

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Comparing serial reproduction and serial prediction of random walk

Permalink

<https://escholarship.org/uc/item/0bf0x655>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 45(45)

Authors

Huang, Jiaqi
Busemeyer, Jerome

Publication Date

2023

Peer reviewed

Comparing serial reproduction and serial prediction of random walk

Jiaqi Huang (huangajq@iu.edu)¹, Jerome Busemeyer (jbusemey@indiana.edu)¹

¹ Department of Cognitive Science, Indiana University, 1001 E. 10th Street, Bloomington, IN 47405 USA

Abstract

Current studies of the serial reproduction paradigm focused on stimuli that were statistically independent of each other. We explored serial reproductions of random walk series and examined whether Bayesian models previously built for independent stimulus could be adapted to autocorrelated stimulus. We found that Bayesian models captured most of the empirical results qualitatively, but could be further improved by incorporating recency effects. Besides, given that the optimal strategy of iterative prediction of random walk was to reproduce the current stimuli, we also compared serial prediction of random walk to serial reproduction. We found that serially reproduced and predicted series both decorrelate as a function of chain position and that the means of the series increase in both tasks, which matched qualitative predictions of the Bayesian models.

Keywords random walk; memory; Bayesian cognition; serial reproduction;

Introduction

Serial reproduction is an experimental paradigm where a set of stimuli is recursively reproduced by participants in a chain: the first person in the chain reproduces the original stimuli, and the reproduced stimuli will then be recursively reproduced by the next person in the chain until the end of the chain (Bartlett, 1932). Serial reproduction represents an information path of social networks (Zhang & Busemeyer, 2021), and thus reveals how memory biases may affect information propagation (Lyons & Kashima, 2003; Kashima, 2000; Lee, Gelfand, & Kashima, 2014). Many computational frameworks had been proposed for serial reproduction (Huang, Zhang, Busemeyer, & Breithaupt, 2022; Hemmer & Steyvers, 2009), and among them the most influential was the Bayesian model developed by Xu and Griffiths (2010). The Bayesian model was supported by numerous empirical studies (Xu & Griffiths, 2010; Langlois, Jacoby, Suchow, & Griffiths, 2021; Jacoby & McDermott, 2017), and had been also recently applied to explain deep learning models (Yamakoshi, Hawkins, & Griffiths, 2022).

However, empirical studies that support the Bayesian model had only explored serial reproductions where stimuli were independent of each other. It raises the question of whether the Bayesian model can be applied to serial reproductions of stimuli that contain systematic autocorrelations. In this work, we attempted to answer this question by applying the Bayesian model to serial reproductions of random

walk series. Random walk series were chosen over other autocorrelated stimuli because (1) random walk is a real-world process that we encounter on a daily basis (eg. weather, stock markets) (2) random walks were widely applied to explain both cognitive processes (Busemeyer & Townsend, 1993; Usher & McClelland, 2004) and neural processes (Ashby & Waldron, 2000; Gold & Shadlen, 2007). Another important reason for exploring random walk series is that the optimal strategy for predicting the next stimuli in a random walk series is to simply reproduce the current stimuli. That said, if humans used the optimal strategy for making predictions, the outcomes of serial reproductions and serial predictions of random walk series would have been similar. In fact, there was good evidence that humans' predictions of random walk series deviate from this optimal strategy (Spicer, Zhu, Chater, & Sanborn, 2022; Zhu, Spicer, Sanborn, & Chater, 2021), but none of them so far focused on the serial effects of random walk predictions.

In this work, we empirically compared serial reproductions and serial predictions of random walk series and examined whether the empirical results matched the qualitative predictions of the Bayesian models. Besides, we proposed an extension of the Bayesian model in Xu and Griffiths (2010) to account for potential recency effects induced by the autocorrelations of the random walks.

Bayesian Model of Serial Reproduction

In Xu and Griffiths (2010), serial reproduction can be viewed as a sequence of memory reconstructions in a chain. At chain position n , the model assumed that participant A_n 's previous experience established a prior of the true state of the world $\eta(n)$, with $\eta(n) \sim N(\mu(n), \sigma(n)^2)$, and that the noisy observation followed $x(n)|\eta(n) \sim N(\eta(n), \sigma_x^2)$. The reconstructed true state given the noisy observation $\eta(n)|x(n)$ then followed the Gaussian distribution $N(\lambda(n)x(n) + (1 - \lambda(n))\mu(n), \lambda(n)\sigma_x^2)$, where $\lambda(n) = 1/(1 + \sigma_x^2/\sigma(n)^2)$ (Gelman, Carlin, Stern, & Rubin, 1995). Since the reproduced stimulus $x(n)$ only depended on the stimuli in the previous chain position $x(n-1)$, serial reproduction was a Markov chain with transition probability

$$p(x(n+1)|x(n)) = \int p(x(n+1)|\eta(n))p(\eta(n)|x(n))d\eta(n), \quad (1)$$

where $x(n+1)|\eta(n) \sim N(\eta(n), \sigma_x^2)$. Using the above results, we could write the model as a first-order autoregressive process:

$$x(n+1) = (1 - \lambda(n))\mu(n) + \lambda(n)x(n) + \varepsilon(n+1) \quad (2)$$

where $\varepsilon(n+1) \sim N(0, (1 + \lambda(n))\sigma_x^2)$. In the case when participants had the same prior in the same chain as in Xu and Griffiths (2012), $\mu(n) = \mu_0$ and $\lambda(n) = \lambda$ were constants, and the autoregressive model converged into $N(\mu_0, \sigma_x^2 + \sigma_0^2)$ in the limits.

Since stimuli were assumed to be independent of each other in the same chain position in Xu and Griffiths (2012), we can rewrite Equation 2 as:

$$x(n+1, t) = (1 - \lambda(n))\mu(n) + \lambda(n)x(n, t) + \varepsilon(n+1, t), \quad (3)$$

where t denotes the trial number. Equation 3 thus highlights the fact that $\lambda(n), \mu(n), \varepsilon(n+1)$ were the same across all trials in the same chain position n .

Extending the Bayesian Model

When reproducing random walk series, there could be potential recency influences of the stimuli in the previous trial on reproducing the current stimuli, induced by the trial-by-trial autocorrelations. Thus, the original Bayesian model in Equation 3, which assumed that stimuli were independent, may need modifications. One way to account for potential recency effects is to extend the Bayesian model by incorporating a trial-by-trial autoregressive term:

$$x(n+1, t) = (1 - \lambda(n))\mu(n) + \varepsilon(n+1, t) + \lambda(n)(\beta(n)x(n, t) + (1 - \beta(n))x(n, t-1)), \quad (4)$$

where $0 \leq \beta(n) < 1$ for any n , and $x(n, t-1) = 0$ for $t = 0$. $\beta(n)$ represents an recency coefficient that is constant in trial number t but varies across chain position n . To reduce complexity, only one recency term is included. The term $x(n, t-1)$ represents the $t-1$ th stimuli generated by the n th participants, which is also the $t-1$ th stimuli for the $n+1$ th participants in a serial reproduction paradigm.

Experiment 1: Serial Reproduction

In the first experiment, we explored serial reproductions of random walk series. We examined whether empirical results supported qualitative predictions of the Bayesian model in Xu and Griffiths (2010), and whether incorporating the recency term improved the Bayesian model.

Method

Participants 90 Participants from the United States were recruited through Prolific. Participants were divided into 5 chains, with 18 participants per chain.

Stimuli Stimuli were pictures of missiles with the same width but in varying heights. The starting series of the heights of the missiles for the first participant was generated from a random walk process with drift rate 0 and standard deviation

40. The heights of the missiles were confined to the range of [100, 500] pixel length, and the random walks reflected at the boundaries. The starting heights of the random walk starting series were chosen at random for each of the 5 chains. Stimuli for the subsequent participants were the responses of the previous participant in the chain.

Procedure Participants first went through a set of instructions, before they entered the main phase of the experiment with 200 trials in total. For each trial, participants were shown two screens. In the first screen, they were shown the stimulus whose height would be remembered and reproduced. After they examined the to-be-reproduced missile, they clicked on the “continue” button to enter the second screen where they were shown a second missile whose height was adjustable. The starting height of this adjustable missile was fixed across all 200 trials and across all participants within a chain, but differed between chains. The 5 starting heights for the 5 chains were correspondingly 200 for the first chain, 250 for the second, 350 for the third, 300 for the fourth, and 400 fifth, in pixel length. Participants then adjusted the height of this second missile using their keyboard, until the height of the second missile matched that of the to-be-reproduced missile they remembered. They confirmed their responses by clicking on the “continue” button again on the second screen to enter the next trial. Participants were told in the instruction that they would be awarded bonus for reproducing stimuli as accurate as possible. We paid bonus based on their total sum of square errors across all trials. Participants were initially unaware of the serial dependency of the stimulus, but they could gradually learn this serial dependency as the experiment goes.

Model Fitting

Two versions of the original Bayesian model (OB) in Equation 3 and the extended Bayesian model (EB) in Equation 4 were fitted:

1. $\sigma(n), \mu(n), \sigma_x$ were constant for any chain position n for both models, and $\beta(n)$ was also constant in n for EB. $\sigma(n)$ and $\sigma(x)$ were confined to the range [0,200], and $\mu(n)$ to the range [100,500]. For EB, $\beta(n)$ was also constant and confined to the range [0,1].
2. σ_x was constant in chain position n for both models, but $\sigma(n)$ and $\mu(n)$ which defined the prior were sampled from a mixture of two possible values σ_1, σ_2 and μ_1, μ_2 correspondingly. That said, $\sigma(n)$ had probability of ω to equal σ_1 and $1 - \omega$ to equal σ_2 , and similarly for $\mu(n)$. σ_1 was confined to the range [0,100], σ_2 to the range [100,200], σ_x to the range [0,200], μ_1 to the range [100,300], and μ_2 to the range [300,500]. For EB, the extra parameter $\beta(n)$ was also assumed to be sampled from a Bernoulli mixture of two values β_1, β_2 confined to the range [0,0.5] and [0.5,1] correspondingly.

Conceptually, the first versions of the models assumed that all participants in the same chain had the same prior mean and standard deviation, the same observation noise, and the same

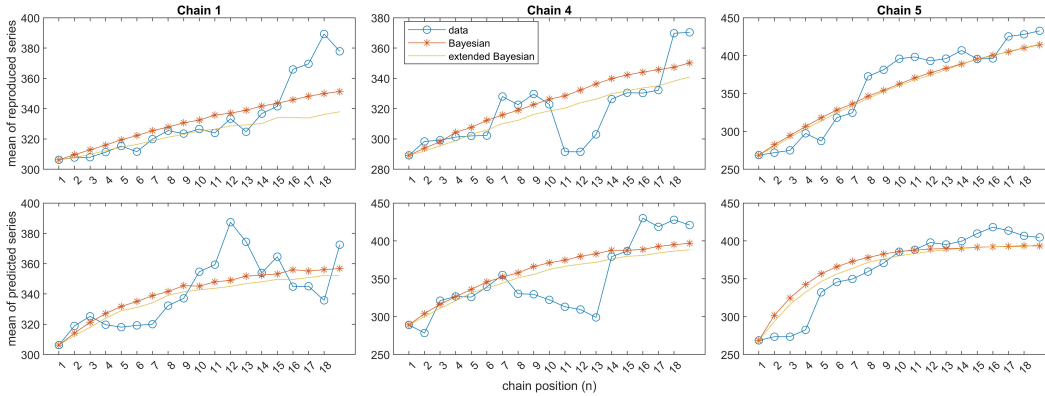


Figure 1: Serial development of the series means for three example chains. The three plots in the upper panel show the result of the serially reproduced series, and the lower panel show that of the serially predicted series. Empirical observations are displayed in blue, the predictions of the best version of the original Bayesian model are displayed in red, and the predictions of the best version of the extended Bayesian model are displayed in yellow. The title indicates the chain number for the three example chains.

recency effect. The second version assumed that prior means, standard deviations, and recency effects were changing as a function of chain position n , but could only change in two different ways. The reason for adopting the second version is that participants in our experiment are not trained to have a common prior as in Xu and Griffiths (2010). We, therefore, assumed that participants may have different priors that could affect their reproduced stimulus. To minimize complexity, we make the simplest assumption that participants' priors can only change in two ways, and we encourage future research to explore more complex hierarchical Bayesian models.

We fitted one set of parameters for each of the five chains by minimizing G^2 errors and compared all four versions of the two models using the Bayesian Information Criterion (BIC). For the first version, the mean BIC over the five chains of EB is $35453 < 35495$ against that of the OB. For the second version, the mean BIC of EB is $34729 < 34878$ against that of the OB. For all of the five chains and for both OB and EB, the second versions of the models were preferred. Furthermore, for all five chains, EB is the better model compared to the same version of OB. Thus, according to the fitting result, EB is the better model overall.

Next, we examined the parameters of the fitted models. Since the second versions of the models involve mixtures of participants' priors, which makes the parameters harder to interpret, we focused on examining parameters of the first versions. Conceptually, in OB, the parameter $\mu(n)$ represents the prior mean, and $1 - \lambda(n)$ represents the relative importance of the prior in reproducing the stimuli. For EB, $\mu(n)$ and $1 - \lambda(n)$ are interpreted in the same way, and there is an additional weight $\lambda(n)(1 - \beta(n))$ that represents the relative importance of the recency term according to Equation 3. The mean $\mu(n)$ over the five chains for first version of OB is 409.19 ± 64.17 , where 64.17 is the standard deviation, and that for EB is 423.92 ± 62.57 . The mean $1 - \lambda(n)$ for OB is 0.070 ± 0.026 , and that for EB is 0.057 ± 0.023 . Since the rel-

ative weights should sum to 1, the result implies that the influence of the prior is presented, but is relatively small compared to the influence of the true stimulus. For EB, the relative importance of the recency term $\lambda(n)(1 - \beta(n))$ is 0.063 ± 0.013 , which is also relatively small compared to the influence of the true stimulus, but are slightly bigger than the influence of the prior, showing that the improvement of the fitting induced by the recency term is robust.

Result and Discussion

Mean and Standard Deviation We first examined empirically whether there was any significant change in the mean and the standard deviation of the reproduced series as a function of the chain position n . We fitted a mixed effect regression model to the means: $\mu \sim 1 + n + k + (1 + n|k)$, where μ denotes the means, n denotes the chain position, and k denotes the chain number of the 5 chains. The same model was applied to standard deviations. The fixed effect of chain position on the mean was 4.579 with p value $0.03 < 0.05$, and that on the standard deviation was -1.064 with p value $0.06 > 0.05$. The significant increase in the means could be a result of converging toward the participants' prior mean, as predicted by the Bayesian model. Such a claim is supported by the fitted first versions of both models, where the $\mu(n)$ parameters are all near the ceilings above the mean of the starting random walk series around 350 for all of the 5 chains.

We next explored whether the Bayesian models predicted the empirical development of means and standard deviations. To compute the predictions, we ran 100 simulations of both Bayesian models using their better versions and best-fitted parameters for each chain. In each simulation, we computed the means and standard deviations of reproduced series for each chain position. We then averaged the means and standard deviations from all simulations for each chain position and took these averaged values as the models' predictions. The results are displayed in Figure 1 and Figure 2. Both models captured

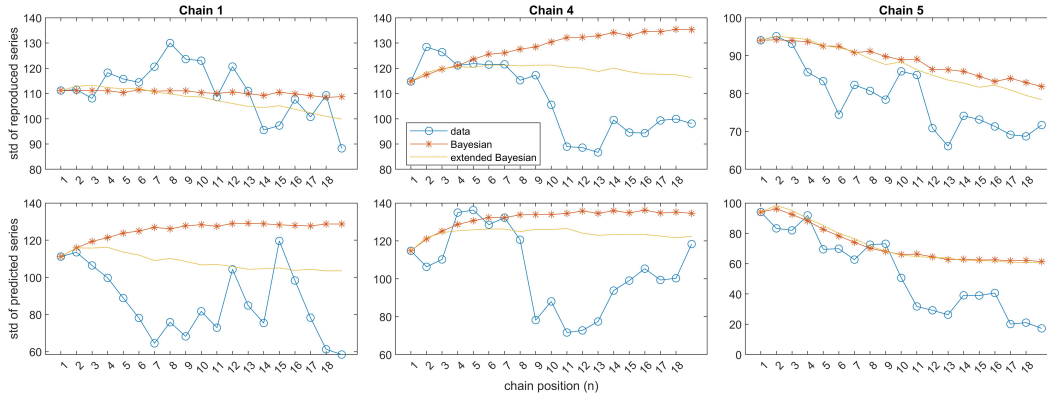


Figure 2: Serial development of the series standard deviations for three example chains. The three plots in the upper panel show the result of serially reproduced series, and the lower panel show that of serially predicted series. The yellow and orange lines are the models' predictions.

the significant increase in the means as a function of chain position but did not capture the changes in standard deviation well.

Autocorrelation of reproduced series An important qualitative prediction of the original Bayesian model in Xu and Griffiths (2010) was that the random walk series will decorrelate as chain position n grew. In particular, when participants in the same chain had a common prior, the distribution for all trials would converge into a Gaussian distribution with the mean being the mean of the common prior (Xu & Griffiths, 2010). However, even if not all participants had the same prior, the series would still decorrelate: consider the most general solution to the difference equation in Equation 3

$$\begin{aligned}
 x(n, t) = & \left(\prod_{p=0}^n \lambda(n) \right) x(0, t) + \sum_{q=1}^n \left(\prod_{l=1}^q \lambda_l \right) \varepsilon_q \\
 & + \sum_{j=0}^{n-1} \left\{ \left(\prod_{k=0}^j (\lambda(n-k))^{\phi(k)} \right) \cdot (1 - \lambda(n-j-1)) \cdot \mu(n-j-1) \right\},
 \end{aligned} \tag{5}$$

where $\phi(k) = 0$ for $k = 0$, and $\phi(k) = 1$ for any other k . Since, $\lambda_n < 1$ for all chain position n by definition, the product term indexed by p in the solution would decrease monotonically, while the sum term indexed by j would increase monotonically, as $0 < \lambda_n < 1$ and $\mu_n > 0$ for all chain positions n . Given that the only term that contained autocorrelation was $x(0, t)$, whose influence became very small as the product term vanished, and that the white noise term indexed by q and the sum term indexed by j combined into a Gaussian distribution independent of trial index t , the model predicted that the starting random walk series $x(0, t)$ decorrelated into a white noise distribution independent of trial number t in the limits. Similarly, for EB in Equation 4, it was also expected that the series would decorrelate as chain position n grew, but because of the additional recency term, the decorrelation can be slower, and the autocorrelation may not completely vanish in the limit.

Empirically, we found evidence supporting this qualitative prediction of the Bayesian models: for all 5 chains, we found

that the autocorrelation for the first and subsequent lags decreased as a function of chain position. A mixed effect regression analysis similar to that for means and standard deviations suggested that the fixed effect of chain position n on the first lag of the autocorrelation function of the reproduced series was -0.018 , with p value $0.006 < 0.05$, showing that the decorrelation was statistically significant.

Figure 3 displays the autocorrelation for the first lag against chain position. As for means and standard deviations, the autocorrelation predictions for each chain position were computed as the averaged result over 100 simulations. Comparing models' predictions with the empirical autocorrelation, the models' predictions seem to decay much faster as a function of chain position. With the additional recency term, EB performed slightly better than OB, but still failed to slow down to the same rate as the decay of empirical autocorrelation. A potential solution to this problem would be to incorporate more recency terms, and we encourage future works to explore this possibility.

Correlation between reproduced series As the serially reproduced series decorrelate, the Bayesian model in Xu and Griffiths (2010), however, predicted that the correlation between the reproduced series in chain position n and the target series in the previous chain position $n - 1$ would not systematically change as a function of chain position. This is because the theoretical correlation between the reproduced series in n and $n - 1$ was λ_n , which was not expected to change systematically as a function n because: (1) participants reproducing the series were randomly assigned to their chain positions n ; (2) participants were unaware of the exact chain position n they were in. In the simplest version of the model in Xu and Griffiths (2010), it was even assumed that λ_n were constant throughout the chain. The same qualitative prediction was true for EB, as $\beta(n)$, the recency coefficient, was also not expected to be a systematic function of chain position n for the same reason as described for $\lambda(n)$.

Empirically, the qualitative prediction of the Bayesian models about correlations between reproduced series of ad-

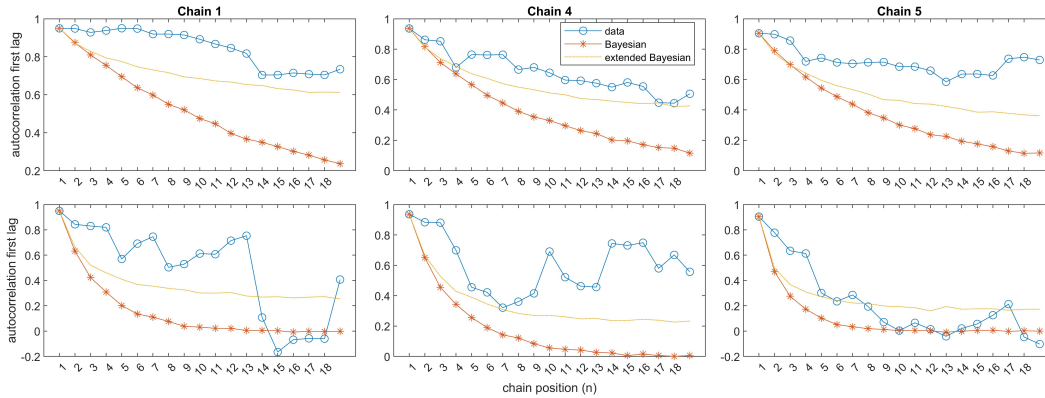


Figure 3: Serial development of the first lag of the autocorrelation function for three example chains. The three plots in the upper panel show the result of serially reproduced series, and the lower panel show that of serially predicted series. The yellow and orange lines are the models’ predictions.

Adjacent chain positions was confirmed by a mixed effect regression analysis, where the fixed effect of chain position n was -0.001 with p value $0.33 > 0.05$. Visually, the result was also evident in the upper panel of Figure 4. The models’ predictions, computed from 100 simulations as for previous statistics, also did not show systematically change as a function of chain position.

Experiment 2: Serial Prediction

The primary purpose of the second experiment was to examine whether the Bayesian model of serial reproduction could be extended to serial predictions of random walk series, given that the optimal strategy for iterative predictions of random walk series was to reproduce the current stimuli.

Method

Participants 90 Participants from the United States were recruited through Prolific. Participants were divided into 5 chains, with 18 participants per chain.

Stimuli The same 5 random walk starting series in experiment 1 was used as the starting series for experiment 2. For subsequent participants, stimuli were the predictions from the previous participants.

Procedure Same as in experiment 1, participants were shown two screens for each trial. They were shown, in the first screen the current stimuli, and in the second screen the missile with adjustable height to express their prediction for height of the next stimuli. As in Experiment 1, participants used their keyboard to adjust their responses, and clicked on the “continue” button to confirm their responses. The starting heights of the adjustable missiles were the same as those specified in experiment 1, and were fixed for each trial in each chain. The same as in experiment 1, the experiment was incentivized and participants’ rewards were based on their sum of square errors of predictions.

Model Fitting

The same four versions of the models in experiment 1 were fitted to the serial prediction data in experiment 2. There were other models, such as the MCMC model (Spicer et al., 2022) that may explain human’s prediction of random walk better stepwise, but our interest was in the serial effects rather than the accuracy of prediction, and thus we did not implement and extend such models. We encourage future work to explore serial extensions of the random walk prediction models. The same as in experiment 1, the better versions of both models were the second versions, with EB performing slightly better than OB in terms of BIC for all five chains. The mean BIC of the second version of EB is $38321 < 38564$ against that of the second version of OB. Regarding the parameters of the first version of the models, the mean $\mu(n)$ for OB is 352.99 ± 67.40 , and that for EB is 357.70 ± 69.62 . The mean $1 - \lambda(n)$ for OB is 0.29 ± 0.12 , and that for EB is 0.23 ± 0.088 . Compared to that for serially reproduced series, the relative importance of the prior increases in estimating serially predicted series. Finally, the mean $\lambda(n)(1 - \beta(n))$ for EB is 0.12 ± 0.030 , which implies that the recency term is also relatively more important for estimating serially predicted series than for estimating reproduced series. Conceptually, the increase in relative importance in priors and recency terms makes sense since participants are no longer instructed to reproduce the series intentionally and thus may adapt to the optimal strategy of predicting the series much slower than that in Experiment 1.

Result and Discussion

Mean and Standard Deviation We ran two mixed effect regressions models as previously to test whether the mean and the standard deviation of the serial prediction series had significant fixed effect related to chain position. The fixed effect of chain position on the mean was 5.206 with p value $0.016 < 0.05$, and on the standard deviation was -2.378 with p value $0.012 < 0.05$. Comparing to the result of experiment 1, the fixed effect of chain position on the mean remained sig-

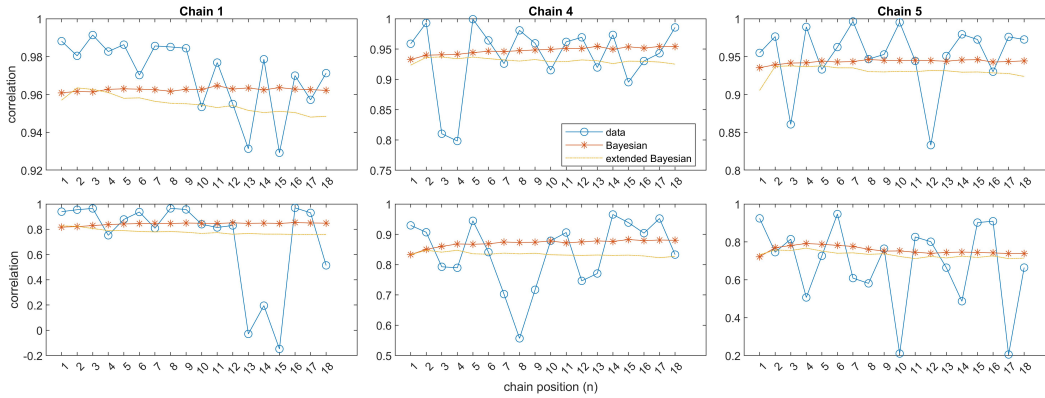


Figure 4: Serial development of the correlations between the adjacent series for three example chains. The three plots in the upper panel show the result of the serially reproduced series, and the lower panel show that of the serially predicted series. The yellow and orange lines are the models' predictions.

nificant and acquired a larger magnitude. Such an increase in magnitude could be explained by an increase in λ_n . Given the significant decrease in standard deviation, the increase in λ_n could only be made possible by a systematic decrease in observation noise σ_x by their definitions in Equation 3. In all four versions of the the two Bayesian models, σ_x was constant, and thus it was expected that all current models failed to simultaneously explain the empirical development of mean and standard deviation found. However, even if we allowed σ_x to change, it was not expected that participants' observation noise would systematically decrease as a function of chain position, as participants were assigned randomly to their chain positions. The present result provided reasonable evidence that the existing Bayesian models for serial reproduction need modifications to account for serial predictions of random walks.

Autocorrelation of predicted series As in experiment 1, we first ran mixed effect regression analysis on the empirical first-lag autocorrelation. The result suggested that the fixed effect of chain position was -0.024 with p value $0.065 > 0.05$, which implied that there was no evidence for a systematic decrease in the first-lag autocorrelation. However, this result did not reject the fact that the predicted series decorrelated: when the autocorrelation reached zero very fast, it remained close to zero for the rest of the chain position, and thus may not had a fixed effect of chain position. As confirmed in Figure 3, it was clear that the predicted series did decorrelate as a function of chain position and indeed in a much faster rate than reproduced series decorrelated in experiment 1. This observation is consistent with the model fitting results, as the relative importance of the prior $1 - \lambda(n)$ is much bigger for predicted series than that for reproduced series, which causes the series to decorrelate faster. The decorrelation was again captured by both models as illustrated by the predictions of the models in the lower panel of Figure 3 computed from 100 simulations. Similar to that in experiment 1, the predictions of the models seem to still decay faster than the empirical autocorrelations, with EB performing slightly better.

Correlation between predicted series mixed effect regression analysis showed that the fixed effect of chain position on correlation between the previous predicted series and the current predicted series was 0.012 with p value $0.96 > 0.05$. The result provided evidence for another qualitative similarity between serial reproduction and serial prediction of the random walk series explainable by the Bayesian frameworks. According to the lower panel of Figure. 4, the predictions of both models again capture this empirical finding qualitatively.

General Discussion

We explored serial prediction and serial reproduction of random walk series, and confirmed several important qualitative predictions of the Bayesian model by Xu and Griffiths (2010) from empirical result, even if stimuli were not independent of each other but were instead autocorrelated. Besides, model fitting to the empirical data show that the Bayesian model could be further improved by incorporating recency influences of reproducing the stimuli in the previous trial on reproducing the stimuli in the current trial. Comparison of serial prediction and serial reproduction experiments show that several serial effects such as decorrelation of the random walk series were consistently presented for both tasks.

Despite that we found serial effects that were the same across serial prediction and serial reproduction of random walk series, it was not yet clear whether participants indeed use the optimal strategy for iterative prediction of random walk series. Thus, future works can explore whether there were certain biases in predicting random walk series that deviate from the optimal strategy, and may intensify in a serial reproduction paradigm, inspired by our work. Besides, the proposed extended Bayesian model in Equation 4 has not yet been derived from a Bayesian probability perspective. We believe that the exact derivation involves Kalman Filters (Welch, Bishop, et al., 1995), and we left this derivation for future research.

Acknowledgments

This material is based upon work supported by the Air Force Office of Scientific Research under award number FA9550-20-1-0027.

References

- Ashby, F. G., & Waldron, E. M. (2000). The neuropsychological bases of category learning. *Current Directions in Psychological Science*, 9(1), 10–14.
- Bartlett, F. C. (1932). *Remembering: A study in experimental and social psychology*. Cambridge university press.
- Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychological review*, 100(3), 432.
- Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. (1995). *Bayesian data analysis*. Chapman and Hall/CRC.
- Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annu. Rev. Neurosci.*, 30, 535–574.
- Hemmer, P., & Steyvers, M. (2009). A bayesian account of reconstructive memory. *Topics in Cognitive Science*, 1(1), 189–202.
- Huang, J., Zhang, Q., Busemeyer, J., & Breithaupt, F. (2022). A quantum walk model for emotion transmission in serial reproduction of narratives. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 44).
- Jacoby, N., & McDermott, J. H. (2017). Integer ratio priors on musical rhythm revealed cross-culturally by iterated reproduction. *Current Biology*, 27(3), 359–370.
- Kashima, Y. (2000). Maintaining cultural stereotypes in the serial reproduction of narratives. *Personality and Social Psychology Bulletin*, 26(5), 594–604.
- Langlois, T. A., Jacoby, N., Suchow, J. W., & Griffiths, T. L. (2021). Serial reproduction reveals the geometry of visuospatial representations. *Proceedings of the National Academy of Sciences*, 118(13), e2012938118.
- Lee, T. L., Gelfand, M. J., & Kashima, Y. (2014). The serial reproduction of conflict: Third parties escalate conflict through communication biases. *Journal of Experimental Social Psychology*, 54, 68–72.
- Lyons, A., & Kashima, Y. (2003). How are stereotypes maintained through communication? the influence of stereotype sharedness. *Journal of personality and social psychology*, 85(6), 989.
- Spicer, J., Zhu, J., Chater, N., & Sanborn, A. (2022). How do people predict a random walk? lessons for models of human cognition.
- Usher, M., & McClelland, J. L. (2004). Loss aversion and inhibition in dynamical models of multialternative choice. *Psychological review*, 111(3), 757.
- Welch, G., Bishop, G., et al. (1995). An introduction to the kalman filter.
- Xu, J., & Griffiths, T. L. (2010). A rational analysis of the effects of memory biases on serial reproduction. *Cognitive psychology*, 60(2), 107–126.
- Yamakoshi, T., Hawkins, R. D., & Griffiths, T. L. (2022). Probing bert’s priors with serial reproduction chains. *arXiv preprint arXiv:2202.12226*.
- Zhang, Q., & Busemeyer, J. (2021). A quantum walk model for idea propagation in social network and group decision making. *Entropy*, 23(5), 622.
- Zhu, J., Spicer, J., Sanborn, A., & Chater, N. (2021). Cognitive variability matches speculative price dynamics.