

**Figuring Out Preference or Balancing Out Effort: Do Inferences From Incentives
Undermine Post-Incentive Motivation?**

Indranil Goswami, University at Buffalo
goswami4@buffalo.edu

Oleg Urminsky, University of Chicago
oleg.urminsky@chicagobooth.edu

January 15, 2018

Please contact authors for an updated version before citing or circulating

ABSTRACT

Do financial rewards undermine motivation by prompting people to infer from their incentivized choices that they only did a task because of the incentive? We examine this prediction of the *overjustification hypothesis* by manipulating the salience and availability of initial (i.e., pre-incentive) task attitudes in repeated choices between tasks. While reminding people about their initial task attitude prior to the incentive arrested post-incentive disengagement, we fail to find persistent disengagement after an incentive ends, when initial attitudes are *not* salient. Contrary to the prediction of overjustification that the negative longer-term effects of incentives will be magnified when all of a person's initial choices to do a task were incentivized, we also find no evidence of any post-incentive reduction in that case. Overall, our results contradict inferential accounts of post-incentive behavior and instead point to the role of perceived effort and justification. We discuss implications for dynamic consumer responses to promotions.

Temporarily available incentives are pervasive, whether in the form of project bonuses at work or price promotions in the marketplace. While managers typically see such incentives as beneficial in promoting the incentivized behavior, influential findings in psychology have suggested that incentives result in an unfavorable attitude toward the incentivized task, resulting in a persistent backlash after the incentive is no longer available (Deci, Koestner, and Ryan 1999). In this paper, we systematically investigate one specific account, overjustification (Lepper, Greene, and Nisbett 1973) that relies on an inferential mechanism for the prediction that financial incentives backfire.

We find very little support for the theory that choosing incentivized tasks leads people to infer that they were only doing the task for the incentive, and subsequently engage less in the task. Instead, we find only a momentary reduction in post-incentive engagement. Even under circumstances predicted to be most detrimental by the overjustification hypothesis, we do not find evidence of persistent disengagement. We discuss how past empirical findings can be reconciled with our results and directions for future research.

THEORETICAL DEVELOPMENT

Research in social psychology has shown that circumstances and situations not only affect behavior but can also affect people's beliefs about their own behavior (see Ross and Nisbett 2011 for a review). According to self-perception theory (Bem 1972), "individuals come to know their own attitudes, emotions, and other internal states partially by inferring them from observations of their own overt behavior and/or the circumstances in which the behavior occurs." Apart from direct tests (Bandler, Madaras, and Bem 1968; Bem 1965, 1966) this theory has been used to interpret several behavioral findings, including the foot-in-the-door effect

(Freedman and Fraser 1966), overjustification effect (Lepper and Greene 1975), and the effect of no-choice option on persistence (Schrift and Parker 2014).

In particular, the overjustification hypothesis emerged as a corollary of the insufficient justification effect (Festinger and Carlsmith 1959), in which lack of external justification for a counter-attitudinal act leads people to infer that they are intrinsically motivated to do the task. Conversely, the overjustification hypothesis predicts that when strong external justifications for an act (e.g., reward contingency, surveillance, etc.) are present, people infer that they were acting because of the external justification and are *not* intrinsically motivated to do the task (Lepper, Greene, and Nisbett 1973; Ross 1975), as they would if observing another person's behavior (Bem 1967). In this paper, we focus on the potential inference-based implications of a particular kind of justification – pre-announced contingent monetary incentives – on post-incentive task motivation and preference.

Overjustification effects of monetary incentives

Prior research on overjustification, much of it with children, found that when people were paid to do a task, they subsequently were less likely to engage in the task when no incentives were present, than if no incentives had been previously available (Greene and Lepper 1974; Kruglanski, Alon, and Lewis 1972; Lepper et al. 1973; Ross 1975). Young children's inferential processes (and tendency to generalize from those inferences) may differ from those of adults. However, proponents of the overjustification hypothesis have maintained that the inferential process and the consequent detriment would happen for adults as well (Lepper, Henderlong, and Gingras 1999; see Tang and Hall 1995 for a review).

Several theoretically important moderators have been identified, which help distinguish overjustification from other effects of incentives (e.g., reduction in autonomy, Deci and Ryan 1985). First, the negative effects of incentives on subsequent behavior does not extend to unexpected rewards, which were announced after the choices had been made (Lepper, Greene, and Nisbett 1973). According to the overjustification account, this is because incentives announced after the choice could not have influenced the choice. While unannounced incentives may also result in less during-incentive effort, this confound has been addressed in other papers (Kruglanski, Alon, and Lewis 1972; Ross 1975). Second, the overjustification effect has been shown to be absent for tasks commonly done for money (e.g., money-intrinsic tasks, Kruglanski et al. 1975), presumably because these tasks are already assumed to be chosen for extrinsic reasons.

The role of pre-incentive choices.

Although several studies in the overjustification literature included a pre-incentive period, generally for baseline measurement, researchers have largely not discussed the possibility that people would draw inferences from their pre-incentive choices. Luyten and Lens (1981) speculated that if an initial free-choice period provides an experience of intrinsic motivation, it may buffer the effects of self-perception. Fazio (1981) manipulated the salience of pre-incentive behavior, and found no incentive-based reduction in engagement, when young children were shown pictures of themselves engaging in the target task in an earlier session. In Study 1, we investigate the effect of measuring pre-incentive attitudes and reminding participants of these initial attitudes.

More generally, research on overjustification has failed to test the strong implication that overjustification effects should be stronger when all of a person's prior choices of a task

occurred when the task was incentivized (i.e., there was no pre-incentive period), compared to when some of the prior choices of the task were not incentivized (i.e., people experienced a pre-incentive period). According to the theory, participants in the first situation should more readily infer that the incentive dictated their behavior, and consequently should show a stronger post-incentive reduction in engagement. We test this prediction, varying whether the incentivized task is a work task (Study 2) or a fun task (Study 3).

EXPERIMENTAL PARADIGM

Prior research on overjustification tended to measure aggregate behavior and ignore the dynamics of post-incentive decision-making. In our experiments, we observe a series of choices, with and without incentives, in order to measure how behavior shifts after an incentive ends. Participants made repeated choices between doing a 30 second cognitive task, intended to be moderately effortful but interesting, and a 30 second leisure activity, watching a video (adapted from Goswami and Urminsky 2017; see Web Appendix A for additional details). The cognitive task required finding the two numbers in a grid of 12 numbers that added up to 10. Participants had 30 seconds to do the cognitive task, after which the page automatically advanced to the next choice. The video clips were also of 30 seconds duration. Pretests confirmed that both the cognitive tasks as well as the videos had an average rating of at least the mid-point on a scale measuring how “interesting and enjoyable” the task was ($1=Low$, $9=High$; $M_{cognitive\ task} = 6.02$; $M_{video} = 6.64$), per prior definitions of intrinsically motivating activities (Deci, Koestner, and Ryan 1999).

In the base versions of the experiments, participants made 30 repeated choices, split into three rounds, with incentives introduced in the middle round in some conditions. Participants in the experiment were not informed about how many choices they would make or how these

choices were divided into rounds, reducing the possibility of strategic forward-looking behavior. This research design enables us to assess initial preferences without an incentive, the effect of the incentive on choices and, in the final round, how preferences are affected by the termination of the incentive.

Target sample sizes were decided in advance in all the experiments. Participants with duplicate IP addresses, who reported technical problems, who failed an attention check question, or who did not complete Round 1 were removed prior to analysis. Participants who did complete Round 1 but then dropped-out before the end of study (between 2% and 7%; see Web Appendix C) were coded as not doing the target task during the rest of the experiment, and were included in the analysis, to minimize potential differential selection effects from experimental conditions.

STUDY 1: REMINDERS OF PRE-INCENTIVE TASK ATTITUDES

Method

Adult participants (N= 367 valid completes) were recruited from Amazon MTurk and randomly assigned to either the incentive-only condition, the incentive-with-reminder condition, or a control (i.e., a no incentive and no reminder) condition.

In the control (i.e., a no-incentive) condition, participants completed three rounds of choices between the two types of tasks, totaling 30 task choices. In the incentive-only condition, after participants completed Round 1, they were informed that they could earn a real 5 cents bonus for each cognitive task they solved correctly during the next round only, and confirmed their understanding. At the end of Round 2, they were told how much bonus they had earned, to be paid at the end of the experiment.

The incentive-with-reminder condition was the same as the incentive-only condition, except that just before reading about the incentive, participants were told how many cognitive tasks they had chosen to do in Round 1 and rated their liking for the cognitive task (1=Like a great deal, 6=Dislike a great deal). Subsequently, after learning about their total bonus earnings at the end of Round 2, participants were reminded how many cognitive tasks they had chosen to do in Round 1, and were asked to recall their earlier rating of the cognitive task. This condition was designed as a conceptual replication of Fazio's (1981) manipulation, to capture pre-incentive attitudes towards the task and then remind people of those initial attitudes once the incentive ended.

Results

The cognitive task was chosen around half of the time during Round 1, the baseline period, across conditions (control, 55%; incentive-only, 57%; incentive-with-reminder, 58%). Consistent with the pre-tests, this suggests that the task was intrinsically motivating and participants were willing to engage in the task without an incentive. Furthermore, choices of the cognitive task in the control condition held constant during the subsequent rounds (56% in both Round 2 and Round 3; see Figure 1¹).

In the incentive-only condition, the incentive boosted task engagement during Round 2. The proportion of cognitive task choices increased to 83% when the incentive was available – significantly higher than in the control condition (56%) during the same period, controlling for the baseline proportion of cognitive task choices in Round 1² ($\beta = 0.26, t = 7.81, p < .001$).

¹ The figures in the manuscript show aggregated means at the round level. Web Appendix B shows the means at every choice level with lowess smoothers.

² In all studies we report linear regression t-tests that control for the individual level engagement in the target task in round 1, when measured.

Therefore, the incentive was successful at increasing participants' task engagement when it was available.

The experiment was designed to test how the termination of an incentive affected choices of whether to do the previously incentivized task, both immediately and over a longer horizon. Immediately after the incentive ended (i.e., in choice 19, the first choice in Round 3) significantly fewer participants in the incentive-only condition chose to do cognitive task than in the corresponding choice in the control condition (39% vs. 52%; $\beta = 0.15$, $t=2.76$, $p=.006$). This decrease in engagement in the previously incentivized task once the incentive is no longer available is consistent with prior findings in the overjustification literature.

However, contrary to the prediction of the overjustification hypothesis, the decrease in engagement was only momentary. Averaging across the entire post-incentive period (i.e., Round 3), people chose to do a similar number of cognitive tasks in the incentive condition as in the control condition during the same period (53% vs. 56%; $\beta = 0.04$, $t=1.24$, $p=.215$). Since the decrease in post-incentive engagement was short-lived, the net effect of temporary incentives on target task choices across the entire experiment was significantly positive (incentive: 66% vs. control: 55%; $\beta = 0.09$, $t=3.28$, $p=.001$). These results replicate prior momentary-reduction findings (Goswami and Urminsky 2017) that challenged autonomy-based accounts of intrinsic motivation (Deci and Ryan, 1985).

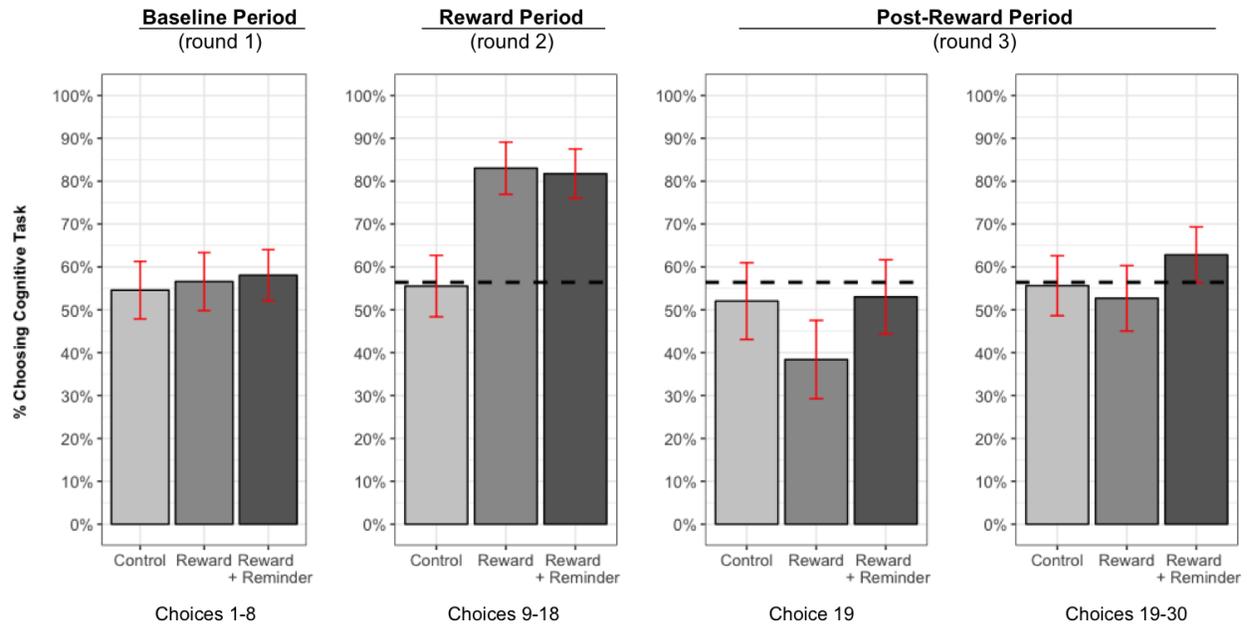


Figure 1. Study 1: Mean proportion of cognitive task choices in the control and the two incentive groups during each round of the experiment. Dotted lines represent the baseline effort (i.e., average effort level in Round 1 across conditions), and the vertical red lines are 95% CIs.

In the incentive-with-reminder condition, participants chose the cognitive task significantly more often during Round 2 (i.e., when incentivized) than in the control condition (82% vs. 56%; $\beta = 0.24$, $t=7.62$, $p<.001$), and at a rate nearly identical to the incentive-only condition (82% vs. 83%; $\beta = 0.02$, $t<1$, NS). However, reminding participants of how many cognitive tasks they had freely chosen to do and having them recall their prior task attitude successfully arrested the decrease in initial post-incentive engagement with the target task. In choice 19, the first choice made after the incentive ended, the likelihood of selecting the cognitive task in the incentive-with-reminder condition was comparable to the control condition (53% vs. 52%; $\beta = 0.01$, $t<1$, NS) and significantly higher than in the incentive-only condition (53% vs. 39%; $\beta = 0.14$, $t=2.42$, $p=.016$).

In the longer run, considering all choices in Round 3, participants in the incentive-with-reminder condition chose to do the cognitive task directionally more than in the control condition (63% vs. 56%; $\beta = 0.05$, $t=1.32$, $p=.19$), and significantly more than in the incentive-only condition (63% vs. 53%; $\beta = 0.09$, $t=2.26$, $p=.024$).

Discussion

Study 1 demonstrates a momentary reduction in choices of the incentivized cognitive task that was eliminated by reminders of pre-incentive engagement and attitude. This finding parallels Fazio (1981), consistent with over-justification, in which participants infer low intrinsic motivation from the incentive, unless reminded of cues about pre-incentive (intrinsic) motivation.

However, the full pattern of results is difficult to reconcile with the overjustification hypothesis. The inferential mechanism underlying overjustification theory predicts a persistent decrease in engagement with the temporarily-incentivized task unless participants are prompted to revise their beliefs again, by new information. In contrast, we find that choices of the previously incentivized task recovers quickly, and both the incentive-only and incentive-with-reminder conditions quickly converge to the same level of cognitive task choices as in the control condition.

These results suggest a reinterpretation of some of the prior evidence of overjustification, and of specifically Fazio (1981). Our findings suggest that the inferences people make after an incentive ends may be more about their momentary assessment of their current state, relative to their desired balance between work and leisure (Boksem and Tops 2008; Goswami and

Urminsky 2017; Inzlicht, Schmeichel, and Macrae 2014), than about the nature of the task itself or about their own enduring preferences.

Our findings are unlikely to be explained by other factors, such as fatigue, depletion or need for variety. Since participants in the two incentive conditions exerted similar level of effort and self-control (i.e. chose the cognitive task at similar rates) during Rounds 1 and 2, they were presumably in similar states of psychological depletion, physiological fatigue, and motivation to seek variety, and those process-accounts would predict similar behavior in Round 3. However, the results differed significantly between the incentive-only and incentive-with-reminder conditions.

The overjustification theory predicts that the downstream effects of incentives will depend on whether conditions favored or precluded the formation of initial attitudes towards the target task, prior to the introduction of the incentive. When people's initial choices of an activity occur when the activity is already incentivized, people should be more likely to infer that their task engagement happened because of incentives, and not because of preference for the task. As a result, the overjustification hypothesis predict a stronger post-incentive drop in motivation after rewards end when the person has not had prior task experience without an incentive. We test this prediction of the overjustification hypothesis in the next two studies.

STUDY 2: INTRODUCING TASK WITH REWARDS

Method

Adult MTurk participants (N=197 valid completes) were randomly assigned to a 2 (Incentive, Control) x 2 (Pre-incentive Round: Yes, No) between subjects design. For half the participants in the incentive group, the incentive was introduced in the second of three rounds

(baseline, reward and post-reward, as in Study 1) whereas for the other half, there were only two rounds (reward and post-reward) and thus the cognitive task was incentivized from the very beginning of the experiment. Correspondingly, half the participants in the control group had three rounds, and the other half had two rounds. Otherwise, the procedures were the same as Study 1.

Results

The pre-incentive-round conditions were identical to the control and incentive-only conditions in Study 1, and we replicated the results of Study 1 (see left panel of Figure 2). Participants did significantly more of the cognitive task in Round 2 in the incentive condition than control (85% vs. 58%; $\beta = 0.24$, $t=4.59$, $p<.001$), but significantly less in the first choice after the incentive ended (48% vs. 64%; $\beta = 0.18$, $t=2.08$, $p=.040$). However, this disengagement was only momentary, and there was no overall difference between incentive and control during Round 3 (58% vs. 54%; $\beta = 0.007$, $t<1$, NS). The net effect of incentives (summing Rounds 2 and 3) was positive compared to the no-incentive control condition (70% vs. 56%; $\beta = 0.11$, $t=2.27$, $p=.025$).

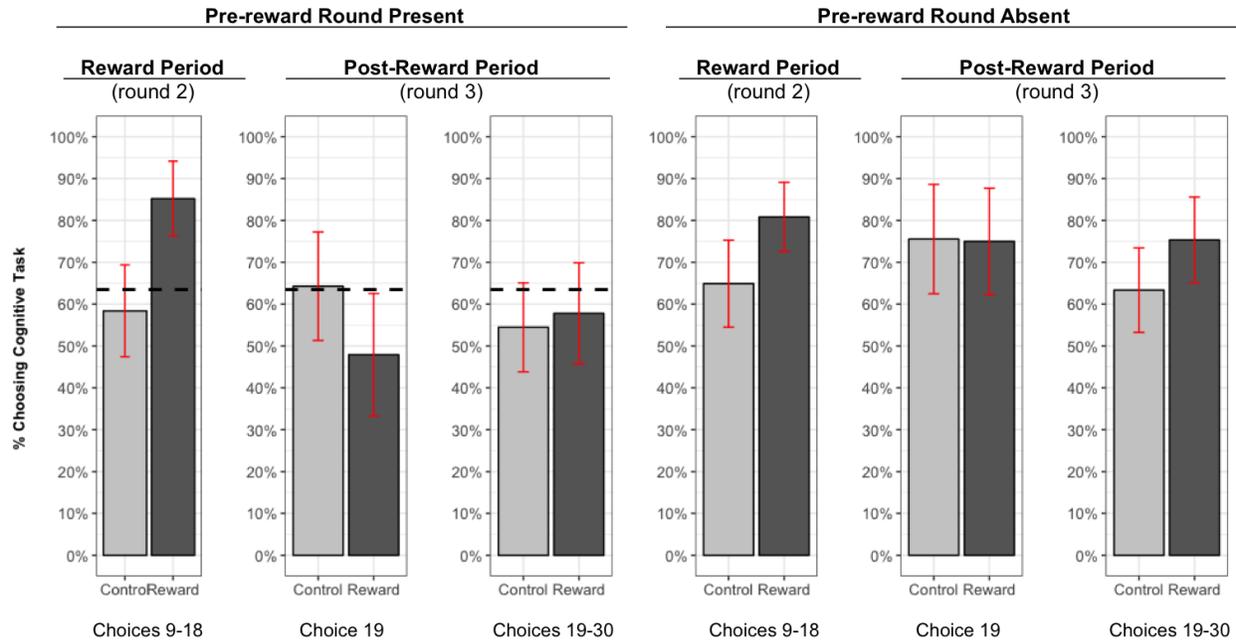


Figure 2: Study 2: Mean proportion of cognitive task choices in the control and incentive conditions. The left panel shows the results when the experiment had a pre-incentive period. The right panel shows the results when the experiment had *no* pre-incentive period and started directly with the incentive round. Dotted lines on the left-side charts represent the baseline effort (average effort level in Round 1 across conditions). The vertical red lines are 95% CIs.

In the no-pre-incentive-rounds conditions, participants were incentivized to choose the cognitive task over the video task from the very beginning of the experiment. Per the overjustification hypothesis, this should strengthen the inferential link between the incentive and beliefs about task preference as the participants could not make inferences from their unincentivized choices of the task. Participants in the incentive condition chose to do the target task more when the incentive was available, than in the corresponding control condition (81% vs. 65%; $\beta = 0.16, t=2.43, p=.017$). However, after the incentive ended, in the first choice of Round 2, participants in the incentive group were as likely as the control group to choose the cognitive task (75% vs. 76%; $\beta = 0.005, t<1, NS$). Summing over the post-incentive period, those in the

incentive condition were in fact marginally more likely to choose to do the cognitive task compared to participants in the control condition (75% vs. 63%; $\beta = 0.12$, $t=1.68$, $p=.096$).

Discussion

The results of Study 2 do not support the prediction of the overjustification hypothesis. Instead of finding a stronger or more persistent reduction in participants' post-incentive motivation to do the cognitive task when the task had been associated with an incentive from the beginning (and therefore, inferences about task liking would only be based on incentivized choices), we no longer find even a momentary post-incentive reduction in the no-pre-incentive-rounds conditions.

The cognitive task was chosen based on pre-tests indicating that people freely chose to do it, and that they found it interesting and enjoyable, satisfying the definition of an intrinsically motivating task. However, it is possible that what motivates people to freely do many "work" tasks, including the cognitive task, is future benefits, such as improved skills. If this is the case, then the cognitive task would be considered extrinsically motivating, under the broader definition of intrinsic motivation as benefits arising inside the task (e.g., in-the-moment enjoyment) and extrinsic motivation as benefits occurring outside the task (e.g., future higher level of skill; see Kruglanski et al. 2002; Woolley and Fishbach 2015). Under this interpretation, if participants were already thinking of the cognitive task as extrinsically motivating, introducing the incentive might not have substantially changed their inferences. This requires an additional assumption that beliefs about one's own non-monetary extrinsic motives are stronger than beliefs about one's own intrinsic motivation. According to this version of the overjustification hypothesis, a leisure task, for which engagement itself is the motivation, should show large and

persistent negative effects of incentives, particularly if the participant did not make initial un-incentivized choices.

We test this in the next study using the same design as Study 2, except that participants in the incentive condition were paid for watching videos. This provides a strong test of an inference-driven overjustification account based on a broader definition of extrinsic motivation.

STUDY 3: INTRODUCING A LEISURE TASK WITH INCENTIVES

Method

Adult MTurk participants (N=475 valid completes) were recruited for the study that employed the same 2 (Incentive, Control) x 2 (Pre-incentive Round: Yes, No) between subjects design, except that participants were incentivized to watch video instead of doing the cognitive task.

Results

We first analyze the pre-incentive-round conditions, in which participants first made choices between tasks, before encountering the incentive. In Round 2, when participants were incentivized for watching videos, participants chose significantly more videos than in the control condition (90% vs. 52%; $\beta = 0.38$, $t=10.83$, $p<.001$; See left panel of Figure 3).

However, unlike the findings for cognitive tasks in Studies 1 and 2, after the rewards ended, there was no decrease in engagement with watching videos – either momentarily or overall during Round 3. In the first choice after the incentive ended, 54% of the participants in the incentive condition chose to watch a video compared to 48% in the control ($\beta = 0.06$, $t=1.26$, $p=.210$). Likewise, summing over the post-incentive period, participants chose to watch videos

55% of the time in the incentive group compared to 50% in the control ($\beta = 0.06, t=1.46, p=.146$). Overall, the temporary incentive for the leisure task resulted in a significantly more choices of the activity compared to the no-incentive condition (71% vs. 51%; $\beta = 0.21, t=6.67, p<.001$).

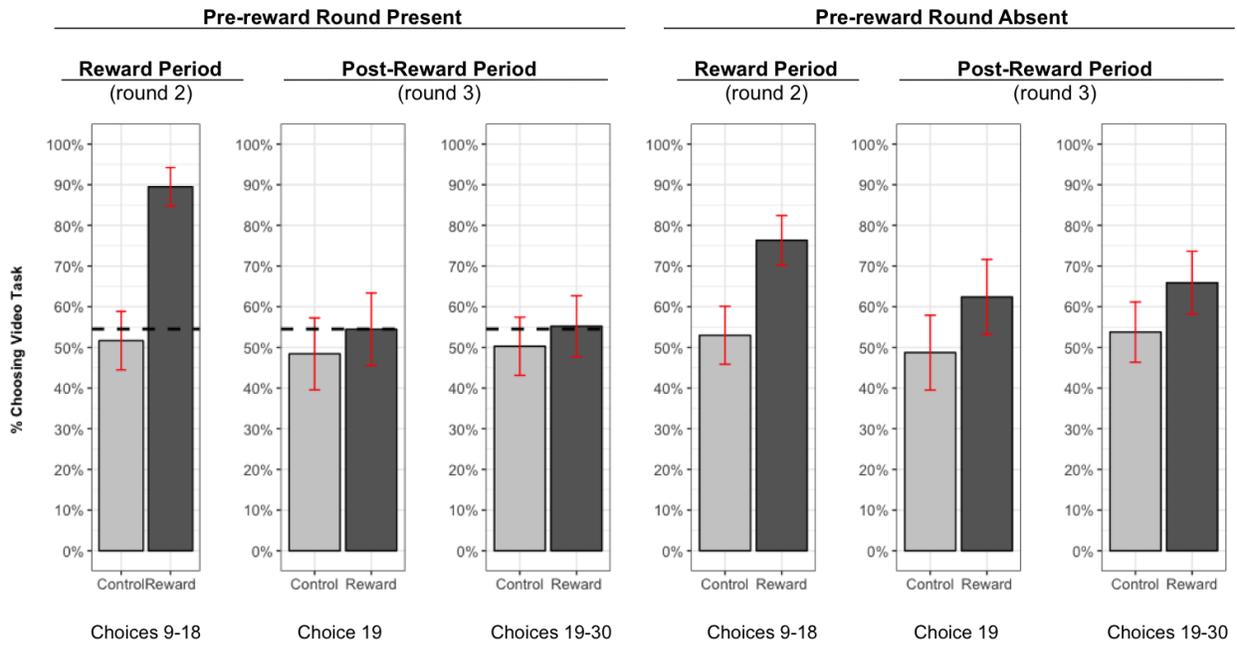


Figure 3: Study 2: Mean proportion of video choices in the control and incentive conditions. The panel on the left shows the results when the experiment had a pre-incentive period. The panel on the right shows the results when the experiment had *no* pre-incentive period and started directly with the incentive round. Dotted lines on the left-side charts represent the baseline effort (average effort level in Round 1 across conditions). The vertical red lines are 95% CIs.

In the no-pre-incentive-round conditions, participants were incentivized to watch videos from the beginning, and we found no indication of reduced motivation to watch videos due to the incentive. Participants in the incentive condition chose to do the target task more when the incentive was available than in the control condition (76% vs. 53%; $\beta = 0.23, t=4.89, p<.001$). More importantly, participants were significantly *more* likely to choose the video in the first

post-incentive choice than control participants (62% vs. 49%; $\beta = 0.14$, $t=2.08$, $p=.039$), and across the entire post-incentive period (66% vs. 54%; $\beta = 0.12$, $t=2.24$, $p=.026$). Overall, the net effect of incentives was strongly positive (71% vs. 53%; $\beta = 0.17$, $t=3.69$, $p<.001$).

Discussion

Study 3 provides the strongest test of the overjustification hypothesis as it applies to contingent monetary incentives. However, contrary to the predictions of the account, we found no persistent decline in choices of the video task after participants were paid an incentive to watch videos. Even more importantly, starting the experiment with incentivized (i.e. omitting initial unincentivized choices between tasks) did not yield a large and persistent reduction in the motivation to watch videos, but instead significantly *increased* both the initial and overall rate of watching videos after the incentive ended.

GENERAL DISCUSSION

Do financial rewards undermine motivation by prompting people to infer from their incentivized choices that they were only making those choices because of the incentive? The overjustification hypothesis postulates such inferences and predicts a persistent decrease in motivation to engage in the incentivized activity, once the incentive ends. Advocates against the use of incentives in the workplace point to the need to avoid “transforming an interesting task into a drudge.” (Pink, 2011). The concern for policy makers is that post-incentive effects may undercut the gains from incentives when available.

Across three studies, we systematically test the immediate and longer-term post-incentive effects, using a dynamic repeated-decision paradigm in which participants choose between two intrinsically motivating activities, one of which is incentivized. Our studies identified a

momentary reduction in choices of the previously incentivized task under certain conditions, but consistently failed to find any evidence for a persistent post-incentive decrease. Furthermore, eliminating the opportunity to initially make unincentivized choices, would be predicted by overjustification to magnify the inferences and subsequent reduction, but no permanent disengagement was found (Studies 2 and 3).

The studies call into question the validity of applying a self-perception framework in predicting the consequences of temporary monetary incentives. However, our experiments with an adult population may tap into different psychological mechanisms than some of the research on overjustification, which involved young children who might still be forming beliefs about money and incentives in general. Future research might systematically test for differences between children and adults in their responses to incentives.

If post-incentive behavior is not driven by over-justification, what explains the observed behavior? Recent work on dynamic post-incentive behavior (Goswami and Urminsky 2017) has suggested that exerting extra effort in response to modest incentives may change people's sense of balance between work and leisure, such that it creates a temporary desire or justification to "take a break" (i.e., engage in leisure). This would explain why we observe a momentary reduction in choices of a previously incentivized "work" task (Studies 1 and 2), but no persistent effect. Furthermore, this view is consistent with our findings that even the momentary reduction in choices of the task is eliminated when people are reminded that they like the task (Study 1), when they have done less of the task because there was no pre-incentive round (Studies 2 and 3), or when the low-effort leisure task is incentivized (Study 3).

Implications for consumer behavior

Marketing promotions, an integral component of the marketing mix, are by design temporary. Researchers have previously proposed that overjustification-based self-perception might result in reduced choices of the promoted brand after certain deals end (Dodson, Tybout, and Sternthal 1978). While such post-promotion patterns have been observed, the reason is not always clear. Frequent promotions can change consumers beliefs about future promotions, in effect “training” them to stockpile and only when the product is on sale, resulting in post-promotion sales dips (Mela, Jedidi, and Bowman 1998; Neslin, Henderson, and Quelch 1985). Instead of consumers drawing inferences about their own tastes and reasons, per overjustification, they may instead draw inferences about the quality of the brand based on how much it resorts to such incentives (Scott 1976), although there is mixed evidence for this account (Blattberg and Neslin 1990; Davis, Inman, and McAslister 1992; Raghurir and Corfman 1999). Our findings call into question the overjustification explanation. We predict, based on our results, that any dip in motivation to buy the promoted brand would be only momentary and occur immediately after the promotion has ended.

REFERENCES

- Bandler Jr, R. J., Madaras, G. R., & Bem, D. J. (1968). Self-observation as a source of pain perception. *Journal of Personality and Social Psychology*, 9(3), 205.
- Bem, D. J. (1965). An experimental analysis of self-persuasion. *Journal of Experimental Social Psychology*, 1(3), 199–218.
- Bem, D. J. (1966). Inducing belief in false confessions. *Journal of Personality and Social Psychology*, 3(6), 707.
- Bem, D. J. (1967). Self-perception: An alternative interpretation of cognitive dissonance phenomena. *Psychological Review*, 74(3), 183.
- Bem, D. J. (1972). Self-perception theory. *Advances in Experimental Social Psychology*, 6, 1–62.
- Blattberg, R. C., & Neslin, S. A. (1990). *Sales promotion: Concepts, methods, and strategies*. Prentice Hall.
- Boksem, M. A., & Tops, M. (2008). Mental fatigue: costs and benefits. *Brain Research Reviews*, 59(1), 125–139.
- Davis, S., Inman, J. J., & McAslister, L. (1992). Promotion has a negative effect on brand evaluations: Or does it? Additional disconfirming evidence. *Journal of Marketing Research*, 143–148.
- Deci, E. L., Koestner, R., & Ryan, R. M. (1999). A meta-analytic review of experiments examining the effects of extrinsic rewards on intrinsic motivation. *Psychological Bulletin*, 125(6), 627.
- Deci, E. L., & Ryan, R. M. (1985). Cognitive Evaluation Theory. In E. L. Deci & R. M. Ryan (Eds.), *Intrinsic Motivation and Self-Determination in Human Behavior* (pp. 43–85).

Boston, MA: Springer US. Retrieved from http://dx.doi.org/10.1007/978-1-4899-2271-7_3

Dodson, J. A., Tybout, A. M., & Sternthal, B. (1978). Impact of deals and deal retraction on brand switching. *Journal of Marketing Research*, 72–81.

Fazio, R. H. (1981). On the self-perception explanation of the overjustification effect: The role of the salience of initial attitude. *Journal of Experimental Social Psychology*, 17(4), 417–426.

Festinger, L., & Carlsmith, J. M. (1959). Cognitive consequences of forced compliance. *The Journal of Abnormal and Social Psychology*, 58(2), 203.

Freedman, J. L., & Fraser, S. C. (1966). Compliance without pressure: the foot-in-the-door technique. *Journal of Personality and Social Psychology*, 4(2), 195.

Goswami, I., & Urminsky, O. (2017). The dynamic effect of incentives on postreward task engagement. *Journal of Experimental Psychology: General*, 146(1), 1.

Greene, D., & Lepper, M. R. (1974). Effects of extrinsic rewards on children's subsequent intrinsic interest. *Child Development*, 1141–1145.

Inzlicht, M., Schmeichel, B. J., & Macrae, C. N. (2014). Why self-control seems (but may not be) limited. *Trends in Cognitive Sciences*, 18(3), 127–133.

Kruglanski, A. W., Alon, S., & Lewis, T. (1972). Retrospective misattribution and task enjoyment. *Journal of Experimental Social Psychology*, 8(6), 493–501.

Kruglanski, A. W., Riter, A., Amitai, A., Margolin, B.-S., Shabtai, L., & Zaksh, D. (1975). Can money enhance intrinsic motivation? A test of the content-consequence hypothesis. *Journal of Personality and Social Psychology*, 31(4), 744.

- Kruglanski, A. W., Shah, J. Y., Fishbach, A., Friedman, R., Chun, W. Y., & Sleeth-Keppler, D. (2002). A theory of goal systems. In M. P. Zanna (Ed.), *Advances in Experimental Social Psychology* (Vol. 34, pp. 331–378). Academic Press.
- Lepper, M. R., Greene, D., & Nisbett, R. E. (1973). Undermining children's intrinsic interest with extrinsic reward: A test of the "overjustification" hypothesis. *Journal of Personality and Social Psychology*, 28(1), 129.
- Lepper, Mark R., & Greene, D. (1975). Turning play into work: Effects of adult surveillance and extrinsic rewards on children's intrinsic motivation. *Journal of Personality and Social Psychology*, 31(3), 479.
- Lepper, Mark R., Henderlong, J., & Gingras, I. (1999). Understanding the effects of extrinsic rewards on intrinsic motivation—Uses and abuses of meta-analysis: Comment on Deci, Koestner, and Ryan (1999).
- Luyten, H., & Lens, W. (1981). The effect of earlier experience and reward contingencies on intrinsic motivation. *Motivation and Emotion*, 5(1), 25–36.
- Mela, C. F., Jedidi, K., & Bowman, D. (1998). The long-term impact of promotions on consumer stockpiling behavior. *Journal of Marketing Research*, 250–262.
- Neslin, S. A., Henderson, C., & Quelch, J. (1985). Consumer promotions and the acceleration of product purchases. *Marketing Science*, 4(2), 147–165.
- Pink, D. H. (2011). *Drive: The Surprising Truth About What Motivates Us*. Penguin.
- Raghubir, P., & Corfman, K. (1999). When do price promotions affect pretrial brand evaluations? *Journal of Marketing Research*, 211–222.
- Ross, L., & Nisbett, R. E. (2011). *The person and the situation: Perspectives of social psychology*. Pinter & Martin Publishers.

- Ross, M. (1975). Salience of reward and intrinsic motivation. *Journal of Personality and Social Psychology*, 32(2), 245.
- Schrift, R. Y., & Parker, J. R. (2014). Staying the course: The option of doing nothing and its impact on postchoice persistence. *Psychological Science*, 25(3), 772–780.
- Scott, C. A. (1976). The effects of trial and incentives on repeat purchase behavior. *Journal of Marketing Research*, 263–269.
- Tang, S.-H., & Hall, V. C. (1995). The overjustification effect: A meta-analysis. *Applied Cognitive Psychology*, 9(5), 365–404.
- Woolley, K., & Fishbach, A. (2015). The experience matters more than you think: People value intrinsic incentives more inside than outside an activity. *Journal of Personality and Social Psychology*, 109(6), 968.

Web Appendix A: Study Stimuli

Exhibit A1: The experimental paradigm and the depiction below is adapted from Goswami and Urmitsky 2017.

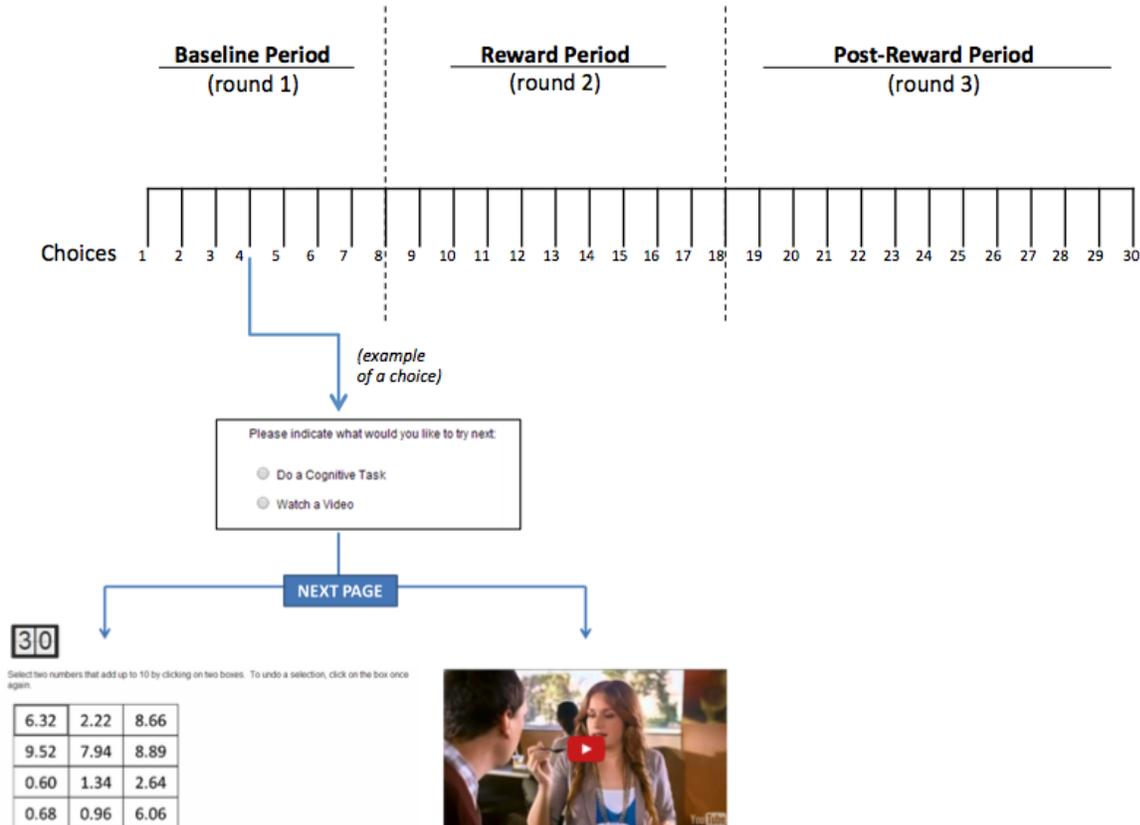


Exhibit A2: Study and tasks introduction in Study 1

PLEASE READ THE INSTRUCTIONS CAREFULLY.

In this survey, you will be given a series of **choices between doing cognitive tasks and watching videos of interesting television advertisements** collected from across the world.

You can do as many of the cognitive tasks as you want, or can just enjoy the videos.

Exhibit A3: Elicitation of task attitude at the end of round 1 Study 1

You chose to do $\{e://Field/R1Num\}$ cognitive tasks (i.e., $\{e://Field/R1Percent\}$ % of the tasks) in Round 1.

Please take a few moments to think about these cognitive tasks. How much did you like the cognitive tasks?

Please do not rush. Take your time to indicate your attitude (there are no right or wrong answers here).

Like a great deal	Like a moderate amount	Like a little	Dislike a little	Dislike a moderate amount	Dislike a great deal
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Exhibit A4: Study and tasks introduction in Study 2

PLEASE READ THE INSTRUCTIONS CAREFULLY.

In this survey you will be given a series of **choices between doing cognitive tasks and watching videos of interesting television advertisements** collected from across the world.

The cognitive task will train your mental reasoning skills, and we will use your results to calibrate and standardize a training test.

You can do as many of them as you want, or can just enjoy the videos.

Exhibit A5: Study and tasks introduction in Study 3

PLEASE READ THE INSTRUCTIONS CAREFULLY.

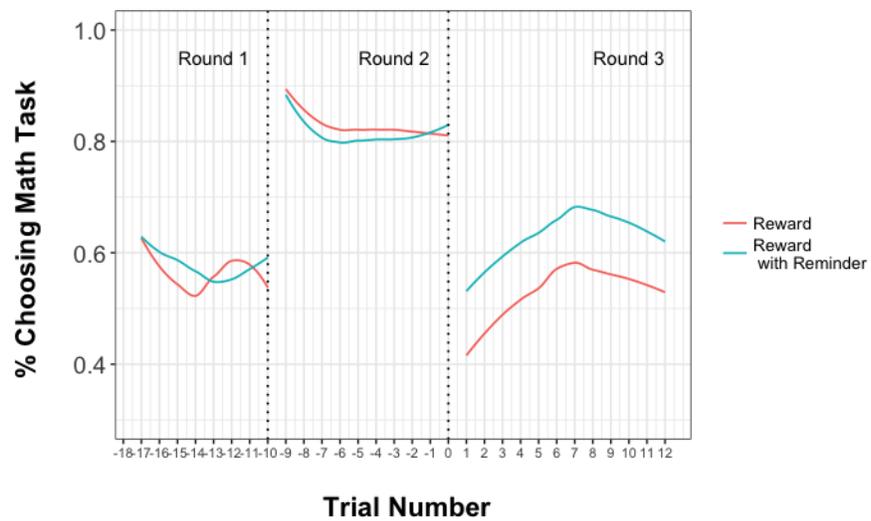
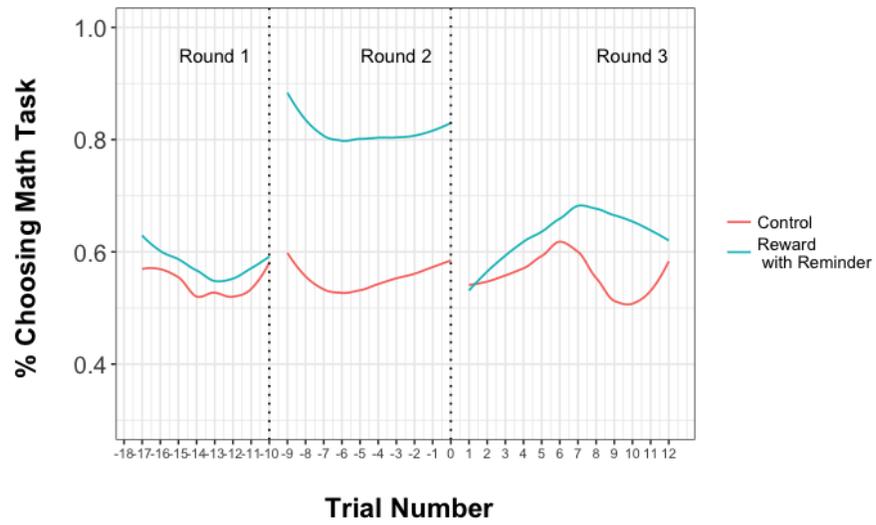
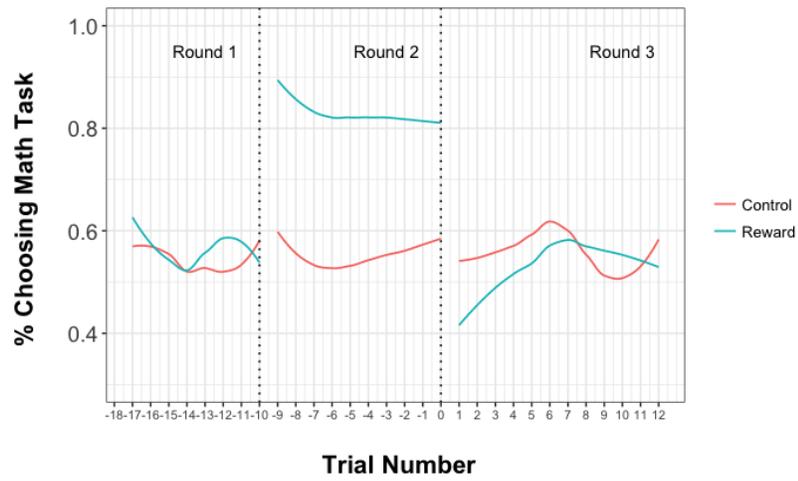
In this survey, you will be asked to do a task. The task is to **watch videos of television advertisements**. We will use your pattern of responses to design experimental stimuli for our future studies on attention and perception.

Since doing the same task can be tiring, you will also have an option of a different task, **solving cognitive math problems**, so that you can take a break.

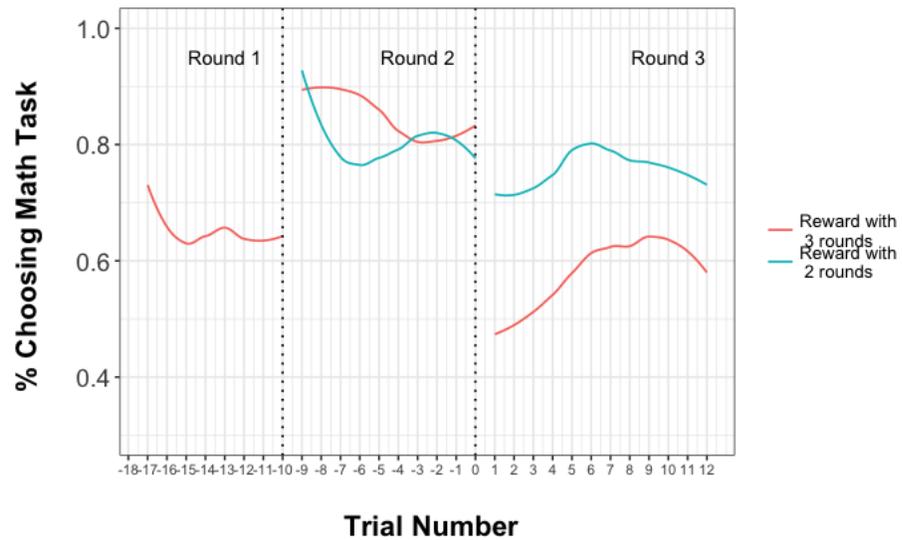
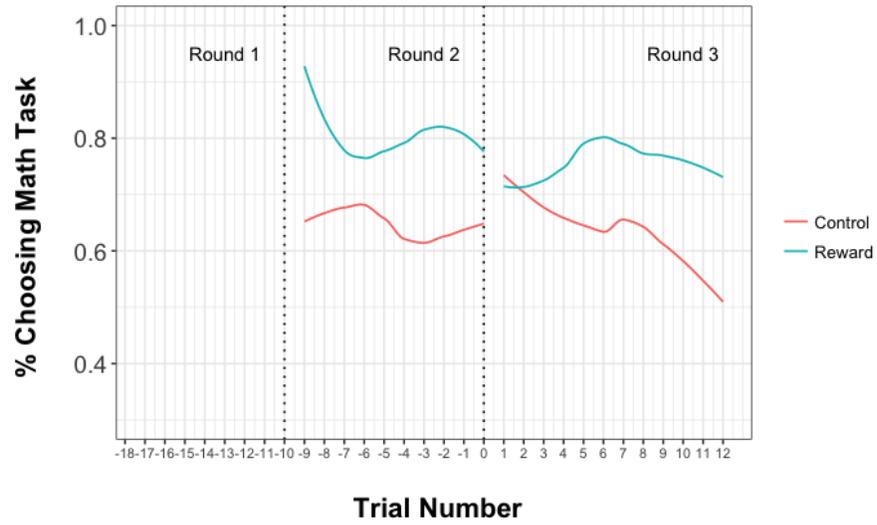
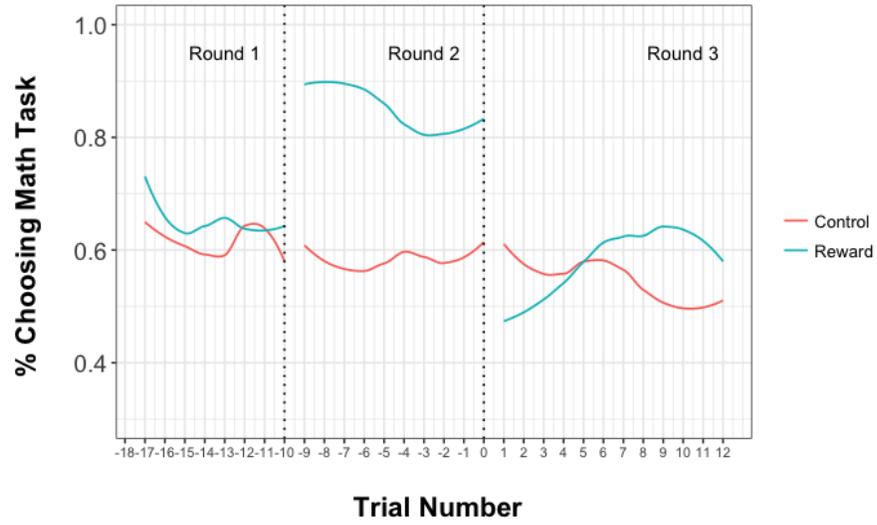
It is completely up to you to choose which task you want to do.

Web Appendix B: Raw Choice Data of All Studies with Lowess Smoother

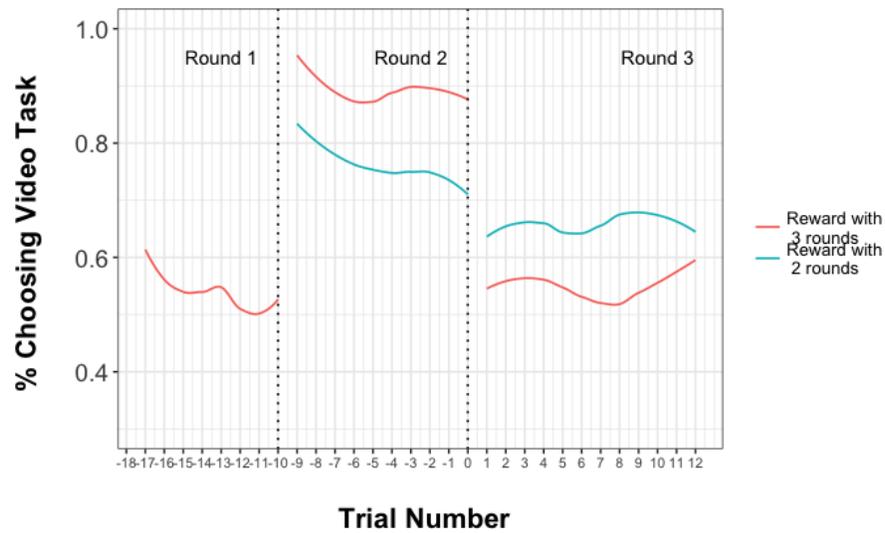
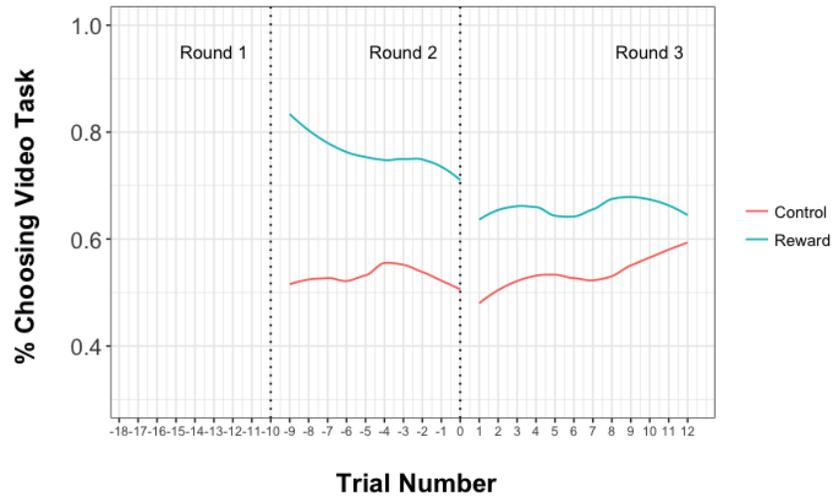
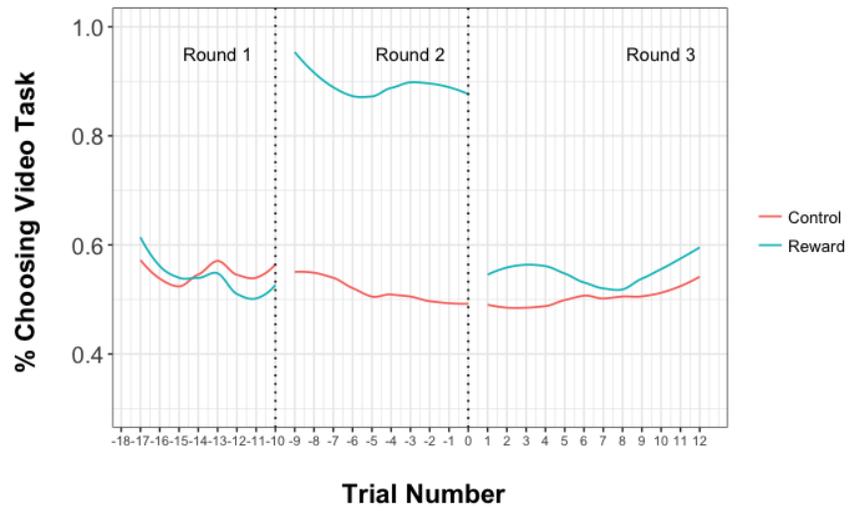
Study 1



Study 2



Study 3



Web Appendix C: Additional Methodological Details

Study compensation. Potential participants in all conditions were told that the study would take around 30 minutes to complete, and in addition to the base participation fee of \$1.45, they had a possibility of earning additional bonuses. Participants in the control condition earned the base payment and at the end of study, a surprise bonus (similar in magnitude to the average bonus earned by the incentive group) was announced and given in order to be consistent with the Mturk task description.

Study 1

Study 1 sample. The sample size (N=367) included participants who dropped out after round 1 (4%) in the analysis, coding their choices as not doing the target task.

Study 1 attitude recall check. Before proceeding to the post-incentive round, participants in the incentive-with-reminder condition successfully recalled their initial rating of the cognitive task (both mean ratings = 3.11, paired $t(110) = 0.16$, NS). The near-perfect recall suggests that participants indeed deliberated on their pre-incentive attitude and indicated their task liking. Consequently, we would expect they had access to strong cues about task attitude before they started their post-incentive round.

Study 2

Study 2 sample. The sample size (N=197) included participants who dropped out after round 1 (6.6%) in the analysis, coding their choices as not doing the target task.

Study 3

Study 3 sample: The sample size (N=475) included participants who dropped out after round 1 (1.9%) in the analysis, coding their choices as not doing the target task.