OPEN

SUBJECT AREAS:

APPLIED PHYSICS

DECISION

DYNAMICAL SYSTEMS

APPLIED MATHEMATICS

Received
12 December 2012Accepted
18 June 2013Published
17 July 2013

Correspondence and
requests for materials
should be addressed to
D.O.H. (denis.ohora@
nuigalway.ie.)

Local dynamics in decision making: The evolution of preference within and across decisions

Denis O'Hora¹, Rick Dale², Petri T. Piiroinen³ & Fionnuala Connolly⁴

¹School of Psychology, National University of Ireland, Galway, University Road, Galway, Ireland, ²School of Social Sciences, Humanities and Arts, University of California Merced 5200 North Lake Rd., Merced, CA 95343, ³School of Mathematics, Statistics and Applied Mathematics, National University of Ireland, Galway University Road, Galway, Ireland, ⁴Harvard School of Engineering and Applied Science, Harvard University 29 Oxford Street, Cambridge, MA 02138.

Within decisions, perceived alternatives compete until one is preferred. Across decisions, the playing field on which these alternatives compete evolves to favor certain alternatives. Mouse cursor trajectories provide rich continuous information related to such cognitive processes during decision making. In three experiments, participants learned to choose symbols to earn points in a discrimination learning paradigm and the cursor trajectories of their responses were recorded. Decisions between two choices that earned equally high-point rewards exhibited far less competition than decisions between choices that earned equally low-point rewards. Using positional coordinates in the trajectories, it was possible to infer a potential field in which the choice locations occupied areas of minimal potential. These *decision spaces* evolved through the experiments, as participants learned which options to choose. This visualisation approach provides a potential framework for the analysis of local dynamics in decision-making that could help mitigate both theoretical disputes and disparate empirical results.

Choosing between two or more available options is an activity extended in time, and during that time, decisions evolve. The empirical literature on decision making distinguishes between perceptual and value-based decisions. Perceptual decisions (e.g., “Did I just smell coffee?”) typically become more accurate the longer they take, up to some threshold (for reviews, see^{1,2,3}) and approximate optimality^{4,5,6}. Value-based decisions (e.g., “Would I like a cup of coffee?”), in contrast, exhibit a number of well-known idiosyncrasies^{7–11}. To date, these effects have, in the main, been explained by theories concerned with the distribution of outcomes based on utility^{12,13}. Recently, however, there has been a shift towards the development of models, some informed by perceptual choice models, that attempt to characterize processes occurring during decision making^{14–24}. Evidence of these processes has been derived from patterns of accuracy and response time, but there has been relatively little research on the behavioral correlates of these processes. The current paper reports an investigation of sub-decision dynamics and a new visualisation method to facilitate theoretical development in this area.

Important features of decision dynamics can be unveiled by tracking the motor execution that accompanies it. Specifically, researchers have examined characteristics of response trajectories (the path from initiation of a response to its completion) for decisions under varying conditions. Studies have employed a variety of response equipment: the computer mouse^{25–31}, the Nintendo Wii remote^{32,33} and a motion-tracking system fixed to digits³⁴ see^{35–37} for reviews). For example, Dale *et al.*²⁶ found that participants showed slower and longer response trajectories when they had to choose between “MAMMAL” and “FISH” to categorize “sea lion” than to categorize “salmon” or “lion”. In social categorization, participants take longer and show more divergent trajectories when categorizing relatively androgynous faces (with ambiguous sex features) as “MALE” or “FEMALE,” than when categorizing more typical male or female faces²⁹. In these cases, categorizing more typical exemplars was easier (faster and more direct) than categorizing exemplars that were atypical or more similar to the incorrect category, a characteristic clearly observable from the response trajectories. These findings are not limited to categorization. Spivey *et al.*²⁵ demonstrated that phonological competition between available choices was observed in response trajectories. When asked to “Click the Candle”, participants’ responses were slower and less direct when required to choose between a picture of a candle and a candy, than when asked to choose between a picture of a candle and a pickle.

In line with some models of perceptual decision making, the foregoing studies suggest that early in a decision, distributed neural representations may be partially consistent with multiple outcomes^{35,38,39}. As time elapses, information accumulates until these unstable patterns of neural activation dynamically evolve into more stable patterns associated with one of the available outcomes. Competition between outcomes increases the duration of the early unstable phase during which multiple responses remain possible. This gives rise to deflection in response trajectories from the shortest direct path. For example, in a study of attitudes, Wojnowicz *et al.*³⁹ found that participants responded with greater deflection when responding in opposition to social stereotypes. In addition, in these trials, velocity exhibited early instability followed by a compensatory increase in velocity, which was predicted by a computational model of decision making^{40,19}. Indeed, the gradual accumulation of information (sequential sampling) that may underlie the action dynamics effect is a common feature of many current decision making models⁴¹.

The foregoing literature suggests that great insights may be gained by tracking decisions as they unfold. To date, researchers have investigated neural activity during decision making in an attempt to characterize the evolution of perceptual^{42,43} and value-based^{44–46} decisions, but few studies have investigated changes in behavior during value-based decisions^{47,48,12}. Koop and Johnson examined response trajectories during a version of the Iowa Gambling Task and found that response trajectories towards Good decks (rewards greater than losses) gradually became more direct across blocks of decisions, whereas response trajectories towards Bad decks (losses greater than rewards) did not.

One way to conceptualize how we make decisions is to consider the available outcomes as attractors in a *decision space*^{49,35}. We make a decision when our behavior reaches the vicinity of one of these attractors. Killeen⁴⁹ provided an account of conditioning as basic physics within which he proposed that behavior could be understood as movement in a behavior space in which incentives or reinforcers lie at centers of basins of lowered potential. Responses that are reinforced become more probable and faster over time as the trajectories in behavior space are warped through contact with these basins. In a similar vein, Spivey and Dale³⁵ provided a model of decision dynamics that proposed that response trajectories in a binary choice may be understood as a path through a two-well attractor landscape. This landscape, the decision space, is an expression of the cognitive evaluation of the available outcomes in a particular decision and the motor execution of the response influenced by this evaluation. More recently, dynamic models of perceptual decision making inspired by neural models also emphasise the role of attractor dynamics in decision making^{50,24,51}. From this perspective, competing outcomes (i.e., valued choices) induce divergent, slow, and longer action dynamics when those outcomes are near equal in their potentiality.

In the experiments described here, participants played a simple game to gain points, which constituted a forced-choice simple discrimination task. In each experiment, some choices earned more points than others and participants gradually learned which choices earned them the most points. Each decision provided two choices and, in the majority of decisions, one choice earned more points than the other (High/Low, e.g., 20 points/5 points). In some decisions, participants were required to decide between two choices that were both worth the same value, either a specific higher-point value (High/High, e.g., 20 points/20 points) or a specific lower-point value (Low/Low, e.g., 5 points/5 points). The relative reward in both of these decision situations is 1, because the available choices in both cases earn precisely the same number of points ($20/20 = 5/5 = 1.0$). In light of the research summarized thus far, one might expect that, with equipotential response options, (a) there ought to be competition in both situations, and (b) the degree of competition should be similar.

Figure 1 | Experimental task and learning data. The top panel provides a schematic of the experimental procedure. The lower left panel depicts mean probability of choosing a High option in a High/Low decision across participants on consecutive blocks of 6 decisions in each of the three experiments (error bars denote standard errors). The shaded area indicates probability below 80%, the threshold used to infer that participants learned to choose the High option, and the dashed line indicates 50%, the choice probability expected by chance. The lower right panel is a bubble plot of sensitivity to relative reward measured by the $\log_2$ probability of choosing a High option. As relative reward increased (the horizontal axis), more participants reliably chose the High choice in High/Low decisions. The size of each point is determined by the number of participants that obtained that value controlled for the number of participants in that experiment (i.e., probability density of that value within each experiment). The shaded area and black dashed line correspond to the same values as in the left panel. The blue dashed line indicates the mean $\log_2$ probability at each level of relative reward.

In actuality, the decision dynamics, both within and across decisions, were markedly different in these two conditions.

Results

Across three experiments, the ratio of the high-point reward to the low point reward was increased: Experiment 1 (7/5; $n = 34$), Experiment 2 (10/5; $n = 37$) and Experiment 3 (20/5; $n = 55$). Learning was faster and more reliable as the high-point value increased. On average, participants learned more quickly to choose the higher of the two options in the 20/5 experiment than in the other two experiments; after just 12 decisions, participants chose the high-value symbol on over 80 percent of High/Low decisions. By halfway through the experimental session (18 decisions), the high-value symbol was chosen reliably in High/Low decisions in all three experiments (see Fig. 1, bottom-left panel). The high-point value also affected the proportion of participants who learned to consistently choose the High options. The bottom-right panel of Fig. 1 provides the distribution of high-point/low-point choice ratio across participants in each experiment. The $\log_2$ of the ratio of the probability of choosing the high choice to the probability of choosing the low choice on any High/Low decision (i.e., $\log_2(p_{High}/p_{Low})$) is employed as a measure of the degree to which the participant reliably chose High options in High/Low decisions (in the literature on basic learning principles, the $\log_2$ ratio is employed as an index of relative allocation of behavior to two independent responses and it has been shown to be sensitive to relative magnitude of reinforcement⁵²).

Table 1 | Comparison of fits of binomial linear mixed effects models to predict High choices in High/Low decisions

Model	No.	Df	AIC	BIC	logLik	Chisq	Chi Df	Pr(>Chisq)
Intercept only	1	2	2757.19	2769.20	-1376.60			
Model 1 + Expt	2	4	2754.05	2778.08	-1373.03	7.14	2	0.0281
Model 1 + Expt*Trial	3	7	2463.53	2505.58	-1224.77	296.52	3	0.0000
Model 1 + Expt*log(Trial)	4	7	2405.13	2447.18	-1195.57	58.40	0	0.0000

Mean $\log_2(p_{\text{High}}/p_{\text{Low}})$ increased as relative reward increased across experiments. There was a weak but significant positive correlation ($r(126) = 0.22, p = .012$) between the log proportion of High to Low responses, $\log_2(p_{\text{High}}/p_{\text{Low}})$, and the log proportion of reward magnitudes, $\log_2(M_{\text{High}}/M_{\text{Low}})$. We employed a series of binomial linear mixed effects models to further analyse this effect (see Table 1). In line with the previous observation, the addition of Experiment as a predictor improved the intercept-only model. In this model, accuracy was significantly lower in the 7/5 experiment than the 20/5 experiment, $b = -0.5800, z = -2.589, p = .0096$, and it was marginally significantly lower in the 10/5 experiment than the 20/5 experiment, $b = -0.4086, z = -1.842, p = .0655$. The model was improved by adding the learning effect across decisions and log transforming the effect of trials/decisions better fit the ceiling effect on accuracy that can be seen in Fig. 1. In the final model, $\log(\text{Trial})$ was a strong significant predictor, $b = 1.1373, z = 11.9905, p < .0001$, but the interaction effect of Experiment and $\log(\text{Trial})$ was not significant. There no significant difference in the change in accuracy across trials between the 20/5 experiment and the 10/5 experiment, $b = -0.1454, z = -1.0232, p = .3062$, or between 20/5 and 7/5, $b = -0.0666, z = -0.4635, p = .6430$.

Data selection for trajectory analysis. During decisions, the positional (pixel) coordinates of the response trajectories were recorded. In line with previous work on action dynamics (e.g.²⁵), we then assumed a coordinate system in which the initiation of each decision trajectory constituted the origin. The horizontal (x) and vertical (y) axes of the computer screen constituted the dimensions of this 'decision space'.

In the analyses of the response characteristics, we filtered out participants who did not appear to learn. Only participants who chose the high-value symbol in High/Low decisions on 80% of the last 12 decisions ($n_{7/5} = 26, n_{10/5} = 26, n_{20/5} = 48$) in each experiment were included. These participants were deemed to have demonstrated that they had learned the crucial distinction in the experiment. The remaining participants were considered not to have learned the High/Low distinction well enough for their responses to be comparable to those who had. The greatest per-experiment proportion of participants satisfied this response requirement in the 20/5 experiment (88%), followed by the 7/5 experiment (77%) and then the 10/5 experiment (71%). The fact that a greater proportion of participants reliably learned the distinction in the 7/5 experiment than the 10/5 experiment was unexpected as overall learning (mean p_{High}) was better on average in the 10/5 experiment. In the following analyses, we also excluded trajectories in which participants chose a low-value symbol in a High/Low decision, as we considered such decisions to be errors and errors are known to exhibit different

response characteristics (e.g., reaction times¹⁶) from expected or correct decisions.

Reaction time. Reaction time patterns were not affected by the increase in relative reward across experiments. In the first 12 decisions, reaction times were relatively similar across decision types. In High/Low and High/High decisions, reaction times gradually decreased throughout the experiments, but, in Low/Low decisions, reaction times remained consistently high. To examine the effects of decision type on reaction times across trials, we fit a series of linear mixed effects models, using maximum likelihood estimation on log-transformed reaction times (see Table 2). There was no significant effect of relative reward (Experiment) on log reaction time. Reaction time decreased significantly across decisions, $b = -0.0087, t = -14.98, p < .0001$, but there was no main effect of decision type. The change in reaction times across Low/Low decisions was significantly different, $b = 0.0081, t = 6.75, p < .0001$, from the change across High/Low decisions, the change across High/High decisions was not $b = -0.0004, t = -0.38, p = .7063$. That is, as can be seen in Fig. 2, exposure to the learning context reduced the reaction times of High/Low and High/High decisions, but not Low/Low decisions.

Maximum deviation. For each trajectory, we extracted how much a trajectory deviated from the endpoint response option. This is an additional measure of cognitive conflict: If a trajectory has high maximum deviation, it means participants are moving their computer mouse closer to the alternative response. With a lower maximum deviation, it reflects a more direct movement towards their choice (exhibiting less conflict, and more confidence). Similarly to the reaction times, maximum deviation was relatively similar across decision types in the first 12 decisions and then decreased in High/Low and High/High decisions in the remaining decisions. Maximum deviation during Low/Low decisions remained consistently high throughout the experiment with a slight increasing trend. To statistically investigate these patterns, we also used a series of linear mixed effects models (see Table 3). Findings were similar to those previously found for reaction times. There was no significant effect of relative reward (Experiment) or main effect of Decision type, but maximum deviation decreased significantly across decisions, $b = -0.7594, t = -3.26, p = .0011$. The change in maximum deviation across Low/Low decisions was significantly different, $b = 1.6263, t = 3.37, p = .0008$, from the change across High/Low decisions, the change across High/High decisions was not, $b = 0.1921, t = 0.40, p = .6901$. As participants progressed through the experiment, the deflection towards the unchosen symbol decreased in High/Low and High/High decisions suggesting reduced conflict between choices. This did not occur in Low/Low decisions, suggesting

Table 2 | Comparison of fits of linear mixed effects models used to predict log transformed reaction time

Model	No.	df	AIC	BIC	logLik	Test	L.Ratio	p-value
Intercept only	1	4	1283.34	1307.69	-637.67			
Model 1 + Decision	2	6	1158.34	1194.88	-573.17	1 vs 2	128.99	0.0000
Model 1 + Decision * Trial	3	9	877.59	932.39	-429.79	2 vs 3	286.75	0.0000
Model 1 + Decision * Trial * Expt	4	21	893.04	1020.91	-425.52	3 vs 4	8.55	0.7411

persistent conflict in these decisions throughout the experiment. Results, by experiment, are displayed in Fig. 2.

Horizontal velocity. Within a decision, the available outcomes were presented at the extreme left and right of the experimental display. As a consequence, competition between the available outcomes may be expressed in the horizontal component of response velocity ($\frac{dx(t)}{dt}$).

To further analyze the competition between outcomes, point to point velocity in the x direction was calculated and then interpolated to 101 time steps to normalize response time and provide an average velocity profile in the last 18 decisions for each condition in all three experiments (see upper panel of Fig. 3). In all three experiments, horizontal velocity towards the eventual choice was slower to increase in the Low/Low condition than in the other two conditions, providing evidence of persistent early conflict during these decisions. It is worth remembering that the mean RT of Low/Low decisions was significantly greater than High/Low or High/High decisions, so the evolution of velocity seen in this figure underestimates the differences in real time (Low/Low decisions took approximately 300 ms longer in real time on average). In binned segments of trajectory (Fig. 3, lower panel), horizontal velocity during the third quintile (40–60%) was significantly lower ($p < .05$) in Low/Low than both High/Low and High/High in all three experiments. Otherwise, the evolution of x velocity within decisions was quite similar for equivalent conditions across the three experiments.

Response trajectory shape. If it is harder to make a given response because participants are experiencing conflict, then the trajectories of those responses should be more divergent: They should spend more time between choices, and in some cases may even move towards the alternative choice before making a response. We interpolated each trajectory into 101 time steps so they could be overlaid into an average plot (see^{25,26}). Trajectories of decisions to the left-hand choice were reflected in the vertical axis at the origin so that all trajectories ended at the right-hand choice for ease of comparison.

The final 18 decisions were used because, by this point in all three experiments, participants were choosing the high point choice on 80% of all High/Low decisions. This is shown separately for each experiment in the lefthand column of Fig. 4. In all three experiments, trajectories during Low/Low decisions were very different to those observed in High/Low and High/High decisions. In decisions that included a high-point choice, trajectories were relatively direct, whereas trajectories in Low/Low decisions exhibited considerably larger deflections towards the unchosen symbol. These deflections suggest stronger and more persistent competition between the available response options during Low/Low decisions.

The most conspicuous difference was between High/High and Low/Low decisions, in which the relative magnitude of the points available for each choice was the same. In a statistical comparison of x coordinates (paired t tests) at each interpolated time step²⁶, Low/Low trajectories travelled significantly closer to the unchosen option than High/Low trajectories for considerable portions of the trajectories (7/5: steps 52 to 101; 10/5: steps 55 to 82; 20/5: steps 57 to 93) and closer than High/High trajectories for similar portions (7/5: steps 41 to 90; 10/5: steps 55 to 75; 20/5: steps 50 to 94). In addition, in Experiments 1 (7/5) and 3 (20/5), small portions of the High/High decision trajectories were significantly closer to the eventual choice than High/Low trajectories (7/5: steps 29 to 53; 20/5: steps 44 to 52; departures of less than 8 consecutive points were not considered significant to preserve family-wise error-rate), offering some weak but significant suggestion that High/High decisions may have been easier than High/Low decisions. It is also worth noting that this facilitation of High/High over High/Low was earlier in the trajectories than the conflict effects observed in the Low/Low trajectories.

Modeling decision space. By connecting empirical action dynamics with systems of differential equations, insightful new visualizations of decision spaces are possible. We sought to characterize the cognitive landscape from which decisions emerge. Using momentary velocities and accelerations derived from the positional coordinates in the decision trajectories, we inferred a potential field that captured

Figure 2 | Measures of choice conflict across blocks of 12 decisions. L/L denotes Low/Low, H/L denotes High/Low and H/H denotes High/High decision types. (A) The top panel depicts the mean reaction time and bootstrapped confidence intervals (calculated using the R package ggplot2⁵⁹) in each block of 12 trials in each experiment. (B) The lower panel depicts the same values for maximum deviation, the furthest point in a trajectory from the straight line from the initiation of a response to its completion.

Table 3 | Comparison of fits of linear mixed effects models used to predict maximum deviation

Model	No.	df	AIC	BIC	logLik	Test	L.Ratio	p-value
Intercept only	1	4	39931.56	39955.91	-19961.78			
Model 1 + Decision	2	6	39870.32	39906.86	-19929.16	1 vs 2	65.23	0.0000
Model 1 + Decision * Trial	3	9	39859.73	39914.53	-19920.87	2 vs 3	16.59	0.0008
Model 1 + Decision * Trial * Expt	4	21	39872.18	40000.04	-19915.09	3 vs 4	11.56	0.4817

characteristics of this *decision space*. Low choices in High/Low decisions were included in these analyses.

Using the computer-mouse coordinates collected during choices, it was assumed that motion during a choice can be described by a potential field given by the function $V(x, y) = V_x(x) + V_y(y)$, where x and y are the screen coordinates. As a first approximate, it was assumed that motion in the x -direction was independent of motion in the y -direction, hence $V(x, y)$ could be separated into the two functions $V_x(x)$ and $V_y(y)$. Thus, the most simple system of second-order ordinary differential equations describing the motion could be written

$$\frac{d^2x(t)}{dt^2} + \frac{dx(t)}{dt} + \frac{dV_x(x)}{dx} = 0, \quad (1)$$

$$\frac{d^2y(t)}{dt^2} + \frac{dy(t)}{dt} + \frac{dV_y(y)}{dy} = 0. \quad (2)$$

The functions $V_x(x)$ and $V_y(y)$ were approximated using the experimental data in two steps. First, the screen was discretized into a mesh $(\hat{x}, \hat{y})$ and, using the mouse-cursor position data, approximations of the velocities $(\frac{d\hat{x}}{dt})$ and $(\frac{d\hat{y}}{dt})$ and accelerations $(\frac{d^2\hat{x}}{dt^2})$ and $(\frac{d^2\hat{y}}{dt^2})$ were calculated at each mesh point by interpolation and averaging over a chosen set of the participants' motion data. Second, using eqs. (1) and (2), the averaged velocities and accelerations were fed back into the integrals

$$V_x(x) = - \int_x \frac{d^2\hat{x}}{dt^2}(\hat{x}) + \frac{d\hat{x}}{dt}(\hat{x})d\hat{x}, \quad (3)$$

$$V_y(y) = - \int_y \frac{d^2\hat{y}}{dt^2}(\hat{y}) + \frac{d\hat{y}}{dt}(\hat{y})d\hat{y}, \quad (4)$$

from which the overall potential function $V(x, y)$ was calculated. Note that this was done independently of the time when the mouse cursor reached a specific position on the screen.

Across the three blocks of 12 choices in each experiment, the induced decision spaces evolved to match the relative values of the available choices. This is most clearly seen in the increase in potential at the location of the low value choice in the High/Low decision space (see Fig. 4, second column from left). For High/Low decisions, in early decisions, these spaces showed relatively equal potential at the location of low- and high-reward choices and across decisions, the location of the high-reward choice decreased in potential and the location of the low-reward choice increased in potential. Visually, the decision space tipped towards the high-reward choice demonstrating the facilitation of high-reward choices across decisions.

Further evidence of the change in dynamics across decisions is observed in the gradients of the potential functions $V(x, y)$ corresponding to decision dynamics during the first and final block of 20/5 decisions (Fig. 5). On the high value side of the figure (right-hand side), both early (red arrows; first block) and late (green arrows; third block) point towards the stimulus on that side of the figure (i.e., 20).

Figure 3 | Mean horizontal velocity across 100 time steps in the final 18 response trajectories. (A) The top panel shows the evolution of horizontal component of velocity towards the ultimate choice in the response trajectory. Velocities in the x direction were calculated between consecutive points in each trajectory. (B) The velocity time series were interpolated into 100 time steps and then the mean velocity at each time step calculated to depict the time-normalised evolution of velocity. The lower panel shows mean horizontal velocity and bootstrapped confidence intervals during consecutive quintiles of 20 time steps.

Figure 4 | Interpolated mean trajectories and decision spaces for blocks of 12 trials within each experiment. (A) The left column depicts the 20% trimmed mean trajectories for the final 18 trials in each experiment. (B) Surface plots depict inferred potential fields based on momentary velocities and accelerations derived from positional coordinates within response trajectories (see text for details). The positions of the circles depicted on the decision spaces indicate the approximate starting point of trajectories (i.e., 0,0) and the colors of the circles denote the decision type.

Figure 5 | The gradients of the potential functions $V(x, y)$ corresponding to decision dynamics during the first and third block of 20/5 decisions, with arrows pointing “down hill”, and where the high value stimulus is located on the right-hand side of the figure. The first block is highlighted with red arrows and the third blk with green arrows. There is a marked change in the arrow directions on the left-hand side (low value side) between the first and third block, indication of a positive bias towards the high value in the third block, where no such bias can be seen in the first block. Shaded circles denote approximate locations of choice symbols.

However, there was a marked change on the low value side of the figure, as participants learned the values of the available choices. In the first block (red arrows), motion towards the low value stimulus was likely to persist; red vectors on the low value side of the figure point toward the low value stimulus (i.e., 5). In contrast, in the final block (green arrows), the vector of motion changed towards the middle of screen indicating a reduction in the x component of velocity in the direction of the low value stimulus; green vectors on the low value side of the figure no longer point toward the low value stimulus.

The decision spaces also highlighted differences between High/High and Low/Low decision dynamics. During High/High decisions, stronger attraction (steeper slopes) indicated easier decisions (faster, more direct) and potential at the choice locations was stable or decreased across blocks indicating evolution of stronger attraction towards high-point choices across decisions. In contrast, during Low/Low decisions, weaker attraction (shallower slopes) towards the available choices indicated greater persistence of indecision (slower, less direct choices). The potential at 4 of the 6 low-point choice locations increased across Low/Low blocks indicating weakening of attraction to these choices. Finally, in the middle of the decision space, the slope towards choices plateaued, providing tentative evidence of a late 'saddle point' prior to eventual movement towards the available choices. These decision space models provide support for Killeen's⁴⁹ and Spivey and Dale's³⁵ position that a choice between two alternatives can be understood as movement in a two-well attractor landscape.

Discussion

Across three experiments, participants were sensitive to the relative reward for outcomes in a binary decision task. As the ratio of the higher reward to the lower increased across experiments, participants were more likely to reliably choose the outcomes with the higher rewards. Within experiments, for those who learned to reliably choose the high reward outcomes, the evolution of preference within and across decisions was similar across experiments. When one outcome was more favorable than the other (High/Low), decisions were fast and direct. When both outcomes earned equal rewards, decision dynamics depended upon the value of the available rewards. If both choices earned high rewards (High/High), the decisions were fast and direct, like High/Low decisions. If both choices earned low rewards (Low/Low), decisions were slow and indirect, with persistent early instability. Potential surfaces derived from positional coordinates fit with theoretical characterizations of decision making (e.g.,^{49,35,24}) in which available choices are attractors in a decision space.

To date, preferential decision making has, in the main, been investigated by analyzing patterns of choice allocation and response times. Our analyses of decision dynamics potentially provide much greater detail on the evolution of decisions. In the current experiments, reaction times were sensitive to the greater competition during Low/Low decisions. However, when we controlled for reaction time (Fig. 3), the evolution of horizontal velocity towards the available choices remained markedly different across conditions. Low/Low decisions were not simply slower versions of High/Low and High/High decisions; they exhibited a distinct temporal evolution characterized by persistent early competition that was not observed in the other decisions.

Low/Low decisions induced greater competition than High/High decisions. If one considers that the rewards available for choices determined the attraction towards the available choices, then one might expect persistent competition between two strong attractors in High/High decisions. Instead, participants chose one of the available high value options quickly and directly. Greater competition between similar low value choices than similar high value choices has been seen in some recent experimental data and in basal ganglia

and diffusion models of decision conflict⁵³. We can see two possible explanations for these differences in the current study. The first is that participants may have set a reference point based on some function of the mean points per decision^{10,19,13}. Since low value symbols were below this reference point, they may have been experienced as a loss and aversive, even though they remained nominally rewarding. Research on framing effects¹⁷ provides convincing evidence that repulsion from perceived bad choices (loss aversion) is a powerful determinant of eventual choice. In this case, in addition to attraction towards the available choices, low value choices may have exerted a repulsion on response trajectories, in effect a dynamic expression of approach-avoidance conflict⁵⁴. In High/High decisions, there would have been no repulsion to counter the attraction towards the eventual choice resulting in faster more direct trajectories. A weakness with this account is that repulsion from Low value options would encourage faster and more direct choices in High/Low decisions than in High/High decisions, which was not observed (in fact there was some evidence of the reverse effect). Within a binary choice framework, it is difficult to distinguish attraction towards the terminal choice from repulsion from alternative choices, but analyses of decision dynamics provide a fresh approach to distinguishing between these different possibilities.

An alternative explanation of greater competition during Low/Low decisions is that the slower uncommitted responding in Low/Low choices was due to a negative incentive contrast effect. A negative incentive contrast effect occurs when following experience with a high value reward, the incentive effect of a low value reward is reduced. For instance, animals are known to run slower and less directly towards a less preferred reward following training with a preferred reward (see⁵⁵ for a review). The fact that responding to High/High and High/Low decisions exhibited similar characteristics suggests that the highest available stimulus on a trial dominated decision dynamics. Further research under tighter laboratory controls will facilitate more thorough investigation of these findings including testing the stability, in longer protocols, of the differences identified herein. In particular, the seeming absence of conflict during High/High decisions relative to High/Low decisions invites further empirical research. In an attractor model, Low/Low decisions might be understood to constitute a decision space with two weak attractors that fail to capture response trajectories until later in the decision. In High/High and High/Low decisions, stronger attractors captured response trajectories earlier and pulled trajectories faster towards completion.

As mentioned previously, some dynamical models of decision-making predict greater competition between similar low value choices than between similar high value choices. In particular, basal ganglia and diffusion models predict these effects, but so, arguably, do models that include lateral inhibition, when inhibition is high^{4,19,21}. Indeed, the very low conflict observed during High/High decisions might constitute evidence of strong inhibition of the High alternative in these cases. The foregoing models constitute a subset of sequential-sampling models of decision making, in which information gradually accrues over time to bias the eventual choice. The current method of collecting detailed movement data on decisions seems particularly appropriate for the analysis of such models, since they propose processes that unfold over time and eventually arrive at one of the available options. Dynamic competition between graded continuous representations of choices may account for gradual shifts in trajectory and 'changes of mind' (e.g.²⁴). Finally, a number of these models demonstrate attractor dynamics, which may allow researchers to develop model decision spaces to compare with decision spaces generated from experimental data.

Within a broader learning context, Schöner and Kelso^{56,57} proposed a dynamical account of learning as the integration of *intrinsic dynamics* and *behavioural information*. In our experiments, intrinsic dynamics refer to the attractor landscape in the space of possible

responses (including physical, biological and learned constraints) prior to exposure to a new decision and behavioural information both mitigates the expression of those dynamics in that decision and updates the intrinsic dynamics for future decisions. Average decision spaces, such as those generated in the current analyses, characterise the gradual evolution of preference in different subsets of environmental contingencies. In this way, we approximate the average change in the intrinsic dynamics of preference and the average impact of behavioural information. Metaphorically, the decision space is the 'playing field' on which decisions compete. However, Schöner and Kelso emphasise that the greatest contribution of their account is to enable modelling of intrinsic dynamics at the individual level. In contrast, our average decision spaces necessarily occluded inter-individual differences in learning and, thus, differences in the structure of decision spaces at the level of the individual participant (with his/her particular physical, biological and learning constraints). At present, many trajectories are required to establish stable decision spaces and, thus, it was not possible to compare decision spaces at the level of the individual participant or trajectory. As we learn more about the expression of decisions in two-dimensional space, then it might be possible to generate decision spaces for individual participants or decisions. For instance, one might hypothesise a baseline space (e.g., for a condition or an individual), from which a specific trajectory might provide deviations that would allow us to generate an hypothetical decision space for that trajectory. It is our hope that analyses of decision dynamics and the depiction of decision spaces will provide a flexible testbed for future theoretical development.

Methods

Participants were drawn from Amazon Mechanical Turk (mturk.com), restricted to the United States, with the following inclusion criteria. Participants were required to have completed a minimum of 100 Mechanical Turk tasks prior to participation and they were required to have an approval rate 95% or above (i.e., in 95% of cases they were paid for the work they had done), a measure of how effectively they typically dealt with Mechanical Turk tasks. Participation in the experiments earned US\$0.50. At the beginning of each participant's performance, they answered a series of demographic questions (which could be filled in before, or after, but not during, their performance). This included simple non-identifying information such as age and first language. From Amazon, participants were forwarded to an external link that presented a "game-like" interface using Adobe Flash®.

At the beginning of the experiment, the following instructions were presented:

Thank you for your work on this brief HIT.

In the next screen, you will see a small green button. Click it to begin the first round. There are 36 rounds and in each round you will choose between two shapes to earn points. Different shapes give different points. The more points you earn the better! Are you ready?

Click the bottom-right corner of this instruction box to continue...

On each decision, participants clicked on a button at the center of the bottom of the screen to begin their response, which was made to one of two shapes from the Bodoni font set presented on the top-left or top-right of the computer screen. This produced, on each decision, a computer mouse-cursor trajectory from the bottom center to the left or to the right (see Fig. 1).

Across three experiments, the ratio of the high-value choice to the low-value choice was systematically manipulated. Participants were assigned to either Experiment 1 (7/5; $n = 34$), the high-point reward was 7, and the low was 5, in Experiment 2 (10/5; $n = 37$), high was 10 and low was 5, and in Experiment 3 (20/5; $n = 55$), high was 20 and low was 5. To control for any potential effect of the Bodoni font, at the start of each experiment, two shapes were randomly (across participants) assigned to the experiment-specific high value, and the two others to the low value (these stimulus-reward pairings remained consistent within each participant's set of decisions, so that they could learn these values). In each decision, two of the four symbols (High 1, High 2, Low 1 and Low 2) were presented as choices, creating three possible choice situations: a high-value choice versus high value choice (High/High), a high-value choice versus low-value choice (High/Low), and a low-value choice versus low-value choice (Low/Low). The Flash game presented 36 of these decisions in a random sequence, 24 in the High/Low condition, and 6 in each of High/High and Low/Low conditions. On clicking a choice, the points for that choice were presented in green text (e.g., +20) in the middle of the screen and a point counter in black text at the top of the screen was incremented by that value (see Fig. 1).

Amazon Mechanical Turk provides access to a demographically diverse population⁵⁸. However, though arguably providing increased ecological validity, the online environment is relatively uncontrolled compared to typical decision-making

experimental contexts. We employed a number of exclusion criteria designed to exclude participants who did not pay due attention to the experimental task. From a total of 140 participants, participants who satisfied the following criteria were excluded: those whose mean reaction time was below 500 ms ($n = 9$), those with a reaction time on any trial of greater than 10 s ($n = 6$), those who completed fewer than 36 trials ($n = 5$) and those who chose the left or right stimulus on over 75% of trials ($n = 2$). Overall, 11 participants were excluded (8%). Following these exclusions, trajectories of the remaining 129 participants were visually analysed and three further participants were removed because of unusual patterns (e.g., "swooping", moving directly left or right before moving upwards towards the choices). 126 participants were assigned to the three experiments as described above.

1. Luce, R. D. *Response times: Their role in inferring elementary mental organization*. (Oxford University Press, 1986).
2. Link, S. W. *The wave theory of difference and similarity*. (Psychology Press, 1992).
3. Gold, J. I. & Shadlen, M. N. The neural basis of decision making. *Ann. Rev. Neurosci.* **30**, 535–574 (2007).
4. Bogacz, R., Usher, M., Zhang, J. & McClelland, J. L. Extending a biologically inspired model of choice: Multi-alternatives, nonlinearity and value-based multidimensional choice. *Philos. Trans. R. Soc. Lond., Ser. B* **362**, 1655–70 (2007).
5. Stafford, T. & Gurney, K. N. Biologically constrained action selection improves cognitive control in a model of the Stroop task. *Philos. Trans. R. Soc. Lond., Ser. B* **362**, 1671–84 (2007).
6. Usher, M., Elhalal, A. & McClelland, J. The neurodynamics of choice, value-based decisions, and preference reversal. In Chater, N. & Oaksford, M. (eds.) *The Probabilistic Mind: Prospects for Bayesian cognitive science*, 277–300 (Oxford University Press, Oxford, UK, 2008).
7. Herrnstein, R. Relative and absolute strength of response as a function of frequency of reinforcement. *J. Exp. Anal. Behav.* **4**, 267–272 (1961).
8. Herrnstein, R. J. On the law of effect. *J. Exp. Anal. Behav.* **13**, 243–266 (1970).
9. Tversky, A. Elimination by aspects: A theory of choice. *Psychol. Rev.* **79**, 281–299 (1972).
10. Kahneman, D. & Tversky, A. Prospect theory: An analysis of decision under risk. *Econometrica* **47**, 263–292 (1979).
11. Pierce, W. D. & Epling, W. F. Choice, matching, and human behavior: A review of the literature. *Behav. Anal.* **6**, 57–76 (1983).
12. Koop, G. J. & Johnson, J. G. Response dynamics: A new window on the decision process. *Judgm. Decis. Mak.* **6**, 750–758 (2011).
13. Vlaev, I., Chater, N., Stewart, N. & Brown, G. D. A. Does the brain calculate value? *Trends Cogn. Sci.* **15**, 546–54 (2011).
14. Busemeyer, J. R. & Townsend, J. T. Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychol. Rev.* **100**, 432 (1993).
15. Townsend, J. T. & Busemeyer, J. R. Dynamic representation of decision-making. In Port, R. F. & Van Gelder, T. (eds.) *Mind as Motion: Explorations in the Dynamics of Cognition*, 101–120 (MIT Press, 1995).
16. Ratcliff, R., Van Zandt, T. & McKoon, G. Connectionist and diffusion models of reaction time. *Psychol. Rev.* **106**, 261–300 (1999).
17. Kahneman, D. & Tversky, A. *Choices, values, and frames*. (Cambridge University Press, 2000).
18. Van Zandt, T. ROC curves and confidence judgments in recognition memory. *J. Exp. Psychol.: Learn. Mem. Cogn.* **26**, 582–600 (2000).
19. Usher, M. & McClelland, J. L. Loss aversion and inhibition in dynamical models of multialternative choice. *Psychol. Rev.* **111**, 757–69 (2004).
20. Stewart, N., Chater, N. & Brown, G. D. A. Decision by sampling. *Cogn. Psychol.* **53**, 1–26 (2006).
21. Wang, X.-J. Decision making in recurrent neuronal circuits. *Neuron* **60**, 215–234 (2008).
22. Dickhaut, J., Rustichini, A. & Smith, V. A neuroeconomic theory of the decision process. *Proc. Natl. Acad. Sci. USA* **106**, 22145–50 (2009).
23. Gigerenzer, G. & Gaissmaier, W. Heuristic decision making. *Ann. Rev. Psychol.* **62**, 451–82 (2011).
24. Albantakis, L. & Deco, G. Changes of mind in an attractor network of decision-making. *PLoS Comput. Biol.* **7**, e1002086 (2011).
25. Spivey, M. J., Grosjean, M. & Knoblich, G. Continuous attraction toward phonological competitors. *Proc. Natl. Acad. Sci. USA* **102**, 10393–10398 (2005).
26. Dale, R., Kehoe, C. & Spivey, M. J. Graded motor responses in the time course of categorizing atypical exemplars. *Mem. Cogn.* **35**, 15–28 (2007).
27. Farmer, T. A., Anderson, S. E. & Spivey, M. J. Gradiency and visual context in syntactic garden-paths. *J. Mem. Lang.* **57**, 570–595 (2007).
28. Farmer, T. A., Cargill, S. A., Hindy, N. C., Dale, R. & Spivey, M. J. Tracking the continuity of language comprehension: Computer mouse trajectories suggest parallel syntactic processing. *Cogn. Sci.* **31**, 889–909 (2007).
29. Freeman, J. B., Ambady, N., Rule, N. O. & Johnson, K. L. Will a category cue attract you? Motor output reveals dynamic competition across person construal. *J. Exp. Psychol.: Gen.* **137**, 673 (2008).
30. Freeman, J. B. & Ambady, N. Motions of the hand expose the partial and parallel activation of stereotypes. *Psychol. Sci.* **20**, 1183–1188 (2009).
31. Barca, L. & Pezzulo, G. Unfolding visual lexical decision in time. *PLoS One* **7**, e35932 (2012).

32. Dale, R., Roche, J., Snyder, K. & McCall, R. Exploring action dynamics as an index of paired-associate learning. *PLoS ONE* **3**, e1728 (2008).
33. Duran, N. D., Dale, R. & McNamara, D. S. The action dynamics of overcoming the truth. *Psychon. Bull. Rev.* **17**, 486–91 (2010).
34. Song, J.-H. & Nakayama, K. Numeric comparison in a visually-guided manual reaching task. *Cogn.* **106**, 994–1003 (2008).
35. Spivey, M. J. & Dale, R. Continuous dynamics in real-time cognition. *Curr. Dir. Psychol. Sci.* **15**, 207–211 (2006).
36. Song, J.-H. & Nakayama, K. Hidden cognitive states revealed in choice reaching tasks. *Trends Cogn. Sci.* **13**, 360–6 (2009).
37. Freeman, J. B., Dale, R. & Farmer, T. A. Hand in motion reveals mind in motion. *Front. Psychol.* **2**, 59 (2011).
38. Spivey, J. M. *The Continuity of Mind* (Oxford University Press, 2007).
39. Wojnowicz, M. T., Ferguson, M. J., Dale, R. & Spivey, M. J. The self-organization of explicit attitudes. *Psychol. Sci.* **20**, 1428–35 (2009).
40. Usher, M. & McClelland, J. L. The time course of perceptual choice: The leaky, competing accumulator model. *Psychol. Rev.* **108**, 550–592 (2001).
41. Teodorescu, A. R. & Usher, M. Disentangling decision models: From independence to competition. *Psychol. Rev.* (in press).
42. Vanrullen, R. & Thorpe, S. J. The time course of visual processing: From early perception to decision-making. *J. Cogn. Neurosci.* **13**, 454–461 (2001).
43. Ratcliff, R., Philastides, M. G. & Sajda, P. Quality of evidence for perceptual decision making is indexed by trial-to-trial variability of the EEG. *Proc. Natl. Acad. Sci. USA* **106**, 6539–44 (2009).
44. Hare, T. A., Schultz, W., Camerer, C. F., O'Doherty, J. P. & Rangel, A. Transformation of stimulus value signals into motor commands during simple choice. *Proc. Natl. Acad. Sci. USA* **108**, 18120–18125 (2011).
45. Chib, V. S., De Martino, B., Shimojo, S. & O'Doherty, J. P. Neural mechanisms underlying paradoxical performance for monetary incentives are driven by loss aversion. *Neuron* **74**, 582–94 (2012).
46. Niv, Y., Edlund, J. A., Dayan, P. & O'Doherty, J. P. Neural prediction errors reveal a risk-sensitive reinforcement-learning process in the human brain. *J. Neurosci.* **32**, 551–62 (2012).
47. Glöckner, A. & Herbold, A. K. An eye-tracking study on information processing in risky decisions: Evidence for compensatory strategies based on automatic processes. *J. Behav. Dec. Mak.* **24**, 71–98 (2011).
48. Franco-Watkins, A. M. & Johnson, J. G. Applying the decision moving window to risky choice: Comparison of eye-tracking and mousetracing methods. *Judgm. Decis. Mak.* **6**, 740–749 (2011).
49. Killeen, P. R. Mechanics of the animate. *J. Exp. Anal. Behav.* **57**, 429–463 (1992).
50. Wang, X.-J. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* **36**, 955–968 (2002).
51. Wong, K.-F. & Huk, A. C. Temporal dynamics underlying perceptual decision making: Insights from the interplay between an attractor model and parietal neurophysiology. *Front. Neurosci.* **2**, 245–54 (2008).
52. Aparicio, C. F. & Baum, W. M. Dynamics of choice: relative rate and amount affect local preference at three different time scales. *J. Exp. Anal. Behav.* **91**, 293–317 (2009).
53. Ratcliff, R. & Frank, M. J. Reinforcement-based decision making in corticostriatal circuits: Mutual constraints by neurocomputational and diffusion models. *Neural Comput.* **24**, 1186–229 (2012).
54. Hovland, C. I. & Sears, R. R. Experiments on motor conflict. I. Types of conflict and their modes of resolution. *J. Exp. Psychol.* 477–493 (1938).
55. Flaherty, C. F. Incentive contrast: A review of behavioral changes following shifts in reward. *Anim. Learn. Behav.* **10**, 409–440 (1982).
56. Schöner, G. & Kelso, J. A. S. Dynamic pattern generation in behavioral and neural systems. *Science* **239**, 1513–20 (1988).
57. Kelso, J. A. S. *Dynamic Patterns: The Self-Organization of Brain and Behavior* (The MIT Press, Cambridge, MA, 1995).
58. Buhrmester, M., Kwang, T. & Gosling, S. D. Amazon's Mechanical Turk: A new source of inexpensive, yet high-quality, data? *Persp. Psychol. Sci.* **6**, 3–5 (2011).
59. Wickham, H. *ggplot2: Elegant graphics for data analysis* (Springer Publishing Company, Incorporated, 2009).

Acknowledgements

Procedures were approved by the University of Memphis Committee on the Protection of Human Research Participants. Informed consent was obtained from all participants. This research was supported by NSF grant BCS-0720322 to RD and by a Summer Internship Grant from the College of Science, NUI Galway to FC. Support for Open Access publication was provided by the UC Merced Open Access Fund Pilot. We wish to thank two anonymous reviewers for their detailed appreciation of the manuscript and their constructive and supportive suggestions.

Author contributions

D.O.H., R.D. and P.P. wrote the main manuscript text. D.O.H. prepared figures 1 to 3 and P.P. and D.O.H. prepared figures 4 and 5. D.O.H. and R.D. designed the experiment and R.D. developed the experimental programs. P.P. and F.C. developed the mathematical model. All authors reviewed the manuscript.

Additional information

Competing financial interests: The authors declare no competing financial interests.

How to cite this article: O'Hara, D., Dale, R., Piironen, P.T. & Connolly, F. Local dynamics in decision making: The evolution of preference within and across decisions. *Sci. Rep.* **3**, 2210; DOI:10.1038/srep02210 (2013).

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 Unported license. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/3.0>