

Forgone but not forgotten: the effects of partial and full feedback in “harsh” and “kind” environments

Tim Rakow, Ben R. Newell & Louise Wright

Psychonomic Bulletin & Review

ISSN 1069-9384

Psychon Bull Rev

DOI 10.3758/s13423-015-0848-x



Your article is protected by copyright and all rights are held exclusively by Psychonomic Society, Inc.. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your article, please use the accepted manuscript version for posting on your own website. You may further deposit the accepted manuscript version in any repository, provided it is only made publicly available 12 months after official publication or later and provided acknowledgement is given to the original source of publication and a link is inserted to the published article on Springer's website. The link must be accompanied by the following text: "The final publication is available at link.springer.com".

Forgone but not forgotten: the effects of partial and full feedback in “harsh” and “kind” environments

Tim Rakow^{1,2} · Ben R. Newell³ · Louise Wright¹

© Psychonomic Society, Inc. 2015

Abstract In a perfect world, the choice of any course of action would lead to a satisfactory outcome, and we would obtain feedback about both our chosen course and those we have chosen to forgo. In reality, however, we often face harsh environments in which we can only minimize losses, and we receive impoverished feedback. In these studies, we examined how decision makers dealt with these challenges in a simple task in which we manipulated three features of the decision: The outcomes from the available options were either mostly positive or mostly negative (kind or harsh environment); feedback was either full or partial (outcomes revealed for all options or only for the chosen option); and for the final 20 trials in a sequence, participants either chose on each trial or set an “advance-directive” policy. The propensity to choose the better option was explained by several factors: Full feedback was more beneficial in harsh than in kind environments; policy decisions encouraged better decisions and ameliorated the adverse impact of a harsh environment; and beliefs about the value of strategy diversification predicted switch rates and choice quality. The results suggest a subtle interplay between bottom-up and top-down processes: Although harsh environments encourage poor choices, and some decision makers

choose less well than others, this need not imply that the decision maker has failed to identify the better option.

Keywords Repeated choice · Risky choice · Feedback · Forgone outcomes · Switch rates · Strategy diversification · Probability matching · Maximizing strategy

Decisions are often evaluated according to the outcomes obtained—even when these could not reasonably have been foreseen (Yates, Veinott, & Patalano, 2003). For example, investors in a bear market may be dissatisfied with their investment decisions, even though few (if any) better options existed, whereas any investment in a bull market may look “smart,” even though making a profit was like taking candy from a baby. A more principled evaluation of decision quality considers whether the best choice was made *given the information available at the point of decision*. For choice under uncertainty (e.g., investment decisions), feedback from previous similar decisions can be important for identifying the best (current) course of action—whether one seeks to maximize gains (e.g., the bull market) or minimize losses (e.g., the bear market). At best, one learns the outcomes of all options—even those one did not select. Such “full feedback” is the norm for some kinds of decision (e.g., you can learn how well the shares you did not buy performed). At worst, one receives no feedback on a decision, such as the physician who does not know whether their patient failed to return because a treatment killed them or cured them. Arguably, such situations are rare, though often feedback is only “partial”—most commonly because one *only* learns the outcome of the option that one did choose (e.g., you will never know how good last night’s restaurant meal would have been elsewhere). Such partial feedback can create a dilemma: Taking a risk on a new course

Electronic supplementary material The online version of this article (doi:10.3758/s13423-015-0848-x) contains supplementary material, which is available to authorized users.

✉ Tim Rakow
tim.rakow@kcl.ac.uk

¹ University of Essex, Colchester, UK

² Department of Psychology, Institute of Psychiatry, Psychology and Neuroscience, King’s College London, Box PO77, Henry Welcome Building, De Crespigny Park, Denmark Hill, SE5 8AF London, UK

³ University of New South Wales, Sydney, Australia

of action may be the only way to learn whether it is actually the best option.

In two studies, we examined these two components that make for dissatisfying decisions: What happens when the decision *seems* hard because the likely outcome is negative whatever one chooses, and what happens when the decision *is* hard because one has incomplete information about the options. We did so by manipulating the harshness of the environment (“harsh” choices between “bad” options vs. “kind” choices between “good” options) and by manipulating the amount of the outcome feedback available (“full” feedback for all options vs. partial feedback, only for the chosen option).

Economic decision theory predicts no interaction between these factors: Full feedback will always help one maximize the expected value (EV) at least as effectively as partial feedback (irrespective of the kindness of the environment), because full feedback delivers information at a faster rate (each “packet” of feedback contains more data), and perhaps also—in some instances—because it supports direct comparison between the options (e.g., making it easier to spot when one option dominates the others).¹ However, behavioral decision theory, and some behavioral data, might lead us to expect that the natural advantage of full (over partial) feedback might play out differently, depending on whether the world is harsh or kind. Some researchers have shown that the (experienced) regret that can arise from comparing obtained and forgone outcomes is an important driver of choice, with obtained and forgone outcomes receiving similar weights in subsequent decisions (Grosskopf, Erev, & Yechiam, 2006). In contrast, in their analysis of repeated play in strategic games, Camerer and Ho (1999) concluded that “players generally weight foregone payoffs about half as much as actual payoffs” (p. 869). In other words, decision makers may *not* always extract (or use) the full value of the information that full feedback provides.

Despite their general conclusion, Camerer and Ho’s (1999) estimates for the discount rate applied to forgone outcomes varied widely between games, but Camerer and Ho offered little discussion of the source of this variability.² Simon’s (1955) classic analysis of boundedly rational choice may offer one explanation for such variability. If one’s current pattern of choice offers “good enough” outcomes (i.e., outcomes exceeding one’s aspiration level—the performance that would meet one’s personal target for this decision), one need not

consider whether other (nonchosen) options are superior. Presumably, such down-weighting, or ignoring, of forgone outcomes will be more common in a “kind” world in which most outcomes are positive. Therefore, the harshness of the decision environment might determine the extent to which decision makers use information about forgone outcomes.

To examine this possible interaction, we conducted two studies in which participants made repeated choices, and we manipulated (1) whether the decision environment was harsh or kind, and (2) whether participants obtained full or partial feedback. We evaluated the impacts of these features by examining the quality of participants’ choices (i.e., proportions of optimal responses) and their tendencies to alternate between options (switch rate). Switching between options on successive choices may reflect indifference (or high uncertainty) over the best option, exploration of different options, or a (usually misguided) strategy for maximizing profit (Vulkan, 2000). We adopted two approaches to disentangle these possibilities. First, we changed the manner in which participants signaled their preference (trial-by-trial choices vs. a one-off “policy” choice, in which participants indicated, in advance, their preference for a run of 20 consecutive trials; cf. Camilleri & Newell, 2013). Second, we elicited participants’ beliefs about the best strategies when making repeated choices (in a series of simple described problems).

Method: studies 1 and 2

Participants

Study 1 included 74 participants, and Study 2, 77 participants. The participants received a fee (UK£4 in Study 1, UK£4.50 in Study 2) or course credit, plus additional payment contingent on the outcome of their choices. (Power exceeded 90 % to detect small-to-medium effects of $d = 0.35$ within subjects.)

Design and tasks

The experiment comprised a series of two-option choice tasks—described to participants as a series of games—each involving repeated choices, with feedback given after each choice. The options were represented by identical buttons displayed (left–right) on a computer screen, and they were selected by clicking the button. Participants were informed that they could win or lose three pence on each choice. No further information about the payoff distribution of the options was provided.

In Study 1, we employed eight choice sets, and in Study 2 we used only four of these choice sets (Table 1). The sets were either harsh or kind (i.e., $EV \leq 0$ or $EV \geq 0$, respectively), as determined by their predefined probabilities of a gain (p_{win}). The harsh and kind sets were matched for the magnitudes of

¹ We suggest “in some instances” because adding forgone feedback can actually make an inferior option *more* attractive when that option is risky and the worst possible outcome is rare (Yechiam, Rakow, & Newell, 2015).

² Likewise, Yechiam and Rakow (2012) found that this discount rate varies dependably between individuals, but they also did not consider why.

Table 1 Structure of the choice sets

Set	Type of Choice Environment	p_{win} Inferior Alternative	p_{win} Superior Alternative	Cohen's d (Separation)	Used in Study Number
1	Harsh	.1	.3	0.52	1, 2
2	Harsh	.1	.256	0.42	1
3	Harsh	.258	.5	0.52	1, 2
4	Harsh	.3	.5	0.42	1
5	Kind	.7	.9	0.52	1, 2
6	Kind	.744	.9	0.42	1
7	Kind	.5	.742	0.52	1, 2
8	Kind	.5	.7	0.42	1

their EVs (e.g., Set 1 was the reflection of Set 5) and for the standardized separations of the options, as measured by Cohen's d (e.g., Sets 1, 3, 5, and 7 had identical separations between the options). In this way, the choice environment (kind vs. harsh) was manipulated between subjects³ in Study 1, and within subjects in Study 2.

Figure 1 provides an overview of the study designs. Each game had two phases, the first comprising 20 or 40 trial-by-trial choices⁴ (number of trials denoted in the condition name) in which the type of feedback (full vs. partial) varied (within subjects) between games. Thus, in the full-20 condition (Studies 1 and 2) and the full-40 condition (Study 2), participants saw the outcome ("win" or "lose") for both the chosen and nonchosen options. In the partial-20 condition (Study 1) and the partial-40 condition (Studies 1 and 2), participants saw only the outcome for their chosen option. These manipulations created three conditions per study that varied according to the amount of information provided to the participant (Study 1, full-20 vs. partial-20 vs. partial-40; Study 2, full-20 vs. full-40 vs. partial-40). Note that each study had a pair of conditions that were equated for the number of trials (full-20 and partial-20 in Study 1; full-40 and partial-40 in Study 2). Also, each study had another pair of conditions that provided participants with identical numbers of observations (full-20 and partial-40 in both studies).

The second phase of each game had a further 20 choices, for which participants either continued with trial-by-trial choices or set a policy by choosing one of the two options to

be selected for all 20 choices (with feedback being given for the numbers of wins and losses after all the choices had been played out). In Study 1, the type of feedback (partial vs. full) matched that of the preceding phase, whereas in Study 2 no trial-by-trial feedback was provided in this second phase. This mode of choice factor (trial-by-trial vs. policy) was manipulated within subjects; thus, each choice set was encountered six times by each participant (three initial-phase conditions crossed with two final-phase modes).

Procedure

The participants completed the experiment in individual testing booths. Participants read on-screen task instructions (see the [supplemental materials](#)) that highlighted that the games would vary according to their length, the feedback provided, the chances of winning, and the manner in which the final 20 choices would be made. The instructions indicated that wins (vs. losses) paid +3 (vs. -3) pence, that each game began with an initial balance amount, and that participants would be paid their earnings for a game (Study 1) or games (Study 2) chosen at random. In Study 1, the initial balance was £1.60 for the harsh choice sets and £0.40 for the kind sets. This was done to ensure that participants' additional payments did not differ substantially according to their random allocation to choice environment, and it fits with the approach taken in the loss-gain framing literature (e.g., Tversky & Kahneman, 1981) to ensure that choices framed as involving either losses or gains had equivalent sets of possible final wealth states. However, this introduced a confound between initial endowment and choice environment. This confound was absent in Study 2, in which the initial game balance was always £0.00. The participants in Study 2 were paid for one randomly chosen game for each choice set (explained to participants as a random selection "balanced across the different levels of difficulty" of the games). Participants' additional payments ranged from £0.04 to £2.08 in Study 1, and from £0.00 to £2.00 in Study 2.

When ready, participants proceeded to the choice tasks ("games"). The game order was randomized separately for each participant, and the left-right position of the options was randomly determined for each new game. The option buttons were labeled identically ("Choose this option"). The outcomes were randomly determined (independently) for each option (according to their p_{win}) and (independently) for each choice, and feedback was presented as "Win 3p" or "Lose 3p." This feedback remained on screen until the next choice was made. The minimum time allowed between choices was 0.5 s. No running total of past outcomes or balance amount was shown. A pop-up window signaled the onset of the final 20 choices. When a policy choice was required, the pop-up text prompted participants to "Pick ONE option to play out for the remaining 20 choices." This message was also displayed on the screen (above the options), and the option button labels

³ Manipulated between subjects because, with eight choice sets, a fully within-subjects design would have required an unreasonable time commitment from the Study 1 participants.

⁴ "Short runs" of repeated choices were used because (a) over many studies, we have seen only small changes in preferences after the 40th repeated choice, and (b) given the option to freely sample options to determine which is best, few participants sample more than 30 times (Rakow, Demes, & Newell, 2008).

Manipulation	Study 1			Study 2		
Phases 1 & 2						
Environment	Kind environment (4 sets; Option EVs > 0)	Harsh environment (4 sets; Option EVs < 0)		Kind environment (2 sets; Option EVs > 0)	Harsh environment (2 sets; Option EVs < 0)	
	Manipulated between-subjects			Manipulated within-subjects		
Phase 1						
Feedback type (+ length)	Full feedback for 20 trials	Partial feedback for 20 trials	Partial feedback for 40 trials	Full feedback for 20 trials	Full feedback for 20 trials	Partial feedback for 40 trials
	Manipulated within-subjects			Manipulated within-subjects		
Phase 2 (20 trials)						
Mode of choice	Trial-by-trial choices (With the same type of feedback as in Phase 1)	Policy choice (Choose a single option for 20 consecutive trials)		Trial-by-trial choices (No trial-by-trial feedback)	Policy choice (Choose a single option for 20 consecutive trials)	
	Manipulated within-subjects			Manipulated within-subjects		

Fig. 1 Overview of the study designs

altered to “Choose this option for 20 choices.” The durations of the games varied with their designs, typically 30–35 s for the shortest games and 65–70 s for the longest games.

On finishing their final (24th) game, participants completed a questionnaire designed to measure their tendency for strategy diversification (at the expense of EV maximization). This was measured in order to examine whether beliefs about “good” strategies for well-defined decision scenarios (the questionnaire) predicted trial-by-trial choices between ill-defined options (the games), allowing us to explore the interplay between top-down control and the bottom-up influence of outcomes in experiential choice (see Newell, Koehler, James, Rakow, & Van Ravenzwaaij, 2013). Participants specified a preferred strategy for each of four strategy-selection scenarios (i.e., “word problems”), adapted from Gal and Baron (1996) and Rubinstein (2002). The responses were subsequently coded as a *maximizing strategy* (repeatedly and exclusively choosing the EV-maximizing option), a *matching strategy* (selecting different options across repeated choices, in proportion to the probability of a favorable outcome for each option), or an *other strategy*. A scale created by summing the numbers of maximizing strategies adopted (maximum = 4) was found to be reliable ($\alpha = .74$ in Study 1; $\alpha = .71$ in Study 2). Thus, *high scores* indicated a tendency to *avoid* (disadvantageous) strategy diversification. The data were excluded for six participants who did not answer all four problems.

Results and discussion

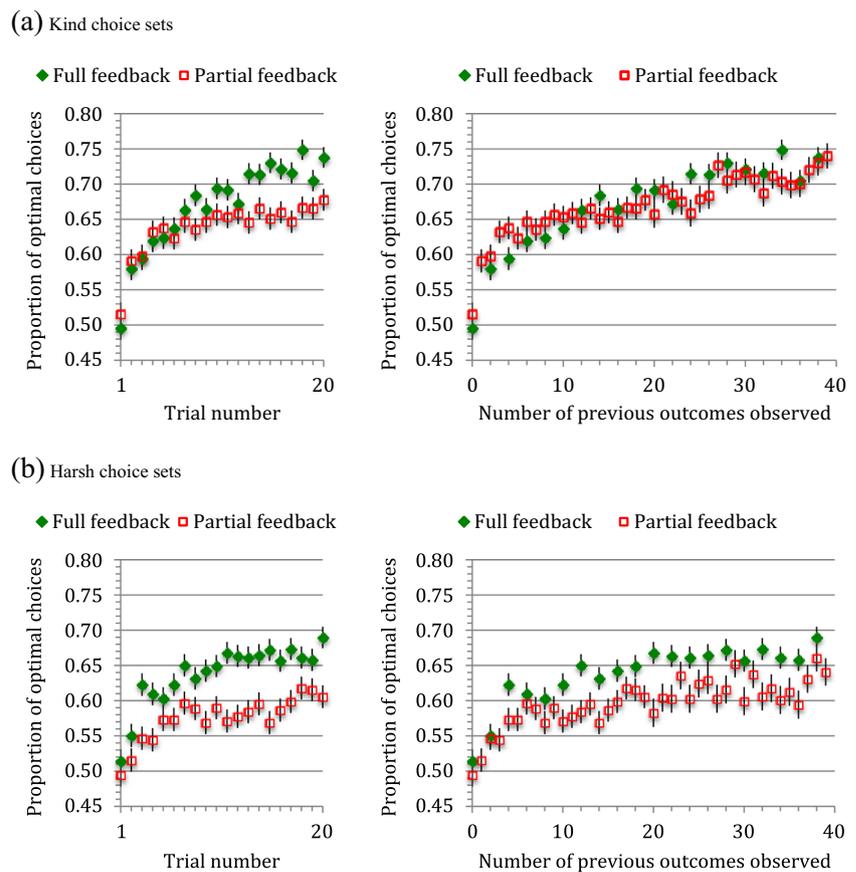
Choice data (proportions of optimal choices)

The tasks faced by participants in the initial phase of each game were identical for both studies; therefore, we combined

the initial-phase data for Studies 1 and 2. Figure 2 plots the proportions of optimal choices (i.e., selecting the higher-EV option) by type of feedback, shown separately for the kind (a) and harsh (b) choice sets. The panels on the left show the data for the first 20 trials of the initial phase, with each data point representing approximately 900 choices. A highly powered analysis of these data across trials revealed significant effects of choice environment, $F(1, 19) = 145.3, p < .001, \eta^2 = .884$, and type of feedback, $F(1, 19) = 76.7, p < .001, \eta^2 = .801$, reflecting higher proportions of optimal choices in kind choice sets and when full feedback was available. Importantly, we observed a substantial and significant environment-by-feedback interaction, $F(1, 19) = 20.9, p < .001, \eta^2 = .524$, reflecting a larger effect of the type of feedback in harsh choice sets ($\eta^2 = .890$) than in kind ones ($\eta^2 = .490$).

The right-hand panels of Fig. 2 illuminate why full feedback helped participants identify the superior option, and how this differed between the harsh and kind environments. Reflecting the fact that full feedback revealed two outcomes per trial, whereas partial feedback only showed the obtained outcome, these figures replot the data according to the numbers of observations seen (e.g., after seven trials, the groups had observed 14 outcomes with full feedback, but only seven outcomes with partial feedback). Thus, this presentation of the data controls for the amount of information delivered in each “packet” of feedback. Here one sees little difference between full and partial feedback in the kind choice sets (Fig. 2a, right), whereas the superiority of full over partial feedback remains in the harsh choice sets (Fig. 2b, right). This apparent pattern in the data is confirmed when an analysis across trials (as above) is conducted in which the number of previously observed outcomes defines how the data points are paired between feedback conditions. In this analysis, the environment-by-feedback interaction is again significant, $F(1, 19) = 47.6, p < .001, \eta^2 = .715$, but this time the effect of type of feedback is

Fig. 2 Studies 1 and 2 combined: Effects of feedback type (full vs. partial) on the proportions of optimal choices for kind choice sets (a; upper panels) and harsh choice sets (b; lower panels). The data are plotted by trial number (left) and by the number of outcomes previously observed (right). Error bars show ± 1 standard error



substantial and significant in the harsh choice sets, $F(1, 19) = 86.9, p < .001, \eta^2 = .821$, but small and nonsignificant in the kind choice sets, $F < 1, \eta^2 = .047$.

Because the advantage of full over partial feedback almost disappears in kind environments when the number of observations provided is controlled for (Fig. 2a, right), we inferred that when the advantage for full feedback is seen (i.e., in the nonequated data—Fig. 2a, left), it arises simply because full feedback delivers information at twice the rate. In contrast, full feedback retains its substantial superiority over partial feedback in harsh environments even when the number of previous observations is controlled for. This pattern suggests that its provision of a larger sample of observations is not the only reason why full feedback is superior to partial feedback in harsh environments.

Switch rates

Figure 3 further illustrates the interaction between the environment and the type of feedback provided, and suggests a possible mechanism for the accuracy data above. Consistent with fewer optimal choices being made in harsh choice sets, participants switch from one option to the next on successive trials more often in harsh sets (in which losses are frequent)

and—importantly—more so when only partial feedback is provided. In contrast, the switch rates are similar for partial and full feedback in kind choice sets (in which gains are frequent).

This pattern is confirmed by analyses of the participants' mean switch rates across Trials 2–20 of Study 1, and across Trials 2–40 of Study 2. Study 1 had a significant main effect of environment, $F(1, 72) = 21.05, p < .001, \eta^2 = .226$, and a significant main effect of type of feedback, $F(1, 72) = 8.52, p = .005, \eta^2 = .106$, which was qualified by a near-significant interaction, $F(1, 72) = 3.86, p = .053, \eta^2 = .051$. In harsh choice sets, participants switched significantly more often with partial than with full feedback ($M = .42$ vs. $.34$), $t(36) = 2.92, p = .006$, whereas in kind choice sets these two switch rates did not differ significantly ($M = .23$ vs. $.21$), $t < 1$. In Study 2 (within-subjects design), switching was again more common in harsh choice sets, $F(1, 76) = 112.74, p < .001, \eta^2 = .597$, but there was no main effect of type of feedback. We did observe, however, a significant and substantial environment-by-feedback interaction, $F(1, 76) = 24.82, p = .017, \eta^2 = .072$, reflecting higher switch rates with partial feedback ($M = .40$) than with full feedback ($M = .34$) in harsh choices sets, $t(77) = 3.78, p < .001$, but the reverse pattern in kind choice sets ($M = .22$ vs. $.26$), $t(77) = 2.69, p = .009$. We note that even though

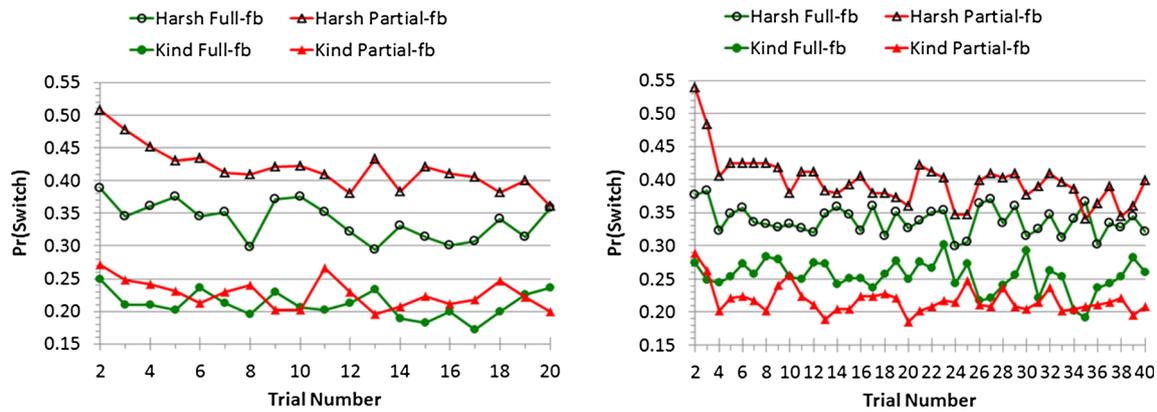


Fig. 3 Switch rates over the first 20 trials of Study 1 (left) and over the first 40 trials of Study 2 (right), by kindness of the choice set and type of feedback (full vs. partial). A switch occurs on trial n when the options selected on trial $n - 1$ and trial n differ

Studies 1 and 2 differed in their designs (between- vs. within-subjects manipulations of choice environment), which introduced a possible confound between initial endowment and environment in Study 1 (see the Method section), the patterns of the data are similar for the two studies (Fig. 3).

Notably, with partial feedback—for which one must switch at least some of the time to learn about the other option—the switch rates remain fairly constant after the 5th trial, and switch rates also remain constant across trials when full feedback is provided (Fig. 3). From this, we infer that most switching represents an (often misguided) attempt to exploit the environment, and that only a small proportion of switches represent genuine exploration.

Closer analysis of the Study 2 switch rates, when they were conditioned on the previous trial outcome, throws valuable light on this pattern of behavior.⁵ Counter to the pattern for the overall switch rates, participants were actually *more likely* to switch *following a loss* in kind choices sets than in harsh sets. In kind choice sets, these conditional switch rates were similar for the full ($M = .646$) and partial ($M = .645$) feedback conditions. In harsh choice sets, switches after a loss were less common with full feedback ($M = .47$) than with partial feedback ($M = .55$). Unsurprisingly, switching after obtaining a gain was rare. The probability of switching after a gain was uniformly low in the harsh choice sets with full and with partial feedback ($M_s = .077$ and $.069$, respectively), and in the kind choice sets with partial feedback ($M = .074$). Intriguingly, however, the probability of a switch after a gain almost doubled in the kind choice sets when full feedback was provided ($M = .128$). Thus, again, the impact of full feedback depended on the nature of the choice set, generating greater benefit in harsh environments. When losses were common, full feedback reduced the probability of switching after a

loss—thereby ameliorating the tendency for inconsistent, and suboptimal, choices in harsh environments. In contrast, when gains were common, full feedback seemingly encouraged choice inconsistency: Switching after a loss was as common as when partial feedback was provided, and there was a slight elevation in the rate of switching after a gain.

Strategic beliefs about strategy diversification

We used correlation to examine whether the maximizing scores compiled from the four scenario problems predicted trial-by-trial choices. In Study 1, the maximizing scores significantly predicted the proportions of optimal choices in the final phase, $r(70) = .40, p < .001$, but not in the initial phase, $r(70) = .18, p = .137$. In Study 2, the maximizing scores predicted the proportions of optimal choices both in the initial phase (in which feedback was provided), $r(71) = .24, p = .040$, and to a slightly greater extent in the final phase trial-by-trial choices (without feedback), $r(71) = .34, p = .004$. The maximizing score also significantly predicted participants' switch rates in both the initial phase, $r(70) = -.33, p = .005$, and the final phase, $r(70) = -.30, p = .010$, of Study 1, and in the initial phase, $r(71) = -.28, p = .015$, and the final phase, $r(71) = -.42, p < .001$, of Study 2.

Thus, participants who more frequently adopted a maximizing strategy in the described (scenario) problems made higher proportions of optimal choices in the repeated choice task and alternated less frequently between options. This suggests that higher-order beliefs about strategic responding may influence repeated choices with feedback (see Newell et al., 2013).

Converging evidence from the second phase of the task also suggests that poor choices rather than (exclusively) poor learning contributed to the suboptimal responding in repeated choices. In Study 2, participants' propensity to identify the best option (signaled by their policy choices) exceeded their propensity to choose that option in trial-by-trial choices. Thus, we found a substantial and significant main effect of mode of

⁵ The Study 2 data were preferred to the Study 1 data for this analysis because the within-subjects manipulation of environment provided a control for dependable individual differences in switch rates (see the main text).

choice, $F(1, 76) = 21.59$, $p < .001$, $\eta^2 = .221$, with a higher proportion of optimal choices for policy decisions ($M = .73$) than for trial-by-trial choices ($M = .64$).⁶ This mimics what occurs in choices between nonindependent options: Decision makers know enough (from observation or prior information) to be confident as to the best option, yet often “flip-flop” between the options when choosing trial by trial (i.e., *probability matching* in sequential binary choice tasks; Vulkan, 2000).

Conclusions

For some decisions, we find out what would have happened had a different course of action been chosen. For many others, we never know. Our studies operationalized these two situations in a simple task by either providing or withholding feedback about the forgone options. We found that although the choices were forgone, this feedback was not forgotten, especially in environments in which participants were resigned to minimizing losses rather than maximizing gains. Overall, the results were consistent with an aspiration-level account (cf. Simon, 1955): When most outcomes are favorable (kind environment), knowing what “might have been” has relatively little influence on choice, but when the world is harsh, the provision of forgone feedback ameliorates ultimately futile attempts to minimize losses via switching between options. Taken together, the results shed light on previous inconsistencies in conclusions about the weight placed on forgone information (Camerer & Ho, 1999; Yechiam & Rakow, 2012) and highlight the importance of considering the interaction between the environment, the information available to the agent, and the agent’s beliefs about how best to exploit that environment.

Author Note Both studies reported in this article were supported by a Leverhulme Trust Research Project Grant (No. RPG-384) to the first

author. We thank Michael Aitken for suggestions concerning the data analysis.

References

- Camerer, C., & Ho, T.-H. (1999). Experience-weighted attraction learning in normal form games. *Econometrica*, *67*, 827–874.
- Camilleri, A. R., & Newell, B. R. (2013). The long and the short of it: Closing the description–experience “gap” by taking the long-run view. *Cognition*, *126*, 54–71. doi:10.1016/j.cognition.2012.09.001
- Gal, I., & Baron, J. (1996). Understanding repeated simple choices. *Thinking & Reasoning*, *2*, 81–98.
- Grosskopf, B., Erev, I., & Yechiam, E. (2006). Foregone with the wind: Indirect payoff information and its implications for choice. *International Journal of Game Theory*, *34*, 285–302. doi:10.1007/s00182-006-0015-8
- Newell, B. R., Koehler, D. J., James, G., Rakow, T., & van Ravenzwaaij, D. (2013). Probability matching in risky choice: The interplay of feedback and strategy availability. *Memory & Cognition*, *41*, 329–338. doi:10.3758/s13421-012-0268-3
- Rakow, T., Demes, K. A., & Newell, B. R. (2008). Biased samples not mode of presentation: Re-examining the apparent underweighting of rare events in experience-based choice. *Organizational Behavior and Human Decision Processes*, *106*, 168–179.
- Rubinstein, A. (2002). Irrational diversification in multiple decision problems. *European Economic Review*, *46*, 1369–1378.
- Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, *69*, 99–118.
- Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, *211*, 453–458. doi:10.1126/science.7455683
- Vulkan, N. (2000). An economist’s perspective on probability matching. *Journal of Economic Surveys*, *14*, 101–118.
- Yates, J. F., Veinott, E. S., & Patalano, A. L. (2003). Hard decisions, bad decisions: On decision quality and decision aiding. In S. L. Schneider & J. Shanteau (Eds.), *Emerging perspectives on judgment and decision research* (pp. 13–63). New York, NY: Cambridge University Press.
- Yechiam, E., & Rakow, T. (2012). The effect of foregone outcomes on choices from experience: An individual level modeling analysis. *Experimental Psychology*, *59*, 55–67.
- Yechiam, E., Rakow, T., & Newell, B. R. (2015). Super-underweighting of rare events with repeated descriptive summaries. *Journal of Behavioral Decision Making*, *28*, 67–75.

⁶ This effect appeared in a $2 \times 3 \times 2$ (Environment \times Feedback \times Mode) analysis of variance; none of the interactions involving mode were significant (all $ps > .11$). The second-phase data of Study 1 were not informative on this issue, because participants received feedback during trial-by-trial choices; this confound between mode and feedback was eliminated by the design of Study 2.