by the unhabituated strong drive for that study’s stimulus mapping. It allows a stimulus whose effective drive has fallen during training to produce a smaller state increment. J is dimensionless, not a measured force, current, calcium concentration, or receptor count. The variable C is the cell’s realized binary contraction at the present event. Contraction-associated activation is a proxy for intracellular processes accompanying the response, such as calcium entry, rather than a claim that the visible shortening itself is the sole cause of habituation. Response-dependent electrophysiological habituation motivates this choice [9]. C multiplies the updates of B , L_w , L_s , and Q , while weak-stimulus priming of T remains possible without contraction. P is used to generate a cell’s response and to predict the observed population fraction; it is not supplied to a cell as an update factor. A simulated event history can differ among cells even when their intrinsic parameters are shared.

The state Q addresses a different feature of the data: cells can recover responsiveness during rest and nevertheless habituate faster during a second bout. Ramdas and colleagues interpreted this dissociation in terms of a persistent second-order memory affecting subsequent learning [6]. Q is this model’s proposed implementation, rather than a molecular variable measured or specifically proposed by that study. Repeated response-associated stimulation increases Q , while its slower relaxation permits training history to remain after the loads have partly recovered. $H = 1 + q_{\text{gain}}Q$ translates that history into larger future habituation increments. H is an algebraic multiplier, not an additional independent state. The denominator 60 fixes the normalization of the Q update relative to the experimental bout length; it is not a biological constant. Setting $q_{\text{gain}} = 0$ removes its influence on the response trajectory.

Possible substrates of Q include persistent kinase/phosphatase regulation, altered abundance or turnover of receptor-regulatory proteins, a maintained transcriptional state, or RNA-mediated and chromatin-associated regulation. These candidates operate at different levels and should not be treated as interchangeable. The molecular experiments in *Stentor coeruleus* provide a reason to examine protein regulation: translation inhibition accelerated habituation and prolonged retention, and recovery was associated with protein synthesis [5]. This does not support simply importing the familiar claim that forming a long-term memory must require new protein synthesis. RNA-mediated epigenetic inheritance is experimentally established in *Paramecium* during developmental genome rearrangement [13]; it provides a ciliate precedent for persistent regulatory information, but no direct evidence that mechanical potentiation uses this process. Neither RNA expression changes, RNA

interference, nor retention through cell division would alone identify RNA or chromatin as the storage substrate of Q .

The role assigned to Q suggests a discriminating experiment. After training, responsiveness could be re-stored while leaving prior training history otherwise intact, followed by a second bout. Faster re-habitation despite that restoration would separate a history-dependent change in learning rate from incomplete response recovery [6]. Combining this behavioral test with time-resolved RNA, protein, phosphorylation, and chromatin measurements could then ask which changes persist with potentiation and which disappear with it. Selective adjustment of those candidates would be needed to assign a causal storage role. The current model provides quantitative hypotheses for such measurements; its behavioral parameter estimates alone do not establish their molecular identity.

The weak response-drive unit was fixed by $R_{w0} = 1$. For Doan protocols, the weak and strong drives are R_{w0} and R_{s0} . Both Ramdas drives are multiplied by a fitted `ramdas_scale`. In the Rajan protocols, the lower-force drive is `rajan_scale` R_{w0} and the higher-force drive is `rajan_ratio3` times the lower-force drive. Thus the force mapping is estimated from behavior and Fig. 4F; the model does not assert that nominal tap labels measure the same physical impulse across laboratories. A high initial contraction probability can result from a larger stimulus, a different baseline, or both, so these scales have to be interpreted together with V_0 .

For each Doan or Ramdas cohort, V_0 was calculated from its first observed response probability p_0 and the first-stimulus drive, $V_0 = 1 - R_{\text{first}} + \text{logit}(p_0)/\gamma$. The first observed tap was used for this calibration and omitted from the trajectory error. The Rajan plots did not provide a comparable pretraining response count for every recovery cohort. Instead, fitted starting probabilities were shared within specified figure groups: one for the Fig. 5C/7E/7F protocols, one for Fig. 5E, and one for the separate Fig. 4F calibration. In each primary protocol simulation the dynamic states began at zero; these V_0 calibrations are assumptions about distinct experimental populations, not observed initial states. The recovery probes themselves were simulated sequentially and can affect later probe responses.

2.4 Parameter fitting and evaluation

The response trajectory objective used squared errors divided by an observation-scale uncertainty with a minimum standard deviation of 0.045 for Doan, 0.05 for Ramdas, and 0.08 for digitized Rajan recovery fractions. Each fitted Doan WS schedule, together with the matched WW1/45 schedule, gave equal weight to taps 2–10 and all remaining observed taps; the first tap was excluded after baseline calibration. This targets both acquisition and later habituation. The weighting was selected during model development and is reported as an analysis choice. Other fitted Doan schedules used their ordinary within-schedule average error, and the two Ramdas bouts contributed equally within each protocol. The fitted family scores were then balanced across Doan, Ramdas, and Rajan, with a lower weight of 0.15 for the six approximate Fig. 4F calibration readings. Small soft penalties discouraged $\tau_P > \tau_B$ and disproportionately strong $k_{sw}R_{s0}$ relative to k_{ww} . These bounds, weights, and penalties are modeling decisions, not independent measurements or a unified observation likelihood.

Fitting used bounded multistart Powell optimization with sequential passes through pairing, controls, force calibration, recovery, and repeated-bout history, followed by differential-evolution reconciliation and local refinement. A persistent candidate bank retained feasible improvements to the joint fitting objective across stages. Earlier V3 solutions served as starting candidates only; all free parameters were re-estimated from the V3 fitting data. The final candidate was checked against those starts under the same objective and simulation settings. These checks do not establish a global optimum.

The search used 128 simulated cells initially, then rescored candidates and refined with 512 cells. A fixed seed within the search provided reproducible event draws, but the resulting objective remains nonsmooth. Evaluation used 2048 cells and an independent seed. Ensemble-size and independent-seed checks quantified simulation variability; these are numerical diagnostics, not confidence intervals for biological parameters.

Held-out schedules did not enter parameter estimation. They had, however, been inspected during model development, so their reported errors are developmental checks of schedules withheld from fitting, rather than

an untouched confirmatory test. The revised weighting was applied only to fitting schedules. It did not use held-out peaks to select weights or introduce a new floor on strong-response probability. Reported root-mean-square errors (RMSEs) are unweighted errors on observed contraction fractions, first calculated within each protocol and then averaged over protocols; they are not the weighted fitting scores. Consequently, a protocol with many more measurements does not automatically dominate the reported family mean.

Three fixed-parameter ablations copied the final fit, set $B_{\text{update}} = 0$, set both cross-habituation rates to zero, or set $q_{\text{gain}} = 0$, and then recomputed trajectories. No other parameters were refitted. These calculations measure the conditional contribution of a mechanism at this particular fitted parameter set, rather than the best possible performance of an alternative reduced model. A separate conditional evaluation recalibrated the Doan and Ramdas starting state with their first three taps and scored only subsequent taps. Because this uses more observations from each held-out cohort, it is not the same prospective prediction as the primary first-tap calibration.

3 Results

3.1 Parameter values and response scales

Table 2 gives the dynamical parameters of the joint V3 fit. These numbers were estimated from the V3 fitting protocols, rather than carried over from V2. The response-drive scale was set by $R_{w0} = 1$ and the threshold in Eq. 3 was set to 1, so their absolute values have no independent voltage or force interpretation. The fitted $R_{s0} = 1.050$ reached the lower allowed value, even though the two stimuli can still produce different behavior because of their initial conditions, habituation, and the sigmoid response function. The steepness $\gamma = 20.19$ makes changes near threshold consequential.

Table 2: Dynamical parameters in the selected realized-event V3 fit. Times are in seconds; the full-precision checkpoint is included with the source. Boundary proximity does not establish a biological lifetime.

Parameter	Value	Parameter	Value
R_{w0}	1.000 (fixed)	k_{ws}	0 (fixed)
R_{s0}	1.050 (lower bound)	k_{ss}	0.03424
γ	20.19	$A_{\text{max},w}$	0.99882 (near upper bound)
B_{update}	0.03663	$A_{\text{max},s}$	0.34619
τ_P	54.45	τ_{Lw}	4379
τ_B	262.65	τ_{Ls}	171692 (near upper bound)
k_{ww}	0.01804	q_{gain}	6.000 (upper bound)
k_{sw}	0.008204	τ_Q	159498

The trace time constant was $\tau_P = 54.45$ s, while the enhancement state decayed with $\tau_B = 262.65$ s. In the absence of intervening events, the trace retains approximately 98%, 83%, and 69% of its value after 1, 10, and 20 s, respectively. The enhancement state retains 84% after 45 s but only 11% after 590 s. These percentages describe the fitted equations; they are not direct measurements of a persisting membrane potential. The Q history state decays much more slowly: with $\tau_Q = 159498$ s, its between-event retention is approximately 98% after one hour and 89% after five hours. In contrast, L_w with $\tau_{Lw} = 4379$ s retains approximately 44% and 2% over those same gaps. This combination can let contraction recover while prior training continues to affect the rate of habituation in the second bout.

3.2 The WS and WW trajectories separate early, then both habituate

The 1/45 WS and WW conditions formed the central pairing contrast. In the observed data, weak-response probability under WS reached 0.497 at trial 3, compared with 0.374 under WW, a difference of 0.123. The V3 curves also placed the WS peak at trial 3, but predicted 0.474 under WS and 0.416 under WW, a difference of

only 0.058. Thus the model reproduces the timing of the peak while underestimating the maximum separation between protocols. It also predicts an early WS advantage at trial 2, where the observed WS and WW fractions were nearly equal (0.396 and 0.405). Figure 2 separates the trajectories and their difference so that this limitation remains visible.

Later, both response fractions declined substantially. At trial 35, the observed WS and WW probabilities were 0.006 and 0.047; the model gave 0.019 and 0.041. The fitted WS and WW protocol RMSEs were 1.63 and 2.51 percentage points, respectively. These whole-trajectory errors summarize 59 scored trials each, so they should be interpreted alongside the smaller predicted peak contrast. The fitting objective balanced early and later segments, yet the peak difference remained underestimated. The weak-only and strong-only fitting controls had RMSEs of 4.56 and 4.09 points. They constrain the self-habituation part of the model, while the WS/WW difference constrains enhancement and strong-to-weak transfer jointly.




Figure 2: The fitted WS and WW trajectories and their contrast. The gray band is a 95% run-level bootstrap interval for the observed contrast (10 WS runs, 14 WW runs); it reflects between-run variation rather than independently resampled taps. The data and model reach the WS peak at trial 3, but the model produces a smaller WS–WW difference.

3.3 Timing and control protocols test the pairing mechanism outside fitting

Across all six fitted Doan protocols, the mean per-protocol RMSE was 2.92 percentage points; across eight held-out complete protocols it was 4.11 points (Table 3). The model predicted WS10/59 and WS20/45 with errors of 2.60 and 2.73 points even though these whole schedules were not used in fitting. It also predicted the much longer 590-s between-pair interval with an error of 4.06 points over the 19 scored observations. Because this condition has fewer observed trials and a different gap, that value should be read as a separate protocol-level result, not pooled with the 59-point curves.

Table 3: Mean per-protocol RMSE for the saved full model. Values are percentage points of contraction probability; fit and held-out protocols are different complete experiments. Rajan values use approximate digitized fractions.

Source	Fit RMSE (protocols)	Held-out RMSE (protocols)
Doan	2.92 (6)	4.11 (8)
Ramdas	4.51 (4)	5.68 (8)
Rajan recovery	8.41 (3)	24.04 (4)

The early part of some held-out curves remains weaker than their aggregate score suggests. In WS20/45, the observed peak was 0.435 at trial 3, while the largest predicted value among the scored trials was 0.367 at

trial 2; the model was on average 3.8 percentage points below observations over trials 2–10. The WW1/59 and weak-only20/45 holdouts each had errors near 5 points. The two multistimulus controls, WSWWW and WSSSW, were treated as secondary checks: the former had an error of 3.58 points and the latter 6.58 points, with a particularly large early underprediction in WSSSW. The remaining single-strong schedule had an error of 3.70 points. Figure 3 shows these held-out trajectories separately, so an overall mean does not hide the less successful controls.




Figure 3: Whole-protocol tests of the fitted Doan dynamics. The additional WS schedules generalize more accurately than some weak-only and multistimulus controls. Their starting probabilities were calibrated from the first observed event in each separate population.

3.4 Two training bouts constrain response recovery and a slow history effect

In the Ramdas protocols, the fitted set crossed 1- and 3-minute stimulus intervals with 1- and 5-hour rests. Its mean protocol RMSE was 4.51 points; the eight unfit combinations averaged 5.68 points. The model predicts both the decline over the first 60 stimuli and the response after the rest before the second series. For example, after 60 one-minute stimuli and a five-hour rest, the observed first response of the second bout was 0.703 and the prediction was 0.764. With the same stimulation interval and a one-hour rest, those values were 0.714 and 0.481. The latter discrepancy shows that even a fitted condition can have a large error at an important transition despite a moderate RMSE over 119 scored taps.

The held-out 1-minute/2-hour and 3-minute/2-hour schedules had errors of 2.76 and 3.29 points. A notable failure was the 2-minute/5-hour schedule, with an error of 9.53 points: at the end of its first bout the observed fraction was 0.067 while the prediction was 0.183, and at the beginning of its second bout they were 0.674 and 0.781. Other two-minute schedules were predicted more accurately, so this is not simply a uniform failure at the middle stimulation frequency. Figure 4 separates the two bouts in each panel and labels the rest duration; no line implies observations during the intervening rest.




Figure 4: Training, rest, and re-habitation in the Ramdas data. The model follows several unseen rest intervals but misses some individual bout transitions and the 2-minute/5-hour schedule. Observed markers are cohort response fractions; neighboring taps often involve the same cells.

3.5 Behavioral force calibration between laboratories

The common response-drive scale requires an explicit mapping between the two laboratories' nominal stimulus levels. The fitted Ramdas multiplier was 0.792 relative to the Doan weak-drive convention. For the Rajan experiments, the force-level-4 drive was $0.252R_{w0}$ and the force-level-3 drive was 1.275 times that value. These are relative *model input drives*; the ratio is neither a force ratio in newtons nor a statement that a Rajan force level is physically identical to a Doan tap. The Rajan multiplier was close to its lower allowed bound, which makes the mapping particularly sensitive to the chosen response scale.

Table 4: Study-level calibration values in the V3 fit. These values define relative input drives and assumed starting probabilities; they are not physical force calibrations.

Parameter	Value	Role
ramdas_scale	0.792	Multiplier for Ramdas response drives
rajan_scale	0.252 (near lower bound)	Rajan level 4 relative to Doan weak drive
rajan_ratio3	1.275	Level 3 / level 4 drive ratio
rajan_p0	0.995 (upper bound)	Baseline for Fig. 5C/7E/7F groups
rajan_p0_5E	0.880	Baseline for Fig. 5E group
rajan_p0_4F	0.876	Separate Fig. 4F calibration baseline

The separate early alternating-force calibration from Fig. 4F had six approximate readings. Under the mixed error rule that treats three near-ceiling level-3 markers as $P \geq 0.97$, their RMSE was 3.49 points. These readings

help distinguish the relative ordering of level 3 and level 4, but a ceiling response cannot determine how far above threshold the stronger level is. In addition, a population with a different V_0 could reproduce a similar initial probability with a different drive. Figure 5 shows the approximate source readings and predictions on the source panel’s nominal force-level scale, without implying an absolute force conversion.




Figure 5: Relative force calibration using Rajan and Marshall’s Fig. 4F. The data are approximate readings of published symbols and do not measure a physical force conversion between laboratories.

3.6 Recovery after force and frequency changes is less well predicted

The three fitted Rajan recovery protocols averaged 8.41 points of error, but the four held-out protocols averaged 24.04 points (Table 3). The fitted 12-hour level-3 training condition began near 0.03 observed contraction probability at the first 2-minute recovery probe and reached 0.43 at 90 minutes; the model predicted 0.086 and 0.429. With level-4 training, the observed values were 0.17 and 0.65, while the model predicted 0.324 and 0.678. The final recovery values are close, but the model rises too early in the level-4 condition. In the fitted 20-minute high-frequency protocol, the observed probability increased from approximately 0.48 at 1.5 minutes to 0.80 at 90 minutes; the prediction increased from 0.63 to 0.88. These comparisons convey information that the mean error alone does not.

The held-out 12-hour slow-frequency Fig. 7E protocol was much more discrepant. At the first 1.5-minute probe, the digitized response fraction was about 0.20, while the model predicted 0.831. At 90 minutes those values were approximately 0.76 and 0.968. Its RMSE was 38.24 points. The other three held-out protocols had errors of 18.55, 18.89, and 20.45 points. Across the 28 held-out digitized observations, predictions exceeded observations by an average of 20.1 points. For probes through 10 minutes, the pooled RMSE was 34.5 points and the mean excess was 30.6 points; for later probes the RMSE was 15.4 points. These diagnostics show that the problem is especially acute in the initial response after training, not only in the later recovery slope.

The source figures compare separate cohorts, but the V3 model lacks a measured pretraining response for every Rajan cohort. Its shared starting probability for the Fig. 5C/7E/7F group was fitted at the upper bound, 0.995, whereas the Fig. 5E group used 0.880. This difference affects the subsequent 12-hour predictions. A counterfactual with one group’s baseline borrowed for another can reduce an error, but the held-out data would then be helping to choose its starting state. It is therefore a diagnostic of confounding, not a validated improvement. The remaining mismatch may also involve the force mapping, the history of the cells before training,

the cumulative effect of recovery probes, or the recovery equation itself. Figure 6 keeps all seven fitted and held-out recovery curves visible.




Figure 6: Force-, duration-, and frequency-dependent recovery. The three fitted curves constrain V3 but do not suffice to predict the four held-out cohorts. Points are digitized published fractions, without raw cell-level counts; lines are predictions with unchanged shared dynamical parameters.

3.7 Fixed-parameter ablations isolate contributions of the fitted states




Figure 7: Conditional contribution of enhancement, strong-to-weak transfer, and slow history. The full and ablated curves use identical starting calibrations and every unchanged parameter. No optimizer is run after a term is removed.

The ablations used exactly the final full-model parameter values except for the specified term. Setting $B_{\text{update}} = 0$ raised the Doan held-out RMSE from 4.11 to 5.47 points. This removes the model's pairing-

dependent enhancement while retaining the habituation updates. Setting both cross rates to zero raised the Doan error to 5.07 points and the Rajan fitted recovery error from 8.41 to 28.24 points. Since k_{ws} was already fixed at zero, this comparison is specifically the effect of removing the fitted strong-to-weak update k_{sw} . It leaves the Ramdas weak-only results unchanged.

Setting $q_{\text{gain}} = 0$ raised the Ramdas held-out error from 5.68 to 12.53 points. It also raised the Rajan held-out error to 30.78 points and the Doan held-out error to 5.70 points. The slow-history factor therefore influences predictions beyond the two-bout experiments in this particular joint fit. Figure 7 shows representative trajectory changes. These are fixed-parameter interventions: an alternative model allowed to refit might compensate through other parameters, and the present comparison does not establish that any state is a uniquely identified biological entity.

3.8 Latent-state trajectories and parameter limits

The event equations give a specific explanation for the shape of the fitted curves. In WS training, a weak event leaves T available for the following strong event. If that event produces contraction, it can raise B before a subsequent weak response is computed. In WW, contraction at a primed weak event can also recruit B ; there is no strong-stimulus gate on Eq. 4. The mean B trajectory reaches approximately 0.065 in WW1/45, although habituation still outweighs enhancement in its net response. A positive enhancement state therefore need not produce a rising observed trajectory. Contraction-associated events increase L_w , and a strong event can increase it through k_{sw} . This makes the later WS response decline even while B remains positive. In the Ramdas schedules, Q accumulates over the first bout and persists during rest; recovered responsiveness can coexist with faster subsequent habituation. Figure 8 shows ensemble means of post-event cell states from the same simulation settings as the exported curves. These are inferred states, not recordings of membrane voltage or receptor availability.

The recovery parameters are much less secure than the trajectory fits. $A_{\text{max},w}$ approached its upper bound, q_{gain} reached its upper bound, R_{s0} reached its lower bound, and τ_{Ls} approached the two-day upper bound. The Rajan response scale also approached its lower bound. The fitted $\tau_{Ls} = 171692$ s is a conditional search result, not evidence for a two-day biological lifetime. All digitized Rajan recovery probes used the weak force level, leaving no direct strong-probe recovery curve to isolate τ_{Ls} . Formal uncertainty intervals have not been established for this realized-event fit. Independent seed and ensemble-size checks produced no fitting-protocol flags under the stated 0.02 maximum probability-change criterion. This checks simulation stability, not biological identifiability. Profiles across plausible digitization and starting-state choices remain necessary before treating fitted recovery constants as measured quantities.

For the Doan and Ramdas held-out populations, the principal evaluation used the first observed event to set V_0 and scored subsequent events. When the first three responses were instead used for a secondary cohort calibration, held-out mean error changed from 4.11 to 3.61 points for Doan and from 5.68 to 5.50 points for Ramdas. This modest improvement shows that part of the apparent trajectory error depends on how much is known about the starting population. It is a conditional evaluation with extra cohort information, not the original held-out test. Figure 9 labels the observations used for calibration and the taps subsequently scored under each rule.




Figure 8: Ensemble-mean post-event states from the selected realized-event simulation. Both WS and WW can recruit enhancement; intrinsic parameters are shared, while realized histories differ. Panels A and B show the paired 1/45 schedules; panel C shows the 1-minute/5-hour repeated-bout schedule. Their shapes explain the model’s behavior but are not independent biological measurements.




Figure 9: Sensitivity of held-out predictions to observed cohort starting responses. Three-tap conditioning uses two additional observations from each held-out population and should not be mistaken for an unchanged prospective prediction.

4 Discussion

The V3 results extend the first version of this work from a pairing-specific explanation to a quantitative comparison across pairing, repeated-bout habituation, and force-dependent recovery. A trace and enhancement state allow an early rise in weak-stimulus responsiveness; habituation states account for the later decline and strong-to-weak transfer; and the slow history state helps describe behavior in the second training series. The relatively low held-out Doan and Ramdas errors are encouraging because whole protocols, including different timing conditions, were excluded from the fitting objective. These schedules were inspected during model development, so this is developmental evaluation rather than untouched confirmatory validation.

The recovery results prevent a stronger conclusion. The high error for several Rajan cohorts shows that fitting representative recovery curves and a behavioral force ratio was insufficient to predict other training durations and frequencies. The current model maps force levels to response drives but does not reproduce the physical actuator force or contact waveform. Matching two stimuli by an initial probability also need not match the hidden response state of the cells. In particular, contraction probabilities close to one carry little information about how far above threshold a stimulus lies. Future cross-laboratory calibration needs the stimulus waveform and a response curve covering unsaturated probabilities, together with baseline and training histories from the same cohorts.

The states L_w and L_s are effective variables for stimulation-associated habituation. They do not establish anatomically separate weak and strong receptor pathways. The receptor-inactivation model of Rajan and

Marshall offers a possible interpretation of history-dependent responsiveness, whereas older electrophysiological experiments emphasized response dependence and changing excitability [4, 8, 9]. Recent molecular work identifies pathways associated with learning in *Stentor coeruleus* [5]. These studies motivate experiments that measure the states more directly, but the present aggregate behavior cannot choose uniquely among channel adaptation, receptor regulation, downstream excitability, or combinations of them.

The model predicts a population response fraction from simulated individual histories. Cells can acquire different states through different realized contractions, while sharing intrinsic parameters. Rajan and colleagues showed that the behavior of individual cells can change sharply even when the average curve is smooth [3]. Calibration to the average does not test whether the simulated distribution of cellular transition times matches the observed one. The present fit does not estimate intrinsic parameter variation among cells. The Rajan data used here add a second problem: published figure fractions carry digitization uncertainty, and their missing pretraining counts leave starting state and recovery dynamics confounded. These limitations do not erase the successful cross-protocol predictions; they define which conclusions they support. The work demonstrates that a common state-space description is useful for several protocols, while its current recovery generalization and parameter interpretation remain incomplete.

More broadly, ciliates provide an opportunity to study how a single cell modifies its behavior through experience. Habituation, response enhancement, recovery, and faster re-habituation ask how stimulus history is retained and how that history changes later responses. These changes can be reproducible at the population level while remaining heterogeneous among cells. In *Stentor coeruleus*, a smooth population habituation curve can arise from abrupt transitions occurring at different times in individual cells [3]. Same goes for the associative learning-like effect seen in the weak-stimulus response enhancement curve [1]. This recalls a distinction in neural systems, where heterogeneous single-neuron responses can participate in more readily interpretable population dynamics, including those reproduced by recurrent neural-network models [15]. The analogy concerns the relationship between individual activity and population behavior. In the present model, the population curve summarizes learning within independently responding cells; it does not require communication between ciliates. Consequently, explaining the average trajectory is one level of explanation, while explaining the distribution of individual learning histories remains another.

The implications extend beyond ciliates. Evidence of habituation in the multinucleate single-celled organism *Physarum polycephalum* also shows that some forms of learning can occur without neurons or synapses [16]. Such findings raise the possibility that nervous systems evolved to expand on cellular capacities for retaining experience and regulating responsiveness, although behavioral similarities alone cannot establish a shared evolutionary origin. At a functional level, an account of learning requires experience to modify an internal state, that modification to persist beyond the immediate stimulus, and the retained state to influence subsequent behavior. The present model provides one candidate implementation through stimulus-dependent updates, state retention, and a history-dependent response drive. Its different timescales also allow responsiveness to recover while prior training continues to modify the rate of later habituation. This does not establish that all five states are necessary, or identify a universal molecular minimum for memory. It instead makes *Stentor coeruleus* a system in which these requirements can be separated experimentally, helping determine which cellular mechanisms are sufficient for particular forms of learning and which additional mechanisms expand their duration, specificity, and flexibility.

4.1 A proposed reporting framework for mechanical stimulation experiments

Table 5: Proposed minimum reporting framework for *Stentor coeruleus* mechanical stimulation experiments. Report cohort-specific values and release the underlying event and response records where available.

Domain	Required information and raw record
Cells and culture	Strain/source, medium, feeding or starvation, temperature, handling, batch, date, and cell identifiers.
Stimulator and geometry	Actuator and probe shape, position, orientation, contact and fluid conditions; device diagram and calibrated waveform.
Stimulus magnitude	Measured force, displacement or impulse with units, duration and variability for each nominal level; unsaturated force–response curve.
Event sequence	Absolute times and magnitudes of every delivered event, including recovery probes; rest periods, missing and delayed events.
Starting state	Prior exposure, training, rest and selection; baseline probe schedule and per-cell or per-cohort pretraining response counts.
Response scoring	Contraction rule and observation window, latency if available, scoring method, ambiguous outcomes, exclusions and missingness.
Replicates	Cells and denominators at each event, replicate and batch IDs, attrition, cohort variation, and linked individual trajectories.
Analysis access	Preprocessing, aggregation, uncertainty calculation, figure source tables, data dictionary, code and persistent repository identifier.

Mechanical stimulation studies can be difficult to compare when the words “weak,” “strong,” and “recovered” refer to different devices or different cell populations. Table 5 is my proposal for a minimum common record, not an established reporting requirement. Preserving the physical event sequence and cell-level denominators alongside plotted averages would let future models reproduce the schedules exactly, calibrate force–response relations, and separate a change in starting responsiveness from a change in learning dynamics.

5 Future Work

The most immediate experiment is a matched calibration across preparations and stimulators. In each laboratory, the same documented culture and handling conditions should be followed by a range of mechanically measured force levels, including values that yield intermediate rather than saturated contraction probabilities. Recording the full stimulus waveform and baseline responses in each cohort would test whether the current Rajan–Doan drive mapping has a stable meaning. Similar initial probabilities alone would not be sufficient: the comparison must include responses after matched training and rest histories.

A second experiment should measure the missing trajectory before and after recovery in the *same* cohorts. Weak and strong training forces could be crossed with short and long training bouts and several rest intervals. Pretraining response counts, every training response, and successive recovery probes should be recorded at the cell level, with an additional group receiving only a single recovery probe at each time. The latter comparison would show whether repeated probes themselves change apparent recovery. Randomized, replicated cohorts would allow the initial state to be estimated without using the held-out recovery outcomes as a calibration target. The contrast between the 12-hour slow-frequency held-out result and the fitted force-level protocols makes this experiment especially informative.

A third test should distinguish a slow history variable from an alternative recovery law. Repeating the same schedule after different intervening rests, and comparing cells with matched end-of-training response but different earlier histories, would test whether the later bout depends on more than the current contraction probability. If a single exponential recovery law cannot explain these experiments, the appropriate extension may involve more than one timescale or a history-dependent transition. That extension should be compared prospectively with the present model on whole held-out protocols, rather than adopted solely because it improves a fit.

Finally, the latent states require independent measurements. Voltage-sensitive recording or intracellular electrophysiology during WS, WW, and recovery schedules could distinguish a short-lived channel-priming state from sustained depolarization, and test whether contraction at either a weak or strong primed event leaves altered channel gating or background inward current, extending the electrophysiological proposal of V1 [2, 8]. Measurements or perturbations of the candidate calcium and phosphorylation pathways identified in recent molecular studies could test whether altered responsiveness accompanies the modeled histories [5]. Time-resolved measurements of RNA, protein turnover, phosphorylation, and chromatin could test candidate substrates of the slower Q state, with selective perturbation and rescue separating causal retention mechanisms from correlated responses. Cortical imaging could test whether a response-facilitating state tracks a maintained membrane–cortex organization. Such experiments should include individual-cell trajectories, because a population average alone may not distinguish graded changes from switches in cell state [3]. Formal profile-likelihood or comparable identifiability analysis, repeated across uncertainty in cohort baselines and digitized points, should precede a biological interpretation of fitted recovery constants [7]. Until those tests are available, the model is best understood as a falsifiable behavioral framework, not a complete cellular account of learning.

Data, code, and figure preparation

The manuscript source archive includes the selected parameter checkpoint, exported trajectories, evaluation tables, and the script used to prepare its figures. Data, code, and figure preparation are also available at: [Archive](#)

References

- [1] Doan, N., Theroux, A., Ramdas, T., & Gershman, S. J. (2026). Associative learning in the protozoan *Stentor coeruleus*. *bioRxiv* preprint. doi:10.64898/2026.02.25.708045.
- [2] Mohamed, H. (2026). Residual depolarization as a model of associative learning-like behavior in the protozoan *Stentor coeruleus*. *Research Square* preprint, version 1. doi:10.21203/rs.3.rs-10346520/v1.
- [3] Rajan, D., Makushok, T., Kalish, A., Acuna, L., Bonville, A., Correa Almanza, K., Garibay, B., Tang, E., Voss, M., Lin, A., Barlow, K., Harrigan, P., Slabodnick, M. M., & Marshall, W. F. (2023). Single-cell analysis of habituation in *Stentor coeruleus*. *Current Biology*, 33(2), 241–251.e4. doi:10.1016/j.cub.2022.11.010.
- [4] Rajan, D. H., & Marshall, W. F. (2025). A receptor-inactivation model for single-celled habituation in *Stentor coeruleus*. *Current Biology*, 35(14), 3327–3340.e4. doi:10.1016/j.cub.2025.05.071.
- [5] Rajan, D. H., Albright, A., Kim, H., Diaz, U., Hudnall, Y., Steube, N., Dey, G., Liu, T., & Marshall, W. F. (2026). Molecular pathways for learning in the single-cell *Stentor coeruleus*. *Current Biology*, 36(9), 2367–2381.e9. doi:10.1016/j.cub.2026.03.080.
- [6] Ramdas, T., Doan, N., Theroux, A., & Gershman, S. J. (2026). A quantitative portrait of habituation in *Stentor coeruleus*. *eLife*, reviewed preprint, version 1. doi:10.7554/eLife.112314.1.
- [7] Raue, A., Kreutz, C., Maiwald, T., Bachmann, J., Schilling, M., Klingmüller, U., & Timmer, J. (2009). Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood. *Bioinformatics*, 25(15), 1923–1929. doi:10.1093/bioinformatics/btp358.
- [8] Wood, D. C. (1969). Electrophysiological studies of the protozoan, *Stentor coeruleus*. *Journal of Neurobiology*, 1, 363–377. doi:10.1002/neu.480010402.
- [9] Wood, D. C. (1988). Habituation in *Stentor*: a response-dependent process. *Journal of Neuroscience*, 8(7), 2248–2253. doi:10.1523/JNEUROSCI.08-07-02248.1988.

- [10] Hennessey, T. M., & Kung, C. (1985). Slow inactivation of the calcium current of *Paramecium* is dependent on voltage and not internal calcium. *Journal of Physiology*, 365, 165–179. doi:10.1113/jphysiol.1985.sp015765.
- [11] Beisson, J., & Sonneborn, T. M. (1965). Cytoplasmic inheritance of the organization of the cell cortex in *Paramecium aurelia*. *Proceedings of the National Academy of Sciences of the United States of America*, 53(2), 275–282. doi:10.1073/pnas.53.2.275.
- [12] Preston, R. R., & Hammond, J. A. (1998). Long-term adaptation of Ca^{2+} -dependent behaviour in *Paramecium tetraurelia*. *Journal of Experimental Biology*, 201(11), 1835–1846. doi:10.1242/jeb.201.11.1835.
- [13] Garnier, O., Serrano, V., Duhaucourt, S., & Meyer, E. (2004). RNA-mediated programming of developmental genome rearrangements in *Paramecium tetraurelia*. *Molecular and Cellular Biology*, 24(17), 7370–7379. doi:10.1128/MCB.24.17.7370-7379.2004.
- [14] Wood, D. C. (1988). Habituation in *Stentor*: produced by mechanoreceptor channel modification. *Journal of Neuroscience*, 8(7), 2254–2258. doi:10.1523/JNEUROSCI.08-07-02254.1988.
- [15] Mante, V., Sussillo, D., Shenoy, K. V., & Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, 503(7474), 78–84. doi:10.1038/nature12742.
- [16] Boisseau, R. P., Vogel, D., & Dussutour, A. (2016). Habituation in non-neural organisms: evidence from slime moulds. *Proceedings of the Royal Society B: Biological Sciences*, 283(1829), 20160446. doi:10.1098/rspb.2016.0446.