Online Appendix

Rapid learning and adaptive search in complex environments: How underestimating noise in performance feedback can leverage and resolve errors of commission

Table of Contents

Online Appendix

Appendix A1. Results using Uniform distribution for landscape calculations.

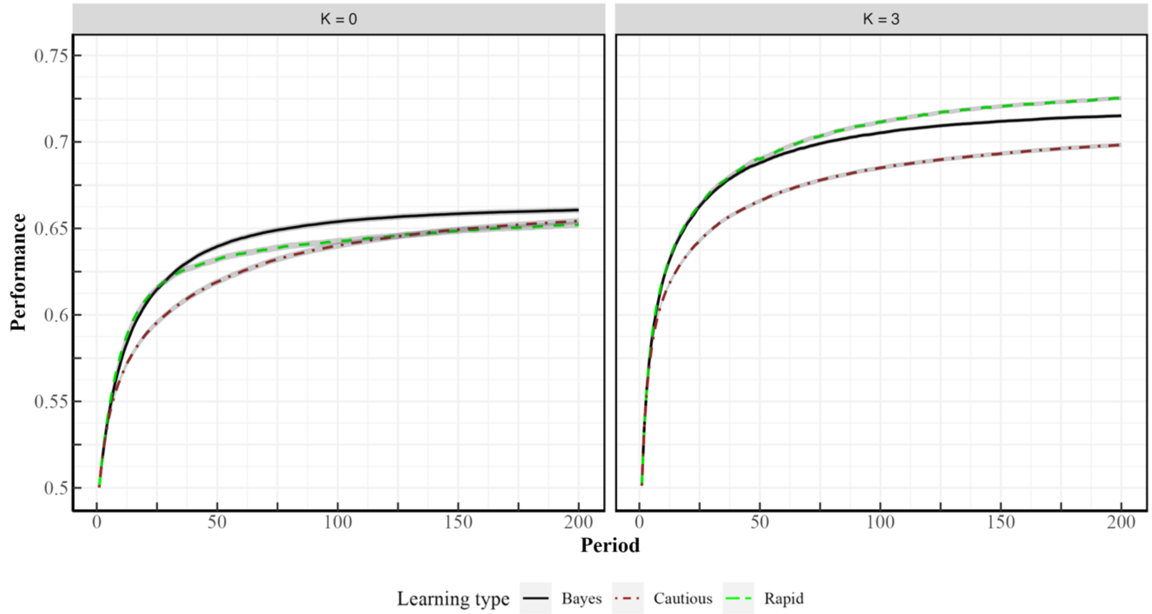

Note. The NK model uses U[0,1] and noise is calculated as a draw from a normal distribution mean = 0 and sd = $\sigma_{Landscape}$.

Appendix A2. N = 6 landscape results.

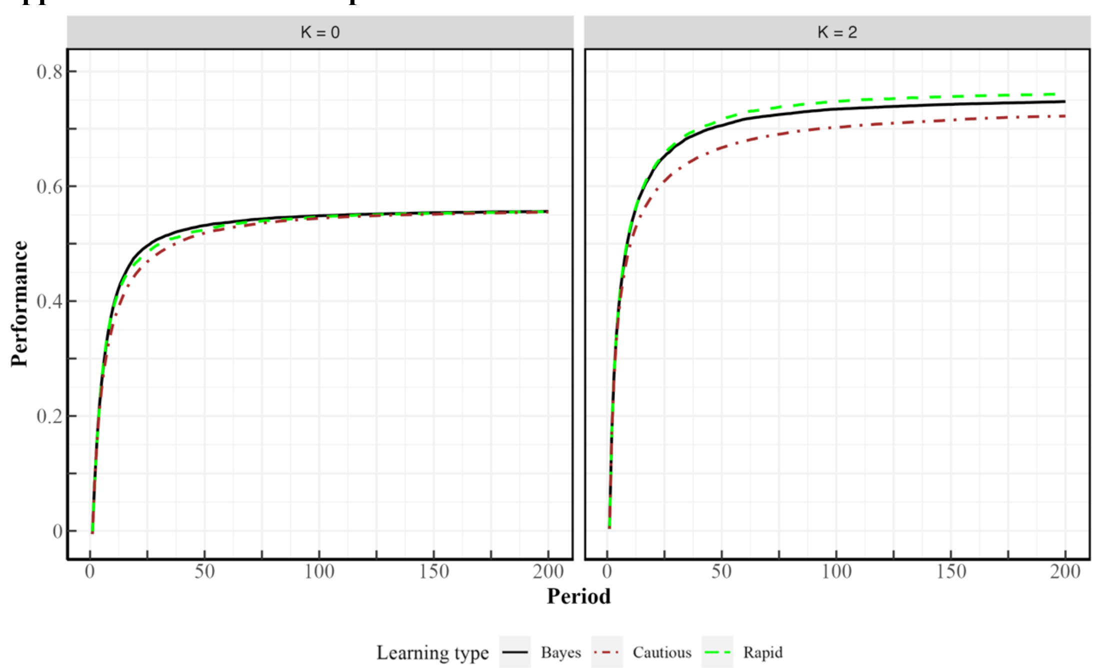

Online Appendix

Appendix A3. Performance results over 200 periods for every level of complexity.

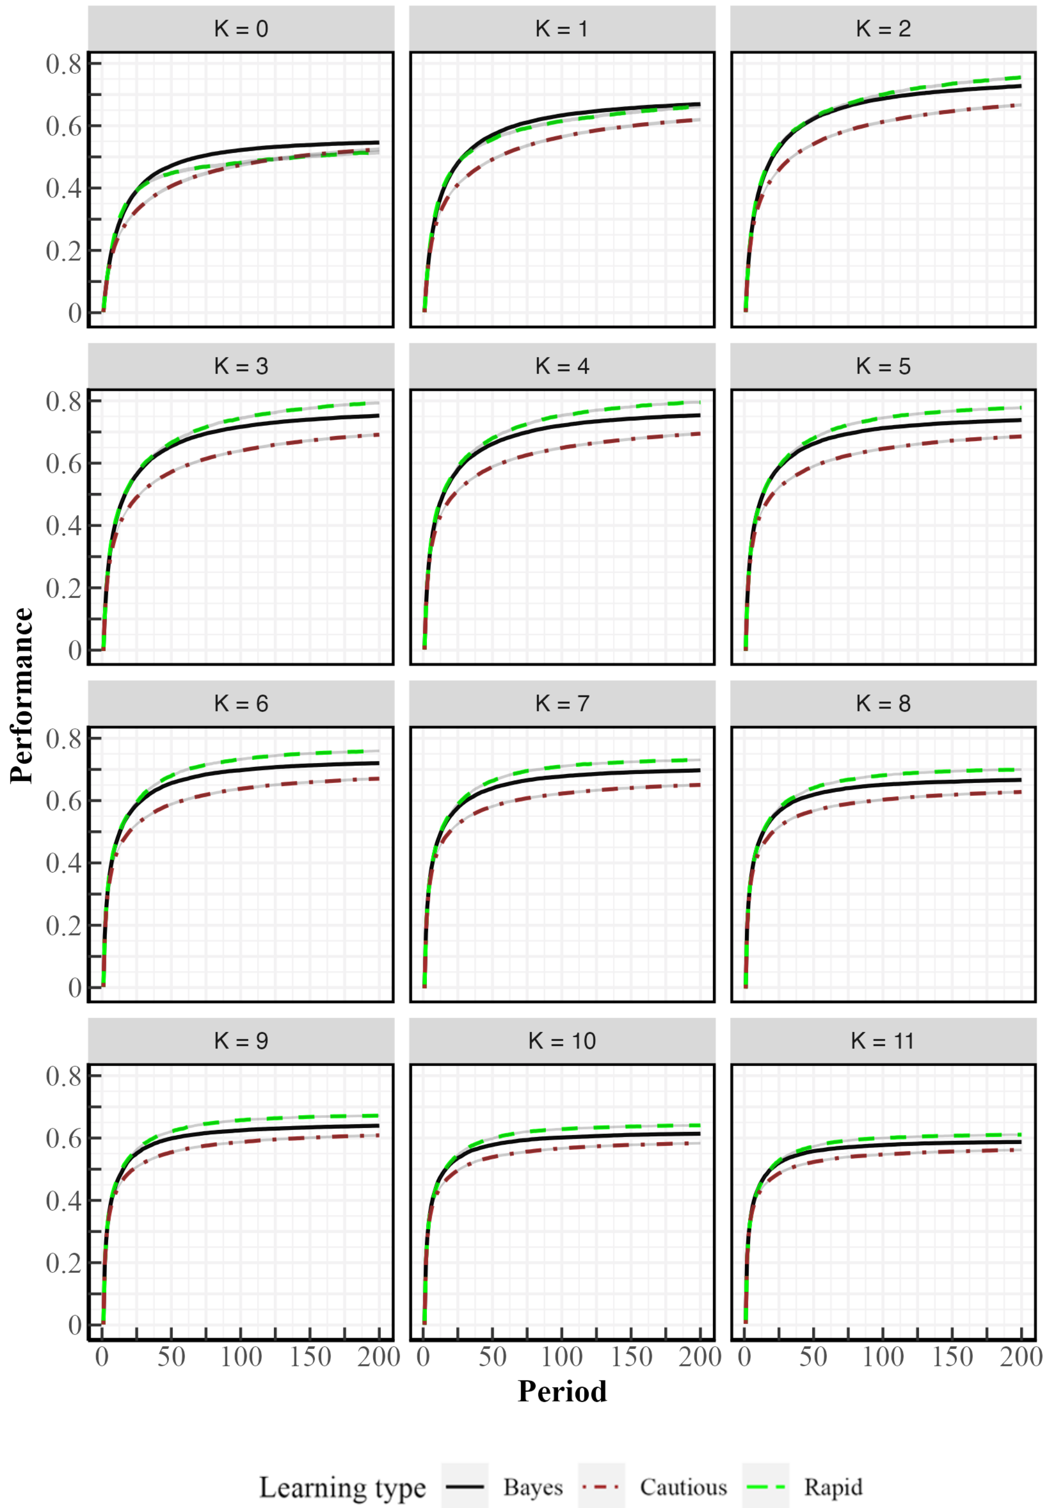

Appendix A4. Results for varying noise-assumption and noise multiplier at t = 200.

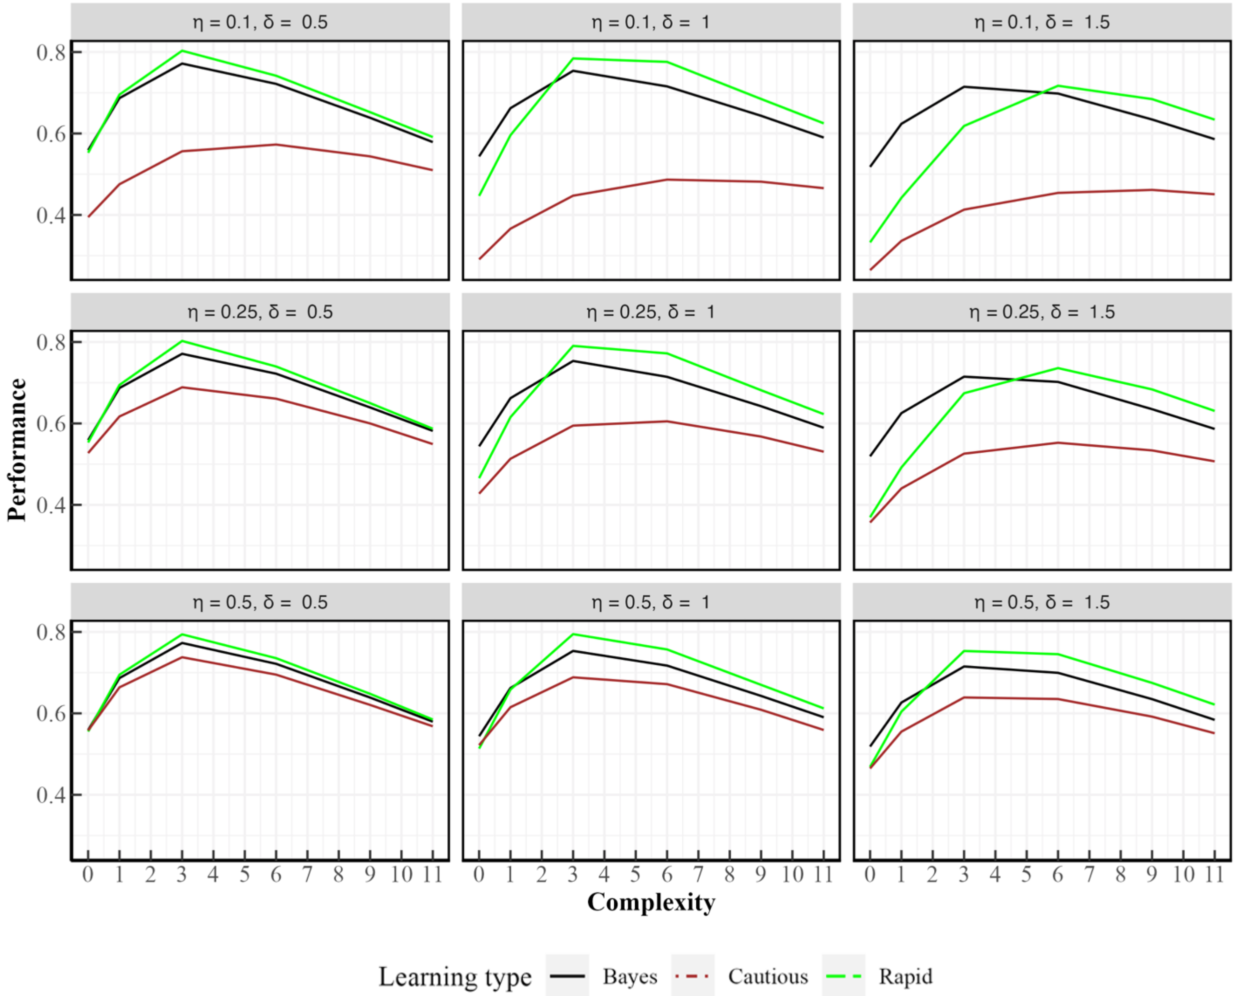

Note. We specify a proportional scale parameter η that translates into the respective alpha values (noise assumption). Specifically, $\alpha_{cautious}$= 1/η; and α_{rapid}= η. For example, when η = 0.1, then $\alpha_{cautious}$ = 10; and α_{rapid} = 0.1. By doing so, we only need to report η in the figure above. Noise multiplier δ = 1, replicates our main results. Values greater than one leads to noisier feedback, and values below one lead to less noisy feedback.

Appendix A5. Long-run results for 2,000 periods (normal and log10 x-axis).

Continues scaled x-axis.

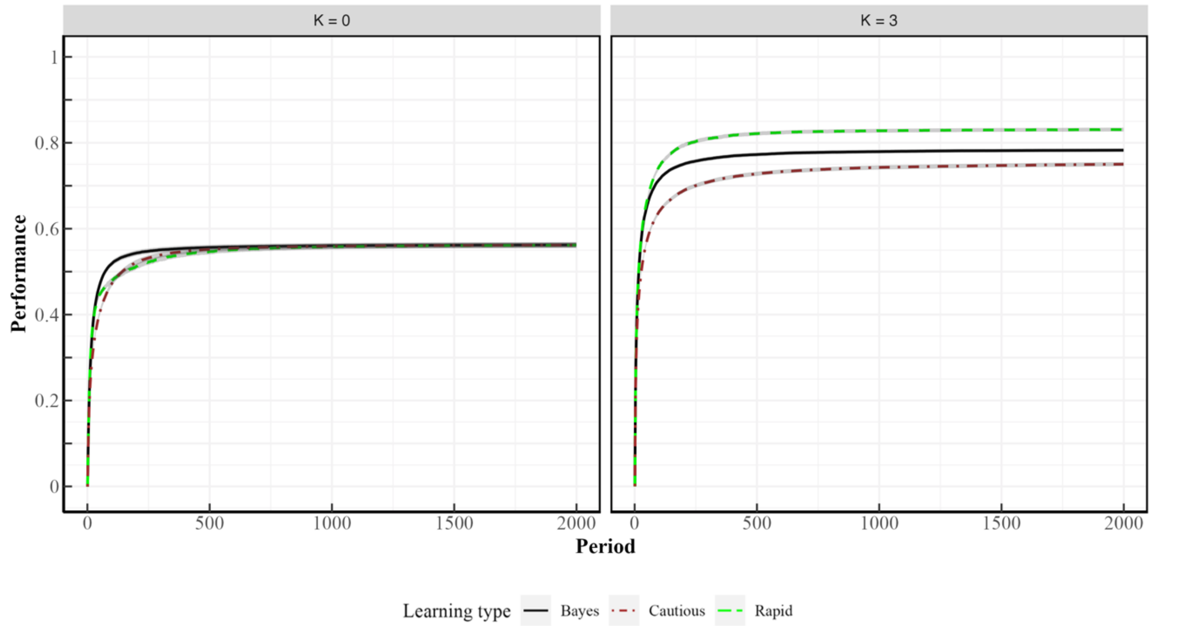

Log10 – scaled x-axis.

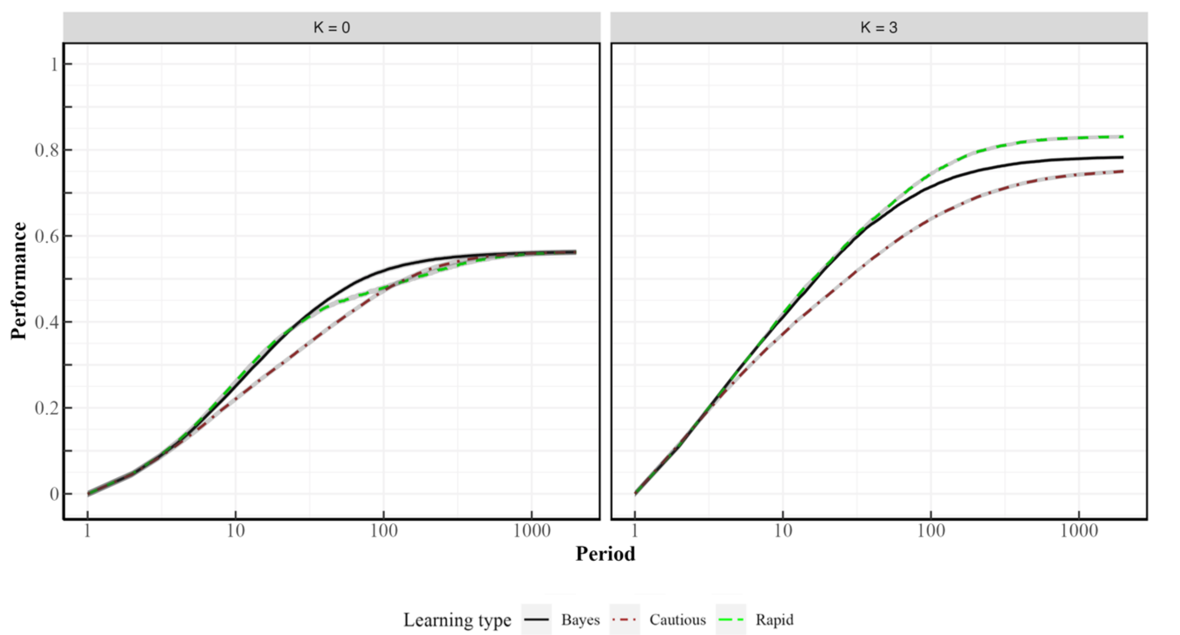

Appendix A6. Varying the search breadth *s*.

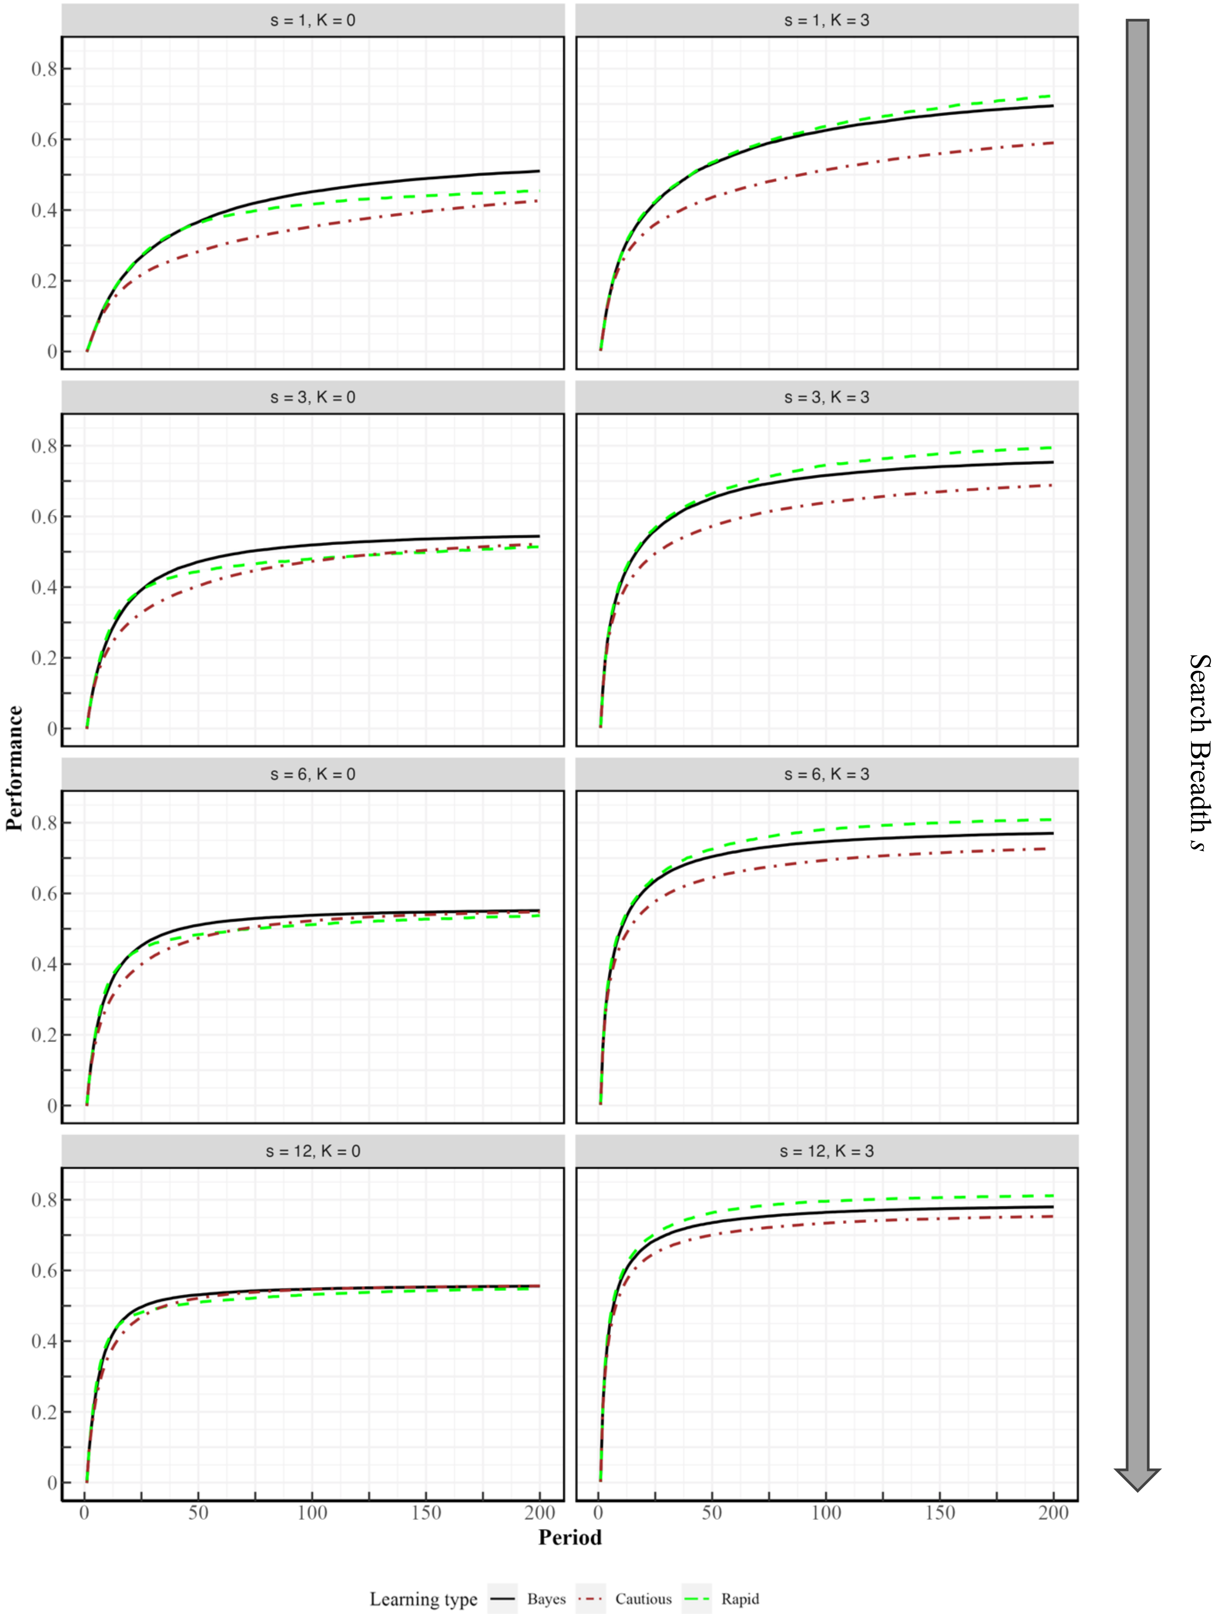

Appendix A7. Bias in prior as a multiple of landscape standard deviation.

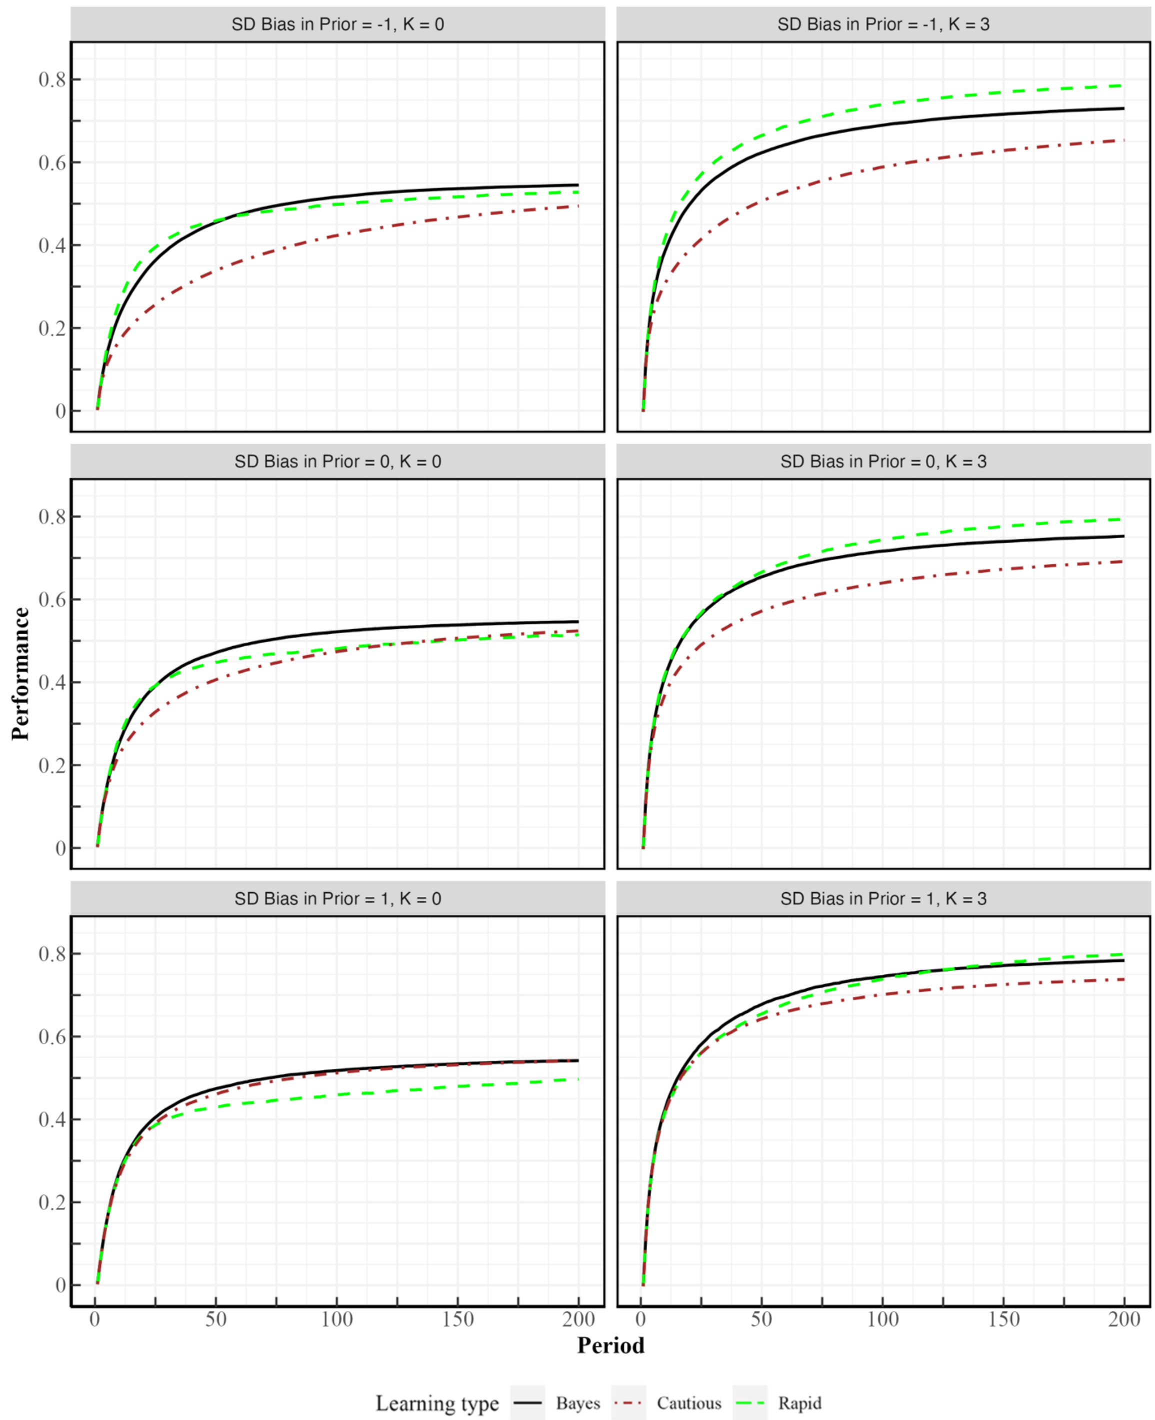

Note. SD Bias in Prior shifts the initial belief in standard deviations of the true landscape standard deviation. For example, when SD Bias = 1, and the true landscape has a mean performance of 0.5 and a standard deviation of 0.2, the biased initial belief is set to (0.5 + 1*0.2) = **0.7**.

Appendix A8. Additional analysis for counterfactual model with static beliefs.

A8a and A8b show adaptation plots. We observe that adaptation rates remain relatively constant over time without belief updating, whereas they decrease in our Bayesian belief updating model, as agents build confidence in their assessments. Overall performance levels are also substantially lower without learning, as agents cannot effectively distinguish signal from noise through repeated sampling. These findings confirm that our main results are specifically driven by the interplay between noise perception and learning (i.e., belief updating), not merely by differences in adaptation thresholds. We also provide the equivalent of our commission and omission error plots (A8c, A8d). We find that without belief updating, error rates are largely persistent over times, except for cautious learners whose total rate drops from the early start but are largely persistent in the long-run.

Figure A8a. Adaptation rate without belief-updating ($K = 0$)

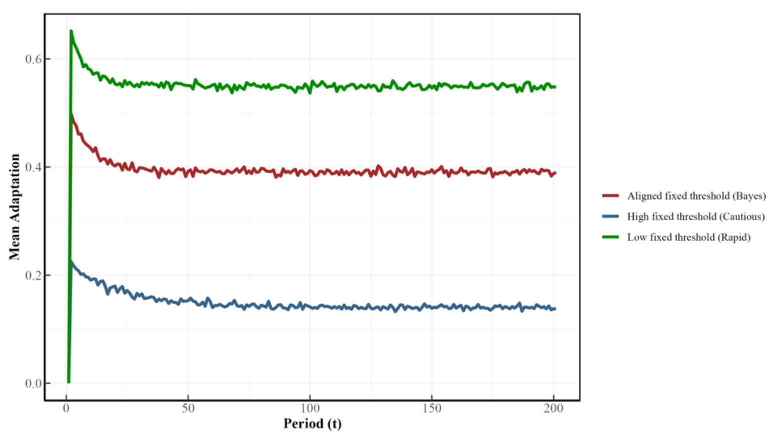

Figure A8b. Adaptation rate without belief-updating ($K = 3$)

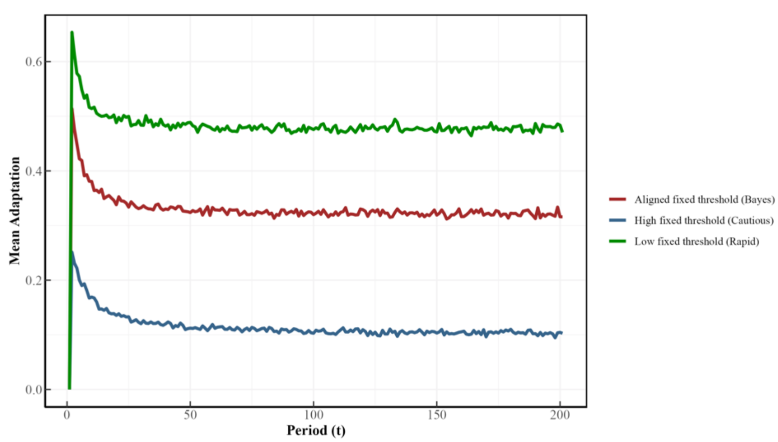

A8c). Omission and commission errors results with no belief-updating (static beliefs) for K = 0.

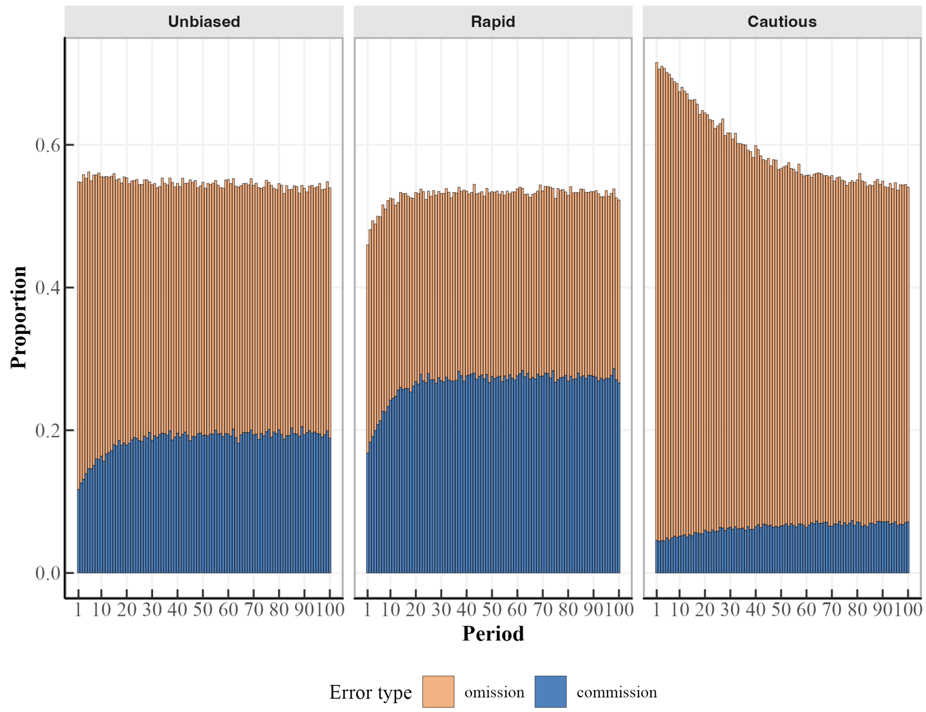

A8d). Omission and commission errors results with no belief-updating (static beliefs) for K = 3.

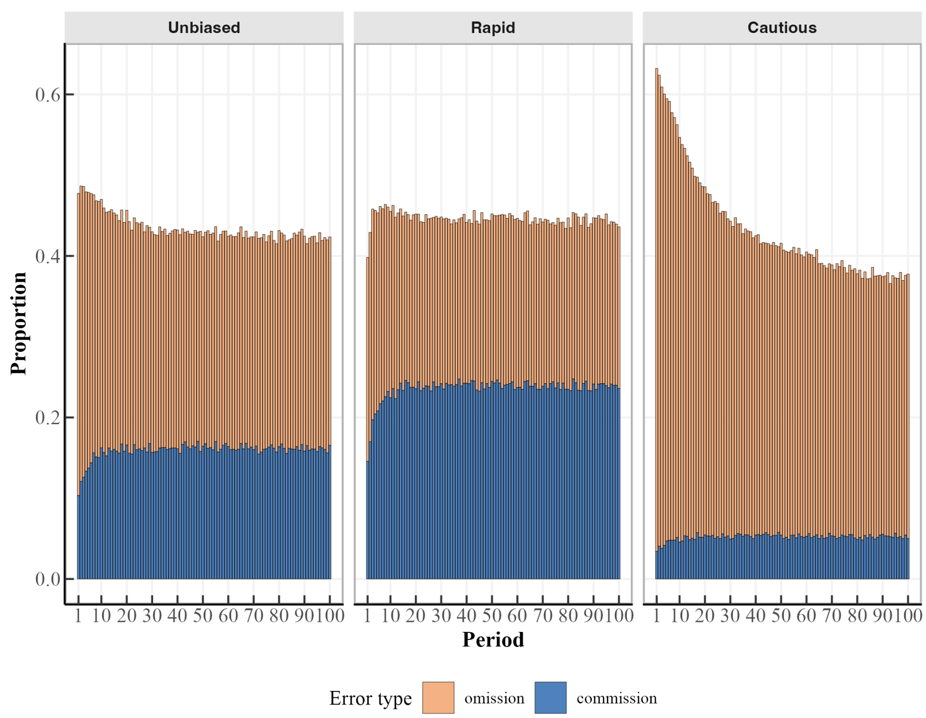

Appendix A9. Number of observations at different adaptation points (K = 3)

In the plot below, we provide the number of feedback signals that decision makers have received for a focal position, when adapting to that position in a K = 3 landscape. We track this for adaptive moves #1, #3, #5, #10, and for Bayes and Rapid learners at move #25 (cautious learners typically don't adapt that frequently). The figure reveals that for the first adaptive move, all three learning types make their initial move with only slightly more than one feedback signal. For adaptive moves #3 and #5, the average number of obtained feedback signals increases for all learners, but most significantly for cautious learners and Bayes, whereas rapid learners adapt based on fewer signals (just over 2 for adaptive move #5).

Importantly, because rapid learners adapt more frequently, we need to inspect later moves, too. We find that the number of observed signals required to "change" their minds increases substantially for later adaptive moves (e.g., for move #25, rapid learners have obtained 11 data points for the position they adopt). This demonstrates that while rapid learners move quickly early on, they eventually zero in on local peaks by requiring more data to change their minds as local priors become more "confident." This occurs because the rapid learner gravitates toward basins with truly high-performing means, and, therefore, where the average feedback is also quite high. In such basins, the rapid learner is more likely to "stay," naturally accumulating more signals for these neighborhood positions—which strengthens the learner's priors. In other words, underestimating noise in feedback can be a beneficial strategy when learning is possible, as this leads to early commission errors that self-correct more easily than omission errors (as we show in the paper). Moreover, with additional feedback, priors eventually become strong enough to counter additional commission errors.

Figure A9: Number of observations at different adaptation points (K = 3)

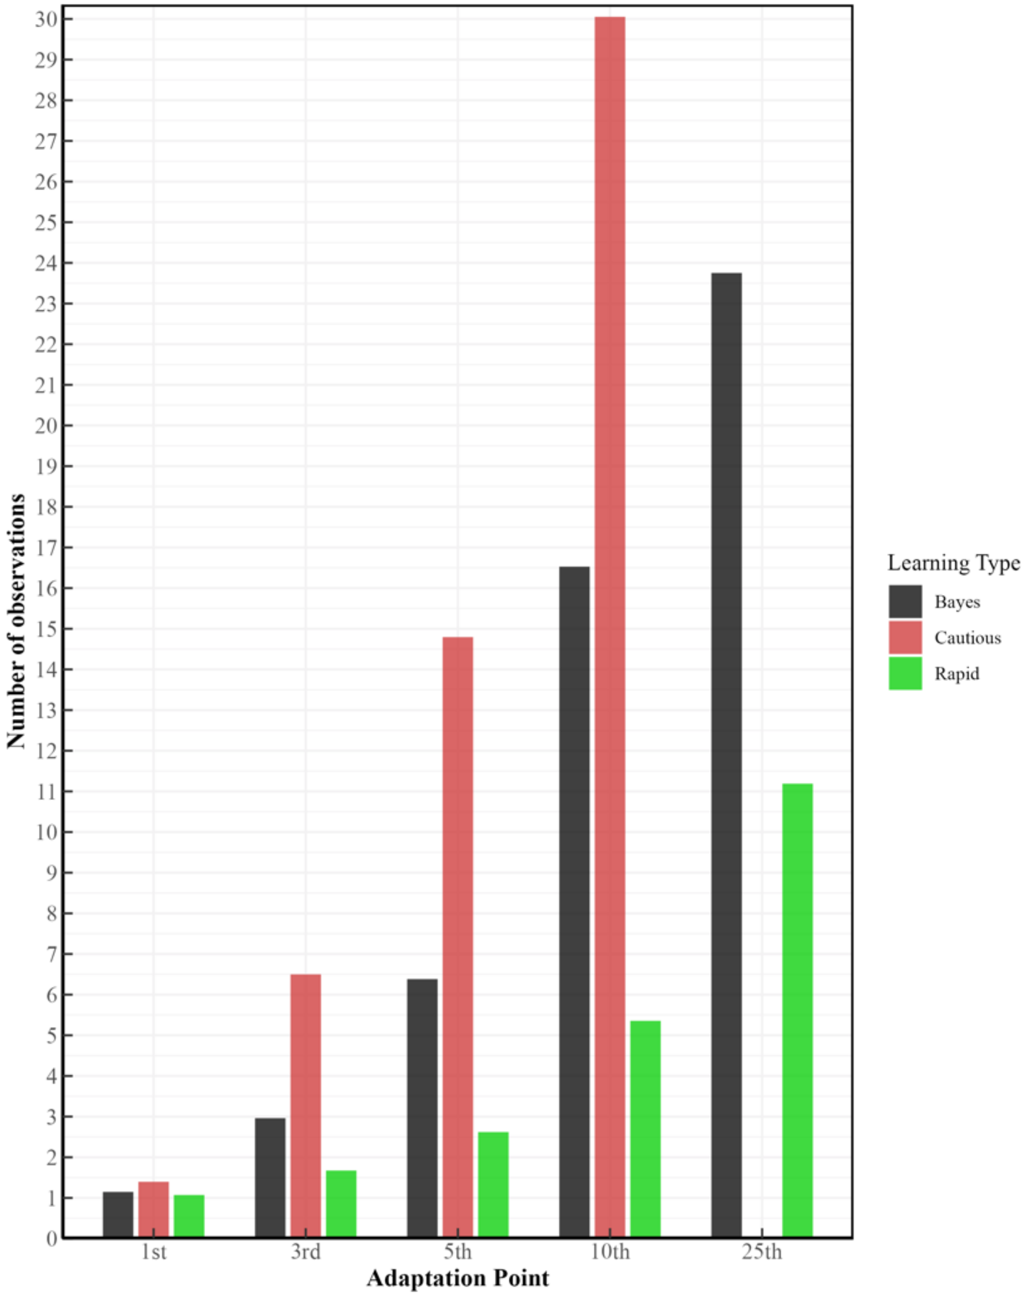

Online Appendix

Appendix A10. Learning rates over time by noise assumption.

Because agents dynamically update beliefs and confidence in their beliefs as shown in Equations 2 and 3 in our manuscript, the learning rate adapts with every performance feedback observation. We can isolate the learning rates for the respective agents as follows.

The learning rate *w* represents the weight given to new information in the belief updating process. From Equation 2, *w* can be written as:

$$\mathrm{w} = \frac{\sigma^2_{m,t-1}}{\hat{\sigma}^2_f + \sigma^2_{m,t-1}}.$$

Further, we can rearrange Equation 3:

$$\frac{1}{\sigma^2_{m,t}} = \frac{1}{\sigma^2_{m,t-1}} + \frac{1}{\hat{\sigma}^2_f}.$$

In other words, the inverse of the posterior variance is the previous inverse posterior variance plus $\frac{1}{\hat{\sigma}^2_f}$ for each additional feedback observation. Consequently, we can state that the inverse of the posterior variance after *n* rounds of feedback is the inverse initial variance plus *n* times $\frac{1}{\hat{\sigma}^2_f}$:

$$\frac{1}{\sigma^2_{m,n}} = \frac{1}{\sigma^2_{initial}} + \frac{n}{\hat{\sigma}^2_f}.$$

Further, as $\hat{\sigma}_f = \sigma_{noise} * \alpha$ and $\sigma^2_{initial}$ in our model is equal to σ_{noise}, we can simplify $\sigma^2_{m,n}$:

$$\sigma^2_{m,n} = \frac{\alpha^2\sigma^2_{noise}}{\alpha^2 + n}.$$

We can then substitute $\sigma^2_{m,n}$ for $\sigma^2_{m,t-1}$ in our learning rate equation:

$$w = \frac{\frac{\alpha^2\sigma^2_{noise}}{\alpha^2 + n}}{\alpha^2\sigma^2_{noise} + \frac{\alpha^2\sigma^2_{noise}}{\alpha^2 + n}},$$

which can be further simplified to:

$$w = \frac{1}{\alpha^2 + n}.$$

This closed-form representation of the learning rate (*w*), therefore, depends on the noise assumption parameter α and the number of observations *n*. Below, we plot the learning rates (y-axis) for the three learning types for the number of observations (*n*) obtained (x-axis). Rapid learners show the largest rates which converge more quickly with the Bayesian rate than cautious learners' rate. With sufficiently many observations, all three types converge to the same rate.

Figure A10: Dynamic learning rates by number of observations and learning type.