

Effects of reward and punishment on action reprogramming: Experiment 2

Frederick Verbruggen

December 9, 2015

Contents

1	Short introduction	2
2	Methods	2
2.1	Subjects	2
2.2	Apparatus and procedure	2
2.3	Statistical analyses	3
3	Analysis of performance on the no-signal trials	4
3.1	Probability of a correct response	4
3.2	Probability of missed response	5
3.3	Reaction times of correct responses	6
4	Analysis of performance on the stop-signal trials	7
4.1	Probability of responding on stop-signal trials	7
4.2	Stop-signal delay	7
4.3	Signal-respond vs. no-signal RT	8
4.4	SSRT	9
5	Conclusions	10

1 Short introduction

The first pilot experiment (UTM-ActR-RW1) suggests that reward and punishment have a similar effect on performance in a stop-change task. In this experiment, we used a stop-signal task.

2 Methods

2.1 Subjects

108 volunteers (36 per condition) from the University of Exeter participated for monetary compensation (£5) or partial course credit, plus money won in the stop-signal task (see below). 16 subjects (3 in the control condition, 7 in the reward condition, and 6 in the punishment condition) were replaced: 13 subjects were replaced because their percentage of correct no-signal trials was below 70%; 2 subjects were replaced because the probability of a correct stop was below 30%; and one subject did not complete the full experiment. The experiment was approved by the local research ethics committee at the School of Psychology, University of Exeter. Written informed consent was obtained after the nature and possible consequences of the studies were explained.

2.2 Apparatus and procedure

The apparatus and materials were the same as in the first pilot experiment (UTM-ActR-RW1).

The sequence of events in a trial is shown in the Instruct pdf, which can be found in the experiment Documentation folder. The trial started with a fixation cross in the center of the screen, flanked by the letters M or W (fixation interval). For half of the subjects, the cross was always flanked by Ms; for the other subjects, it was flanked by Ws. After 500 ms, one of the digits replaced the fixation cross and remained on the screen for 500 ms (stimulus-presentation interval). After that, the fixation cross was shown again for 500 ms (preparation interval). Subjects were told to use this interval to prepare their response to the digit. Finally, the fixation cross was replaced by the up- and downwards pointing arrow cue.

On no-signal trials (2/3 of trials), the arrow cue instructed subjects to execute the go response: they had to press the "down arrow" key for digits smaller than 5, or the "up arrow" key for digits larger than 5. The duration of the go interval (during which subjects could execute their go response on no-signal trials) was 500 ms. Immediately after the response was executed or after 500 ms had elapsed, the stimuli were replaced by a feedback message for 1000 ms. We presented "too soon" when subjects had responded before the arrow cue was shown; "too slow" if they did not respond in time; "correct" if they pressed the correct response key during the go interval; and "incorrect" if they had pressed the incorrect key.

On signal trials (1/3 of trials) the left or right letter changed (from M to W, or from W to M; see Instruct pdf) after a variable delay, instructing subjects to cancel their planned go response to the digit. Left- and right signals occurred with equal probability, and the order was randomized. We used a staircase tracking procedure to obtain a probability of stopping of .50. The stop-signal delay was initially 100 ms and was subsequently adjusted with a 1-up/1-down tracking procedure: when subjects successfully stopped their response, SSD increased by 50 ms; when they erroneously executed the go response, SSD decreased by 50 ms. Immediately after a response was executed or after 500 ms had elapsed since the presentation of the arrow cue, we presented feedback for 1000 ms. The feedback message depended on the group to which the subjects were assigned. In the punishment group, we presented: "failed stop. *You lose 10 points*" if subjects executed the response to the digit (italics are added here to highlight the difference between conditions); or "correct stop", if the go response was successfully cancelled. In the reward group, we presented: "failed incorrect" if subjects executed the response to the digit; or "correct stop. *You win 10 points*", if the go response was successfully cancelled. In the control group, we presented: "failed stop" if subjects executed the response to the digit; or "correct stop", if the go response was successfully cancelled.

Subjects in the punishment and reward group were informed at the beginning of the experiment that at the end of the experiment, the points would be converted into money (100 points = £0.1). They were also told that the money would only be awarded if overall performance on digit trials was also satisfactory (in other words, if they responded correctly and in time on the majority of trials). The start balance in the

punishment group was 2500 points; the start balance in the reward group was 0. There were 256 stop-signal trials. Due to the tracking procedure, both groups were expected to end with 1250 points or £1.25.

The experiment consisted of 8 blocks of 98 trials. Subjects received a break after every block. During the break, we presented as feedback to the subjects their mean RT on no-signal trials, number of no-signal errors, number of missed no-signal responses, and percentage of correctly replaced responses. In the punishment and reward groups, we also presented the current balance. Subjects had to pause for 15 seconds.

At the beginning of the experiment, all subjects completed the short form of the Sensitivity to Punishment and Sensitivity to Reward Questionnaire (see also the first experiment). However, there were not enough subjects in each condition for meaningful correlational analyses. Therefore, the SPSRQ scores were not analyzed.

2.3 Statistical analyses

For each dependent variable of interest, we performed a mixed univariate analysis of variance with Condition (control, punishment, or reward) as between-subjects variable, and Part (first half vs. second half of the experiment) as within-subjects factor. Part was included in the analyses because previous work indicates that learning in response-inhibition tasks can be influenced by reward and punishment. The effect size for the ANOVAs is generalized eta squared.


We also ran a series of pairwise comparisons. For each comparison, we calculated the confidence intervals, t-values, corresponding Bayes factors (using the BayesFactor package in R), Cohen’s d and Hedge’s g. For paired t-tests, we calculated d_{av} and d_z (see Lakens, 2013).

3 Analysis of performance on the no-signal trials

First, we analysed performance on no-signal trials (i.e. trials on which subjects had to respond to the digit).

3.1 Probability of a correct response


Accuracy of $p(\text{correct}) = \text{correct trials} / (\text{correct trials} + \text{incorrect trials})$. Trials without a response or an anticipatory response are omitted here.



	Effect	DFn	DFd	SSn	SSd	F	p	p<.05	ges
1	(Intercept)	1.00	105.00	208.75	0.06	395284.15	0.00	*	1.00
2	condition	2.00	105.00	0.00	0.06	1.86	0.16		0.03
3	part	1.00	105.00	0.00	0.01	0.55	0.46		0.00
4	condition:part	2.00	105.00	0.00	0.01	0.23	0.79		0.00

No significant differences in go accuracy.


3.2 Probability of missed response



	Effect	DFn	DFd	SSn	SSd	F	p	p<.05	ges
1	(Intercept)	1.00	105.00	3.32	0.77	451.54	0.00	*	0.79
2	condition	2.00	105.00	0.01	0.77	0.36	0.70		0.01
3	part	1.00	105.00	0.06	0.12	48.94	0.00	*	0.06
4	condition:part	2.00	105.00	0.00	0.12	0.21	0.81		0.00

The probability of missed responses decreases over time, but there is no main effect of condition.

3.3 Reaction times of correct responses



Effect	DFn	DFd	SSn	SSd	F	p	p<.05	ges
1 (Intercept)	1.00	105.00	32358482.74	284855.29	11927.60	0.00	*	0.99
2 condition	2.00	105.00	11023.06	284855.29	2.03	0.14		0.03
3 part	1.00	105.00	737.17	47706.35	1.62	0.21		0.00
4 condition:part	2.00	105.00	1409.77	47706.35	1.55	0.22		0.00

No main effects or interactions.

	diff	lowerCI	upperCI	df	t	p	BayesF	d	g
control vs punish	-8.64	-26.86	9.58	70.00	-0.95	0.35	0.36	0.22	0.22
control vs reward	-17.50	-34.63	-0.36	70.00	-2.04	0.04	1.41	0.48	0.47
punish vs reward	-8.86	-25.40	7.69	70.00	-1.07	0.29	0.40	0.25	0.25

Uncorrected p-values are reported in the table. The difference between the reward and control condition is no longer significant after Holm-Bonferroni correction for multiple comparisons (we have to apply this correction because the main effect of condition is not reliable). In sum, all numerical go RT differences between the individual conditions are non-significant.

4 Analysis of performance on the stop-signal trials

Next, we analysed performance on stop-trials (i.e. trials on which subjects had to stop their digit response).

4.1 Probability of responding on stop-signal trials


```
## condition part presp.mean presp.sd
## 1 control first 0.5243056 0.05164623
## 2 control second 0.5388455 0.06029845
## 3 reward first 0.5047743 0.02682237
## 4 reward second 0.5099826 0.02578369
## 5 punish first 0.5136719 0.05153121
## 6 punish second 0.5141059 0.04294620
```

P(respond—signal) is close to .50 for all conditions and parts, indicating that the tracking procedure worked well.

4.2 Stop-signal delay


```
## condition part ssd.mean ssd.sd
## 1 control first 106.7595 40.97704
## 2 control second 108.2739 52.88140
## 3 reward first 129.7689 45.44811
## 4 reward second 143.5586 53.87427
## 5 punish first 123.0417 42.68736
## 6 punish second 132.6367 51.08746
```

4.3 Signal-respond vs. no-signal RT



	Effect	DFn	DFd	SSn	SSd	F	p	p<.05	ges
1	(Intercept)	1.00	105.00	56772981.35	806470.83	7391.67	0.00	*	0.98
2	condition	2.00	105.00	30930.31	806470.83	2.01	0.14		0.03
3	part	1.00	105.00	4485.64	120839.80	3.90	0.05		0.00
5	trial	1.00	105.00	151043.31	24212.36	655.02	0.00	*	0.14
4	condition:part	2.00	105.00	2647.22	120839.80	1.15	0.32		0.00
6	condition:trial	2.00	105.00	583.07	24212.36	1.26	0.29		0.00
7	part:trial	1.00	105.00	1.48	15697.65	0.01	0.92		0.00
8	condition:part:trial	2.00	105.00	48.07	15697.65	0.16	0.85		0.00


Is signal-respond RT shorter than no-signal RT for all subjects?



In sum, signal-respond RT is shorter than no-signal RT for all subjects in all conditions.

4.4 SSRT

The probability of a missed response was relatively high in this experiment. This complicates the SSRT estimation because subjects may not have initiated an action to begin with. Therefore, for the SSRT estimation, RT for missed responses was set to the MAX RT (see Verbruggen, Chambers, & Logan, 2013).



Effect	DFn	DFd	F	p	p<.05	ges
2 condition	2.00	105.00	2.78	0.07		0.04
3 part	1.00	105.00	16.75	0.00	*	0.02
4 condition:part	2.00	105.00	0.68	0.51		0.00

The main effect of condition is marginally significant.

	diff	lowerCI	upperCI	df	t	p	BayesF	d	g
control vs punish	13.59	-2.75	29.92	70.00	1.66	0.10	0.79	0.39	0.39
control vs reward	16.84	1.77	31.91	70.00	2.23	0.03	1.98	0.53	0.52
punish vs reward	3.25	-10.60	17.11	70.00	0.47	0.64	0.27	0.11	0.11

The difference between the control and reward condition is significant, but it does not survive Holm-Bonferroni correction for multiple comparisons (again, we have to apply this correction because the main effect of condition is not reliable). Note that the main effect and the differences between the control and incentive conditions are significant when missed no-signal responses are excluded from the SSRT calculation (when missed responses are excluded, the nth RT is lower; consequently, SSRT is also lower, especially in those conditions in which p(miss) is the highest).

5 Conclusions

The numerical differences observed in this pilot are generally consistent with the results observed in the first pilot. Incentives encouraged subjects to make proactive control adjustments: they tended to slow down on no-signal trials, and this may have resulted in (numerically) shorter SSRTs. However, the numerical differences between the control condition and the incentive conditions failed to reach significance and the Bayes factors were inconclusive or provided some support for the null hypothesis. The effect size estimates (see tables with paired t-tests) indicate that we would need at least 64 subjects per condition to have sufficient power.

Consistent with Experiment 1, the Bayesian analyses indicate that reward and punishment have a similar effect on stop-signal performance: