Exploring Dynamic Decision Making Strategies with Recurrence Quantification Analysis

Erin N. McCormick
(enmccorm@andrew.cmu.edu)
Carnegie Mellon University &
ORISE at Air Force Research Laboratory
Pittsburgh, PA 15213 USA

Leslie M. Blaha
(leslie.blaha@us.af.mil)
Air Force Research Laboratory
Pittsburgh, PA 15213 USA

Cleotilde Gonzalez
(coty@cmu.edu)
Dynamic Decision Making Laboratory
Carnegie Mellon University
Pittsburgh, PA 15213 USA

Abstract

Aggregate statistics, such as percentage of choices, drive many insights about sequential behavior in decision making research. However, aggregation leaves usable information and potential insights unexamined. Here, we introduce the use of recurrence plots (RP) and recurrence quantification analysis (RQA) to explore individual choice sequences and determine generalized patterns of decision making strategies in a dynamic decision task. We illustrate the insights that RPs and RQAs reveal in a data set collected in a past study involving a dynamic, binary choice task (McCormick et al., in preparation). Patterns of recurrence reveal multiple, distinguishable, individual choice patterns among participants who were equally successful in adapting to the dynamic environment. We discuss how RQA of choice behavior can augment our understanding of decision strategies when paired with traditional aggregate assessments.

Keywords: Dynamic decision making; Recurrence quantification analysis; Choice sequences; Decision strategies; Visual analytics

Introduction

Sequences of decisions are a core component in the analysis of dynamic decision making (DDM; Brehmer, 1992; Edwards, 1962; Gonzalez et al., 2017). However, common analytic approaches for describing choice behavior (i.e., strategies) often aggregate over these decision sequences, reporting measures like total points, choice proportions, or maximizing rates. This aggregation obscures important variation in choices over time and between individuals that could elucidate patterns of choice in dynamic environments.

Patterns of individual choice sequences can provide evidence for or against theoretical explanations for behavior in dynamic tasks. For example, in common 2-alternative forced choice tasks (2AFC), generalized behaviors such as *Win-Stay-Lose-Shift* (Biele et al., 2009), *Hot-Stove* (Denrell & March, 2001), *stickiness* and *recency* (Rakow & Miler, 2009; Avrahami & Kareev, 2011) have been investigated using fixed choice sequences. Research studying where these patterns appear often uses aggregate group analysis (Rakow & Miler, 2009; Biele et al., 2009) or computational model-fitting at the individual level (Lejarraga et al., 2014) to examine how individual variation suggests decision strategies.

While past studies inform how human choices may fit these standardized behaviors, these analyses do not provide a complete view of the variety of strategies that individuals may exhibit. In DDM tasks, where patterns of choice may vary as a function of the individuals' decisions as well as independently from exogenous events (Gonzalez et al., 2017), a more detailed analysis of sequential patterns is necessary.

In this paper, we use recurrence plots (RPs) and recurrence quantification analysis (RQA) to describe individual decision makers' choice sequences. We analyze a subset of data collected in a previous DDM study that used a 2AFC task (McCormick et al., in preparation). This study tested whether decision makers successfully adapted to a changing decision environment when the type of change was manipulated. The simple task builds on findings that even simple dynamic environments can embody much complexity, but allow clearer analysis of the factors affecting decision making than more complicated, traditional DDM tasks (Gonzalez et al., 2017).

Using RPs to investigate *how* decision makers did or did not adapt reveals multiple, distinguishable individual choice patterns—even among decision makers who were similarly adaptive (optimal) when compared on aggregate measures. These individual choice patterns for adaptation among similarly successful decision makers provide descriptive insight into the variety of ways individuals adapt to dynamic environments. Such descriptive insights provide a first step to assess whether existing theories of decision making successfully account for the systematic patterns uncovered by RPs and RQA.

Analyzing Recurrence: A Short Primer

RPs and RQA statistics comprise a visual analytics approach to studying patterns in complex systems and time series data (Marwan et al., 2007; Webber & Marwan, 2015). Recurrence refers to a point where a system returns to a state value it has previously exhibited. Originated in Physics, recurrence emphasized periodicities and other patterns in a phase space embedding of a complex system (Eckmann et al., 1987; Marwan et al., 2007). In Psychology, recurrence analysis has been applied to motor control, conversation, and interpersonal dynamics but not to decision-making (for review, see Coco & Dale, 2014). We apply this approach to DDM by examining the patterns of recurrence in choice sequences.

In a choice sequence, each choice trial is one point in time, and we define recurrence $\mathbf{R}$ for trials $i, j = 1, \ldots, T$ by the binary function:

$$\mathbf{R}_{i,j} = \begin{cases} 1, & choice_i = choice_j \\ 0, & choice_i \neq choice_j \end{cases}. \tag{1}$$

Note that **R** is agnostic to the specific choice on each trial; it merely codes if there has been a repetition of an earlier choice (event state). Defining the possible event states in a sequence is an important step in RQA that introduces flexibility and necessitates deliberate care. For this demonstration, we start simply by defining the event states as the set of possible choices in the task. In our 2AFC case study, this is whether the decision maker chose Option A or Option B on each trial.[1] Thus, **R** captures when the participant repeats selections of Option A or Option B.

Recurrence can be computed as (1) auto-recurrence of a choice sequence against itself, or (2) cross-recurrence analyzing two choice sequences against each other (Coco & Dale, 2014). In the present work, we focus only on auto-recurrence. We revisit cross-recurrence concepts in the Discussion.

Recurrence plots (RPs) depict the values of **R** in a heatmap-like grid visualization (see Figure 1). Both the horizontal and vertical dimensions represent the same sequence of trials. The first trial is in the lower left corner, and time proceeds forward from left to right on the x-axis and bottom to top on the y-axis, placing the last trial in the upper right corner. The binary values of **R** are colored black if recurrent ($\mathbf{R}_{i,j} = 1$) or white if not ($\mathbf{R}_{i,j} = 0$). Auto-recurrence plots are always symmetric over the diagonal running from lower left to upper right. This diagonal is referred to as the Line of Incidence (LOI) where $\mathbf{R}_{i,i} = 1$. We typically remove the LOI from auto-recurrence RPs (e.g., Figure 1) because the state at time t is always recurrent (and therefore uninformative).

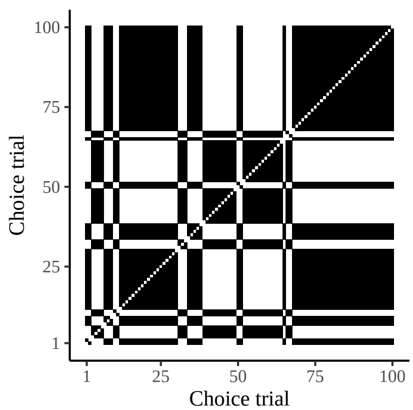

Figure 1: Example auto-recurrence plot where black points mark recurrence and white points mark non-recurrence in a choice sequence of $T = 100$ trials. The first row/column indicates whether each trial is a repeat of the choice made in Trial 1; the second row/column indicates recurrence of the choice made in Trial 2, etc.

It is not hard to see patterns in the black and white blocks of Figure 1. In our 2AFC context, each black point indicates a trial in the choice sequence where either Option A or B is repeated. Diagonal lines indicate repeated sequences of states (e.g., ABAB or AAAA), while horizontal and vertical lines are recurrent choices of the exact same state (e.g., AAAA). Thus, the large, solid blocks (e.g., upper right) indicate a long sequence of repeating the same choice over several trials.

Through interpretation, the recurrence patterns in Figure 1 help us describe behavior in this 2AFC task. Overall, this participant starts by frequently switching between the two choices, then has three subsequent periods where they prefer one choice: their first choice, the alternate choice, then back to their first choice.

Let us consider the strategy in more detail. First, the participant tries one choice in Trial 1 (i.e., Choice 1), and exhibits some switching between the two options in the first few trials (alternating black and white cells); black cells indicate repeating Choice 1 and white off-diagonal cells indicate selecting the other option, Choice 2. Next, around roughly Trials 10–30, the participant repeatedly makes Choice 1 (i.e., the large black block around the LOI indicates the repeated choice), as illustrated by the long horizontal segment for 10–30 on the x-axis at Trial 1 on the y-axis.

Third, around trials 40–70, the participant mostly makes Choice 2 (i.e., given the large black block around the LOI), and we know this is not Choice 1 because the horizontal block at Trial 1 on the y-axis is white for those trials. Finally, around trial 72, the participant returns to making Choice 1 for the remaining trials. Again, the black block around the LOI indicates long-duration recurrence; we know this is Choice 1 because the horizontal bar at Trial 1 on the y-axis during this range of the x-axis is black.

Interpreting RPs is a process of visual inspection, and the meaning of recurrence patterns must be inferred in the context of the data and domain. With this caveat in mind, we can use RPs to see far more detailed patterns than an aggregate summary of this participants' behavior would provide. For example, an aggregate summary might say the participant in Figure 1 preferred Choice 1 on 65% of the trials. However this observation does not capture the dynamic switching between an early preference for Choice 1, a middle preference for Choice 2, and a return to preferring Choice 1 for the last quarter of the trials.

Recurrence Quantification Analysis

Recurrence quantification analysis (RQA) provides well-defined statistics derived from diagonal and vertical structures in the RP (Webber & Zbilut, 1994; Zbilut & Webber, 1992). We define a sample of the core RQA statistics in Table 1, using the formalisms from Webber & Marwan (2015), and give the values for the RP in Figure 1; we briefly summarize RQA statistics here (see Webber & Marwan, 2015, for comprehensive coverage of available statistics). The recurrence rate (RR) represents the proportion of all RP points that are recurrent. In a 2AFC sequence of length $T = 100$, this can range from $RR \approx 50$ (half the trials making Option A and half the trials making Option B) to $RR \approx 100$ if a single choice

[1] Another set of possible event states could represent not just the choice made, but the outcome experienced, which would provide an additional layer of insight into choice behavior.

Table 1: Example Recurrence Quantification Statistics

Statistic	Equation	Fig. 1 Values (Max Possible)
Recurrence rate	$RR(\varepsilon,N) = \frac{1}{N^2 - N} \sum_{i \neq j=1}^{N} R_{i,j}^{m,\varepsilon}$	53.5 (100)
Percent determinism	$DET = \frac{\sum_{\ell=d_{min}}^{N} \ell H_D(\ell)}{\sum_{i,j=1}^{N} R_{i,j}}$	95.3 (100)
Average diagonal line length	$\langle D \rangle = \frac{\sum_{\ell=d_{\min}}^{N} \ell H_D(\ell)}{\sum_{\ell=d_{\min}}^{N} H_D(\ell)}$	5.5 (49)
Maximum diagonal line length	$D_{\max} = \underset{\ell}{\operatorname{argmax}}\, H_D(\ell)$	32 (99)
Laminarity	$LAM = \frac{\sum_{\ell=v_{min}}^{N} \ell H_V(\ell)}{\sum_{i,j=1}^{N} R_{i,j}}$	98.4 (100)
Average vertical line length	$TT = \frac{\sum_{\ell=v_{min}}^{N} v H_V(\ell)}{\sum_{\ell=v_{min}}^{N} v H_V^{\varepsilon}(\ell)}$	8.1 (50)

is made on all trials.[2] For Figure 1, just over half the RP is recurrent behaviors.

Diagonal lines, not including but running parallel to the LOI, indicate repeated sequences over time; this can occur in long blocks of the same choice or a matching repeated sequence. For example, a sequence perfectly alternating between Options A and B has long diagonals but no vertical structures, because the sequence over time matches, but no choice is an immediate repeat of the previous choice. The following statistics are all based on the distribution of diagonal lines: DET (proportion of $\mathbf{R}_{i,j}$ falling onto diagonal lines), $\langle D \rangle$ (average diagonal line length), and $D_{\max}$ (maximum diagonal line length). Figure 1 has a fairly deterministic pattern falling onto many relatively short diagonal lines, given the low $\langle D \rangle$ and medium $D_{\max}$ values.

Vertical[3] structures indicate staying in the exact same state over time. The statistics LAM (proportion of $\mathbf{R}_{i,j}$ falling into vertical structures) and TT (trapping time, or average vertical line length) reflect vertical structures. For Figure 1, LAM indicates a large proportion of recurrent points are in vertical structures, and the average vertical structure is relatively short. Together with the RPs, RQA provides a set of values to help describe the recurrence patterns. It is desirable to use the whole RQA vector because no single statistic uniquely describes an RP.

[2]Approximate values because the LOI is removed.

[3]In auto-recurrence, vertical and horizontal lines are equivalent.

Case Study: A 2AFC DDM Task

To illustrate the use of RPs and RQA in a DDM task, we analyze data from one experimental condition from the study by McCormick et al. (in preparation). In the chosen experimental condition, participants ($n = 101$) performed a 2AFC task where the initially worse (in expected value) choice gradually improved (in expected value) over $T = 100$ choice trials. Both choices were risky gambles with the same high and low outcomes (500 or 0 points, respectively). One option was *stationary*: the probability of receiving the high outcome was 50% over all $T = 100$ trials. The other option was *non-stationary*: the probability of receiving the high outcome was 1% on trial 1 and increased by 1% on each successive trial. Thus, the non-stationary option had a lower expected value than the stationary option during trials 1 to 49; equal expected value for trial 50; and higher expected values than the stationary option during trials 51 to 100.

Participants received no information about the choice outcomes nor their probabilities at the start of the experiment. As shown in Figure 2, they simply saw two buttons to choose between, labeled A and B; these represent the stationary and non-stationary choices, which was not communicated to participants, and the labels were randomly assigned. Once participants chose between option A and B, they received feedback containing the outcome of their last choice, a reminder of that choice, and the total points they had accumulated up to that trial. Participants completed the study online, and were compensated with a base payment plus an individualized bonus that was a fraction of the points the participant had accumulated over the $T = 100$ trials.

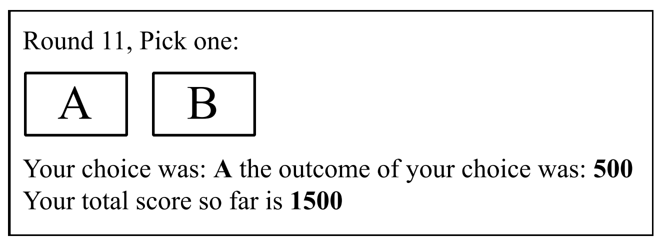

Figure 2: Representation of an example choice trial in the McCormick et al. task. Participants received feedback about their previous choice and current accumulation of points.

A fully informed, optimal decision maker (one who maximized their expected number of points on each trial) would choose the stationary option for trials 1–49, and then adapt to choose the non-stationary option for trials 51–100. Thus, a simple aggregate metric of performance is the rate of maximizing choices (max-rate), which can be calculated over all 100 trials, or separately for trials 1–49 (max-rate before the expected value switch) and trials 51–100 (max-rate after).

As reported in McCormick et al. (in preparation), participants were divided between adaptive and non-adaptive groups using the max-rate before and max-rate after scores. Adaptive individuals were of two types. Agile decision makers made maximizing choices more frequently both before and after the switch, indicating successful adaptation of their choices to the switch in expected value. Although a minority, clumsy decision makers made non-maximizing choices both before and after the switch, suggesting adaptation to expected value shifts but not in a maximizing manner. Non-adaptive participants also fell into two types. Fortunate non-adaptive decision makers made maximizing choices after the switch, but not before, suggesting they did not adapt to changes in the decision environment and happened to choose the ultimately-maximizing option from the start (when it was not maximizing). Conversely, rigid non-adaptive decision makers made maximizing choices before the switch, but not after, suggesting they persistently chose the initially-maximizing option even when it was no longer maximizing.

The authors observed that maximization rates revealed a reasonable number of adaptive decision makers (a majority agile) and an over-influence of initial experiences among non-adaptive decision makers (with many rigid participants). However, further insights about the dynamics of the strategies used by adaptive and non-adaptive decision makers were missing.

RP and RQA Analyses

To apply RP and RQA to the McCormick et al. (in preparation) data, we defined two possible event states: whether participants made a stationary choice or non-stationary choice each trial. Analyses were completed using R for statistical computing (R Core Team, 2013) and the *crqa* library (Coco & Dale, 2014).[4] In the sample data (one experimental condition out of six total), no participants selected the same choice for the entire 100 trials. To aid visual interpretation, the RPs are colored to differentiate stationary and non-stationary choice states. This coloring does not change the RP and RQA computations, which still reflect the total recurrence.

Results

The RPs for all 101 decision makers are plotted as a "recurrence quilt" in Figure 3, revealing a variety of individual choice patterns. These quilts have been ordered by overall maximization rate (provided in parentheses above each individual's RP). A number of characteristics across the quilt are immediately noticeable. The upper few lines of the quilt contain RPs predominantly light blue in color, indicating a preference for the stationary option (which was initially maximizing). A number of plots have a solid (light or dark) appearance; these RPs reflect long periods of recurrence, lasting nearly the whole task (e.g., RPs 23, 39). These participants seem stay with a particular choice. At this level of data, we might call this a static preference behavior pattern, and seek to better understand how static patterns, which seem non-adaptive, correspond to theoretical definitions of adaptation in the DDM task context.

[4] Additional hyper-parameters for the *crqa* function were the following: delay = 1, embed = 1, rescale = 1, radius = 0.0001, normalize = 0, mindiagline = 2, minvertline = 2, tw = 1 (where tw is the Theiler window setting that removes the Line of Incidence).

Other RPs show much more variable patterns, reflecting frequent shifting between choices (e.g., RPs 11, 55, 83). Then throughout there are RPs that show recurrent blocks alternating between dark and light. These participants seem to prefer making one choice for some time, then switch to repeatedly making the other, and then switching back at least once more (e.g., RPs 4, 42, 47, 97). This variability in choice sequences likely captures differences in individual strategies: some sequences perhaps indicate shifts in preference over time, some simply show a tendency to explore, and others capture more adaptive behaviors consistent with optimal or other rational task strategies over time. Specifically, RPs toward the bottom of the quilt (participants with higher maximization rates) show a shift from light to dark blue, indicating an adaptive shifting consistent with the task design and optimal strategy.

Non-Adaptive Decision Makers The RPs quickly illustrate that non-adaptive participants with low maximization rates exhibited choice patterns that vary in the amount of alternation between choice options. Figure 4 shows four recurrence plots from non-adaptive decision makers. All four participants had similar overall maximizing rates (provided in Table 2) and illustrate both rigid (RPs 24 and 15) and fortunate (RPs 39 and 49) choice patterns.

RPs 24 and 39 reflect recurrence of a single choice throughout, while the RPs 15 and 49 capture more frequent switching between options. RQA for RPs 24 and 39 (Table 2) reflect high values for both diagonal ($\langle D \rangle$, $D_{\max}$) and vertical structures (LAM and TT). These RPs show nearly maximal RR and DET values, confirming they are in the same recurrent state nearly the entire time. We can describe these decision makers as engaged in limited exploration of their options at the start of the task, followed by sticking with one option for the remaining trials. They seem non-adaptive at both the aggregate score and choice sequence levels of behavior.

RQA for RPs 15 and 39 show lower proportions of recurrent points (RR in Table 2), the larger amount of white space reflects more choice switching. Their statistics also show short average diagonal ($\langle D \rangle$) and vertical (TT) line lengths. Even with frequent switches, these two decision makers exhibit an overall preference for one of the options—as shown by the dominant colors in each RP—and do not change their preferences with the expected value shift. Thus, these decision makers use non-adaptive strategies, but with far more frequent alternation and stochastic recurrence patterns than the participants in RPs like 24 and 39.

Adaptive Decision Makers Figure 5 shows four RPs from Adaptive decision makers. Three of the four examples (RPs 101, 99, 98) have agile choice patterns: a preference for the stationary option in the first half of trials, and a preference for the non-stationary option in the second half of trials (see Table 2). In addition to meeting the criteria for agile choice patterns, these participants have three of the highest overall max-rates. Thus, these RPs are consistent with the optimal strategy, but are more variable than perfectly optimal.

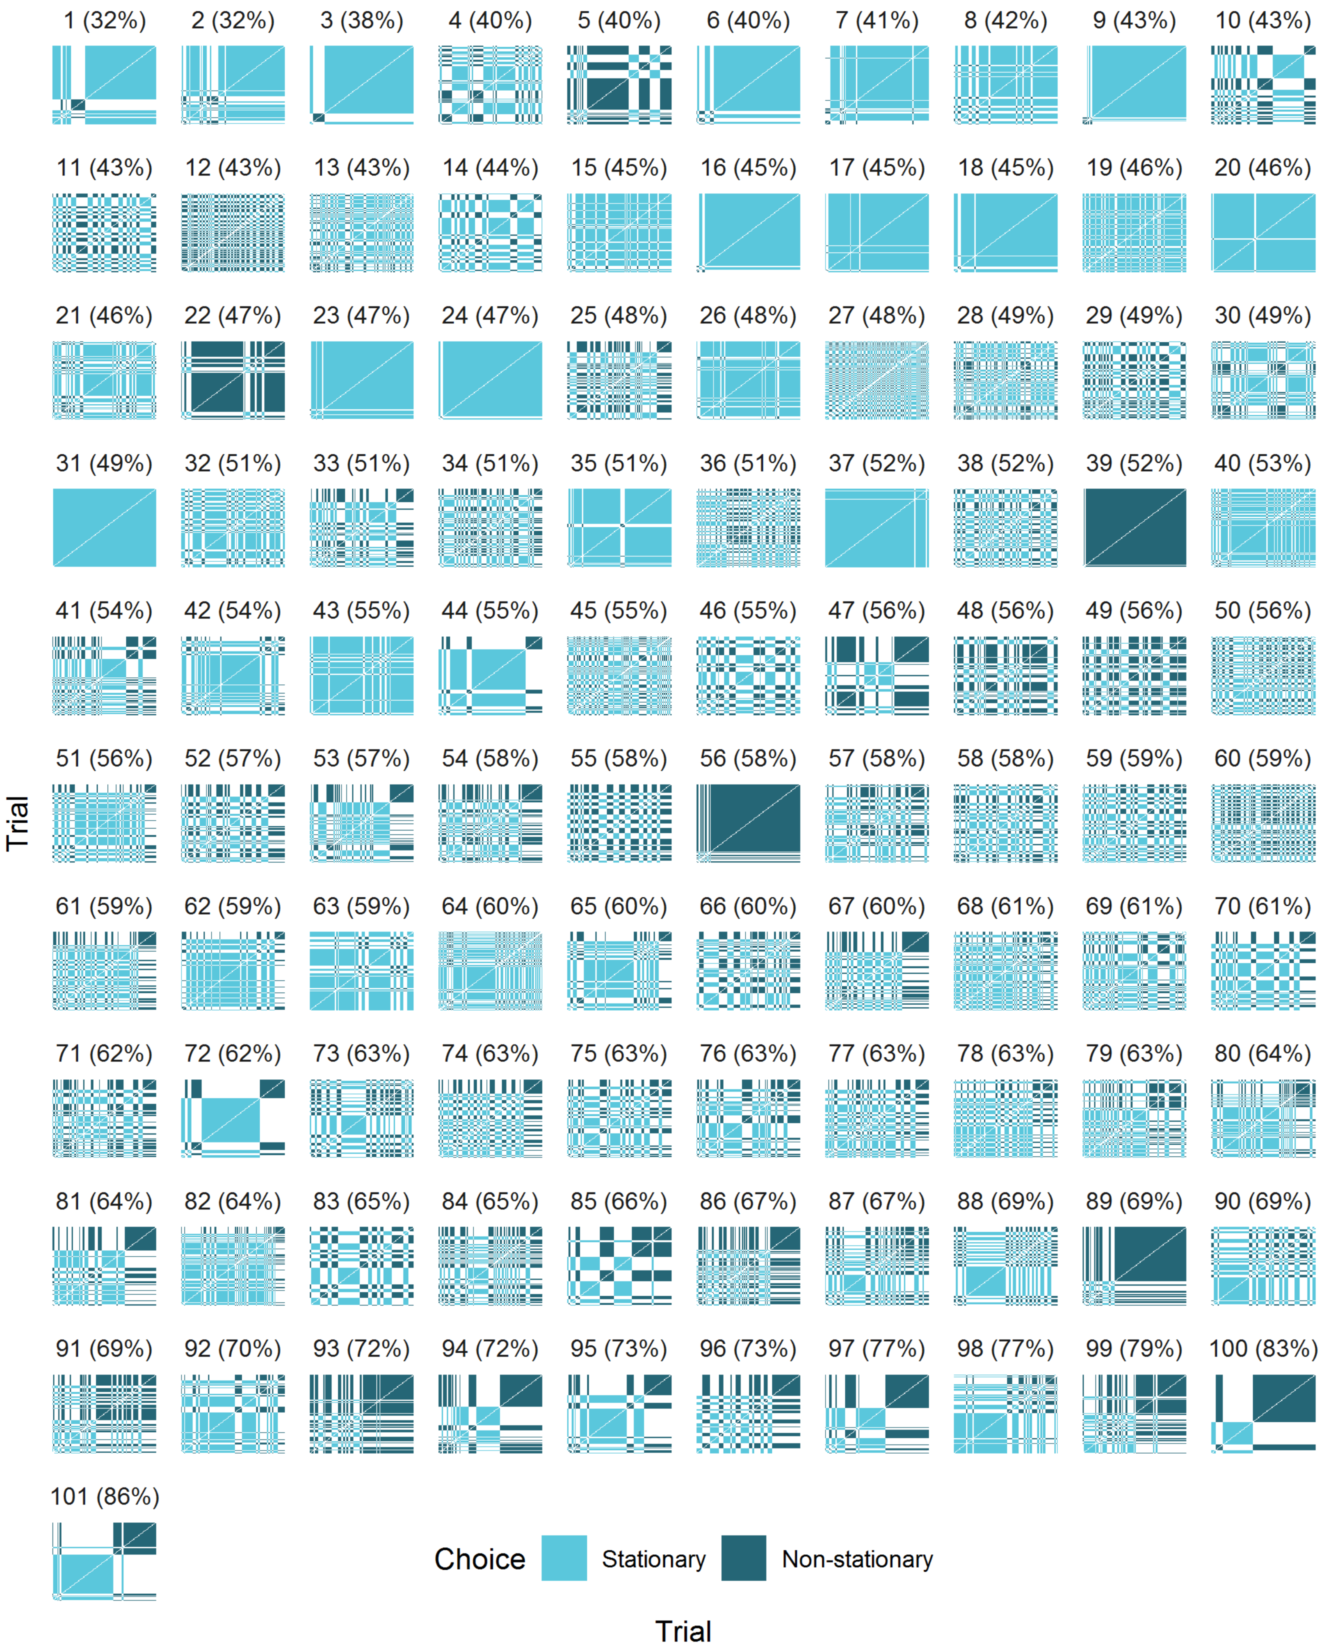

Figure 3: RPs for individual participants ordered by overall maximization rate (provided in parentheses). White points, indicate non-recurrence, light blue indicates recurrence of the stationary choice, and dark blue indicates recurrence of the non-stationary choice.

Table 2: RQA for Exemplar Recurrence Plots

	Maximization Rates			RQA Statistics					
ID	Overall	Before	After	RR	DET	$\langle D \rangle$	$D_{\max}$	LAM	TT
Figure 4: Non-adaptive Decision Makers									
RP24	0.47	0.96	0.00	95.1	99.9	25.1	95	99.9	33.2
RP39	0.52	0.02	1.00	95.1	98	48.5	96	98.9	49.3
RP15	0.45	0.79	0.12	72.1	92.8	5	20	95.7	7.5
RP49	0.56	0.46	0.66	51.4	81.4	3.6	21	87.6	5.2
Figure 5: Adaptive Decision Makers									
RP101	0.86	0.94	0.78	50.0	94.3	9	49	97.0	13
RP99	0.79	0.67	0.92	52.9	82.8	4.3	27	89.6	6.2
RP98	0.77	0.96	0.58	56.2	88.1	4.3	31	93.3	6.2
RP10	0.43	0.48	0.40	49.2	93.2	4.23	29	98.3	5.8
Optimal Strategy									
	1.00	1.00	1.00	49	100	25.5	49	99.9	25.5

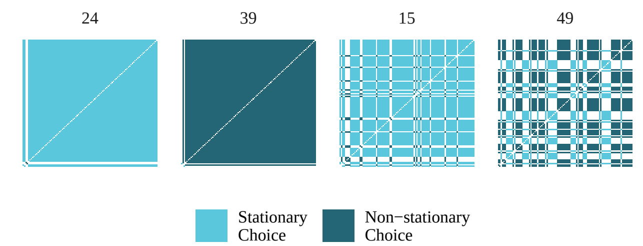

Figure 4: Recurrence plots for the sequence of choices made by non-adaptive decision makers.

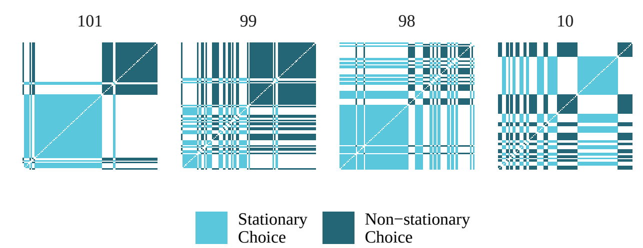

Figure 5: Recurrence plots for the sequence of choices made by adaptive decision makers.

RP 101 shows long periods of recurrence, with higher $D_{\max}$ and DET for this group in Table 2. Conversely, RPs 99 and 98 show more variability, and related smaller recurrence patterns than RP 101, but identical average vertical and diagonal lengths to each other. RP 99 exhibits variability in the early trials before shifting to a longer period of non-stationary recurrence later. RP 98 exhibits the opposite, with variability in the later portion of the experiment and recurrence of longer duration earlier. RP 10 also demonstrates an adaptive strategy, but a non-optimal one–one of the few clumsy choice patterns observed. On average, RP 10 prefers the non-stationary option early, and notably shows a period of recurrence in the stationary option later, when it was not the maximizing option. RP 10's RQA reflects a medium amount of overall RR and low average diagonal and vertical line lengths; the maximum diagonal is of a medium length, reflecting a least one recurrent sequence of medium length.

Defining New Types of Adaptation

We remind readers of the pattern described for Figure 1; this is RP 47 in Figure 3. RP 47 also seems to illustrate a type of adaptive behavior observed in several participants but that was not distinguishable by McCormick et al. (in preparation) at the aggregate score level. We call this a "U-shaped" adaptation strategy: it reflects a move from one preference to another and back to the first. We also see this in RPs 41, 67, 72, 94. While this U-shape clearly reflects shifting preferences, the non-monotonic RP patterns may not be connected to the underlying (monotonic) changes in expected value, or may be mediated by the gamble outcomes together with the changing task dynamics. These and other types of adaptive patterns are of interest for future study of human choice patterns.

Discussion

We introduce recurrence plots and recurrence quantification analyses to gain more insight about human strategies in DDM tasks. Although RPs and RQAs have been used in other fields, they have not previously been used to analyze choice sequences. We demonstrate RPs and RQA in a 2AFC DDM task where the expected value of the options changes over time. Our case study illustrates the benefits of using RQAs in situations where typical aggregate statistics would obscure the diversity of human choice behavior. Visual inspection and quantification of the recurrent patterns in our case study's choice sequences revealed multiple, distinguishable choice patterns among individuals who may be categorized similarly when using aggregate statistics. RPs and RQA highlighted nuances in both adaptive and non-adaptive strategies. These recurrence patterns increase the detail available to researchers, and reveal the diversity of decision strategies that theoretical and computational models of choice seek to explain and simulate. We believe studying this diversity can generate new insights about decision making strategies that are currently being overlooked in our data analysis practices.

Theoretical Implications The variety of choice patterns visualized in the RPs raises a simple question: what recurrence patterns are predicted by existing theories and models of dynamic decision making? We propose that future work should analyze the predictions of computational DDM models using RPs and RQA, including simple decision rules such as *Win-Stay-Lose-Shift* (Biele et al., 2009) and the *Hot-Stove* (Denrell & March, 2001). Differences between models and humans at the level of choice sequence patterns can be used to refine theories accordingly.

Importantly, observed RP patterns provide specific targets for assessing theory predictions that would not be salient in aggregate measures. We observed that individuals exhibiting similar maximization performance could differ in how they switch between the two options. Some decision makers switch rarely and some switch frequently over all trials, even when they show an overall preference for one choice. Some decision makers switch frequently only before the change in expected value, and others do so only after. These patterns raise the question of whether the frequency of switches, and the timing of those switches are choice patterns captured in existing theories and computational models, especially given the importance of exploration in DDM environments. Decision rules such as *Win-Stay-Lose-Shift* might predict that the observed differences are due to experienced outcomes rather than individual differences. Experimentally holding experienced outcomes relatively constant or defining event states to include outcomes would allow researchers to use RPs and RQAs to test whether such individual variation persists.

Applications RP and RQA can be extended from human data to model-simulated choice sequences, testing whether models produce the same richness and variety of behavioral patterns. In addition to auto-recurrence (the focus of the present work), cross-recurrence (CRQA) can be used to examine relationships between two different time series, which extends the study of DDM strategies to multi-player settings. Researchers might compare two people completing the same task (e.g., game theoretic settings), a person to a modeled strategy, a person and computational model completing the same task, or two models to each other. RP and RQA can also extend from 2AFC to any number of choice options, allowing study of multi-choice tasks like rock-paper-scissors or the Iowa Gambling Task. There is rich opportunity to expand applications of recurrence analysis in the decision making domain.

Acknowledgements

This research was supported in part by an appointment to the U.S. Air Force Research Laboratory Research Participation Program by the Oak Ridge Institute for Science and Education, and in part by a Seedling Award from the 711th Human Performance Wing Chief Scientist. Cleotilde Gonzalez was supported by the Army Research Office, Network Science Program, Award Number: W911NF1710431.

References

Avrahami, J., & Kareev, Y. (2011). The role of impulses in shaping decisions. *Journal of Behavioral Decision Making*, *24*(5), 515–529.

Biele, G., Erev, I., & Ert, E. (2009). Learning, risk attitude and hot stoves in restless bandit problems. *Journal of Mathematical Psychology*, *53*, 155–167. doi: 10.1016/j.jmp.2008.05.006

Brehmer, B. (1992). Dynamic decision making: Human control of complex systems. *Acta Psychologica*, *81*(3), 211–241.

Coco, M. I., & Dale, R. (2014). Cross-recurrence quantification analysis of categorical and continuous time series: An R package. *Frontiers in Psychology*, *5*, 510.

Denrell, J., & March, J. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science*, *12*(5), 523–538.

Eckmann, J.-P., Kamphorst, S. O., & Ruelle, D. (1987). Recurrence plots of dynamical systems. *Europhysics Letters*, *4*(9), 973–977.

Edwards, W. (1962). Dynamic decision theory and probabilistic information processings. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *4*(2), 59–74.

Gonzalez, C., Fakhari, P., & Busemeyer, J. (2017, August). Dynamic Decision Making: Learning Processes and New Research Directions. *Human Factors*, *59*(5), 713–721.

Lejarraga, T., Lejarraga, J., & Gonzalez, C. (2014). Decisions from experience: How groups and individuals adapt to change. *Memory & Cognition*, *42*(8), 1384–1397.

Marwan, N., Romano, M. C., Thiel, M., & Kurths, J. (2007). Recurrence plots for the analysis of complex systems. *Physics Reports*, *438*(5-6), 237–329.

McCormick, E. N., Cheyette, S. J., & Gonzalez, C. (in preparation). *Choice Adaptation to Changing Environments: Trends, Feedback and Awareness of Change.*

R Core Team. (2013). R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria. Retrieved from `http://www.R-project.org/`

Rakow, T., & Miler, K. (2009). Doomed to repeat the successes of the past: History is best forgotten for repeated choices with nonstationary payoffs. *Memory & Cognition*, *37*(7), 985–1000.

Webber, C. L., Jr, & Marwan, N. (Eds.). (2015). *Recurrence quantification analysis: Theory and best practices*. Springer.

Webber, C. L., Jr, & Zbilut, J. P. (1994). Dynamical assessment of physiological systems and states using recurrence plot strategies. *Journal of Applied Physiology*, *76*(2), 965–973.

Zbilut, J. P., & Webber, C. L., Jr. (1992). Embeddings and delays as derived from quantification of recurrence plots. *Physics Letters A*, *171*(3-4), 199–203.