1

2

3

4

5 **Beyond Nudging:**

6 **Generalisable and transferable learning in human decision-making**

7 Andreas Jarvstad^{1,2}

8 ¹Department of Experimental Psychology, University of Oxford

9 ²Department of Psychology, City, University of London

10

11

12

13

14

15

16

17

18 Dr Andreas Jarvstad, Department of Psychology, City, University of London, Northampton

19 Square, Clerkenwell, London EC1V 0HB, andreas.jarvstad@city.ac.uk, ORCID: 0000-0002-

20 3175-8733.

21

22 **Keywords:** decision-making, generalisation, transfer, prospect theory, decisions from

23 description, decisions from experience, risky choice, BART

Canonical work on human decision-making demonstrates consistent, stereotyped and sub-optimal risk-reward trade-offs. However, recent work on experienced-based-, perceptual- and sensori-motor choice appear to severely restrict the scope of the canonical work, demonstrating trade-offs that are either differently sub-optimal, or much closer to optimal. Such dissociations may reflect highly specific mechanisms, or other confounds between domains; impossible to tease apart with current observational methods. We develop a method for strong causal inference based on experimentally manipulating risk preferences. Using a double-blind randomized control design, we trained people in a single domain, evaluate the extent to which training *generalizes* beyond the training-set, and more importantly the extent to which training *transfers* to risk domains in which participants were not trained. Participants were trained to maximize their expected earnings (i.e., to be risk-neutral). In total, we tested for transfer to four different risk domains, all of which are known to dissociate. We find that training is necessary and sufficient to cause reliable changes in decision-making. Importantly, training *transfers*, with participants becoming more risk-neutral also in domains for which they had had no training, showing that risk-reward trade-offs in different domains and tasks share common substrate. Shared mechanisms open up the opportunity to for practical general-purpose training programmes to improve human decision-making.

Introduction

Deciding when to cross a busy road, or how to invest ones' savings, involves choosing between potentially valuable, but risky actions (cross now, buy stocks), and less valuable but safer actions (wait for car to pass, buy bonds). For decades, researchers have probed such risk-reward trade-offs with verbal stimuli, such as: "Would you prefer £100 with a .4 probability, or £200 with a .25 probability?", revealing remarkably consistent, stereotyped and sub-optimal trade-offs (1-3).

Recent research has called the generality of these canonical results into question. When risks and rewards are experienced rather than described (4-6), or when making sensori-motor and perceptual decisions (7-11), choices appear radically different; sometimes differently sub-optimal (e.g., 4) and sometimes closer to optimal (e.g., 11, but see 12-13). This variation by task-domain, suggests that the underlying neural mechanisms for risk-reward trade-offs are highly domain-specific, and by extension that the canonical work (1-3) reflects human behaviour only under extremely limited circumstances.

This view is also reflected in recent high-profile ameliorative projects, so called 'Nudges' (14, but see 15, 63). Nudging focusses on narrow interventions, seeking to improve single decisions in highly specific domains (e.g., pension plan saving; 16). Domain-specific mechanisms favour such narrow interventions, because learning will transfer poorly across domains: as one learns to make better decisions in one domain, the learning will have little bearing on ones' decisions in other domains.

However, by other accounts the apparent domain-specificity is surprising. A defining feature of human cognition is said to be our ability to generalise from a limited set of learning stimuli, and to transfer that knowledge across narrow domains. This view is, for example, reflected in research on artificial intelligence (AI), where general-purpose AI is thought to require a human-like ability to generalise and transfer knowledge (17-18). Despite such

strong intuitions, transfer of learning has proven difficult to demonstrate empirically (19-21, but see 22-25). People can of course learn a multitude of things, including to make better decisions (26-28, see also *Discussion*), but improvements are often likely due to stimulus-response learning which generalise and transfer poorly (19, 28).

Here we ask whether people can learn to tune their risk-reward trade-offs such that the learning *generalises* (goes beyond the training set), and whether such learning *transfers* (goes beyond the training domain). The former (generalisation), would indicate that risky choice is more flexible than implied by the replication of key findings across decennia (1-3), and has implications for the feasibility of cognitive training paradigms. The latter (transfer) allows us to directly address mechanism specificity.

We trained participants in one domain, and evaluated the extent to which training *transferred* to several non-trained domains: full transfer implies identical mechanisms, and no transfer implies wholly separate mechanisms. Thus, allowing us to address the extent to which risk-reward trade-offs are underpinned by domain-specific mechanisms – as implied by behavioural dissociations (e.g., 4). As far as we are aware, the application of transfer learning to different risk domains as done here is novel, though the method itself is not (see e.g., 24 for the method applied to working memory and general intelligence).

Testing for generalisation and transfer

We focus on two risk domains: canonical described risk (1-2) and sensori-motor risk (29,10). The latter involves inferring risk from sensori-motor stimuli: much like one might infer the probability of successfully crossing a busy road by estimating time-to-contact (30) with oncoming cars (thus relying on the general idea that risk can be extracted from experienced-based task stimuli, 31).

There were three phases: pre-training, an extensive training phase and a post-training phase (Fig 1A). In the pre- and post-phases participants performed two tasks. First

participants performed a standard pointing task (Fig 1A,B), which involved speeded reaching towards targets (blue rectangles displayed on a touch-screen). On each trial, participants had limited time in which to reach to hit the target (target widths varied across trials). Hits were rewarded, and time-outs were penalized, with cumulative virtual points (yellow bar, Fig1B) later converted to money (see *Methods*). As in previous sensori-motor decision-making studies (e.g., 30), this pre-task was solely for the purpose of participants' gaining experience with the perceptuo-motor stimuli with which they later made decisions – much like people may use their experiences with observing oncoming cars to judge the risk of crossing a busy road.

A choice task followed (Fig 1A,C). It involved choices between pairs of options, composed of value and probability information. Typically, one option was risky and had a low probability of success (but a high potential reward), whilst the other option was relatively safe (but with a low potential reward). Choices thus required a trade-off between risk and reward. At this stage, no feedback was provided, and participants simply indicated which of two options they preferred on each trial (though choices had real economic consequences, *Methods*).

Each participant made 400 choices, 200 of which were with described risk. For example, participant might choose between [$\pounds 6$, $p=.65$] or [$\pounds 10$, $p=.4$]. That is, $\pounds 6$ with a probability of .65 (otherwise nothing), or $\pounds 10$ with a probability of .4 (otherwise nothing). Two-hundred choices were made with matched sensori-motor stimuli. These stimuli were matched on the basis of participants empirical performance (29,10). For example, the described probability in [$\pounds 6$, $p=.65$] can be represented as a pointing target of a width that a given participant would hit with $p=.65$.

Thus, in 'motor' trials, participants traded off risk derived from sensori-motor experience (the probability with which they would hit the target) and money, and for

‘descriptive’ trials, participants traded off described risk (the probability of reward) and money (Fig 1D).

Fig 1: Design and tasks. A) Study design. Day 1: participants completed the pre-training phase, after which they were randomly and double-blindly assigned to one of the four groups. Day 2-7: each

group trained twice daily for six days, for a total of 12 training episodes with 60 trials/episode. Day 8: participants completed the post-training phase repeating the tasks of the pre-training phase. B) Pointing task. To start a trial, participants depressed the space-bar key, a brief random delay followed after which a target appeared. Participants released the space bar and reached rapidly to try and hit the target (using the same hand). Virtual points (yellow bar) accumulated when targets were hit and were deducted when responses were too slow. Virtual points were converted into money upon study completion. The post-training task involved fewer trials but was otherwise identical to the pre-training task (*Methods*). C) Choice task. On each trial, participants indicated whether they preferred the left or the right option. The task was self-paced, and risk domain (description/motor) varied randomly from trial to trial, for a total of 400 trials. No feedback was provided, but at the end of the study participants were paid the average outcome of a selection of their choices (*Methods*). D) Training phase. Participants were allocated to one of four groups (double-blind assignment). Two groups trained with feedback and two groups trained without feedback, with each group training on one task only. Following choices which did not maximize expected earnings, feedback groups experienced a small delay, after which they were required to correct their choice before continuing. Control groups practised identical tasks without feedback.

Following the pre-training session, participants were randomly and blindly divided into four groups (Fig 1A,D). Each group practised risky choice in one, and only one, risk domain twice daily for six days, for a total of 12 training episodes. As in pre/post Choice task, each trial involved a choice between two options. To prevent stimulus-response learning, each training problem was encountered only once. All training groups trained on formally equivalent decision-problems in the same identical order.

Two groups practised with feedback (treatment groups). One treatment group practised with description problems ('Desc-F') and one with motor problems ('Motor-F'). Whenever these participants made a choice that did not maximize their expected earnings (*Methods* for details), a small delay followed, and the choice had to be corrected before continuing. Thus, participants in the treatment groups were trained to maximize their expected earnings. This feedback-scheme resulted in feedback that was sparse in both frequency and information content. The two active control groups ('Desc-C', 'Motor-C') trained with identical

problems, but without feedback. Unlike the pre-and post-phases, training was not incentivized.

Following training, participants completed the post-training phase, which repeated the tasks of the pre-training phase (Fig 1A).

Results

Feedback induces fast and slow adaptation during training

Fig 2A shows choice performance (defined as proportion of choices maximizing expected value, *Methods*) for the feedback and control groups respectively across every session and training episode. Despite near-identical performance prior to training (pre-feedback–pre-control; $t(38)=-.95$, $p=.348$, $CI=-.078$ - $.028$, $D = .3$, $BF_{NULL}=2.26$ [JZS Bayes Factor (31)], un-paired t-test), feedback groups outperformed control groups in the very first training episode (feedback-control; $t(38)=2.856$, $p=.007$, $CI=.024$ - $.139$, $D = .9$, un-paired t-test), indicating very rapid learning. There was also a second trend, whereby feedback groups, but not control groups, continued to improve throughout training; the slopes of linear fits to training data were positive for feedback groups ($t(19)=2.12$, $p=.048$, $CI=0.001$ - $.0052$, $D = .47$, one-sample t-test), but not for control groups ($t(19)=-0.18$, $p=.86$, $CI=-.0028$ - $.0023$, $D = .04$, $BF_{NULL}=4.24$, one-sample t-test). As expected, given that the training stimuli varied in difficulty and were sequenced identically across groups, training performance across episodes 1-12 was correlated across feedback and control groups ($r(10) = .585$, $p = .046$, Pearson's r): when one group did well so did the other.

Fig 2: Feedback and control groups' choice performance as a function of training duration. A) Performance (proportion of optimal choices) in pre-training, training and post-training. Pre- and post-session error bars are bootstrapped 95% CIs. Linear trends are averages of robust linear fits to individual participants' data. B) Smoothed performance over the first 30 trials in Training episode 1. Error surfaces are SEM. Grey bars on the x-axis indicate significant group differences at $p < .05$ by cluster-corrected un-paired t-tests. C) Performance over trials from trial t, when error N was made, up until trial t+n, where n is the trial just before the next error was made. Grey bars on the x-axis indicate significant group differences at $p < .05$ by cluster-corrected un-paired t-tests

Fig 2B-D explores the very rapid initial learning in the first training episode. Fig 2B shows smoothed performance over the very first 30 trials in the first training episode. Performance appears non-monotonic across trials due to data-smoothing and trial-by-trial variation in difficulty. Nonetheless, the key contrast between feedback and control groups shows that the two groups diverge after 18 trials (grey x-axis bars, Fig 2B). This divergence must be driven by feedback following error trials. Plotting performance between the Nth error and the Nth+1 error (Fig 2C), reveals that six instances of feedback is sufficient for divergence between feedback and control participants. Thus, participants learned very rapidly initially and then later more gradually.

Training with feedback improves performance and transfers to non-trained domains

The observed improvement during training might reflect altered risk-reward trade-offs but could also reflect avoidance of aversive feedback. If people's risk-reward trade-offs were fixed (2), they should revert to their true initial preferences when real money is at stake (post-training). This, however, is not what we observe.

Figure 3A shows choice performance differences between the post- and pre-training phases, with positive values indicating an improvement following training. Recall that each group performed both tasks but received training on one task only. The description-feedback group ('Desc-F'), for example, received 'description' training with feedback (dark blue square), but did not receive 'motor' training (light green triangle).

As can be seen from the distance of the 95% confidence intervals to 0 (Fig 3A, see also *SI Additional Analyses*), both feedback groups show large and highly significant improvements in the trained domain (~10% points on average). Importantly, participants who trained on the description task, also improved on the motor task (and vice versa). Thus, training transferred across risk domains (large-medium effect sizes, description $D = 1.26$; motor $D = .66$). Still transfer was not complete (Fig 3A), with performance improving less in the non-trained domains (Desc-F post-pre: $t(16)=-2.41$, $p=.028$, $CI=.006 - .096$, $D = .59$; Motor-F post-pre: $t(17)=-3.91$, $p=.002$, $CI=-.085--.025$, $D = .92$, paired t-tests).

Fig 3: Post-training performance relative to pre-training performance A) Main study: training-induced changes in performance split by participant group and risk domain. Group assignment is indicated on the x-axis (F=feedback, C=control). Blue squares denote description and green triangles denote motor choice problems. Darker shades indicate training and lighter shades indicate no training. Error bars are bootstrapped 95% confidence intervals on the mean. B-C) Independent replication of the motor-feedback condition B) Left plot: mental arithmetic task. The task was to sum four central numbers on each trial. Participants responded on a virtual keypad and received feedback on the distance between their estimate and the true sum (red numbers). If the estimated sum lay within the target interval (green +/-) participants accumulated points. If participants timed-out points were deducted. B) Right

plot: Choice task. Participants made decisions in three risk domains: described risk, sensori-motor risk, and mental arithmetic risk. The added arithmetic choice is analogous to the motor choice, but with the probability represented by arithmetic targets instead of pointing targets C) Choice performance deltas for the replication study. Error bars are bootstrapped 95% confidence intervals on the mean.

There was no evidence of improvement in the control groups, in either the trained or the un-trained domains (Fig 3A). Moreover, the effects of feedback were specific to risk-reward trade-offs. Performance on the pointing task remained unaffected (*SI Non-specific effects*). Feedback was therefore necessary and sufficient to induce substantial and specific changes in risk-reward trade-offs. Importantly, and as expected given the double-blind group assignment, the pre-training performance was highly similar across the four groups (*SI Pre-training performance*).

These analyses show that participants did not revert to their pre-training risk-reward trade-offs but used the feedback to tune their risk-reward trade-offs to make more optimal decisions. Thus, people's risk-reward trade-offs are not fixed but adaptive. Crucially, the transfer to un-trained domains suggest that risk-reward trade-offs rely on (at least partially) shared mechanisms.

Emergence of a transfer gradient with a cognitive experience-based task

Although decision problems were exactly matched across the description and motor tasks, the tasks differed along more than one dimension. Not only are motor decisions less 'cognitive' than decisions from description, they also involve second-order uncertainty (uncertainty about the probability). We therefore performed an independent replication of the 'motor-feedback' condition, the condition with the weakest training effects (Fig 3A), including also a more 'cognitive' task involving second-order uncertainty: mental arithmetic (10).

In the arithmetic task (top panel, Fig. 3B), participants summed four central numbers, keyed in their response on a virtual keypad, and received feedback on how far their estimated

sum was from the true sum (red numbers, 3B). As in the motor task, participants gained points if their response lay within a target interval (green +-, 3B), and lost points if they responded too slowly. And just as before, participants used their experience with the task(s) to trade-off risk and rewards (bottom panel, Fig. 3B), which now involved three different risk domains: description, arithmetic and motor.

The confidence intervals in Fig 3C show that both the within-domain improvement, and the across-domain transfer replicated. In other words, training induced highly significant improvements across all three tasks. Additionally, with a third task, a marginal transfer gradient emerges (repeated measures ANOVA, $F(2,36) = 3.07, p = .059, \eta_p^2 = .146$; linear trend $F(1,18) = 6.082, p = .024, \eta_p^2 = .253$); suggesting that decision-making is relatively blind to whether trade-offs involve more or less cognitive domains (motor and arithmetic are relatively similar, c.f., 8), and more affected by whether risk is from ‘description’ or from ‘experience’ (i.e. involve second order uncertainty, e.g., 4).

Training and the nature of risk-reward trade-offs

The previous analyses show that risk-reward trade-offs improved, but not how. To determine how, we fit Prospect Theory as a measurement model (32, *Methods*) separately to each participant, risk domain and pre/post-training phases. For gambles involving potential gains, Prospect Theory describes the role of risk and reward in peoples’ decision-making with two parameters (α, γ). Risk-reward trade-offs can thus be visualised in a two-dimensional risk-preference space.

Vectors in this space (Fig 4A), describe the average pre-training preference (black circle), the extent and direction of training effects (vector amplitude and angle), and post-training preferences (red circle). The blob at the base of each vector shows the distribution of vectors expected under the null hypothesis of no learning (*Methods*), and thus allows

inference and demonstrates parameter recoverability (see also *SI Model recovery of changes in risk-reward trade-offs*).

Each quadrant in this parameter space (Fig 4A), implies qualitatively different preferences. The left-half of space ($\log \alpha < 0$), implies that the subjective utility of money increases at a slower rate than actual money (diminishing marginal utility). Conversely, the right-hand space ($\log \alpha > 0$) implies increasing marginal utility. The bottom-half of the space ($\log \gamma < 0$), implies overweighting of small probabilities, that is willingness to bet on very low probability outcomes (e.g., lottery tickets). Conversely, the top-half of this space ($\log \gamma > 0$) implies ‘avoidance’ of low probability outcomes. For participant data, there were three broad trends.

Fig 4: Training-induced shifts in preferences for money and risk. A) Hypothetical example with labels. The vectors describe how participants preferences, for money (x-axis) and probability (y-axis), changed from pre-training (black circle) to post-training (red circle). The four quadrants imply qualitatively different preferences. The background contours indicate the proportion of optimal choices expected given any combination of α and γ , for a hypothetical noise-less decision-maker. The expected proportion at $\log(\alpha)=0$ and $\log(\gamma)=0$ is 100%. Expanding out from the centre, each additional contour-line represents a reduction in performance by 2%. The blobs at the base of each vector show the distribution of vectors expected under the null hypothesis that participants remain unaffected by training. Vectors reflect robust estimators of the mean (50% trimmed means). B) Main study, C) Replication of the motor-feedback condition.

First, prior to training, and consistent with previous work (4,6,10,33), preferences were markedly different across domains (black circles, Fig 4B,C). Experience-based preferences ('motor', 'arithmetic') lay in the upper-left quadrant and description-based preferences lay in or near the lower-left quadrant. Thus, participants over-weighted small probabilities for description, and under-weighted small probabilities for experience-based choice (motor & arithmetic). Thus, the dissociations across domains noted in the introduction, which easily lead to the conclusion that different mechanisms are at play, replicated. Preferences for money, were highly similar across domains; all lay in the left half of the parameter space. This similarity was evident also at the level of individual participants; preferences for money were highly correlated across domains ($r(65)=.51$, $p=.00001$, data from Fig 4B).

Second, participants risk-reward trade-offs were highly adaptive: feedback groups showed marked changes in trained domains (top row Fig. 4B,C). The vectors extend well beyond the null-distributions and towards neutral risk preferences (towards $\log \alpha = 0$ and $\log \gamma = 0$). Parameters for value- and probability weighting were closer to neutral following training, both in the main study (α $Z(35)=-4.63$, $p<.0001$, $r=.78$; γ $Z(35)=-3.19$, $p=.0014$, $r=.54$, signed-rank tests), and the replication (α : $Z(19)=-3.78$, $p=.0002$, $r=.88$; γ $Z(19)=-2.01$, $p=.04$, $r=.46$) - with mostly large effect sizes ($r \geq .5$). Thus, improvements in performance

(Fig 3A,C) were underpinned by an effective linearization in preferences for both money and probability (Fig 4A,B).

This model, in which feedback affects preferences for both money and risk, is superior to models that assume that only one of the two change (e.g., only preferences for money being affected by training, controlling for degrees of freedom, see *SI Model comparison* for details, and also *SI Model recovery of changes in risk-reward trade-offs*).

Control groups, on the other hand, were largely unaffected by training. Their post-training preferences were similar to their pre-training preferences, and lay inside, or near, the vector null-distributions (Fig 4A). Moreover, control-participants preferences were highly correlated across sessions, suggesting temporally stable base-line preferences (see *SI Temporal Stability of risk-reward trade-offs* for details, though evidence for longer-term stability is mixed, 49 and references therein). The lack of an improvement in the control groups, mirrors the performance-based analyses, and shows that feedback is necessary and sufficient for inducing near-optimal risk-reward trade-offs.

Third, feedback altered preferences in similar ways in non-trained domains. This can be seen by the highly similar vector angles for trained and un-trained tasks (top row Fig. 4B, Fig 4C). For money (α , x-axis), the vectors approach optimal weighting, but by half the magnitude of the trained domains. If preference for money were independent across domains, training-induced changes should be identical across domains, but they were not. Instead, preferences for money may be partially domain-dependent and/or, contrary to standard assumptions (1,3-4), risk and reward may interact during choice (see e.g., 35-36 and 28).

For probability (γ , y-axis, Fig 4B,C), the vectors for the un-trained tasks have similar direction to the trained tasks, but the vectors originated from different parts of parameter-space, making changes in un-trained domains less optimal than in trained domains.

Nonetheless, training improved preferences for risk, as becomes clear when considering the costs associated with the whole gamut of possible risk-preferences.

The grey contours in Fig 4 show the expected performance (proportion optimal choices), for different combinations of α and γ for a hypothetical noise-less decision-maker, given the stimulus distribution from which choice options were sampled. The innermost contour centred on [0,0] represents optimal choice proportions between .98 and 1, with each additional contour line representing a reduction of .02 (or 2%). As can be seen, the gradients along the major diagonal are steep, and gradients along the minor diagonal are shallow. This asymmetry holds for a range of different assumptions about hypothetical decision-makers (see *SI The cost of deviating from linear preferences* for details). The asymmetry is important because the cost of a given deviation from optimality is proportional to the steepness of the gradient.

For the groups who trained with feedback, the vectors for the un-trained domains approach the low-cost minor diagonals. Participants who trained with feedback were significantly closer to the minor diagonal in the un-trained domain following training than they were before training (across all feedback groups $Z(73)=4.99, p<.000001, r=.58$; split by group: ‘desc-F’ to motor $Z(17)=2.25, p=.025, r=.55$; ‘motor-F’ to description $Z(18)=1.85, p=.064, r=.44$; ‘motor-F replication’ to description $Z(19)=2.86, p=.004, r=.66$, ‘motor-F replication’ to arithmetic $Z(19)=3.02, p=.0025, r=.69$, signed-rank tests). Thus, preferences in the un-trained domain approached the low-cost minor diagonal following training - with mostly large effect sizes ($r \geq .5$).

The Prospect Theory model also includes a third parameter, which captures participants ability to distinguish between risky prospects: κ (not shown in Fig 4, see *Methods* for details). For Kappa the clearest improvement was in the description-feedback group ($Z(17)=2.39, p=.017, r=.58$, signed-rank), but there were signs of improvements in the motor-

feedback group ($Z(18)=1.72$, $p=.085$, marginal, $r=.41$), and in the in the replication group
albeit less clearly ($Z(19)=1.04$, $p=.295$, $r=.24$). Crucially, the improvements in precision were
all in the trained domain. In un-trained domains no, significant changes were observed (all
 p 's $> .523$). Thus, although tuning of risk-preferences transferred across domains,
improvements in the ability to discriminate between prospects were comparatively weak and
no reliable transfer was observed.

Beyond canonical risk-reward trade-offs with explicit risk and reward

Thus far, the data show that training-induced improvements generalise and transfer.
Well-matched tasks are crucial for exploring cognition across domains: lest task differences
be mistaken for domain differences. Here, the stimuli were formally identical (Fig 1, *Methods*
Main Study), and the surface properties were highly similar across domains (all tasks
involved canonical explicit risk-reward trade-offs).

However, when transfer is observed (Fig 3-4), questions of transfer beyond similar
tasks arises (i.e., 'far' transfer). To this end, we ran a partial replication of the 'Description-
feedback' condition (Fig 1), within a single session for a shorter training regime, and with an
additional decision-making task: the Balloon Analogue Risk Task (BART, 64-66, based on
pilot data, see *Methods BART Study*).

BART shares none of the surface properties of canonical tasks (c.f., Fig 1, Fig 5A), but
like the latter involves trading off risk and reward. In BART, participants' goal in is to
maximize earnings by inflating virtual balloons (Fig 5A). With each pump of the balloon,
trial-wise earnings increase. However, like real balloons, the balloon will eventually break. If
the balloon breaks, trial-wise earnings are forfeit and participants move on to the next trial.
Thus, like canonical risky choice, the task captures risk-reward trade-offs: risk another pump
for more reward or play it safe and cash earnings.

Based on pilot data (*Methods BART Study*), we made three predictions: 1) the transfer of training-induced improvements (Fig 3) was expected to replicate: even with a shorter training regime; 2) *before training* there would be no detectable association between level of riskiness across the Choice task and BART (broadly consistent with previous work, see Table S2 in 49). That is, the proportion of choices favouring the riskier option ('Description' / 'Motor') would be unrelated to the number of balloon pumps; 3) finally, we expected changes in riskiness expressed in the choice task (*after training*) to be positively associated with changes in BART. That is, increased risk-seeking in the choice task, for example, should be associated with increased risk-seeking in BART.

Fig 5. BART Study. A) BART paradigm. Each trial involves decisions about whether to risk another pump of the balloon, or whether to cash trial-wise points. If the balloon inflates without breaking, trial-wise points increase (yellow numbers). If the balloon breaks, trial-wise points are forfeit. If a cash decision is made, trial-wise points are added to the total number of points (green numbers). B) Choice task performance. Post/pre-training choice performance deltas for the non-trained domain ('Motor') for participants who improved on the trained domain ('Description'). For comparison, the plot also shows identically computed deltas for the 'Description – Feedback' condition in the Main Study (in which participants were trained on the same task, but training was more extensive). Error bars are bootstrapped 95% confidence intervals on the mean. The p -value denotes an un-paired t-test comparing performance improvements across the two studies. C) Behaviour at pre-training for the BART task as a function of behaviour in the choice task for 'Description' (blue symbols, top-plot) and 'Motor' (green symbols, bottom-plot) respectively. Each data point represents the proportion of adjusted mean pumps (balloon pumps) as a function of the proportion of risky choices for a single participant. Black lines reflect least-squares and red lines reflect robust linear fits. D) As in C, but on differences between post-training and pre-training behaviour. Positive values denote an increase in risk-seeking in post-training for both tasks. D) As in C, but pooling data from the BART Study data and the Pilot Study (*Methods: BART Study*).

As can be seen in Fig 5B, despite the shorter training, there was transfer to the un-trained domain ('Motor'). Those who trained on 'Description' made better decisions in the 'Motor' domain ($t(26) = 2.30, p = .015, CI=.003-.053, D=.44$, one-tailed paired-samples t-test). However, when compared to the transfer in the Main Study, the effect was markedly weaker (Fig 5B, $t(41) = 3.90, p = .0003, CI=.037-.127, D=1.23$, un-paired t-test). We observed no significant reduction in the trained domain relative to the Main Study ('Description', $t(41) = 1.08, p = .29, CI=-.019-.064, D=.34, BF_{NULL} = 2.05$, unpaired t-test, not shown), which suggests that the reduced transfer seen in Fig 5B is due to the reduced training regime (in frequency & duration, *Methods*), and not due to differences between participant sample populations.

As expected, we found no reliable evidence of a pre-training association between behaviour across the two tasks for either risk domain (Fig 5C). Specifically, there was no evidence for an association between proportion risky 'Description' choices and the mean

number of balloon pumps ($r(25) = .005$, $p = .981$, $BF_{\text{NULL}} = 2.79$), nor was there evidence for an association between proportion risky ‘Motor’ choices and balloon pumps ($r(25) = -.064$, $p = .751$, $BF_{\text{NULL}} = 2.69$). These least-squares results (red lines, Fig 5C) were not markedly different from robust linear fits (‘Description’ $t(25) = .015$, $p = .988$; ‘Motor’ $t(25) = -.312$, $p = .754$).

There was, however, a significant association between training-induced changes in risky choice and changes in BART (Fig 5D) for ‘Description’ (Pearson’s $r(25) = .435$, $p = .012$; robust: $t(25) = 2.63$, $p = .023$, both one-tailed), with only a marginal association between changes in the ‘Motor’ domains and in BART (Pearson’s $r(25) = .304$, $p = .062$; robust: $t(25) = 2.57$, $p = .056$, both one-tailed). When pooling the data in the BART Study with that of the Pilot Study (Fig 5E), the strength of the positive association increases somewhat and is highly significant (‘Description’: Pearson’s $r(39) = .525$, $p = .0004$; robust: $t(39) = 3.64$, $p = .0008$; ‘Motor’: Pearson’s $r(39) = .414$, $p = .0072$; robust: $t(39) = 2.84$, $p = .0071$, all two-tailed).

In summary, the Main Study and its partial replication shows that training is *sufficient* and *necessary* for generalisation and transfer (no effects in control groups, Fig 2-4), and that training-induced improvements are *specific* to risk-reward trade-offs (no transfer to the pre-choice pointing task, *SI Non-specific effects*). The BART study, additionally, suggests that transfer is not limited to tasks with highly matched stimuli, and/or tasks with similar surface properties, and is evidence for far transfer.

Discussion

Utilising a double-blind randomised controlled design, we trained participants to maximize their gains in tasks involving canonical risk-reward trade-offs (1-3). We assessed the extent to which training alters choice out-of-sample (i.e., *generalisation*), and importantly how training impacts choice in domains for which participants were *not* trained (i.e.,

transfer). Generalisation and transfer are key to robust cognitive training paradigms and provides direct experimental tests of cognitive and neural trade-off mechanisms.

Participants who received training, tuned trade-offs to maximize expected reward, and became increasingly risk-neutral (Fig 3-4). This would not have been possible, had participants simply learnt stimulus-response associations, and shows generalisable learning in risky choice. The learning also shows that risk-reward trade-offs are highly adaptive (37) and not fixed, as implied by the replication of canonical results across decades (e.g., 2-3), and by seemingly consistent preferences within people across time (as observed here, see also 10, but see 49 and *SI Temporal consistency of risk-reward trade-offs* for discussion). The absence of consistent changes in active control groups shows that the training was necessary and sufficient.

Crucially, the effects partially transferred to domains in which participants had not received training, with large to medium effect-sizes (Fig 3-5, see also *SI Model comparison, SI Model recovery of changes in risk-reward trade-offs*). This highly significant transfer rules out strong domain-specificity. Thus, risk-reward trade-offs targeted by canonical tasks (1-2, 10, 29) are underpinned by (at least partially) shared mechanisms. Moreover, transfer of learning is not limited to similar tasks, but impacts decision-making also in tasks with dissimilar surface properties. That is, we found evidence for *far transfer* of learning (Fig 5D).

Analyses of pre-training data replicated previous dissociations across domains (e.g., 4,10,33, see also *SI Objective task performance predicts decisions*). Specifically, participants were risk-seeking for described risk, and risk-avoiding for experienced risk (Fig 4). Had we, as is commonly done, only run the first part of the study, we might have concluded that the domains must be underpinned by radically different mechanisms. However, once trade-offs are experimentally manipulated through training, it is clear that they share underlying substrate.

The most exciting implication of the current results is the potential for general-purpose interventions. Currently, cognitive training is a multi-billion-dollar industry (38) with controversial efficacy (20), largely attributed to methodological weaknesses (21). Active control groups are rare, and double-blind assignment, the gold standard for evaluating interventions, is rarer still. Perhaps partly due to this, and perhaps due to easier implementation, recent ameliorative work has focused on ‘Nudges’: interventions to improve specific decisions in specific domains (14, but see ‘Boosting’, 63). The observed transfer provides proof of principle that broader cognitive training programmes may be feasible. Of course, proof-of-principle is not implementation, and further work is needed to establish boundary conditions.

One outstanding question is the extent to which training transfers to tasks that, unlike the canonical task, do not require explicit representation of risk and reward. Ultimately, this is a question of the extent to which the same cognitive mechanisms are employed in different tasks (and by extension - the extent to which lab tasks capture real-world demands). The BART study (Fig 5) provides initial evidence that tuning of risk-reward trade-offs can transfer across dissimilar tasks.

Furthermore, we explored different risk domains, but not different value domains. There is evidence for both shared and non-shared value codes for both primary (food, sex) and secondary reinforcers (money, trinkets) (39). It is therefore likely that the learning of different value representations will also transfer. Time is another important value domain; both time invested to solve a particular problem (40), and the willingness to forego immediate reward for future reward (41). Time is also important in terms of training retention. Although we found no evidence of short-term decay (*SI no post-session decay*), longer term retention important for practical applications should be addressed. Finally, we trained people to become more risk-neutral in both probability and money. Future work could

explore the extent to which participants can be trained to tune only one of the two dimensions (e.g., to maximize expected utility).

The observation that risk-reward trade-offs, targeted by canonical tasks (1-3), are achieved by shared substrate does not mean that all risk-related cognition shares this same substrate. For example, training might not transfer to behaviour underpinned by reinforcement learning-like mechanisms (28,42). Similarly, the observed flexibility of risk-reward trade-offs in the feedback groups (Fig 4), and the observed stability in the control groups (*SI Temporary stability*), does not preclude variations in risk-reward trade-offs across development (43) nor in later life (44), nor does it preclude state-dependent preferences (e.g., the value of water depending on levels of thirst, 45), nor does it preclude adaptation to the local context (e.g., 46). The results are also not inconsistent with recent suggestions that preferences are constructed on-the-fly (47), as long as preference-construal has a strong stimulus-related component. Finally, the results do not preclude different risk-reward trade-offs in different domains (48-49), which is indeed what we observe in the pre-training data (Fig 4).

In summary, the empirically observed dissociations in risk-reward trade-offs (1-11) do not reflect independent risk-reward trade-off mechanisms. The observed transfer of learning across tasks – with well matched-stimuli and with the same surface properties (Fig 3-4) – but also to those with dissimilar surface properties (Fig 5) – show that choices are underpinned by shared risk-reward trade-off mechanisms. In other words, risky choice is – at least in part – domain-general. The fact that people apply their learning outside of the training context, improving even in non-trained domains, opens up the opportunity for going beyond Nudging to general-purpose training programmes for broad improvements in human decision-making.

Methods – Main study and partial replication

Participants: The study was approved by the local ethics committee (CUREC R43461/RE001) at Oxford University. Participants provided written consent and were fully informed and debriefed. Standard inclusion criteria were applied, participants had normal, or corrected-to-normal, vision, and were fully mobile, fluent in English, healthy (no known physical or psychological conditions), and between 18-40 years old. Access to a modern Apple or Windows computer was required. In total, 99 participants were recruited from local participation panels. Erroneous self-reported eligibility, random responding, and participant computer compatibility issues meant that the final sample contained 86 participants (Main study $N=67$ distributed thus: $N_{\text{description-F}} = 17$, $N_{\text{motor-F}}=18$, $N_{\text{description-C}} = 16$, $N_{\text{motor-C}} = 16$ with the partial replication, $N=19$).

To address potential concerns associated with different rates of exclusion in the feedback and control groups, we compared the exclusion rates for those participants who trained with feedback to the rates for those who trained without feedback. 56% of excluded participants were from feedback groups (and 44% from control groups). A Bayesian comparison of exclusion rates is marginally supportive of the null hypothesis of no difference between groups ($BF_{\text{NULL}} = 2.12$, $15/27=\text{feedback}$, $12/27=\text{control}$). Thus, there is no evidence to suggest that exclusion depended on group allocation.

In the main study, participants received a base pay of £15 and could earn between £0 and £30 extra depending on performance. The minimum and maximum were theoretical limits, and the typical total reimbursement was ~£32 for ~3hrs of total participation. In the replication study, the base pay was £20 and participants could earn between £0 and £50 extra, with typical earnings being ~£45 for ~4hrs of total participation.

Procedure: The pre- and post-training sessions were run at the Department of Experimental Psychology and involved two tasks: pointing and choice (Fig 1). First, participants were familiarized with the response time requirements of the pointing task, after

which they completed the pointing task, and thereafter the choice task (Fig 1). Upon completion of the pre-training session (Day 1), participants were sent training software. They were instructed to train twice daily for six days, with a minimum of four hours between training episodes (Day 2-7). The training data was saved in an encrypted format inaccessible to participants. Following training, participants completed the post-training session at the Department of Experimental Psychology (Day 8).

Assignment of participants to groups was double-blind and performed by computer. Due to a bug, the first 14 participants were all assigned to feedback groups. There was no statistically detectable difference in post-training improvements between these initial participants, and those later assigned to feedback groups (description $t(33)=-.97$, $p=.33$, $BF_{\text{null}}=2.11$; motor $t(33)=0.64$, $p=.68$, $BF_{\text{null}}=2.58$).

The procedure in the replication study was identical to the ‘Motor-Feedback’ condition in the main study, except that it included an additional mental arithmetic task (Fig 3B) and the training time was doubled to 10 minutes twice daily.

Apparatus. Stimuli were displayed on a touchscreen (Iiyama T2336MSC-B2) and experiments were written in Matlab (Mathworks) using Psychtoolbox (56) on Linux (Ubuntu 16.04). Home practice was performed on modern Apple or Windows computers, with the physical size of the stimuli matched to the physical size in the laboratory.

Pointing task: The pointing task (Fig 1) involved rapid reaching to hit targets (blue rectangles) displayed on a touch screen. Target widths varied across trials (from .07 cm to 4.5 cm). Participants initiated a trial by depressing, and holding down, the space bar. Following this a target appeared after .5, .75 or 1 seconds (uniformly sampled), after which participants had .7 seconds respond.

A dot marking where the screen was hit was shown allowing participants to where they had pointed. Visual and auditory feedback signaled hits, misses and late responses

582 respectively. Hits were rewarded with 1 point and late responses penalized with -3 points.

583 Cumulative points were converted to money at the end of the study.

584 The pre-training pointing task involved at least 300 trials (with up to 60 extra trials to
585 compensate for time-outs), and the post-training task involved at least 150 trials (with up to
586 30 extra trials to compensate for time-outs). In the replication study, pre-training involved at
587 least 200 trials (with up to 40 extra trials), and the post-training involved at least 100 trials
588 (with up to 20 extra trials).

589 *Mental arithmetic task:* The task involved rapidly summing four numbers (Fig 3B). The
590 four numbers were sampled uniformly in the range 1 to 25, with the constraint that the
591 difference between the sum of the four numbers and the expected value of their sum, be
592 smaller than 30. On each trial, participants therefore saw four numbers in the range 1 – 25
593 whose sum was approximately uniform in the ~22-80 range (across trials).

594 Each trial begun by the presentation of 4 central numbers along with a target interval
595 (e.g., ± 12 , Fig 3B). This screen was on for 1 second, after which a virtual keypad was
596 displayed, and a response was required within 2 seconds. A feedback screen showed the
597 target interval and the difference between the true and the estimated sum. If the estimated
598 sum lay within the target interval a hit was scored. Hits were rewarded with 1 point and late
599 responses penalized with -3 points. The number of trials was identical to the pointing task.

600 *Choice task stimuli* Pairs of choice options were generated by sampling uniformly from
601 probabilities in the range ($p=.05$ to $p=.95$ in steps of $.05$, including also $p=.99$) and values in
602 the £1-20 range (in integer steps). This resulted in choice problems with a wide range of
603 differences in expected value between the two options ($|\Delta EV| \sim \text{£}0\text{-}14$). Pairs were selected
604 such that only a fraction were dominated (i.e., both the probability and value was larger for
605 one option). Participants were generally attentive. Out of 5 possible times, participants chose

the dominating option 4.88 times (description) and 4.9 times (motor), with no participant making more than a single error (Main study, pre-training session).

Choice task Two-hundred problems for each risk domain (2x200) were included in the main study. For the replication 120 were included for each domain (3x120). The problems in the pre- and post-sessions were identical but presented in a different random order. On each trial, participants indicated whether they preferred the left or the right option. Their choice was briefly highlighted after which the next trial begun. At this stage, no feedback was provided, but participants knew that a sub-set of their choices would be executed at the end of the study and that they would receive the average outcome of these executions. Specifically, participants were instructed that: ” ... Your choices have real financial implications. At the end of the experiment, the computer will use its random number generator to execute some of your chosen gambles. ... You will receive the average outcome of the ... options the computer executes. ...”.

Training task. The choice problems in the training phase (Main study N=720, Replication study N=1200) were selected independently of the pre/post problems. Each training problem occurred exactly once (i.e., each training stimulus was unique). Due to a bug a fraction of the training problems was also included in the pre/post sessions (Main: 0.69%; Replication: 0%).

All participants and groups practiced identical choice problems presented in an identical order. Whenever feedback groups chose an option that did not maximize expected monetary returns, there was a two second delay (time penalty), after which the correct choice was highlighted, and participants were required to select it before they could continue (i.e., participants were required to approach the better option). Control groups did not receive feedback, but otherwise practiced identical and identically sequences choice problems.

Maximizing expected value is the optimal strategy for maximizing income in this study, and a frequently used standard in learning (28). For the purposes of demonstrating generalization and/or transfer it does not matter what participants learn – as long as it impacts their decisions for non-trained stimuli (generalization), and/or their decisions for non-trained stimuli in different risk domains (transfer). Thus, in principle we could also have, for example, trained participants to become less optimal by increased over- (or under) weighting of small probabilities. For ethical reasons, as well as ease-of-implementation, we choose to give expected value feedback, thus providing *training to maximize participants take-home bonus-pay*. Though future studies could explore, for example, training people become more linear only in how they weight probabilities when making trade-offs (i.e., to become expected utility maximizers). Though in this context, note that observed diminishing marginal utility for stakes in most risky choice tasks is not consistent with expected utility theory (“... expected-utility theory tells us that people will be virtually risk neutral in decisions on the scale of laboratory stakes”, p1286, 48).

Model fitting We use x and p to denote values and probabilities respectively, for the value function (1, SI Eq 1), we estimate α (only positive outcomes),

$$v(x) = \begin{cases} x^\alpha & \text{if } x \geq 0 \\ -\lambda(-x)^\beta & \text{if } x < 0 \end{cases} \quad \text{Eq 1}$$

We use Prelec’s (58) one-parameter probability weighting function,

$$w(p) = e^{(-(-\ln p)^\gamma)} \quad \text{Eq 2}$$

The value of a single choice option ($\psi(A)$), its prospect value, is the product of the weighted value and the weighted probability. A logistic choice function maps differences in

prospects onto choice probabilities, taking into account peoples limited ability to discriminate between prospects (59), whilst potentially capturing other sources of noise:

$$p_B = \frac{1}{1 + e^{k(\varphi(A) - \varphi(B))}}, \quad \text{Eq 3}$$

Differences in prospects were standardized prior to fitting allowing contrasts between precision parameters across a range of value and probability weighting parameters (N.B., this linear transform does not affect probability and value weight parameter estimation).

The model was fit to individual participants' choices separately for each risk domain by maximum likelihood methods using Matlab's Multistart solver. Denoting the choice of prospect A as r and a choice of prospect B as $1-r$, the negative log-likelihood of the value weight (α), the probability weight (γ) and the noise weight (k) given participants choices was minimized:

$$-L(\alpha, \gamma, k) = - \sum_{i=1}^n (r_i \log(P_{A_i}) + (1 - r_i) \log(P_{B_i})) \quad \text{Eq 4}$$

Model parameters were constrained, corresponding to uniform Bayesian priors on the parameters. The probability and value parameters were constrained to lie between .01 and 100; a range sufficient to capture both extreme under and overweighting. The precision parameter was constrained to lie between the smallest possible positive floating-point number and 50; a range sufficient to capture both extremely low and extremely high precision.

Note that Prospect Theory was fit as a measurement model, to capture the effect training had on people's preferences for the two key attributes: money and probability. The fitting should not be considered an empirical test of Prospect Theory for which there exists known deviations (e.g., 60), nor should the fitting of Prospect Theory be taken as a

commitment to a specific choice mechanism. Instead, Prospect Theory was fit because it accounts for behaviour such as that studied here very well (2,10) with highly recoverable parameters (*SI Model recovery of changes in risk-reward trade-offs*).

Analyses All analyses and model fits were carried out in Matlab, except for ANOVAs which were performed in R. A mixture of robust non-parametric and parametric statistics, and bootstrap (61) methods were used. Raw performance analyses were performed on the proportion of choices maximizing expected value. Robust regression and t-tests were used to analyze performance across training episodes (Fig 2). To analyse training-induced changes in performance (Fig 3), post-training performance was subtracted from pre-training performance separately for each participant and risk domain. Group means on the performance differences and their 95% confidence intervals allow inference about group-level learning and were supplemented by paired t-tests and repeated measures ANOVA.

Whenever the null hypothesis of no change was of theoretical interest, JZS Bayes Factors (30) were used for inference. We express BF's as evidence in favour of the null hypothesis. BF= 2, for example, means that the null hypothesis of no difference between means (for a t-test) is twice as likely as the alternative hypothesis. Although BF's indicate evidence for one hypothesis over another continuously, it is common to use arbitrary significance criteria (equivalent to $p < .05$ for NHST), with a cut-off of BF=3 for 'significance', with less evidence interpreted as inconclusive (62).

For inferences with respect to average Prospect Theory parameters describing (group) risk-reward trade-offs (Fig 4), we used trimmed means (50%) with parametrically bootstrapped null distributions, generated as follows: Prospect theory was fit to pre-training data separately for each participant and risk domain. We then simulate agents performing the pre- and post-training session with the best-fit preferences. The group-vectors describing the change from pre- to post for these simulated agents are the vectors one would expect if

participants maintained their risk-reward trade-offs following training (i.e., if training had no effect). The method makes no assumptions about the distribution of parameters and is robust to outlying parameter estimates. These analyses were supported by non-parametric Wilcoxon signed rank tests on the best-fit parameters. Effect sizes for t-tests reflect Cohen's D , for non-parametric tests of differences rank-biserial correlation r and for ANOVA's partial-eta squared η_p^2 .

Missing data Some participants had missing training data. Data was missing for one of two reasons. Firstly, some participants experienced file corruption during the training phase. These participants were instructed to delete the corrupted file and began training anew. Secondly, some participants missed some training episodes. All analyses, except analysis of the training data, also includes participants with missing training data. This allows a principled analysis of the training data (analysing only participants with identical training experience) and provides powerful (using all available data) yet conservative estimates of training-induced effects in the post-training session (assuming that less training results in the same or weaker effects).

Methods – BART Study

Participants. We conducted a pilot study with first-year Psychology Undergraduates participating for course credit ($N=21$). Based on this study, we ran a full study with incentivized participants. In both studies, participants provided written consent, were fully informed and debriefed. In the full study, participants were paid £15 for ~2 hours and an additional performance-related bonus pay (average total pay £21, or ~£10.5/hr). Inclusion criteria matched that of the Main study (see *Main Study: Participants*). An administrative error resulted in the recruitment of some participants age 40 and over. To match the age constraint in the Main study these were excluded. Additionally, one participant attempted to participate twice (using two different participant panel accounts), and a data file for one

participant was accidentally overwritten. These participants were also excluded. The final sample for the full study was N=31. Both studies were performed at City, University of London and approved by the local ethics committee (ETH1819-0270; ETH1920-0778).

Procedure: Pre-training, training and post-training were run in a single experimental session at the Department of Psychology lasting approximately 2 hours. The study involved three tasks: pointing (Fig 1), choice (Fig 1) and BART (Fig 5A, and see below). Participants were familiarized with the response time requirements for the pointing task, after which they completed the pointing task proper. After this, participants completed the choice task and the BART task (order counterbalanced across participants). This was followed by the training part ('Description' - feedback, Fig 1), which was followed by the post-training part (a repeat of the choice task and the BART task). The pilot study was also run in a single session, but with fewer trials (~1.5 hrs), involved the mental arithmetic task (instead of the pointing task), and did not counterbalance task order.

Apparatus. Stimuli were displayed on a touchscreen (Iiyama T2245MSC) and experiments were written in MATLAB (Mathworks) using PsychToolbox (56) on Linux with a soft real-time kernel (Xubuntu 18.04). The pilot study was run in single-user computer booths on computers running Windows.

Pointing task, Choice task and training task. The pointing task and choice task mirrored those of the Main study (Fig 1, Fig 3). For the pointing task, there were a minimum of 200 trials (with up to 40 extra trials to compensate for time-outs). The choice task involved 125 choice trials per risk domain (Description & Motor) for a total of 250 choice trials in the pre- and post-training parts respectively. The training task involved 180 trials.

BART. The BART task (64-66) involved three blocks of 30 trials for a total of 90 trials. At the beginning of each trial, a single small red balloon was displayed centrally (Fig 5A). Trial-wise accumulated points were displayed to the left of the balloon and block-wise

accumulated points were displayed to the right of the balloon. Each trial involved choices between an attempt to inflate the balloon or cashing the current trial-wise earnings. The BART task was modelled after 66, with the probability of the balloon breaking, given a balloon pump, constant across trials and blocks: $p(\text{break}|\text{pump}) = .15$ (with $N=90$ to reduce measurement noise).

A successful balloon pump (the balloon did not break) was rewarded with 1 trial-wise point and resulted in the virtual balloon increasing in size (Fig 5A). If the balloon broke, this was illustrated with a cartoon-like drawing with the text “BANG” replacing the balloon, and the trial-wise point counter was reset to 0 (Fig 5A). If participants chose to cash their current trial-wise earnings, these were added to their block-wise earnings, and the trial-wise counter was reset to 0. The overall goal in this task was to earn as many points as possible, and participants were aware that their final monetary bonus payment was monotonically related to their virtual earnings.

Analyses. The purpose of the BART Study was not to demonstrate generalization and transfer *in general* (Fig 2-4), but that learning in the canonical task can transfer to a very different task. The analyses therefore focus on those participants who showed evidence of learning, for whom transfer is expected (*if* training transfers). Thus, analyses exclude participants whose proportion of optimal choices did not increase in the trained-on risky domain (‘Description’, BART Study $N=4$, Pilot Study $N=7$). Note, this restriction does not guarantee transfer, which is defined as improvement following training on the un-trained domain (‘Motor’), nor does it guarantee a relationship between the trained-on domain and BART.

We first verified that learning to make better choices for ‘Description’ transferred to ‘Motor’ choices (Fig 5A) for the shorter training regime. This analysis replicates that in Fig 3 (and was significant in the Pilot Study: $t(13) = 3.50$, $p = .0039$). We also compared the

magnitude of transfer in the Main study to that in the BART study to more directly evaluate the impact of the length of training. To enable a fair comparison, only the stimuli in the Main Study which were identical to that in the BART Study were selected for analysis, and like the BART analyses include only those participants who showed an improvement (i.e., Fig 5B is not identical to ‘Description-Feedback’ – ‘Motor’ in Fig 3). We then compared, by an unpaired t-test, the size of the transfer effects in the Main and the BART study (Fig 5B).

The focus for the BART study was on the association between behaviour in the two key tasks: risky choice and BART. For this analysis, the dependent variable in the choice task (‘Description’, ‘Motor’) was the proportion of risky choices made in each domain by each participant (i.e., the probability of choosing the higher-reward lower-probability option). The dependent variable in the BART task was the mean number of balloon inflations on successfully terminated trials (64).

BART has a potential learning component and the first 10 trials were excluded as recommended (64). However, we did not find that different criteria for removing trials affected the results. Specifically, including all trials, or removing the whole first block (30 trials), resulted in broadly consistent results (with the exception of the ‘Motor’ choice effect which went from marginal to significant for one analysis). The robustness is consistent with previous work (64) and the absence of a reliable trend across trials (*SI Figure 9 - BART study supplementary analyses*).

Based on pilot data (and 49), we expected behaviour at base-line (pre-training) to be unrelated across the two tasks. Specifically, we expected the probability of choosing a risky option - $p(\text{risky choice})_{\text{pre-training}}$ – to be unrelated to the mean number of balloon inflations on successfully terminated trials - $\text{balloon pumps}_{\text{pre-training}}$. For the Pilot study we observed no significant association with BART at pre-training for either the ‘Description’ ($r(12) = .284, p = .325$) or the ‘Motor’ choice task ($r(12) = .393, p = .164$).

On the other hand, we expected a positive association between behavioural changes from pre-training to the post-training. In the Pilot, the association between ‘Description’ and BART was significant (Description’ ($r(12) = .56, p = .036$), and there was a weak trend towards an association between ‘Motor’ and BART ($r(12) = .432, p = .122$). Thus, we expected training-induced increases in risk-taking (choice task) to be predictive of an increased risk-taking in the BART task. Accordingly, we computed the difference in risky choice across post/pre-training ($\Delta risky\ choice = p(risky\ choice)_{post} - p(risky\ choice)_{pre}$), and the difference in BART across post/pre-training $\Delta BART = balloon\ pumps_{post} - balloon\ pumps_{pre}$.

For these analyses of association, we used least-squares (Pearson’s r). However, for robustness, we also report and plot equivalent robust linear regression analyses.

Data availability

All data will be made publicly available upon publication (through the Open Science Framework). Analysis code is available upon request.

Acknowledgments

Thanks to Chris Summerfield and his research group for discussions and comments, especially Dan Bang, Santiago Hernandez and Keno Juechems. Andreas Jarvstad was supported by a British Academy Postdoctoral Fellowship (D-MAD, PF150005). The data will be made publicly available upon publication (Open Science Framework). Analysis code is available upon request. The author declares no competing interest.

- 1 M. Allais, The foundations of a positive theory of choice involving risk and a criticism of the postulates and axioms of the American school. In M. Allais and O. Hagen (eds.), *Expected utility and the Allais Paradox*, 27 - 145. New York. (1952/1979, Springer).
- 2 D. Kahneman, A. Tversky, Prospect theory: An analysis of decision under risk. *Econometrica*, 47, 263–292 (1979).
- 3 J. Rieskamp, The probabilistic nature of preferential choice. *J. Exp. Psychol. Learn. Mem. Cogn.*, 34(6), 1446–1465 (2008).
- 4 R. Hertwig, G. Barron, E.U. Weber, I Erev. Decisions from experience and the effect of rare events in risky choice. *Psych. Sci.*, 15, 534–539 (2004)
- 5 A. Camilleri, B. Newell, When and why rare events are underweighted: A direct comparison of the sampling, partial feedback, full feedback and description choice paradigms. *Psych. Bull. Rev.*, 18, 377– 384 (2011).
- 6 D.U. Wulff, M. Mergenthaler-Canseco, R. Hertwig. A meta-analytic review of two modes of learning and the description-experience gap. *Psychol Bull*, 144(2), 140 (2018)
- 7 V. Navalpakkam, C. Koch, A. Range, P Perona, Optimal reward harvesting in complex perceptual environments, *Proc. Natl. Acad. Sci. U.S.A.*, 107, 5232–5237 (2010).
- 8 J. Trommershäuser, L.T. Maloney, M.S. Landy, Decision making, movement planning and statistical decision theory. *Trends Cogn. Sci.*, 12, 291–297 (2008).
- 9 L. Whiteley, M. Sahani, Implicit knowledge of visual uncertainty guides decisions with asymmetric outcomes. *Journal of Vision*, 8, 1–15 (2008).
- 10 A. Jarvstad, U. Hahn, S.K. Rushton, P.A. Warren, Perceptuo-motor, cognitive, and description-based decision-making seem equally good. *Proc. Natl. Acad. Sci. U.S.A.*, 110(40), 16271–16276 (2013).
- 11 J. Trommershäuser, L.T. Maloney, M.S. Landy, (Statistical decision theory and the selection of rapid, goal-directed movements. *J. Opt. Soc. Am. A. Opt. A.*, 20, 1419–1432 (2003)
- 12 A. Jarvstad, U. Hahn, P.A. Warren, S.K. Rushton, Are perceptuo-motor decisions really more optimal than cognitive decisions?. *Cognition*, 130(3), 397–416 (2014).
- 13 D. Rahnev, R Denison, Suboptimality in Perceptual Decision Making. *Behav. Brain Sci.*, 1–107 (2018).
- 14 R.H. Thaler, C.R. Sunstein, *Nudge: Improving decisions about health, wealth, and happiness*. New Haven, (1999, CT Yales University Press).
- 15 R.P. Larrick. Debiasing. *Blackwell handbook of judgment and decision making*, 316–338 (2004)
- 16 R.H. Thaler, S. Benartzi, Save more tomorrow™: Using behavioral economics to increase employee saving. *J. Polit. Econ.*, 112(S1), S164–S187 (2004).
- 17 L. Torrey, J. Shavlik, Transfer learning. *Handbook of Research on Machine Learning Applications and Trends: Algorithms, Methods, and Techniques*, 1, 242 (2009).
- 18 B.M. Lake, T.D. Ullman, J.B. Tenenbaum, S.J. Gershman, Building machines that learn and think like people. *Behav. Brain Sci.*, 40 (2017).
- 19 E.L. Thorndike, *Animal Intelligence*. New York (1911, Macmillan)
- 20 J.W. Kable, M.K. Caulfield, M. Falcone, M. McConnell, L. Bernardo, T. Parthasarathi, P. Diefenbach, No Effect of Commercial Cognitive Training on Brain Activity, Choice Behavior, or Cognitive Performance. *J Neurosci.*, 37(31), 7390–7402 (2017).
- 21 Foroughi, C. K., Monfort, S. S., Paczynski, M., McKnight, P. E., & Greenwood, P. M. (2016). Placebo effects in cognitive training. *Proc. Natl. Acad. Sci. U.S.A.*, 113(27), 7470–7474.
- 22 C.H. Judd, The relation of special training and general intelligence. *Educ. Rev.*, 36, 28–42 (1908).
- 23 R.E. Nisbett, D.H. Krantz, C. Jepson, Z. Kunda, The use of statistical heuristics in everyday inductive reasoning. *Psych. Rev.*, 90(4), 339 (1983).

- 24 S.M. Jaeggi, M. Buschkuhl, J. Jonides, W.J. Perrig, Improving fluid intelligence with training on working memory. *Proc. Natl. Acad. Sci. U.S.A.*, 105(19), 6829-6833 (2008).
- 25 S.J. Goodbody, D.M. Wolpert, Temporal and amplitude generalization in motor learning. *Journal of Neurophysiology*, 79(4), 1825-1838 (1998).
- 26 R.K. Jessup, A.J. Bishara, J.R. Busemeyer, Feedback produces divergence from prospect theory in descriptive choice. *Psych. Sci.*, 19, 1015-1022 (2008).
- 27 D.R. Shanks, R.J. Tunney, J.D. McCarthy, A re-examination of probability matching and rational choice. *J. Behav. Decis. Mak.*, 15, 233 – 250 (2002).
- 28 I. Erev, E. Ert, O. Plonsky, D. Cohen & O. Cohen, From anomalies to forecasts: Toward a descriptive model of decisions under risk, under ambiguity, and from experience. *Psychological review*, 124(4), 369 (2017).
- 29 R.S. Sutton, A.G. Barto. *Reinforcement learning: An introduction*. (1998, MIT Press: Cambridge, MA).
- 30 S-W, Wu, M.R. Delgado, L.T. Maloney, Economic decision-making compared with an equivalent motor task. *Proc. Natl. Acad. Sci. U.S.A.*, 106, 6088-6093 (2009)
- 31 D.N. Lee, A theory of visual control of braking based on information about time-to-collision. *Perception*, 5(4), 437-459 (1976).
- 32 A. Tversky, C.R. Fox. Weighting risk and uncertainty. *Psych. Rev.* 102, 269-283, (1995).
- 33 R.E. Kass, A.E. Raftery, Bayes Factors. *J. Am. Stat. Assoc.*, 90, 773-795, (1995).
- 34 A. Tversky, D. Kahneman, Advances in prospect theory: Cumulative representation of uncertainty. *J. Risk Uncertain.*, 5, 297-323 (1992).
- 35 S-W Wu, M.R. Delgado, L.T. Maloney, The neural correlates of subjective utility of monetary outcome and probability weight in economic and in motor decision under risk. *J. Neurosci.*, 31, 8822-8831 ((2011).
- 36 C.F. Camerer, Recent tests of generalizations of expected utility theory. In *Utility theories: Measurements and applications* (pp. 207-251). Springer, Dordrecht. (1992).
- 37 S.J. Kachelmeier, M. Shehata, Examining risk preferences under high monetary incentives: Experimental evidence from the People's Republic of China. *Am. Econ. Rev.* 1:1120-41 (1992).
- 38 J.R. Anderson, *The adaptive character of thought*. Hillsdale, NJ: (Lawrence Erlbaum Associates,1990)
- 39 J. Selk, Amidst billion-dollar brain fitness industry, a free way to train your brain. *Forbes*. Available at www.forbes.com/sites/jasonselk/2013/08/13/amidst-billion-dollar-brain-fitness-industry-a-free-way-to-train-your-brain/#6b3457647d41. Downloaded 2018/03/24 (2013).
- 40 J.A. Clithero, A. Rangel, Informatic parcellation of the network involved in the computation of subjective value. *Soc. Cogn. Affect.*, 9(9), 1289-1302 (2013).
- 41 R. Bogacz, P.T. Hu, P.J. Holmes, J.D. Cohen, Do humans produce the speed-accuracy trade-off that maximizes reward rate? *Q. J. Exp. Psychol. A.*, 63,863-891 (2010).
- 42 S. Frederick, G. Loewenstein, T. O'Donoghue, Time discounting and time preference: A critical review. *J. Econ. Lit.*, 40(2), 351-401 (2002).
- 43 J.M. Hotaling, A. Jarvstad, C. Donkin, B.R. Newell. How to change the weight of rare events in decisions from experience. *Psych. Sci.* (In Press).
- 44 I.P. Levin, J.A. Weller, A.A. Pederson, L.A. Harshman, Age-related differences in adaptive decision making: Sensitivity to expected value in risky choice. *Judg. Decis. Mak.* 2, 225-233 (2007).
- 45 R. Mata, A.K. Josef, G.R. Smanez-Larkin, R. Hertwig. Age differences in risky choice: a meta analysis. *Annal. N.Y. Acad. Sci.* 1235, 18-29. (2011).
- 46 E. Karni, D. Schmeidler, K. Vind, On state dependent preferences and subjective probabilities. *Econometrica*, 1021-1031 (1983).
- 47 M.W. Khaw, P.W. Glimcher, K. Louie, Normalized value coding explains dynamic adaptation in the human valuation process. *Proc. Natl. Acad. Sci. U.S.A.* 114(48), 12696-12701 (2017).
- 48 N. Chater. *The mind is flat: The illusion of mental depth and the improvised mind*. Penguin, UK (2018).

- 49 E.U. Weber, A.R. Blais, N.E. Betz, A domain-specific risk-attitude scale: Measuring risk perceptions and risk behaviors. *J Behav. Desc.-Mak*, 15(4), 263-290 (2002).
- 50 R. Frey, A. Pedroni, R. Mata, J. Rieskamp, R. Hertwig, Risk preference shares the psychometric structure of major psychological traits. *Sci. Adv.*, 3(10), e1701381 (2017).
- 51 L.L. Lopes, The rhetoric of irrationality. *Theory Psychol.*, 1(1), 65-82 (1991).
- 52 J. Dickhaut, A. Rustichini, V. Smith, A neuroeconomic theory of decision process. *Proc. Natl. Acad. Sci. U.S.A.*, 106, 22145–22150 (2009).
- 53 F. Lieder, T.L. Griffiths, M. Hsu, Overrepresentation of extreme events in decision making reflects rational use of cognitive resources. *Psych. Rev.*, 125(1), 1, (2018).
- 54 J.M. Beck, W.J. Ma, X. Pitkow, P.E. Latham, A. Pouget, Not noisy, just wrong: the role of suboptimal inference in behavioral variability. *Neuron*, 74(1), 30-39, (2012).
- 55 M. Kleiner, D. Brainard, D.Pelli, “What’s new in Psychtoolbox-3?” Perception 36 ECV Abstract Supplement (2007).
- 56 M. Rabin, Risk aversion and expected-utility theory: A calibration theorem, *Econometrica*, 68, 1281-1292 (2000).
- 57 D. Prelec, The probability weighting function. *Econometrica*, 66, 497–527 (1998).
- 58 F. Mosteller, F. Nogee, An Experimental Measurement of Utility. *J. Polit. Econ.* 59, 371–404 (1951).
- 59 M.H. Birnbaum, New paradoxes of risky decision making. *Psych. Rev.* 2008 Apr;115(2):463.
- 60 B. Efron, R.J. Tibshirani, *An introduction to the bootstrap*. (CRC press,1994).
- 61 H. Jeffreys, *Theory of Probability*, (Oxford University Press, 1998)
- 62 R. Hertwig, T. Gruene-Yanoff, Nudging and Boosting: Steering or Empowering Good Decisions, *Persp. Psych. Sci.*, 12(6), 973-986 (2017).
- 63 C.W. Lejuez, J.P. Read, C.W. Kahler, J.B. Richards, S.E. Ramsey, G.L. Stuart, D.R. Strong, R.A. Brown. Evaluation of a behavioral measure of risk taking: the Balloon Analogue Risk Task (BART). *J. Exp. Psych. Applied*, 8(2), 75, (2002).
- 64 T.L. White, C.W. Lejuez, & H. de Wit, Test-retest characteristics of the Balloon Analogue Risk Task (BART), *Exp. Clin. Psychopharmacol.*, 16(6), 565–570, (2008).
- 65 D. van Ravenzwaaij, G. Dutilh, E-J. Wagenmakers, Cognitive model decomposition of the BART: Assessment and application, *J. Math. Psych.*, 94–105, (2011)

828

829