The status of human goal-directed actions after manipulations of the action-outcome contingency

Omar David Perez[1,3,4,*], Sarah Oh[2,3,*], Anthony Dickinson[5], Jose Arenas[1], & Emiliano Merlo[6]

[1] Department of Industrial Engineering, University of Chile, Santiago, Chile.

[2] Department of Psychology, University of California, Berkeley, CA, USA.

[3] Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, California, USA.

[4] Instituto Sistemas Complejos de Ingeniería (ISCI), Chile.

[5] Department of Psychology and Behavioural and Clinical Neuroscience Institute, University of Cambridge, Cambridge, UK.

[6] School of Psychology, University of Sussex, Falmer, UK.

Author Note

Correspondence regarding this article should be sent to Omar D. Perez (omar.perez.r@uchile.cl) or Emiliano Merlo (e.merlo@sussex.ac.uk). * Authors contributed equally to this work.

Abstract

Goal-directed actions are those performed with the expectation of producing a specific outcome and are therefore sensitive to changes in the outcome's current value without re-experiencing the action-outcome contingency. In this study, we examined whether human action remains value-sensitive after experiencing the removal of the action-outcome contingency. Across two experiments, participants were trained to perform different actions for distinct outcomes before the outcomes were omitted for the target actions in the extinction experiment and presented independently of responding in the non-contingent experiment. One outcome was then devalued before responding was tested in the absence of outcome presentations. An overall analysis revealed that outcome devaluation decreased responding in this test to a degree that was not significantly reduced by either the extinction or non-contingent training. These experiments replicate with human participants the insensitivity of outcome devaluation to response extinction initially reported by Rescorla (1993) with rats and extend it to non-contingent training.

Keywords: goal-directed, extinction, non-contingent learning, humans

The status of goal-directed actions after manipulations of the action-outcome contingency

Introduction

A wealth of evidence in the last 40 years or so has demonstrated how actions are controlled by either goal-directed or habitual systems, which can be probed behaviorally, computationally, and neurally. Under goal-directed control, actions are directed towards obtaining currently valuable outcomes, whereas habitual responding is simply strengthened or reinforced by the outcome without representing it as a goal (Daw & O'Doherty, 2013; Dickinson, 1985; Dolan & Dayan, 2013; Perez & Dickinson, 2020).

The goal-directed status of instrumental responding is evident by its sensitivity to the current value of the outcome in an outcome revaluation test (Adams & Dickinson, 1981). After training an animal to associate an action with an outcome (e.g., pressing a lever for a food outcome), the value of the outcome is changed, usually by satiating the animal with the outcome or by pairing it with an aversive stimulus. If the action is goal-directed, the animal will reduce the action associated with the now devalued outcome even in the absence of the outcome thereby reflecting sensitivity to the current value of the outcome as a goal. Conversely, if the action becomes habitual, the animal will continue to perform the action despite outcome devaluation, indicating that the behavior is driven by past reinforcement rather than by the current desirability of the outcome. This paradigm is critical in distinguishing between these two types of behavioral control and has been extensively used to explore the neural circuits underlying goal-directed and habitual behaviors (Daw et al., 2005; Dolan & Dayan, 2013; Lee et al., 2014; Tricomi et al., 2009; Yin & Knowlton, 2006).

Table 1. **Design of experiments (based on Rescorla, 1993).** R and O denote responses and outcomes, respectively.

Training	Extinction or degradation	Retraining	Devaluation	Choice test
S1:	**Extinction experiment**	R1 → O3	Devalue O1	S1: R1 vs R2
R1 → O1				
R2 → O2	R1 → nothing	R2 → O3		
S2:	R2 → nothing			S2: R3 vs R4
R3 → O1		R3 → O3		
R4 → O2	**Degradation experiment**	R4 → O4		
	R1 → nothing / O1 free			
	R2 → nothing / O2 free			

Rescorla and colleagues used outcome revaluation experiments in a sophisticated analysis of the associative representations mediating such goal-directed responding (Rescorla, 1993). One of Rescorla's most surprising discoveries is that extinction of the response by omitting the outcome it was trained with does not appear to impact the strength of an established response–outcome (R-O) association mediating goal-directed control (Rescorla, 1993). The design of one of his experiments is illustrated in the extinction component of Table 1. Initially, he trained his hungry rats to perform four different instrumental responses with two of the responses (R1 and R3) being reinforced with one food outcome (O1), and the other two responses (R2 and R4) with a second food outcome (O2). Having established responding, two of the responses (R1 and R2) were extinguished by omitting their respective outcomes.

At issue was whether this extinction had weakened the R–O association mediating goal-directed control. However, in order to assess such control with an outcome revaluation test following extinction, responding had to be re-established using a third outcome (O3). In this way, the level of responding was comparable between extinguished and non-extinguished responses at the outset of the final test. In this test, Rescorla devalued O1 by conditioning a food aversion to it in the absence of the opportunity to perform any of the responses. Finally, the animals were given a choice between R1 and R2 and between R3 and R4 in the absence of any outcomes. Testing in absence of a contingent outcome is critical if responding is to reflect the extent to which devaluation impacts behavior via the R–O association, as any decrease in responding could otherwise be attributed to punishment by experiencing the aversiveness of the outcome after devaluation.

The fact that R3, originally trained with the now-devalued O1, was performed less on test than R4, trained with the still-valued O2, confirmed that the initial instrumental training established goal-directed control over responding. Surprisingly, however, the magnitude of the outcome devaluation effect revealed that the contrast between R1 and R2 was similar to that observed for R3 versus R4. Extinction of R1 and R2 by the total removal of their training outcomes had no effect on the extent to which the current value of these outcomes impacted performance. The theoretical implication is that extinction reduces responding by establishing a new R-no-outcome association, without decreasing the strength of the R–O association established during training.

More recent experiments by Crimmins and colleagues (2022) examined the impact of removing the action-outcome contingency by presenting the outcome independently of responding rather than by simple extinction. They reported that, at least under certain circumstances (see Discussion), such non-contingent training also had no effect on the magnitude of goal-directed control as assessed by outcome devaluation by specific satiety in rats.

Given the theoretical importance of these findings, and the possibility that humans may employ more complex causal action–outcome representations and assign value to actions differently than other species, our experiments examined the impact of simple extinction and non-contingent training on goal-directed control in human participants. To do so, we assessed the effects of these contingency manipulations using a paradigm in which instrumental actions consisted of key presses and rewards were delivered as token outcomes.

Methods

Participants and apparatus

We recruited participants from the Prolific online platform (www.prolific.com). The total number of participants for each experiment was 117 for the extinction experiment and 109 for the non-contingent experiment. As there were no previous experiments testing the impact of outcome devaluation on human free-operant responding after manipulations of the action-outcome contingency, we did not perform a power analysis before data collection. However, we aimed to reach 100 participants per experiment, which is considerably higher than similar designs on human action learning. Participants were required to be located in the US or UK, fluent in English, and were paid 8 GBP per hour for their participation. The study was approved by the University of Sussex Sciences & Technology Cross-Schools Research Ethics Committee (SCITEC) (approval ER/EM540/1).

Procedure

Both experiments implemented the same experimental design outlined in Table 1 except for the fact that the instrumental contingency was degraded by omitting all outcomes in the extinction experiment and by presenting outcome independent of responding in the non-contingent experiment. At

the beginning of each experiment, participants were told that their goal was to press keys to obtain coin outcomes of two different types, red and blue, and that their goal was to obtain as many coins as they could. They were also told that there was a third type of coin, green, which could also be earned in some blocks of the experiment, and that this coin was worth the same as the other two. Participants were also informed that the collected coins would be deposited in different fictitious piggy banks of limited size, one for each type of coin.

Each experiment implemented the design illustrated in Table 1 in which two actions were exposed to the degraded contingency (R1 and R2), whereas the other two actions (R3 and R4) did not undergo any manipulation and functioned as within-subject controls. In each block of training, participants had the opportunity to press two keys, either the two top keys, "q" or "p", or the two bottom keys, "z" or "m". The location of the keys on the screen respected their relative position on the keyboard (e.g., "z" was shown on the bottom left side of the screen whereas "p" was shown on the top right location). For each participant, one of the top keys delivered red coins and the other one to blue coins. To avoid participants attributing the type of coin with the side of the response, which would have associated with the two outcomes to different hands, the opposite assignment was programmed for the bottom keys. Therefore, actions performed with one of the hands always led to different outcomes across different blocks of training. For example, if for one participant the assignment for top keys was q → red and p → blue, then for the next block the assignment was z → blue and m → red.

The sequence of stages in the procedure is illustrated in Figure 1. Initially, participants underwent three 120-s blocks of concurrent training for each pair of actions. The starting pair of actions was randomized across participants and the blocks were presented sequentially once this first assignment was set per each participant. Each action delivered its outcome under a random-interval (RI) schedule, so that on average the appropriate coin for each response became available to be collected

every T seconds. This was implemented by programming two independent clocks, each generating a reward with a probability of 1/T per second. The schedule for each action was capped so that a reward was guaranteed to be available for collection by a response if three times the scheduled inter-outcome interval T had elapsed. To minimize accidental reinforcement of switching between responses, we introduced a changeover count: only the second response after a switch could produce a coin, even if the schedule on the new action had already made a reward available while the participant was engaged in the other response. The parameter T was set to 12 s for the extinction experiment and 15 s for the degradation experiment.

Once a total of three blocks of training for each pair of actions finished, participants were presented with a message informing them that the piggy banks were filling up (see “Devaluation” screenshot in Figure 1). This information was added to remind participants that their coins were being deposited and counted during the experiment and to familiarize them with the stimuli that would be presented at the end of the experiment after devaluation. After this screen, participants were presented with 30 coins located in random positions on the screen, 10 of each type, and asked to collect a total of 10 coins that would be deposited into their respective piggy banks. This free consumption phase (“Free coin collection” in Figure 1) was followed by the 15-min extinction or contingency degradation stage wherein two actions (either "q" and "p", or else "z" and "m") were presented and outcomes were omitted in the extinction experiment or followed by the non-contingent presentation of both outcomes each on an independent random time 15-s schedule in the non-contingent experiment. On this latter schedule, participants could achieve the same outcome rate irrespective of whether they responded or not, thereby generating a zero contingency between action and outcome[1].

[1] The use of a random time schedule to generate non-contingent outcomes differs from that employed by Crimmins et al (2022). They trained their rats on a ratio rather than interval schedule and therefore rendered the outcomes non-contingent by arranging that the outcome probability in the absence of a response was the same as that following a response.

After the manipulation phase, participants were given the opportunity to perform a 120-s block of retraining of all actions, which was similar to those in the training phase except for the fact that the outcome was the green coins for all responses. Participants were informed that these coins were worth the same as the blue and red ones, and that they would also be deposited into their respective piggy banks. In this manner, we ensured that participants were performing responses at comparable rates at the end of training given that extinction or degradation should have reduced responding during the manipulation phase. This retraining phase was followed by devaluation of one of the outcomes. In this phase, participants were presented with a screen showing one of the piggy banks being closed so that no more coins of that type could be deposited in the piggy bank, effectively devaluing that outcome with respect to the others. To test if the devaluation phase succeeded in achieving this goal, we then presented participants with another consumption test where they were asked to collect 10 coins using mouse clicks in the same manner as the previous coin collection phase (see Figure 1; Gillan et al., 2015). During this free collection phase, we expected that participants would collect less of the devalued coin than the two other still-valued coins.

Participants were then given the critical outcome devaluation test at the start of which they were told that they could respond to obtain coins in the same manner as in the training phase for 90 s, but that this time the outcomes obtained would be covered by a curtain. They were informed that, apart from that, nothing about the experiment had changed. The coins obtained would be still deposited into their respective piggy banks (provided they were not closed). If human goal-directed behavior survives contingency degradation by extinction (Rescorla, 1993) and non-contingent training (Crimmins et al., 2022), we expect similar devaluation effects in both experiments.

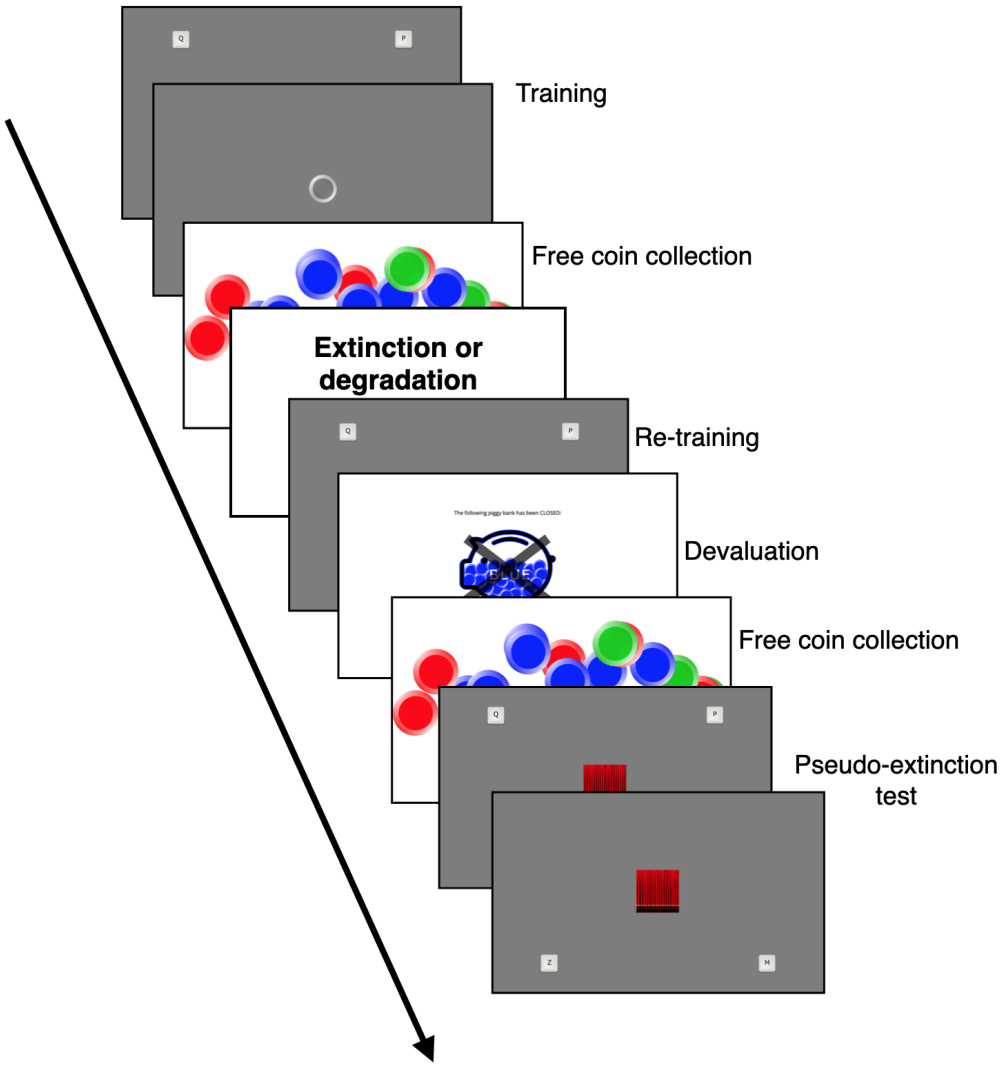

Figure 1. **Sequential presentation of stimuli in the task.** During training, participants were given 120 seconds of concurrent training to press either two top keys ("q" and "p") in one block of training, or two bottom keys ("z" and "m") in another block, followed by a re-training phase where they had the opportunity to press the two top or bottom keys concurrently. One of the two top keys led to red coins and the other to blue coins. The same applied to bottom keys.

Data and Materials Availability Statement

All data, study materials, and analysis code used in this study are openly available at the following OSF repository: https://osf.io/6xg37/. This includes raw data, preprocessing scripts, experimental materials, and R code for the statistical analyses reported in the paper. These resources are provided to support full transparency and reproducibility of the findings.

Results

All our analyses were performed in the R programming language under the RStudio IDE (RStudio Team, 2020). We did not discard any participants based on their performance, unless they stopped responding in more than one block during the training blocks. No participants met this criterion.

A mixed-effects analysis was conducted using the *lmer* function from the lme4 package (Bates et al., 2015), with a random intercept and by-subject random slope for value (valued vs. devalued). The between-participant variable of experiment distinguished performance when the action–outcome contingency was degraded either by extinction or by non-contingent outcomes. Two within-participant variables were included: the contingency variable, contrasting degraded responses (R1 and R2) with control responses (R3 and R4), and the value variable, contrasting responses associated with the devalued outcome (R1 and R3) versus those associated with the still-valued outcome (R2 and R4).

Figure 2A shows the mean response rates during the devaluation choice test for each condition. In both experiments, participants responded less for actions associated with the devalued outcome than for those associated with a still-valued outcome. However, neither the numerical means nor the statistical interactions indicated that the contingency manipulation reduced the magnitude of this devaluation effect. The analysis revealed a robust main effect of value ($b = 29.5$, $SE = 7.19$, $p < .001$), indicating that participants responded more for valued than for devalued outcomes. However, there was no main effect of contingency ($b = 3.81$, $SE = 3.84$, $p = .340$), and no significant interaction between

value and contingency ($b = -5.69$, $SE = 5.43$, $p = .313$). The interaction between value and experiment was also not significant ($b = -18.3$, $SE = 9.99$, $p = .068$), nor did the three-way interaction among value, contingency, and experiment ($b = 7.06$, $SE = 7.83$, $p = .368$). These results indicate that the magnitude of the outcome devaluation effect did not reliably differ between extinction and non-contingent training conditions.

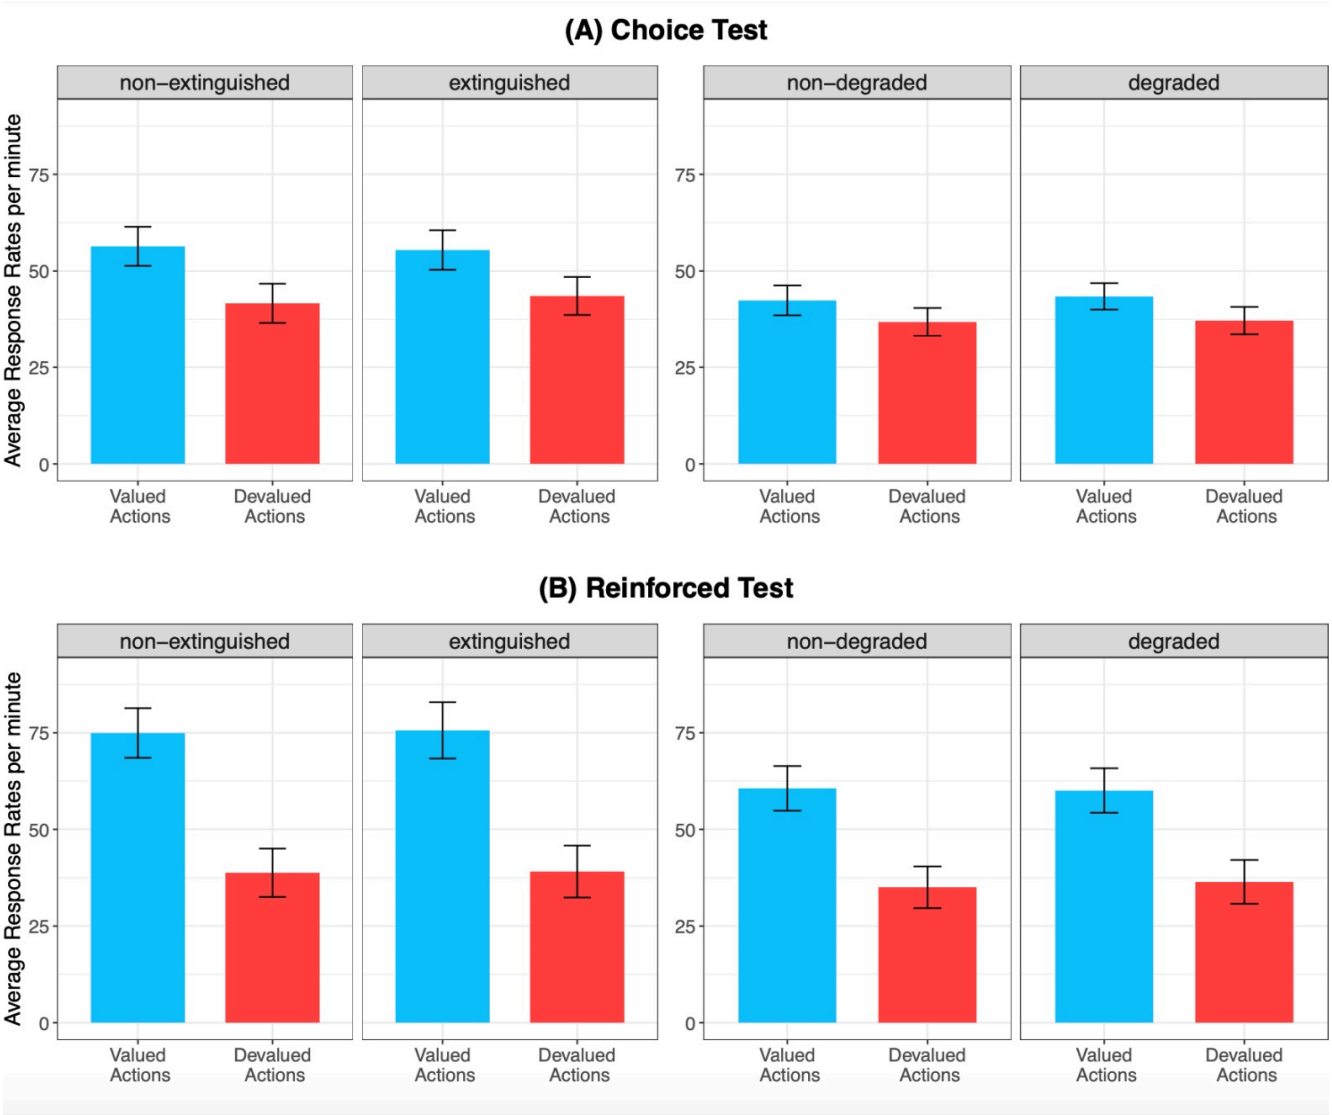

Figure 2. **Response rates in the tests for extinction and degradation experiments. A)** Average responses rates during the pseudo-extinction choice tests. **B)** Average response rates during the reinforced tests. Error bars represent 95% within-subject confidence intervals (Morey, 2008).

The panels of Figure 2B show the response rates during the reinforced test, in which participants received feedback on their actions (the outcomes were reintroduced contingent on responding, as during training). The same value-based pattern was observed: participants responded more for valued than for devalued outcomes. The value effect was numerically enhanced relative to the choice test, as indicated by a significant main effect of test phase ($b = 42.8$, $SE = 7.56$, $p < .001$) and a value × phase interaction ($b = 42.8$, $SE = 7.56$, $p < .001$), confirming that the devaluation effect was stronger in the reinforced than in the extinction-like test. However, this effect did not interact significantly with experiment ($b = 2.15$, $SE = 7.43$, $p = .775$), suggesting that the pattern was similar across extinction and degradation manipulations. Likewise, there was no reliable three-way interaction among value, phase, and experiment ($b = –11.7$, $SE = 14.9$, $p = .433$), indicating that the enhancement of the devaluation effect during the reinforced test generalized across both experimental conditions.

Figure 3A shows the final mean response rates during the last 2 min of each phase of the extinction and degradation experiments. An experiment × value × contingency analysis of these final training blocks revealed no significant main effects or interactions. This indicates that performance prior to the manipulation phase was comparable across conditions.

Figure 3B shows response rates across the contingency manipulation phase, separated by experiment. As expected, participants progressively reduced their response rate for the degraded actions over time. However, degraded response rates remained higher under non-contingent training than under extinction by the end of the manipulation phase. This description was supported by a significant

Experiment × Time-block interaction (b = -.252, SE = .037, p <.001), confirming that response suppression during the manipulation was stronger in the extinction condition.

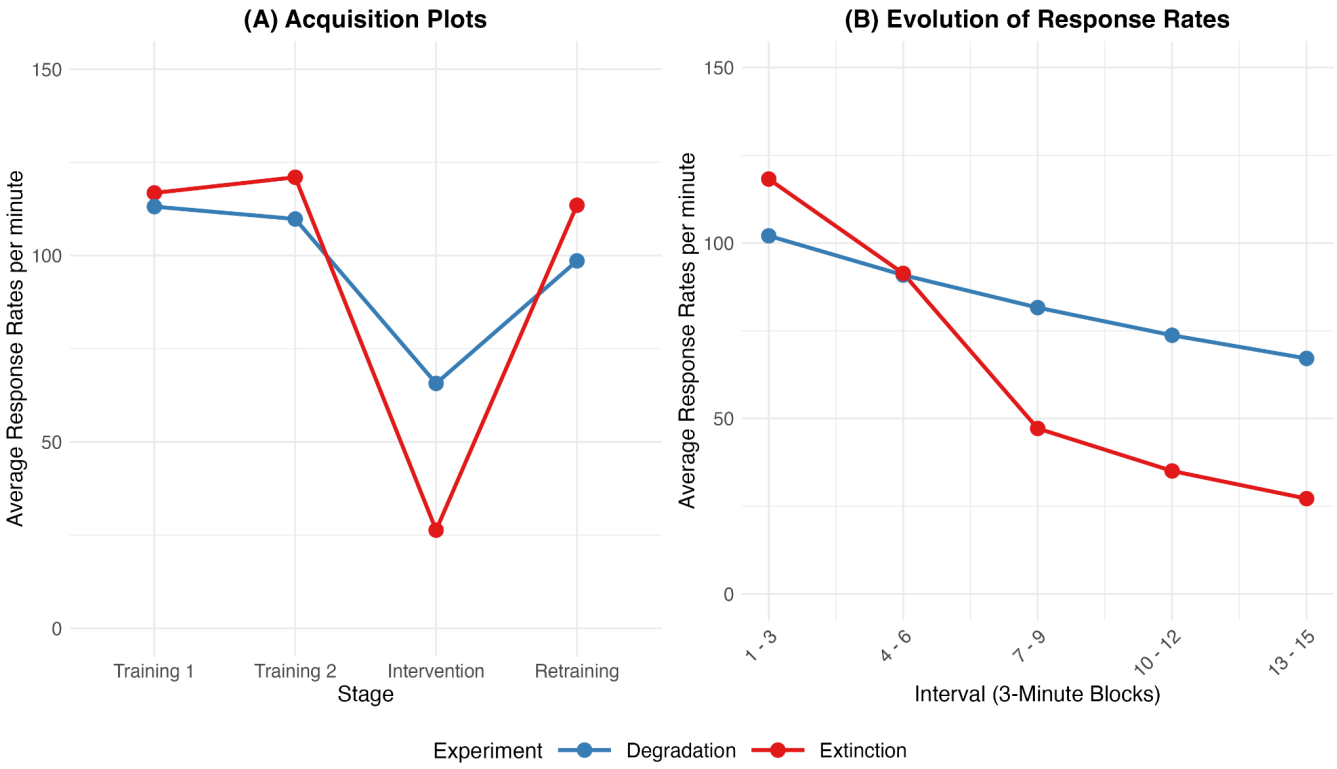

Figure 3. **Response rates for training and during the intervention phase. A)** Response rates during training, intervention phase and retraining phases. Means are calculated for the last 2 minutes of each phase. **B)** Response rates during the manipulation phases. Means are calculated for each of the 5 3-min blocks.

To explore whether individual differences in behavioral suppression during the manipulation phase predicted value sensitivity, we included the response suppression score (response rate in the last 2 min of the intervention phase minus response rate in the last 2 min of the last block of training) in the model and tested its interaction with outcome value. This analysis revealed a significant value ×

suppression interaction (b = –0.171, SE = 0.076, p = .025), indicating that participants who showed greater reductions in responding during the manipulation phase also showed a stronger devaluation effect. However, this effect did not interact reliably with experiment or contingency conditions. Thus, while the degree of suppression covaried with goal-directed performance across individuals, this relationship was consistent across both extinction and non-contingent training.

Discussion

In two experiments, we tested in humans how manipulations of action-outcome contingency affect sensitivity to devaluation of the outcomes produced by the actions. In the extinction experiment, both extinguished and non-extinguished actions were devaluation sensitive, which is consistent with the results obtained by Rescorla (1993) in rodents. Similarly, the second experiment replicated with human participants the Crimmins et al (2022) finding that the degradation of the instrumental contingency by training with non-contingent outcomes can also have no impact on the rodent outcome devaluation effect, at least when the actions involve responses on different manipulanda (see below).

The theoretical interpretation of these findings is complex. The theory of goal-directed learning offered by Perez and Dickinson (2020) is unique in providing an account of why such learning is resistant to extinction. Under this theory, it is the mean local correlation between response and outcome rates that determines the strength of goal-directed control. The local correlation is estimated across time samples in a working memory which records the number of responses and outcomes in each sample. Every time the memory recycles by starting a new sample and dropping an old one, the local correlation in the memory immediately prior to the recycle is estimated and contributes to a running average of the local rate correlation experienced by the agent.

The primary motivation for developing this theory was the fact that ratio schedules of reinforcement, under which this correlation is positive, tend to produce higher response rates and

devaluation sensitivity than interval schedules, which generate relatively low correlations. However, an incidental prediction of the theory is that extinction should have only a limited impact on goal-directed control. The rate correlation can only be computed when both responses and outcomes are represented in memory and, of necessity, the computation is halted when the memory is cleared of representations of at least one of the event types, either the response or the outcome. Given that this system only computes the local rate correlation, this suspension should happen relatively early in extinction with loss of any outcome representation and thereafter the strength of goal-directed control will reflect the mean running average rate correlation established earlier during training. Perez and Dickinson (2020) verified this incidental prediction of the theoretical model by formal simulation[2].

More problematic for the Perez and Dickinson (2020) theory is that response reduction by non-contingent outcomes had no detectable effect on goal-directed control, as demonstrated by the Crimmins et al (2022) rodent experiments and our non-contingent human experiment. In fact, Perez and Dickinson (2020) explicitly offered the prediction that non-contingent training, in contrast to extinction, would reduce goal-directed control as a novel prediction of their theory and therefore these results may well challenge rate correlation theory. However, as Crimmins et al (2022) point out, classic two-process theory (Rescorla & Solomon, 1967) has long argued that an instrumental contingency often arranges not only an action-outcome relationship but also a Pavlovian stimulus-outcome association that endows the stimulus with the capacity to modulate instrumental performance. As a result, what appears to be an instrumental outcome revaluation effect may not be mediated by action-outcome learning as required for goal-directed control but rather by Pavlovian stimulus-outcome learning. Although Jonkman et al. (2010) provided evidence that Pavlovian contextual conditioning can contribute to apparent

[2]The reader may wonder why responding is reduced by extinction if goal-directed strength is maintained. However, Perez and Dickinson (2020) embedded the goal-directed system within a dual-system theory that includes a habitual system. Within this dual-system account, extinction is primarily due to the development of response inhibition by the habitual system, which accords with the explanation offered by Rescorla (1993).

goal-directed control, Crimmins et al. (2022) focus their analysis on the role of their manipulanda, right- and left-hand levers.

This focus reflected the fact that the procedures in the experiments in which they found that non-contingent training had no effect on the outcome devaluation effect were ones in which each response was trained in separate sessions with only the appropriate lever present so that presence of a specific lever predicted the type of outcome. Evidence that the presence of a lever can act as a Pavlovian stimulus comes from a study by Amaya, Stott, and Smith (2020) in which the insertion of a lever signaled the delivery of a food outcome. Although there was no instrumental contingency in that the rats did not have to press the lever to receive the outcome, they regularly learned to approach the lever and contact it and thus often press it. Critically, this Pavlovian lever pressing was mediated by a lever stimulus-outcome association in that conditioning an aversion to the food outcome yielded an outcome devaluation effect on test.

When Crimmins et al (2022) removed this confound by using a single lever which when pressed to the left yielded one outcome with the other outcome depending on right presses. With this bi-directional lever that equally predicted the two outcomes, the results were those anticipated by rate correlation theory, namely that non-contingent training reduced the outcome devaluation effect and thus evidence for goal-directed control. As our online participants did not have access to a bidirectional manipulandum, we ensured that the stimuli accompanying each pair of responses (R1/R2 and R3/R4) had the same Pavlovian associations with both outcomes by training the two responses concurrently so that all stimuli were equally associated with both outcomes. This procedure ensured that it is unlikely that the outcome devaluation effect was mediated by Pavlovian associations with the outcome.

Although both extinction and non-contingent training significantly reduced responding, it is notable that extinction produced a greater reduction (Rescorla & Skucy, 1969). However, it is also

notable that there was still considerable responding after both contingency manipulations, but we think that it is unlikely that greater reductions in responding would have decreased the outcome devaluation effects. Further analyses that included the magnitude of response reduction during extinction and non-contingent training as variables revealed that the size of the outcome devaluation effect increased with greater response reductions. This effect probably reflects individual differences in the cognitive resources applied to the task—the greater the resource allocation to encoding the relative contingencies, the more likely it is that a participant will have reduced responding under extinction and non-contingent training and learned the response–outcome associations and outcome values necessary for the devaluation effect. However, the correlation between suppression and contingency knowledge was not significant ($r = .01$, $p = .667$), suggesting that these variables reflect different processes. When both were included in the model, each independently predicted sensitivity to outcome devaluation: suppression ($b = –0.171$, $SE = 0.076$, $p = .025$) and contingency score ($b = –2.32$, $SE = 0.83$, $p = .005$). This suggests that while suppression may capture attentional engagement or strategic adaptation during the manipulation phase, contingency scores appear to reflect explicit knowledge of instrumental structure. However, they both appear to contribute independently to goal-directed control.

Our paradigm is the first one in humans including a reinforced test at the end of training. Previous studies have ignored this fundamental test in assessing the effect of value on responding after manipulations aimed to prompt habit formation, and relied on consumption tests similar to our coin collection procedure or specific satiety (Gillan et al., 2015; Pool et al., 2019, 2022). Although this test is necessary to show the effectiveness of devaluation, it is not sufficient: what is more critical is showing that devaluation is effective in the same training context and contingent on the actions performed during training. Our data show that the devaluation method was effective in doing so.

In conclusion, the present experiments replicated with human participants the finding with rodents that response extinction does not impact on the magnitude of outcome devaluation effect, a unique prediction of the rate correlation theory of goal-directed control . By contrast, the prediction of this theory that non-contingent training should reduce the outcome devaluation effect was not substantiated.

Preregistration Statement

This study was not preregistered.

Author note

Omar David Perez is supported by Instituto de Sistemas Complejos de Ingeniería (ISCI) ANID PIA/PUENTE AFB220002, ANID-SIA 85220023 and ANID-FONDECYT 1231027. Emiliano Merlo is supported by the University of Sussex School of Psychology start-up funds.

Author contributions

ODP, SO, AD and EM designed the study. SO programmed the tasks. ODP and JA analyzed the data. ODP and AD drafted the manuscript, with critical editions by all authors.

References

Adams, C. D., & Dickinson, A. (1981). Instrumental responding following reinforcer devaluation. *The Quarterly Journal of Experimental Psychology Section B : Comparative and Physiological Psychology*, *33*(2), 109–121. https://doi.org/10.1080/14640748108400816

Amaya, K. A., Stott, J. J., & Smith, K. S. (2020). Sign-tracking behavior is sensitive to outcome devaluation in a devaluation context-dependent manner: Implications for analyzing habitual behavior. *Learning & Memory*, *27*(4), 136–149.

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, *67*(1), 1–48. https://doi.org/10.18637/jss.v067.i01

Crimmins, B., Burton, T. J., McNulty, M., Laurent, V., Hart, G., & Balleine, B. W. (2022). Response-independent outcome presentations weaken the instrumental response-outcome association. *Journal of Experimental Psychology: Animal Learning and Cognition*, *48*(4), 396.

Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, *8*(12), 1704–1711. https://doi.org/10.1038/nn1560

Daw, N. D., & O'Doherty, J. P. (2013). Multiple Systems for Value Learning. In *Neuroeconomics: Decision Making and the Brain: Second Edition* (pp. 393–410). https://doi.org/10.1016/B978-0-12-416008-8.00021-8

Dickinson, A. (1985). Actions and habits: The development of behavioural autonomy. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *308*(1135), 67–78.

Dolan, R. J., & Dayan, P. (2013). Goals and habits in the brain. *Neuron*, *80*(2), 312–325. https://doi.org/10.1016/j.neuron.2013.09.007

Gillan, C. M. M., Otto, A. R., Phelps, E. A., & Daw, N. D. (2015). Model-based learning protects against forming habits. *Cognitive, Affective and Behavioral Neuroscience*, *15*(3), 523–536.

https://doi.org/10.3758/s13415-015-0347-6

Jonkman, S., Kosaki, Y., Everitt, B. J., & Dickinson, A. (2010). The role of contextual conditioning in the effect of reinforcer devaluation on instrumental performance by rats. *Behavioural Processes*, *83*(3), 276–281. https://doi.org/10.1016/j.beproc.2009.12.017

Lee, S. W., Shimojo, S., & O'Doherty, J. P. (2014). Neural computations underlying arbitration between model-based and model-free learning. *Neuron*, *81*(3), 687–699. https://doi.org/10.1016/j.neuron.2013.11.028

Morey, R. D. (2008). Confidence intervals from normalized data: A correction to Cousineau (2005). *Tutorials in Quantitative Methods for Psychology*, *4*(2), 61–64.

Perez, O. D., & Dickinson, A. (2020). A Theory of Actions and Habits: The Interaction of Rate Correlation and Contiguity Systems in Free-Operant Behavior. *Psychological Review*. https://doi.org/10.1037/rev0000201

Pool, E. R., Gera, R., Fransen, A., Perez, O. D., Cremer, A., Aleksic, M., Tanwisuth, S., Quail, S., Ceceli, A. O., & Manfredi, D. A. (2022). Determining the effects of training duration on the behavioral expression of habitual control in humans: A multilaboratory investigation. *Learning & Memory*, *29*(1), 16–28.

Pool, E. R., Pauli, W. M., Kress, C. S., & O'Doherty, J. (2019). Behavioural evidence for parallel outcome-sensitive and outcome-insensitive Pavlovian learning systems in humans. *Nature Human Behaviour*, *3*(March). https://doi.org/10.1038/s41562-018-0527-9

Rescorla, R. A. (1993). Preservation of response-outcome associations through extinction. *Animal Learning & Behavior*, *21*(3), 238–245. https://doi.org/10.3758/BF03197988

Rescorla, R. A., & Skucy, J. C. (1969). Effect of response-independent reinforcers during extinction. *Journal of Comparative and Physiological Psychology*, *67*(3), 381.

Rescorla, R. A., & Solomon, R. L. (1967). Two-process learning theory: Relationships between

Pavlovian conditioning and instrumental learning. *Psychological Review*, *74*(3), 151.

RStudio Team. (2020). *RStudio: Integrated Development Environment for R*. RStudio, PBC. http://www.rstudio.com/

Tricomi, E., Balleine, B. W., & O'Doherty, J. P. (2009). A specific role for posterior dorsolateral striatum in human habit learning. *The European Journal of Neuroscience*, *29*(11), 2225–2232. https://doi.org/10.1111/j.1460-9568.2009.06796.x

Yin, H. H., & Knowlton, B. J. (2006). The role of the basal ganglia in habit formation. *Nature Reviews. Neuroscience*, *7*(6), 464–476. https://doi.org/10.1038/nrn1919