

THE UNIVERSITY OF CHICAGO

THE DYNAMIC EFFECT OF INCENTIVES ON POST-REWARD TASK
ENGAGEMENT

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE UNIVERSITY OF CHICAGO
BOOTH SCHOOL OF BUSINESS
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

BY
INDRANIL GOSWAMI

CHICAGO, ILLINOIS
JUNE 2016

Copyright © 2016 by Indranil Goswami
All Rights Reserved

To my parents who did not live to see this day

TABLE OF CONTENTS

LIST OF FIGURES	vi
LIST OF TABLES	ix
ACKNOWLEDGMENTS	x
ABSTRACT	xi
1 INTRODUCTION	1
2 POST-REWARD TASK ENGAGEMENT	4
2.1 Existing Theories of Post-Reward Engagement Reduction	5
3 EXPERIMENTAL PARADIGM	9
4 EMPIRICAL EXAMINATION OF POST-REWARD BEHAVIOR	12
4.1 Study 1: Dynamics of Post-Reward Behavior	12
4.2 Internal Meta-Analysis: Dynamics of Post-Reward Behavior	17
5 EFFORT-BALANCING PSYCHOLOGY OF POST-REWARD BEHAVIOR	20
6 EXPERIMENTAL INVESTIGATION OF EFFORT-BALANCING ACCOUNT	23
6.1 Study 2: Role of an External ‘Break’	23
6.2 Study 3: Effect of Large Rewards	27
6.3 Study 4: Incentives for Leisure	32
6.4 Study 5: Locking-in Engagement	36
7 ROBUSTNESS OF EFFORT-BALANCING ACCOUNT	39
7.1 Study 6: Investigating Signaling and Desirability	40
7.2 Study 7: Isolating Effort from Rewards	43
7.3 Study 8: Scope of Effort-Balancing	47
8 GENERAL DISCUSSION	51
REFERENCES	56
A STUDY STIMULI, PRE-TESTS, AND DATA CLEANING PROTOCOL	62
B HIERARCHICAL REGRESSION SPECIFICATIONS	72
C STUDIES USED IN INTERNAL META-ANALYSIS	75
D ROBUSTNESS CHECKS	76
E HETEROGENEITY IN POST REWARD BEHAVIOR	79

F	EFFECT ON ACCURACY AND NET OUTCOME	83
G	FIELD STUDIES USING TEMPORARY CONTINGENT INCENTIVES	85

LIST OF FIGURES

2.1	Illustration of persistent post-reward engagement reduction due to temporary incentives, as predicted by prior theories.	4
3.1	Participants chose between doing a math task or a video task in each trial. The specific task (i.e., math problem or video) was revealed on the subsequent page, after their choice.	9
4.1	Proportion of math task choices in the control and reward group in each phase of the experiment in Study 1. Dotted lines represent the baseline (average effort level Round 1), and the vertical lines are 95% CIs.	13
4.2	Raw percentage of participants choosing the math task in each trial in the internal meta-analysis. The average 95% margin of error in the control condition was +/- 4% and the average margin of error in the reward condition was +/- 3%.	17
6.1	Comparison of effortful unrelated activities (choice) and low-effort break (no-choice) conditions to control in Study 2. Dotted lines represent the baseline (average effort level Round 1), and the vertical lines are 95% CIs.	25
6.2	Average math task choices in the control, replication (5 cents), low-reward (1 cent) and high-reward (50) conditions in Study 3. Dotted lines represent the baseline (average effort level Round 1), and the vertical lines are 95% CIs.	28
6.3	Math task used in Study 4. The task required participants to count the number of 1s in the grid (of 150 numbers) in 30 seconds.	33
6.4	Post-reward engagement when the math task was incentivized (left-panel) or the video task was incentivized (right-panel) in Study 4. Dotted lines represent the baseline (average effort level Round 1), and the vertical lines are 95% CIs.	34
6.5	Percentage of participants locking-in math tasks for the post-reward period in Study 5. The vertical lines are 95% CIs.	37
7.1	The post-reward results in the replication condition and when both math and video were said to be equally important in Study 6. Dotted lines represent the average effort level across the conditions in the baseline (i.e., pre-incentive) period and the vertical bars are 95% CIs.	41
7.2	The post-reward results in low required effort and the high required effort conditions in Study 7. Dotted lines represent the average effort level across the conditions in the baseline (i.e., pre-incentive) period and the vertical bars are 95% CIs.	45
7.3	Choice of the effortful task in Study 8. Round 3 employed an effortful spelling task for all the experimental groups. Dotted lines represent the average effort level across the conditions in the baseline (i.e., pre-incentive) period and the vertical bars are 95% CIs.	49
A.1	The exhibit shows how the series of 30 repeated choices (or trials) were organized in the studies. The trials were divided into three rounds for all the experimental groups, with 8, 10, and 12 trials in the pre-, during-, and post-reward rounds respectively. The reward group(s) learnt about the total rewards earned before the start of the third round. However, all payments were made at the very end of the study, and not after the immediate end of the second round.	62

A.2	Four-stage instructions (shown as separated by dotted lines) provided to all participants before the start of all the studies. After the instructions, all participants sampled a task and a video and were asked if everything worked well. The experiment was aborted for those who reported having a problem.	63
A.3	The exhibit shows the types of tasks used for unrelated activities break in Study 2. Participants did all the three tasks of a particular type (Writing: top panel or Brand-name Matching: bottom panel), and each one of these tasks had a time limit of 30 seconds. The exhibit shows the choice condition. In the no-choice condition, one of the options was randomly pre-selected for the participant. . . .	64
A.4	The exhibit shows the manipulations used to make the perceived reward magnitude salient in Study 3. The top panel shows the low-reward condition, and the bottom panel shows the high-reward condition.	65
A.5	The exhibit shows the two types of framing used for the math (work: top panel) and the video (leisure: bottom panel) task in Study 4.	65
A.6	The exhibit shows a typical video task used in Study 4. Unlike other studies (Studies 1-3), in this study the participants in the target-task = video condition were asked to rate the video in order to get their rewards.	66
A.7	The exhibit shows the instructions (two stages, shown separated by a dotted line) given to participants at the end of Round 2 in Study 5. Participants were asked to choose which task they would like to do during the entire ensuing period of Round 3.	67
A.8	The exhibit shows the instructions given to participants at the beginning of Study 6. Participants were told that both tasks are equally important to the researchers.	67
A.9	The exhibit shows the instructions given to participants at the beginning of Round 2 (reward-round for treatment group) in Study 7. Participants in both the control (top panel) and the reward group (bottom panel) were told that the computer will choose the task (math or video) in each trial of the next round. Participants returned to making their own choices after the end of Round 2. . .	68
A.10	Top panel shows the instruction given to participants at the very beginning of Study 8. This was followed by a sample of each task-type. Participants in both the control and the reward group made multiple choices between doing a spelling task or watching a video in Round 3. The bottom panel shows an actual example of a spelling task used in Round 3.	69
D.1	Figure shows the predicted probability of choosing the math task in the reward group (blue triangular dot) and the control group (red circular dot) in Round 3 based on a typical run of the non-parametric model.	78
E.1	The figure shows reward and the control groups further divided into two sub-groups based on initial post-reward behavior (initial post reward behavior lower or higher than individual-level pre-reward average). For example, ‘Reward - Initially Low’ are reward-group participants who showed an initial dip in engagement (compared to their own Round 1 baseline) at the end of Round 2.	79
E.2	Post-reward behavior of participants with high initial task interest using the internal meta-analysis data. The figure shows a significant post-reward decrease in engagement after incentives ended.	81

E.3	Post-reward behavior of participants who low initial task interest using the internal meta-analysis data. The figure shows a significant post-reward decrease in engagement after incentives ended.	81
F.1	Raw data showing the percentage of participants accurately solving the math task, conditional on choosing to attempt it. Incentives did not affect accuracy in our experiments.	83
F.2	Raw data showing proportion of correct answers for every trial in each round. Incentives had a net positive effect on the total number of correct answers, driven by the reward period, despite a significant post-reward decrease after incentives ended.	83

LIST OF TABLES

C.1	List of Studies included in the internal meta-analysis. Not all studies are included in the manuscript as indicated in the Remarks column.	75
D.1	Estimates for Momentary Reduction in Engagement $\beta_{MOMENTARY}$ in meta-analysis with various specifications, SEs are in parenthesis	76
D.2	The table shows the bootstrapped 95% CI from the non-parametric model for the difference in the predicted probability of choosing the math task between the reward and the control condition.	78
E.1	Regression results showing that longer-term post-reward behavior (i.e., return to baseline) not different between control and reward group as a function of their initial post-reward behavior	80
G.1	List of past field studies that have used temporary incentives and looked at post-reward behavior. Only studies that (a) have measured actual post-reward behavior (and not self-report), (b) have studied adults (including a Study with 11th and 12th graders), and (c) where the target task is not a pro-social activity (e.g., contribution to charity, blood donation etc.) are included.	85

ACKNOWLEDGMENTS

I would like to express my heartfelt thanks to my advisor Oleg Urminsky for his continuous support during this and other related research. Words are too feeble to describe how much I have benefited, both personally as well as a researcher, from his endless patience, deep insights, and uncomplaining availability despite his busy schedule. I am also indebted to my other committee members: George Wu, Jean-Pierre Dubé, and Dan Bartels for their valuable feedback throughout my dissertation research. My sincere gratitude to Richard Hahn for his active interest in this work and his statistical insights. I am thankful to Abigail Sussman, Aparna Labroo, Reid Hastie, Pradeep Chintagunta, and Sanjay Dhar for their useful advice at various stages of the research. Finally, I would like to acknowledge the seed grant from the University of Chicago BIG Ideas, and grant from the Initiative on Global Markets at the Booth School of Business for providing financial support for conducting this research.

ABSTRACT

Although incentives can be a powerful motivator of behavior when they are available, an influential body of research has suggested that rewards can persistently reduce engagement *after* they end. This research has resulted in widespread skepticism among practitioners and academics alike about using incentives to motivate behavior change. However, recent field studies looking at the longer-term effects of temporary incentives have not found such detrimental behavior. We design an experimental framework to study the dynamic behavior under temporary rewards, and find that although there is a robust decrease in engagement immediately after the incentive ends, engagement returns to a post-reward baseline that is equal to or exceeds the initial baseline. As a result, the net effect of temporary incentives on behavior is strongly positive. The decrease in post-reward engagement is not on account of a reduction in intrinsic motivation, but is instead driven by a desire to take a ‘break’, consistent with maintaining a balance between goals with primarily immediate and primarily delayed benefits. Further supporting this interpretation, the decrease in post-reward engagement is reduced by contextual factors (such as less task difficulty and higher magnitude incentives) that reduce the imbalance between effort and leisure. These findings are contrary to the predictions of major established accounts and have important implications for designing effective incentive policies to motivate behavior change.

CHAPTER 1

INTRODUCTION

Incentives can be a powerful tool to influence behavior in settings as diverse as health, education, employment, and marketing promotions (Gneezy et al., 2011; Prendergast, 1999; DelVecchio et al., 2006) where people’s decisions often involve a tradeoff between immediate and delayed benefits (Loewenstein et al., 2007). Providing immediate temporary incentives can help motivate people to take beneficial action, countering the effects of hyperbolic discounting and present bias (Ainslie, 1975; Urminsky and Zauberman, 2016). However, policy makers are often skeptical about using incentives based on theories of motivation, signaling, and information. Even when incentives do not reveal additional information, or the task being incentivized has no signaling benefits, critics contend that economic incentives will “smother people’s enthusiasm for activities they might otherwise enjoy” (Kohn, 1999). As a result, people are predicted to engage in *less* of the incentivized behavior after an incentive ends than if the incentive had not been introduced in the first place (Pink, 2011).

This negative view of incentives stems from a large and influential academic literature on Cognitive Evaluation Theory (Deci and Ryan, 1985) and the Overjustification Hypothesis (Greene and Lepper, 1978) which found that adults did less of a task in lab studies (compared to a non-incentivized control group) immediately after an incentive ended. However, these studies only measured post-reward behavior immediately after the incentives were withdrawn, and did not examine the dynamics of longer-term behavior. Nevertheless, a highly influential meta-analysis of this literature generalized beyond the findings to warn that “if people use tangible [i.e., monetary] rewards, it is necessary that they be extremely careful about the intrinsic motivation and task persistence of the people they are rewarding” (Deci et al., 1999, p. 656, emphasis added). In contrast with the prediction from these theories that incentives permanently reduce intrinsic motivation (Dickinson, 1989; Tang and Hall, 1995), recent program-evaluation studies that measured long-term effects after incentives ended have not found evidence of post-reward reduction in incentivized behaviors, with

some even reporting positive long-term effects (e.g., workplace smoking cessation, Halpern et al. (2015); academic performance, Jackson (2010)).

In this research, we investigate the question raised by these conflicting findings: what are the dynamic effects of incentives on people’s task engagement once the incentive ends? To answer this question, we develop an experimental framework that facilitates dynamic measurement of post-incentive task engagement over a series of choices. We find initial reduction in engagement with the incentivized task after the rewards end, consistent with the findings of prior lab experiments. However, this reduction is *momentary* in duration and the incentive yields no persistent negative effect or even a positive longer-term effect over multiple rounds, consistent with field studies of post-incentive behavior. Our results suggest that prior theories of intrinsic motivation were incomplete for predicting post-incentive behavior.

Extending the literature on goal-balancing (Fishbach and Dhar, 2005) and the effect of justification on choices (Lerner and Tetlock, 1999; Shafir et al., 1993), we suggest that people attempt to maintain a balance between investing effort (which often yields delayed benefits) and the immediate enjoyment of leisure (the effort-balancing account). Investing more effort in response to modest incentives can disturb this perceived sense of balance, justifying more leisure immediately after the incentives end. Our account makes testable predictions about how the effect of temporary incentives on post-reward behavior would vary with contextual factors and provides different recommendations for how to design effective temporary incentive policies.

The rest of this dissertation is divided into two parts. The first part comprises three chapters where we systematically examine the longer-term post-reward effects of incentivizing an interesting task in the presence of an outside alternative. In Chapter 2, we define post-reward task engagement and briefly review prior theoretical accounts that make predictions about post-reward task engagement behavior. In Chapter 3, we introduce an experimental framework which enables us to empirically examine the existence and persistence of a post-incentive reduction in task engagement. In Chapter 4, we test the effect of ending an incentive

(Study 1), as well as report an internal meta-analysis of all data we have collected on post-reward behavior.

The second part comprises four chapters where we investigate the psychology of post-reward behavior. In Chapter 5, we introduce an Effort-Balancing account to explain the dynamic effect of incentives on post-reward task engagement. In Chapter 6, we document four additional studies to investigate boundary conditions of our account as a means to understand how the incentive affects task engagement. We vary the presence and type of post-incentive break from the task (Study 2), the perceived magnitude of the incentive (Study 3), the type of target task (effortful vs. leisure, Study 4) and how decisions about engagement are made (up-front vs. sequentially, Study 5). In Chapter 7, we investigate the robustness of our proposed account by framing of the target task and outside alternative as equally important to the experimenter (Study 6), isolating the effect of rewards alone by exogenously fixing effort during the reward period (Study 7), and examining the scope of effort-balancing account by changing the target task *after* the reward period (Study 8). In Chapter 8, we conclude with a general discussion and future research directions.

CHAPTER 2

POST-REWARD TASK ENGAGEMENT

Consider a situation in which people repeatedly choose which of two tasks to do for a set period of time (e.g., 30 seconds) - a ‘target’ task and an alternative task, as illustrated in Figure 2.1. People’s choices of how much of the target task to do without any incentive (in Round 1, the baseline period) are based on their intrinsic motivation for the task. Subsequently, when some people receive a previously unannounced temporary incentive for choosing to do the target task (Round 2), they are likely to do more of this incentivized task¹.




Figure 2.1: Illustration of persistent post-reward engagement reduction due to temporary incentives, as predicted by prior theories.

The key question in our research is how much of the target task people will do in Round 3, after the incentive has ended. The established theories predict a persistent reduction in engagement in the incentivized task after contingent incentives have ended, with fewer choices

1. There are exceptions to this behavior, and in this project we avoid situations where such exceptions can arise. We discuss this later in the text.

of the incentivized task than if the incentive had never been introduced (e.g., vs. control, as shown in Figure 2.1). This reduction in engagement is the behavior associated with a reduction in intrinsic motivation, which is predicted when people had been interested in the target task before the incentive and when the incentive does not increase perceived competence (Deci et al., 1999, 2001; Lepper et al., 1999). The magnitude and duration of this post-incentive reduction, relative to any increase in choosing the task during the incentive period, determines the *net effect* of a temporary incentive policy on behavior change, but this has not been studied in the prior literature.

2.1 Existing Theories of Post-Reward Engagement Reduction

Prior research includes multiple accounts of how incentives will affect post-incentive behavior. While these accounts similarly predict post-incentive decline in engagement, the underlying processes differ, and the accounts therefore make differing predictions about the circumstances in which the decline will be stronger or weaker.

Undermining Autonomy

Cognitive Evaluation Theory argues that people engage in tasks to fulfill innate needs of competence, exploration, and autonomy (White, 1959; Deci, 1971). An external incentive that is conditional on doing the target task will exert control over people’s behavior, and thereby undermine satisfaction from autonomy and exploration (Deci and Ryan, 1985; Ryan and Deci, 2000). As a result, engagement in the task will decline (e.g., compared to initial engagement or a control group) once an incentive ends, unless the incentives communicates competence, which can countervail the negative effects. This account predicts that because the incentive changes people’s interpretation of the target task and the benefit they experience from doing the task, the post-incentive engagement reduction will be persistent (Deci et al., 1999). In this account, the negative effect on intrinsic motivation is stronger either

when the task was initially more intrinsically motivating or when the conditional incentive is larger.

Overjustification

An alternative account, based on self-perception (Bem, 1972), posits that incentives impact inferences about one’s own preferences. After being paid to work on a target task, people may think that they were doing the task only because of the external rewards (Kruglanski et al., 1971, 1972; Lepper et al., 1973; Greene and Lepper, 1978), discount their own intrinsic motives (Nisbett and Valins, 1971; Kelley, 1973) and infer that they do not like doing the task. Based on this revised belief about their own preferences, people would then do less of the task after the incentive ends than they would have if the incentive was not introduced. This account predicts that the reduction will persist until another influence causes those preferences to change, and the reduction will be stronger for larger incentives.

Task and Social Inferences

Incentives can also influence people’s inferences about unknown aspects of the target task, based on what they think of the incentive provider’s motives. Since incentives are often introduced when tasks are unattractive or difficult, people may conclude that an incentivized task is relatively unappealing (Benabou and Tirole, 2003), reducing subsequent choices to do the task. High incentives should send a stronger signal, resulting in stronger reductions, while prior first-hand experience with the target task would be expected to mitigate the effects.

Expectations

Ending an incentive, without advance notice, can be seen as a violation of norms or expectations, resulting in disappointment (Fehr and Falk, 2002). Relatedly, the incentive may

create a reference point, such that then doing the task without an incentive could be viewed as a relative loss (Kalyanaram and Winer, 1995). In both cases, when people learn that their expectation of continuing incentives will not be fulfilled, engagement in the target task can fall (Esteves-Sorensen et al., 2013). If people cope with disappointment and adjust their reference points over time (Gilbert et al., 1998; Schkade and Kahneman, 1998), these accounts predict that the post-incentive reduction will be temporary, and can be prevented when people have advance knowledge of the temporary nature of the incentive.

Temporal Shift in Effort

When decision makers know that the incentive will be temporary, they can strategically shift effort they would have invested during the post-incentive period, instead doing more of the target task during the incentive period. As a result, there might be little or no change in overall task engagement due to the incentives, but an increase during the incentive period and a decrease afterwards. For example, employees might move up task completion in time to help them earn a bonus (Oyer, 1998) or people might stockpile non-perishable goods while a purchase incentive is available (Mela et al., 1998). Reduction in engagement due to effort shifting should occur primarily when people know that the rewards are temporary and can predict the timing of their withdrawal. Temporal shifting should be larger and persist for longer when the incentive is larger.

These accounts all predict an increase during the incentive period, followed by a persistent decline when the incentive ends, and stronger effects for larger incentives. Other accounts, based on information and signaling (Frey and Oberholzer-Gee, 1997; Gneezy and Rustichini, 2000; Ariely et al., 2009; Dubé et al., 2015), predict situations for which an incentive may reduce engagement while the incentive is available (e.g., engaging in pro-social tasks). We constrain our investigation to the effects of *successful* incentives (i.e., those which increased engagement while available) on post-reward engagement. Furthermore, our experimental framework, which we describe in the next chapter, is designed to minimize the potential for

task and social inferences, disappointment about the incentive ending, or temporal shifts in effort, in order to effectively test the two prominent motivation-based accounts discussed first.

CHAPTER 3

EXPERIMENTAL PARADIGM

In order to examine the longer-run dynamics of post-reward behavior, we develop an experimental paradigm in which we can track repeated decisions about whether to engage in the target task after the incentives end. The “quitting paradigm” often used to study persistence (Deci, 1971; Heyman and Ariely, 2004) uses time spent working on an unsolvable task as the key measure. However, in this approach people make one quitting decision, and therefore the measure cannot distinguish between a temporary reduction in motivation to do the target task and a permanent disengagement. In our experimental setup, participants make repeated choices between doing a moderately effortful but interesting math problem (the target task) or watching a fun video clip (the alternative task), as illustrated in Figure 3.1.

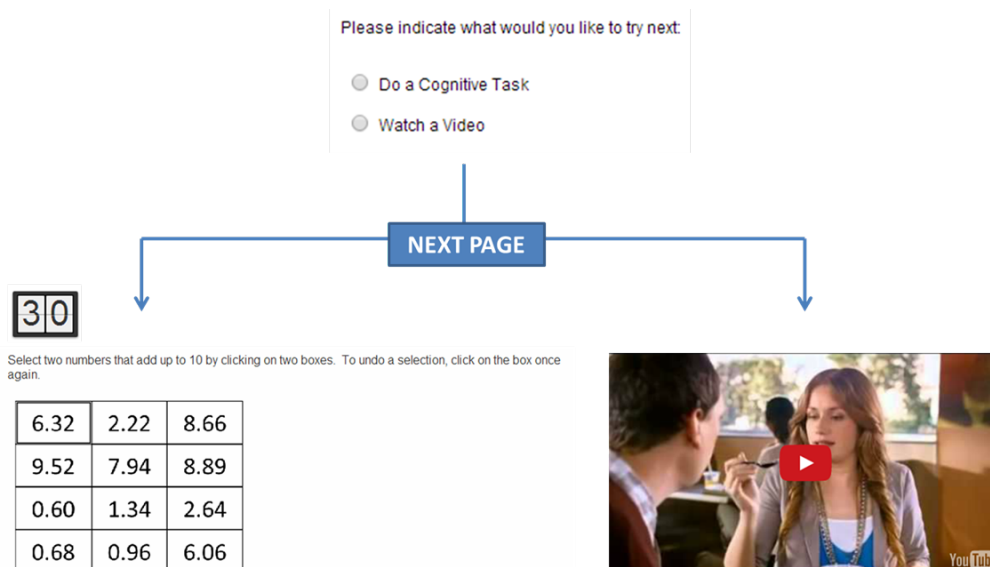


Figure 3.1: Participants chose between doing a math task or a video task in each trial. The specific task (i.e., math problem or video) was revealed on the subsequent page, after their choice.

The 30 trials were divided into three *rounds*: pre-incentive baseline (8 trials), randomized incentive (10 trials), and post-incentive (12 trials). The initial baseline round gave participants first-hand experience with the target task, minimizing the potential for participants

assigned to the reward group to draw inferences from the incentive about unknown task characteristics or preferences. Participants were randomly assigned to reward (incentive during Round 2) or control (no incentive) groups. At the end of each round, participants in both the control and the reward group were informed that the round was over, and were instructed to click “Next” to proceed to the next round. Participants did not know in advance how many trials were in each round, or how many rounds would occur, limiting their ability to strategically shift effort from the post-reward to the reward period.

The math task used in the studies (except for Study 4), was adapted from Experiment 1 of Mazar, Amir and Ariely (2008) and involved finding the two numbers in a grid of 12 numbers that added up to 10. Participants had 30 seconds to solve the problem, the same duration as the video clips in the alternative task¹. All participants were given the following instructions:

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.

To investigate the effect of external incentives on intrinsic motivation, it is important to use tasks that are intrinsically motivating. Deci, Koestner and Ryan (1999) set average ratings of at least the mid-point on a scale measuring how ‘interesting and enjoyable’ the task is ($1 = Low, 9 = High$) as a criterion. A pre-test ($N = 47$) confirmed that both the math task ($M_{math} = 6.02, SD = 2.19$) and the video task ($M_{video} = 6.64, SD = 2.08$) fulfilled this requirement. In addition, our pre-tests confirmed that people would choose to do some math tasks without any incentive, even with the videos as an alternative.

Participants in all the studies were recruited from Amazon Mechanical Turk (AMT),

1. Detailed stimuli, including additional pre-test measures, are available in the Appendix A

an online employment marketplace for brief tasks that is used by companies and academic researchers. Potential participants were told that the study would take around 30 minutes to complete, and in addition to the base participation fee of \$1.75, they had a possibility of earning additional bonuses. At the end of the baseline Round 1, those participants randomly assigned to the reward condition were told that they could earn bonuses by correctly solving the math problems and the amount. These participants were told in advance that the reward opportunity was only available during the next round, so that the end of the incentive would not come as a surprise and cause disappointment (although ultimately that does not seem to matter, see Study 5 in Appendix C).

At the end of Round 2, participants in the reward condition were told how much bonus they had earned, to be paid at the end of the experiment. For the control (no-incentive group) a surprise bonus, similar in magnitude to the average bonus earned by the reward group, was announced at the end of the experiment. Data was included for all participants who completed the first round of the task, to minimize selection effects from the incentive manipulation, with those dropping out after the first round coded as not doing the target task in subsequent trials.

CHAPTER 4

EMPIRICAL EXAMINATION OF POST-REWARD BEHAVIOR

4.1 Study 1: Dynamics of Post-Reward Behavior

Method

A target of 100 participants was requested, yielding 91 surveys. In all studies, the sample size was determined in advance, and we used the same data-cleaning protocol (see Appendix A), implemented prior to data analysis. Records with duplicate IP addresses (3 respondents), or who reported having technical problems with viewing the videos or working on the math task (4 respondents), or who failed the basic attention check (7 respondents) were removed prior to analysis, yielding 77 valid completes. In all studies, we include any participants who completed Round 1 and therefore could have been exposed to an incentive, even if they did not complete the rest of the study. Our final sample included two participants (2.6%) who completed Round 1, but then dropped-out later in the experiment. Immediately after the end of the first round, participants randomly assigned to the incentive condition were informed that they could earn 5 cents for every math task they choose to do and answered correctly.

Results

In the control condition, the math task was chosen 67% of the time during the baseline period, confirming that participants found the task intrinsically motivating. This is a prerequisite for our intended tests, because tasks with have little or no baseline engagement cannot show substantial reduction in engagement due to the floor effect. The average engagement (i.e., choices of the math task) remained at around the same level during the following experimental periods in the control condition, as shown in Figure 4.1.

In the incentive condition, the proportion of math task choices increased to 88% when




Figure 4.1: Proportion of math task choices in the control and reward group in each phase of the experiment in Study 1. Dotted lines represent the baseline (average effort level Round 1), and the vertical lines are 95% CIs.

the incentive was available - a significantly higher level than in the control condition (61%) during the same period, controlling for the baseline proportion of math task choices in Round 1¹ ($\beta = 0.28, t = 5.42, p < .001$). Therefore, the incentive was successful at increasing participants' task engagement while the incentive was available, also a precondition of the intended test.

The key question we designed the experiment to test was whether people's engagement with the focal task after the incentive ended was affected by having previously offered the incentive. For participants in the reward condition, the average percentage of math task choices dropped to 53% in the very first trial in Round 3, immediately after the incentives ended, compared to 72% in the control condition. This initial test yielded a borderline significant difference ($\beta = -0.19, t = -2.02, p = .047$). The drop in people's choices of the math task after the incentivized Round 2 had ended is consistent with the prior findings in lab studies of reductions in the incentivized behavior immediately after rewards ended.

Since we tracked the dynamics of post-incentive task engagement over repeated trials,

1. In all studies we report both linear regression t-tests that control for the individual baseline levels in Round 1 and results from a hierarchical regression model.

we can also distinguish between temporary and permanent disengagement with the target task. Averaging across the entire post-reward period (Round 3), people did a similar number of math tasks in the reward condition as in the control condition during the same period (66% vs. 65%; $\beta = .008, t < 1$). Although we do find an initial decrease in engagement immediately after the incentives ended, this decrease was brief, and the incentive did not result in any overall reduction in post-reward engagement. This result is not consistent with the predictions made by the prior theories of a persistent decrease in engagement due to the incentive.

Because the reduction in engagement was only momentary, the immediate negative post-reward effect of incentives was small in scope relative to the prolonged positive effect during the incentive period. Overall, offering a temporary incentive in the reward condition yielded a significant positive net effect in the study, with people choosing to do more of the math task in the incentive condition than in the control condition during the post-baseline periods, controlling for Round 1 (76% vs. 63%; $\beta = 0.13, t = 2.89, p = .005$).

Hierarchical Regression Specification

In order to more precisely quantify the magnitude and duration of the reduction in engagement, we used a hierarchical non-linear mixed model (Raudenbush and Bryk, 2002), with the per-trial observations nested under individuals. We assume that after the initial decline, the return to the baseline level of engagement is inversely proportional to the number of periods t since the incentive ended. Based on this assumption, we define the probability of individual i choosing to do the math task in the post-reward period (Round 3) during trial t as²:

$$Pr(Y_{ti} = 1) = \text{logit}(\beta_{0i} + \beta_{Mi} * \frac{1}{t}) \quad (4.1)$$

2. See Appendix B for full details. We relax this inverse proportionality assumption in several ways and test robustness to different specifications, including a non-parametric test.

In this model, β_{0i} is a person-specific intercept and β_{Mi} is a person-specific coefficient capturing the momentary post-reward reduction in task engagement. These parameters are in turn a function of time-invariant individual-level covariates \mathbf{X}_i and the experimental condition C_i . The expected proportion of math tasks chosen in each trial t of Round 3 can therefore be written as:

$$Pr(Y_{ti} = 1) = \text{logit}(\beta_{00} + \beta_{10} * \frac{1}{t} + \beta_{01} * C_i + \beta_{02} * \mathbf{X}_i + \beta_{MOMENTARY} * C_i * \frac{1}{t}) \quad (4.2)$$

the coefficient $\beta_{MOMENTARY}$ tests for a difference between the momentary engagement reduction in the experimental condition ($C_i = 1$) compared to the same time periods in the control condition ($C_i = 0$). A significant and negative $\beta_{MOMENTARY}$ indicates a momentary engagement reduction after the incentive ended (assuming no main effect of condition), compared to the corresponding trials in the control condition, controlling for individual differences (e.g., different individual baseline efforts).

We use a similar hierarchical difference-in-difference model to test whether the difference in the overall probability of choosing the math task between two experimental rounds ($R_t = 0$ vs. $R_t = 1$) differs by condition ($C_i = 0$ vs. $C_i = 1$):

$$Pr(Y_{ti} = 1) = \text{logit}(\beta_0 + \beta_1 * R_t + \beta_2 * C_i + \beta_3 * R_t * C_i) \quad (4.3)$$

We can test multiple hypotheses using the key coefficient β_3 , depending on how the rounds (R_t) and conditions (C_i) are coded. Using different coding in this model (see Appendix B), we estimate either β_{REWARD} - the effect of incentives on during-reward performance; β_{POST} - the overall post-incentive engagement level; or the net effect of incentives (combining the incentive and post-incentive periods), β_{NET} . Next, we use this approach to analyze the results of Study 1, followed by an internal-meta analysis of the same two conditions in all the data we have collected.

Regression Results for Study 1

We find a significant momentary engagement reduction after the incentive ended ($\beta_{MOMENTARY} = -2.56, z = -2.59, p < .01$). Moreover, there was no overall reduction in engagement with the math task after incentives ended ($\beta_{POST} = -0.003, z = -0.01, p > .250$). Therefore, choices of the math task returned to the pre-reward baseline level after an initial post-reward momentary decrease. Since the incentive did boost math task choices significantly over the baseline level during Round 2 ($\beta_{REWARD} = +4.60, z = +4.22, p < .001$), the net effect of incentives was a significant increase in the number of math tasks chosen over the course of the study ($\beta_{NET} = +1.21, z = +2.75, p = .006$).

Discussion

In Study 1, consistent with the findings of the existing psychological theories, we find a reduction in engagement with the target task immediately after the incentive ends. However, contrary to the predictions of those theories, this reduction in engagement is not persistent. Instead, tracking trial by trial engagement, we find a momentary reduction in engagement, followed by a return to baseline, and no longer-run reduction in engagement. This absence of a longer-term reduction in engagement is consistent with the empirical findings of several recent field studies. These studies did not track engagement over time, but measured people’s total behavior in the days, weeks or months after the incentive ended and found either no long-term effect (John et al., 2011; Lacetera et al., 2013; Volpp et al., 2006), or a modest increase in engagement (Charness and Gneezy, 2009; Jackson, 2010).

Our findings therefore suggest a potential reconciliation between the immediate post-incentive declines reported in the lab experiments and the lack of long-term post-incentive decline in the field studies. However, it is important to first establish the robustness of the findings beyond this one study, and to explore heterogeneity in the response to incentives across individuals.

4.2 Internal Meta-Analysis: Dynamics of Post-Reward Behavior

In order to test the robustness of the findings, we conducted an internal meta-analysis, combining all data collected in the same reward and control conditions as Study 1 across our studies. The resulting data comes from ten of our studies (see Appendix C for details), totaling 1,098 participants (530 in control, and 568 in the reward group). The large sample size allows us to examine post-incentive behavior more closely, do various robustness checks, and look at heterogeneity in responses. The proportion of math tasks in each trial, plotted separately by experimental condition, is shown in Figure 4.2. The results reveal a significant momentary reduction in engagement in the reward condition relative to the control condition, after the incentive has ended ($\beta_{MOMENTARY} = -3.78, z = -9.56, p < .0001$).




Figure 4.2: Raw percentage of participants choosing the math task in each trial in the internal meta-analysis. The average 95% margin of error in the control condition was $\pm 4\%$ and the average margin of error in the reward condition was $\pm 3\%$.

However, comparing the pre- and post-reward rounds, we find a significant overall increase in choices of the math task post-reward ($\beta_{POST} = +0.36, z = +2.00, p = .045$).

Furthermore, given the momentary nature of post-reward reduction in engagement and the otherwise positive effects of incentives, the combined data also show a strong positive net effect of incentives on the number of math tasks chosen ($\beta_{NET} = +1.21, z = +9.55, p < .0001$). These results are robust to different parameterizations of the post-reward recovery (e.g., t^{-n} with different values of $n > 0$) and also to excluding participants who dropped out before completing the full study. It is important to note that accuracy in correctly solving the math problems did not vary between reward and control across the rounds. As a consequence, the same results hold if we look at the number of successfully solved math problems, rather than the number of attempts (see Appendix F).

Next, we investigate heterogeneity in people’s behavior, quantifying how many participants’ individual behavior is represented by the average patterns shown in Figure 4.2. In the reward group, 41% of the participants showed a decrease in engagement in the first trial of Round 3, relative to their own individual baseline level in Round 1, significantly more than in the no-reward control group (41% vs. 26%; $\chi^2(1) = 28, p < .0001$). The results are similar if we instead compare the average of the first two or even three trials in Round 3. This suggests that, consistent with the average results, many individual participants in the reward group did show a reduction in engagement during the first three post-incentive trials, significantly more than in the control group. This tendency to reduce engagement in the immediate post-reward period was not moderated by initial motivation (i.e., choices of the math task in Round 1, see Appendix E).

We also investigated the possibility that the initial post-incentive decrease and subsequent increase over baseline that we observe in Round 3 could be explained by heterogeneity, resulting from averaging the behavior of some participants who displayed an initial reduction and other participants with a long-term increase. We quantified the degree to which participants had an initial post-reward reduction, by computing the difference in average math tasks in the first three trials of Round 3 versus that person’s Round 1 baseline. We also quantified the participants’ final post-reward behavior, by computing the difference in av-

erage math tasks in the last seven trials of Round 3 versus that person’s Round 1 baseline. Compared to control, we found no relationship between initial and final post-reward behavior in the reward group ($\beta = 0.01, t = 0.41, p > .250$, see Appendix E). Among participants in the reward group who had an initial reduction, 34% were finally above (settled at a higher post-reward baseline), compared to 13% in the corresponding segment of the control group. Among participants in the reward group who were not initially below (i.e., did not show a momentary post-incentive decrease in engagement), 45% were finally above, compared to 31% in the corresponding segment of the control group. Thus, both participants who had an initial reduction and those who did not have an initial reduction showed a similar longer-term increase, relative to control.

Prior theories of intrinsic motivation (Lepper and Greene, 1978; Deci and Ryan, 1985) predict a persistent decrease in engagement after the incentive ends specifically because introducing the incentive would result in a decrease in intrinsic motivation for the target task. Consistent with our finding that the reduction in engagement is only momentary and engagement eventually increases to be higher than baseline, self-reported interest in the task at the end of the study, a measure of intrinsic motivation, was actually higher in the reward group compared to the control in our meta-analysis ($M_{reward} = 5.97, SD = 2.35$ vs. $M_{control} = 5.59, SD = 2.51; t(890) = 2.3, p = .02$). In fact, in Deci, Koestner and Ryan (1999), conclusions about intrinsic motivation were based only on a meta-analysis of task engagement measures, and their analysis likewise did not find the predicted decrease in self-reported task-liking after contingent incentives ended. These results suggest that the momentary reduction in task engagement after rewards end may not be attributable to the changes in intrinsic motivation that they have proposed, but may need to be explained by a different psychological process.

CHAPTER 5

EFFORT-BALANCING PSYCHOLOGY OF POST-REWARD BEHAVIOR

We hypothesize that our findings may be explained by people attempting to strike a balance between investing effort and enjoying leisure. Some intrinsically motivating tasks, such as watching videos, represent leisure and provide immediate benefits with minimal effort. Other intrinsically motivating tasks require immediate effort or cost to gain some combination of immediate and delayed benefits. For example, going to the gym, going to the doctor for a regular health checkup, learning a new skill or practicing an existing skill are intrinsically motivating, in that people are willing to engage in these tasks without any external rewards, but require upfront effort. As a result, people often feel they should be engaging in such activities for their future welfare, but may want to engage in more enjoyable and less effortful activities (e.g., watching TV). We can think of decisions about whether to engage in these kinds of effortful intrinsically motivating tasks as similar to self-control conflicts between goals representing more immediate and more delayed benefits (Ainslie, 1975; Bazerman et al., 1998; Milkman et al., 2008).

People often strive to find a balance between competing goals (Dhar and Simonson, 1999; Fishbach and Dhar, 2005), including between effort and leisure (Kool and Botvinick, 2014), and between effort and expected reward (Kivetz, 2003). In the absence of incentives, we can think of people’s choices between doing tasks associated with more delayed benefits (e.g., solving math problems) or doing tasks with immediate benefits (e.g., watching videos) as representing their preferred level of balance. In our setting, this would be measured by the pre-reward baseline level of engagement in the target task (vs. the leisure alternative).

Introducing a temporary incentive adds an immediate benefit to any pre-existing intrinsic motivation associated with the task. When goals compete for resources, people are more motivated by the goals that are more salient (Shah and Kruglanski, 2002) and more prox-

imal (Brown, 1948; Kivetz et al., 2006; Urminsky and Goswami, 2015), which would favor the incentivized task. When people therefore increase work on the incentivized task, their competing leisure goals are likely to be deferred to avoid interruption (Jhang and Lynch, 2015) until a time deemed more appropriate (Shu and Gneezy, 2010).

As people forego leisure in favor of the compensated task, the leisure goal is likely to persist in the background. When the incentive ends, motivation will shift from the completed incentive-earning goal, and thereby from the effortful task, to the neglected competing goal (Fishbach and Dhar, 2005; Förster et al., 2007; Khan and Dhar, 2006; Kruglanski et al., 2002) of pursuing leisure. Recent research on self-control has indeed suggested that depletion-like regulatory failures can represent a motivated switching between labor and leisure to strike a balance between two goals, without necessarily reflecting exhaustion of limited resources (Inzlicht et al., 2014; Inzlicht and Schmeichel, 2012). Choices of immediately gratifying options are also increased by the availability of justifications (Kivetz and Zheng, 2006; Lerner and Tetlock, 1999), such as having just worked on an effortful task. Therefore, immediately after the incentive ends, people will be more motivated to do a leisure alternative and will find it easier to justify doing so.

Prior research has shown that after some level of goal-attainment, the reduction in motivation to pursue that goal is only temporary (Förster et al., 2005). Even minimal satisfaction of an immediate gratification goal can restore people’s motivation to pursue longer-term goals (Urminsky and Kivetz, 2011). Hence, after engaging in extra post-reward leisure, the perceived balance between competing goals is eventually restored.

However, engagement in the focal task may settle at a higher post-incentive baseline post-incentive (compared to the initial pre-incentive baseline) for multiple reasons. Greater engagement with the target task when the incentive is available may promote habit-formation (Neal et al., 2006), resulting in a higher post-reward baseline. More directly, associating a target task with rewards can foster positive associations, and increase task liking (De Houwer et al., 2001; Razran, 1954), particularly when the rewards are high in magnitude. Findings

in the literature on reinforcement behavior also suggests that rewards can sometimes result in eventual positive long-term effects on behavior (Brooks and Bouton, 1993).

We propose that people prioritize effortful incentivized tasks when a reward is available, but then are looking to take a break from the effortful task when the incentive ends, making the alternative leisure task temporarily more appealing and easier to justify. After having a break by doing more of the leisure task, balance is restored and task engagement settles at a baseline level, which may have been increased relative to the initial baseline by the experience of the incentive. This “effort-balancing” account makes similar predictions as the prior theories about engagement reduction immediately after the incentive, but predicts that the reduction will not persist.

The effort-balancing account also makes specific testable predictions about how the post-incentive reduction in engagement will be affected by contextual factors. Giving people a ‘break’ after the incentive ends should help restore balance and ameliorate the reduction in engagement. Incentivizing a less effortful task or paying a higher incentive (but holding effort constant) should reduce the imbalance arising in the incentive period, likewise precluding a post-incentive reduction in engagement (and potentially fostering a higher post-incentive baseline). These predicted boundary conditions are not predicted by prior theories of intrinsic motivation, enabling us to use these competing predictions to test between accounts in the remaining studies.

CHAPTER 6

EXPERIMENTAL INVESTIGATION OF EFFORT-BALANCING ACCOUNT

6.1 Study 2: Role of an External ‘Break’

Method

We requested 350 participants from Amazon MTurk, yielding 322 surveys. Unusable cases (due to duplicate IPs, technical problems or a failed attention check) were removed prior to analysis, yielding 257 valid completes, or approximately 85 per pair of conditions to be compared. The few participants who completed Round 1 but then dropped-out part way through (5.4%) were coded as not doing the focal task and included in the analysis.

Participants were randomly assigned to one of six conditions. Two of the conditions (no-reward control, 5-cent per correct answer incentive) were the same as in Study 1. In the other four conditions, participants were given the 5 cent conditional incentive in Round 2, and then were asked to do an unrelated activity for 90 seconds immediately after the incentives ended. Ninety seconds was chosen to be equal to the duration of three trials, the approximate duration of post-reward reduction in engagement observed in the internal meta-analysis.

In the four unrelated-activity conditions, we varied task type (writing vs. logo-matching) and whether or not participants needed to make choices about the activity, in a 2 x 2 design. In the no-choice conditions, participants were assigned to an activity (either writing about their opinions on an assigned topic or matching assigned brand logos to brand names). Participants were assigned to three such tasks, each of 30 seconds duration. In the choice conditions, participants had to first choose which unrelated activity they would do (picking their topics in the writing task or selecting product categories in the logo-matching tasks).

The choice manipulation provides a direct test between the effort-balancing account and prior autonomy-based accounts. Consistent with findings from prior research (Vohs et al.,

2014), a pre-test confirmed that requiring people to make choices (in the choice conditions) was more effortful, on the one hand, but that it also provided significantly more of a sense of autonomy. In contrast, the pre-tests confirmed that the no-choice conditions involved little effort (serving as a break), but did not bolster autonomy. If reduction in engagement occurs because of a loss of autonomy, giving people choices about the unrelated activity (i.e., the choice conditions) may counter this loss of autonomy, ameliorating the reduction in engagement. However, if the reduction occurs because people need a break after exerting effort, then the non-effortful break provided in the no-choice conditions will be more effective in eliminating the reduction.

Results

We tested two types of unrelated activities (writing and logos) for generalizability, and found no differences. We therefore collapsed these conditions in our analysis and report results comparing the two no-choice breaks with the two choice activity conditions.

On average, participants chose to do the math task 66% of the time during the baseline period in the control condition, and there were no systematic changes in this level over the trials. In the no-break incentive (i.e., the replication) condition, there was a decrease in the choice of math tasks immediately after the rewards ended, relative to control (58% vs. 63%), replicating the momentary reduction in engagement ($\beta_{MOMENTARY} = -2.49, z = -2.55, p = .01$). We also observed a significant overall post-reward increase in engagement ($\beta_{POST} = +1.11, z = +2.33, p = .02$) consistent with the finding in our internal meta-analysis.

These results were moderated by providing participants a break after the incentive round. When participants were given a non-effortful break (i.e., an unrelated activity that did not require making choices), participants chose the math task 66% of the time in the immediate post-reward trial - similar to the same trial in the control condition (63%; $\beta = -0.008, t < 1$, Figure 6.1). As a result, there was no momentary reduction,

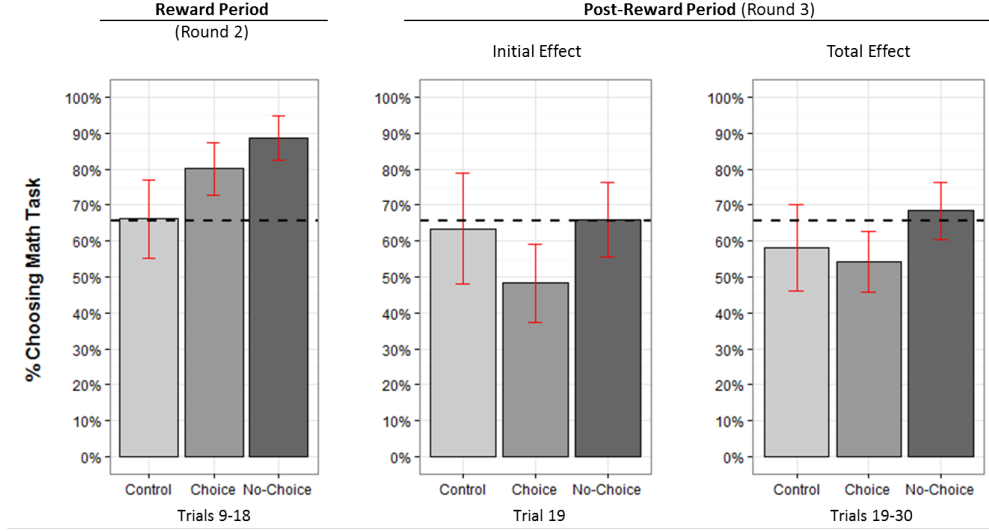


Figure 6.1: Comparison of effortful unrelated activities (choice) and low-effort break (no-choice) conditions to control in Study 2. Dotted lines represent the baseline (average effort level Round 1), and the vertical lines are 95% CIs.

relative to control ($\beta_{MOMENTARY} = -1.14, z = -1.04, p > .250$). The non-effortful break successfully arrested the post-reward engagement reduction found in the replication (no break) incentive condition ($\beta_{MOMENTARY\ interaction} = +1.60, z = +2.13, p = .03$).

These findings are specific to participants taking a break, rather than to doing any unrelated activity. When participants instead engaged in the same activities by making relatively effortful choices, the momentary reduction in engagement found in the incentive condition was not eliminated. Choices of the math task in the first trial of Round 3 were directionally lower, compared to the same trial in the control condition (48% vs. 63%, $\beta = -0.11, t = -1.35, p = .18$), reflecting a significant overall momentary reduction in engagement relative to control ($\beta_{MOMENTARY} = -2.15, z = -2.10, p = .04$). Engaging in the effortful version of the alternative activities yielded a similar post-reward reduction in engagement as in the replication incentive condition ($\beta_{MOMENTARY\ interaction} = +0.62, z = +0.82, p > .250$). Combining the replication and the effortful choices condition, we found that giving people a non-effortful break significantly arrested the immediate post-reward decrease in engagement ($\beta_{MOMENTARY\ interaction} = +1.24, z = +1.97, p = .049$). These findings are consistent

with the proposal that people reduce task engagement after the incentive in order to take a break, and are inconsistent with the autonomy-based prior theories, which would have predicted that giving people choices should be as effective or more effective at eliminating the reduction.

Aside from the initial post-incentive differences, the longer-term engagement in Round 3 was similar overall in both the no-choice break and choice activity conditions. In the choice activity conditions, engagement returned to the pre-reward baseline level (54% vs. 58%, $\beta = 0.004, t < 1; \beta_{POST} = +0.12, z = +0.26, p > .250$). Participants' overall engagement with the math task in the no-choice-break conditions was directionally higher after the incentive (68% vs. 58%; $\beta = 0.07, t = 1.46, p = .15$), reflecting a modest increase in engagement ($\beta_{POST} = +0.48, z = +1.99, p = .046$), similar to the replication incentive condition ($\beta_{POST\ interaction} = -0.44, z = -0.99, p > .250$).

Discussion

In study 2, we tested participants' need to restore balance as an explanation for the findings by providing them with a break after the incentive ended. A non-effortful unrelated activity gave participants a break, potentially restoring their perceived balance between effort and leisure, and successfully preventing an immediate post-reward decrease in task engagement. In contrast, when the same unrelated activities involved making relatively more effortful choices, the activity presumably did not offset the effort exerted during the reward period, and the post-incentive momentary reduction in engagement was observed.

These findings are difficult to reconcile with prior accounts, particularly Cognitive Evaluation Theory, which explains negative effects of incentives in terms of a reduction in perceived autonomy. It is possible that giving participants an unrelated activity could allow them to re-assert their sense of autonomy, countering the theorized effect of incentives. If that were the case, however, an unrelated activity involving more of an opportunity to express autonomy after the incentive ended (i.e., by making choices) should be more effective at countering

the momentary engagement reduction. Instead, we find the opposite, suggesting that it is the lack of required effort that makes the break effective.

In the next study, we test a prediction of the effort-balancing account about how post-incentive engagement is affected by the magnitude of the incentive. In our account, when incentives are higher but during-incentive efforts are similar, people will feel less of a need to balance out the efforts with leisure or a break, and will feel less justified in doing so. In fact, when people find the task incentive particularly rewarding, relative to effort, long-term engagement in the task may actually be bolstered by the incentive. In contrast, the prior accounts involving intrinsic motivation or inferences predict the opposite, a stronger and more persistent post-reward reduction in task engagement from higher incentives, since larger incentives would be experienced as more controlling and a stronger basis for self-perception inferences.

6.2 Study 3: Effect of Large Rewards

Method

Adult participants were recruited from Amazon MTurk. A target of 320 participants was requested, yielding 305 surveys. Unusable cases (duplicate IPs, technical problems, failing the attention check) were removed prior to analysis, yielding 235 valid completes. Participants in this sample who completed Round 1, and then dropped-out part way through (5.5%), were coded as not doing the focal task and included.

Participants were randomly assigned to one of four conditions: two conditions identical to Study 1 (a no-reward control condition and a 5-cent performance-contingent condition), and two other performance contingent-conditions with different relative reward amounts. In the low-reward condition, participants were told that they would earn 1 cent per correct answer in the second round, and that participants in a previous study were paid 5 cents per correct answer, as a basis for comparison. In the high-reward condition, participants were

told they would earn 50 cents per correct answer, and that previous participants had earned 5 cents per correct answer. At the end of the study, participants were asked a few follow-up questions about their experience, including how much they liked the math task.

Results

On average, participants in the control condition chose the math task 60% of the time during Round 1 and there were no systematic trends across trials. The incentives in all three incentive conditions were successful, increasing average choices of the math task during Round 2 compared to control (1 cent: 70%, $t = 1.89$, $p = .06$; 5 cent: 89%, $t = 4.78$, $p < .001$; 50 cent: 89%, $t = 5.51$, $p < .001$). Given that choices of the math task during Round 2 were nearly identical in the 5 cent and 50 cent conditions, participants in the 50 cent condition effectively earned more for the same effort.




Figure 6.2: Average math task choices in the control, replication (5 cents), low-reward (1 cent) and high-reward (50) conditions in Study 3. Dotted lines represent the baseline (average effort level Round 1), and the vertical lines are 95% CIs.

In the low-reward (1 cent) condition, math task choices were directionally lower in the first post-reward trial than in the same trial for the control condition (45% vs. 56%; $\beta = -0.12$, $t = -1.47$, $p = .14$) and there was significant momentary reduction in engagement

($\beta_{MOMENTARY} = -1.94, z = -2.04, p = .04$) after the incentive. However, there was no long-term reduction in engagement (52% vs. 53%; $t < 1; \beta_{POST} = -0.22, z = -0.34, p > .250$).

Similarly, in the replication condition (5 cents), the math task was only chosen 45% of the time in the first trial after the incentive ended, marginally less than in the control condition (56%; $\beta = -0.16, t = -1.93, p = .055$), reflecting significant momentary reduction in engagement ($\beta_{MOMENTARY} = -3.52, z = -3.07, p = .002$) after the incentive. As in the previous studies, engagement returned to baseline, and there was no difference in the average effort level in Round 3 between the replication and control conditions (60% vs. 53%; $\beta = 0.01, t < 1; \beta_{POST} = +0.26, z = +0.44, p > .250$).

Overall, post-reward behavior in Round 3 for the low-reward (1 cent) and replication (5 cent) conditions were very similar. Combining the replication and the low-reward conditions, we find a significant reduction in math choices in the first trial after the incentives were stopped (45% vs. 56%; $\beta = -0.14, t = -1.98, p = .049$) and overall ($\beta_{MOMENTARY} = -2.72, z = -2.98, p = .003$), compared to control. Moreover, there was no long-term reduction of effort in these combined lower-reward conditions (56% vs. 54%; $\beta = 0.007, t < 1; \beta_{POST} = +0.04, z = +0.07, p > .250$).

The key comparison in this study is between the lower-reward (1 cent and 5 cent) conditions and the highly rewarding 50 cent condition. In contrast with the momentary reduction observed in the lower reward conditions, directionally more participants chose the math task immediately after the high incentive ended, compared to the same trial in the control condition (63% vs. 56%; $\beta = 0.07, t < 1$). Participants were significantly more likely to choose the math task in the trial immediately after the incentive ended in the higher-incentive condition than in the lower-incentive conditions (63% vs. 45%; $\beta = 0.20, t = 2.64, p = .009$). Even though participants were doing more of the math task during the reward period in the higher incentive condition, their choices of the math task then increased further to a higher post-incentive baseline level ($\beta_{POST} = +1.97, z = +3.42, p < .001$), resulting in more math

tasks after the incentive in the high reward condition than in the lower reward conditions ($\beta_{POST\ interaction} = +2.25, z = +3.41, p < .001$). Therefore, the high perceived reward did not result in any momentary or persistent post-reward reduction in task engagement.

Discussion

Post-reward Baseline Engagement Level

The proportion of math tasks chosen in all of Round 3, after the incentive ended, was higher for larger incentives (Control: 53%, 1 cent: 52%, 5 cents: 60%, 50 cents: 74%, $F(233) = 4.36, p = .038$). One explanation might be that higher incentives result in more experience with the task during the incentive period, facilitating more habit formation or belief updating about the task, or practice effects. However, since the effort expended during the incentive period was identical in the 50-cent and 5-cent conditions (both 89%), this comparison is particularly instructive.

Contrary to the habit-formation or learning accounts, which would have predicted the same post-reward behavior in both conditions, there was a significantly higher post-reward baseline level in the high-reward (50 cent) condition than in the lower (5 cent) condition ($\beta_{POST\ interaction} = +1.88, z = +2.84, p = .004$). This finding is instead consistent with theories of evaluative conditioning (Razran, 1954; De Houwer et al., 2001), which suggest that highly rewarded activities become more attractive, as people’s attention is spontaneously drawn to previously rewarded activities, potentially mediated by the release of dopamine (Anderson et al., 2016).

Research in social psychology has long suggested that paying people for an interesting target task turns play into work (Lepper and Greene, 1975) because of a reduction in intrinsic motivation in the target task. Consistent with the results of our internal meta-analysis reported earlier, we do not find any evidence for this concern. Post-study self-reports suggested that participants felt that the math task was actually less like work

($M_{50\text{ cents}} = 4.32, SD = 2.57$ vs. $M_{5\text{ cents}} = 5.68, SD = 2.48$; $t(100) = 2.7, p = .007$) when they were paid a higher amount (e.g., 50 cents), compared to being paid a lower amount (e.g., 5 cents). Participants also rated working on math tasks under high rewards to be a significantly better opportunity for them compared to the 5-cent replication condition ($M_{50\text{ cents}} = 8.12, SD = 1.00$ vs. $M_{5\text{ cents}} = 6.55, SD = 2.03$; $t(100) = 4.93, p < .001$).

Distinguishing Between Effort-Balancing and Alternative Accounts

We have proposed that the momentary post-reward reduction in engagement occurs because people want to take a break and engage in leisure to balance out prior efforts, rather than because of a decline in intrinsic motivation. If incentives are seen as controlling and thereby reduce people’s intrinsic motivation, we should have found the strongest and most persistent reductions for the largest incentive. Instead, we find the exact opposite. When participants earned more (50 cents vs. 5 cents) for the same effort, we found no momentary reduction, relative to control. Instead, consistent with our account, participants earning high rewards reduced their engagement after the incentive ended back to the same level as the control condition, and then increased their engagement to a new, higher baseline level.

This study also addresses another interpretation of our findings. Doing more of a focal task (e.g. because of an incentive) could result in fatigue, depletion or satiation with the task and a desire for variety. An increase in effort during the incentive period could then directly result in reduction in engagement, rather than because of perceived balance. If this were the case, then we would expect participants in the 50-cent high-reward condition (who did significantly more of the task than in the 1-cent low-reward condition; 89% vs. 70%; $\beta = 0.18, t = 3.49, p < .001$), to also show a stronger and more persistent reduction in engagement. However, contrary to the fatigue, depletion and satiation accounts, we find the opposite. Participants in the 50 cent condition did more math tasks in Round 2 than participants in the 1 cent condition, but then showed less, rather than more, post-incentive reduction in engagement. In fact, high incentives resulted in a significantly higher post-reward baseline

relative to the low-reward condition ($\beta_{POST\ interaction} = +2.44, z = +3.33, p < .001$), which is also inconsistent with reference point or disappointment accounts of the findings.

In the studies so far, we have incentivized participants for solving math problems, a moderately effortful task, similar to tasks used in prior research. Incentives are generally deemed unnecessary to motivate people to do leisure tasks, and are rarely applied to such tasks in practice (although in-game rewards and gambling may be considered exceptions). While prior research has not tested the effects of incentives on purely leisure tasks, doing so provides a particularly strong test between our account and alternative accounts. Since leisure tasks are typically more intrinsically motivating, prior accounts would predict a stronger post-reward decrease in intrinsic motivation and therefore stronger engagement reduction when people are paid to do the leisure task (Calder and Staw, 1975; Deci et al., 1999). However, since doing the leisure task involves less effort for the same reward, people would feel less of a need to take a break after being paid for leisure, and we would predict little or no momentary reduction in engagement.

6.3 Study 4: Incentives for Leisure

Method

Adult participants were recruited from Amazon MTurk. A target of 320 participants was requested, yielding 305 surveys. Unusable cases (duplicate IP addresses, technical problems, failed attention check) were removed prior to analysis, yielding 246 valid completes. Participants who completed Round 1 but then dropped-out (2.4%) were coded as not doing the focal task and included in the analysis.

Participants were randomly assigned to one of four conditions in a 2 (Target task: Math vs. Video) x 2 (Control vs. Incentive) between-subjects design. For generalizability, we use a different cognitive math task in this study (adapted from Abeler et al. 2011) that required participants to count the number of 1s in a grid of 150 1s and 0s (see Figure 6.3).

0	0	0	0	0	1	1	1	1	0	1	0	0	0	0
0	0	1	1	0	1	0	0	0	0	1	1	1	0	0
0	0	0	0	0	0	1	0	0	1	1	0	0	1	0
1	0	0	0	0	0	0	0	1	1	0	1	1	0	1
1	1	0	0	0	1	1	0	0	0	0	0	1	0	0
1	1	0	0	0	0	0	1	0	1	0	0	0	1	0
0	0	0	1	0	0	1	1	1	0	1	0	0	1	0
0	0	0	0	0	0	1	0	1	1	0	0	1	0	1
1	0	0	0	0	0	0	0	0	1	0	0	0	0	1
1	0	0	0	0	1	1	0	0	0	0	0	1	1	0

Count the number of 1s and enter the count in the box below

Figure 6.3: Math task used in Study 4. The task required participants to count the number of 1s in the grid (of 150 numbers) in 30 seconds.

We pre-tested this math task and found that participants were willing to engage in the task without rewards, even when they had the option of a video task as an alternative choice. In the two incentive conditions, they were either paid 5 cents for correctly completing math tasks, as in the prior studies, or paid 5 cents for each video they watched and rated (1-5 stars). The two control conditions matched the two incentive conditions, highlighting the target task without any incentive.

Results

The math task used in this study was intended to be less intrinsically motivating than the video task, and this is reflected in the relatively low level of baseline attempts both overall (40%) and in particular when video watching was the focal task (30% math in Round 1; see

the dotted lines across charts in Figure 6.4). We have two different control conditions in this study, so we compare each incentive condition to the corresponding control condition.

We replicated the momentary reduction in engagement when people were paid for doing the math task. Fewer people chose the math task in the first trial of Round 3 in the incentive condition after the rewards had ended, compared to in the matching control condition (47% vs. 26%; $\beta = -0.19, t = -2.3, p = .02$; Figure 6.4). There was no long-term reduction of effort due to incentives, relative to control (35% in both). These results were further confirmed in the hierarchical regressions ($\beta_{MOMENTARY} = -4.17, z = -3.75, p < .001$; $\beta_{POST} = -0.15, z = -0.28, p > .250$).




Figure 6.4: Post-reward engagement when the math task was incentivized (left-panel) or the video task was incentivized (right-panel) in Study 4. Dotted lines represent the baseline (average effort level Round 1), and the vertical lines are 95% CIs.

When participants were paid for the leisure task instead, we did not find a reduction in engagement. The proportion of video choices in the first trial after the incentive ended was similar to the corresponding control condition (61% in both, $\beta = 0.02, t < 1$) and we found no long-term reduction of effort (72% in both). These results were further confirmed by the hierarchical regression models ($\beta_{MOMENTARY} = -0.52, z = -0.57, p > .250$; $\beta_{POST} =$

$-0.06, z = -0.10, p > .250$).

In fact, the momentary reduction in engagement observed when incentivizing the math task was completely eliminated when the videos were incentivized instead ($\beta_{MOMENTARY\ interaction} = +4.45, z = +3.62, p = .0003$). There was no difference between the two conditions in terms of the longer-term post-reward baseline level ($\beta_{POST\ interaction} = +0.36, z = +0.52, p > .250$).

Discussion

These findings provide direct support for our effort-balancing interpretation. Prior motivation-based theories would predict stronger and more persistent post-reward reduction in engagement when paying people for a more intrinsically motivating task (e.g., watching videos). Satiation and variety seeking would predict similar reductions for both incentives. In contrast, we found that incentivizing people for the leisure task yielded no reduction in engagement, and was significantly different from incentivizing the math task. Since the leisure task involves little effort, our account suggests that people feel little need to balance out the extra choices of that task after the incentive ended, resulting in no post-reward reduction.

In the studies presented thus far, people made repeated choices about what to engage in during the post-incentive period, allowing us to measure when preferences for the focal task return to (or even exceed) the original baseline level. If people instead make a single binding decision about future engagement immediately after the incentives have ended, then the temporary decrease in motivation for the target task could result in a long-term disengagement. In the next study, we test this possibility, and test a solution predicted by our account, that providing people with a break (as demonstrated in Study 2) would prevent long-term disengagement in locked-in decisions.

6.4 Study 5: Locking-in Engagement

Method

Adult participants were recruited from Amazon MTurk. A target of 250 participants was requested, yielding 235 surveys. Unusable cases (duplicate IPs, technical problems, failed attention check) were removed prior to analysis, yielding 189 valid completes. Power (73%) was slightly less than suggested by Study 1. Participants who completed Round 1 but then dropped-out afterwards (1.0%) were coded as not doing the focal task and included in analyses.

Participants were randomly assigned to one of three conditions in a between-subjects design - a control group, and two incentive condition (5 cents per correct in Round 2), one with a 90 second non-effortful break between Rounds 2 and 3 (as in Study 2), and one without. At the start of Round 3, participants in all conditions were asked to choose one of the two tasks as the only task they would do during the entire ensuing period (i.e., doing only math tasks or only watching videos). In the break condition, this decision was made after the break. All participants were told that they would not be able to change their decision during the third round.

Results

In the control condition, when participants were asked to make a binding future decision immediately after the end of Round 2, 60% of the participants choose to only do math tasks in Round 3. This is similar to the average baseline level of engagement in Round 3 in the control condition of previous studies. In contrast, in the no-break incentive condition, only 43% of the participants choose math tasks for Round 3, marginally less than that in the control condition ($\chi^2(1) = 3.6, p = .057$; Figure 6.5). This is similar to the differences observed in the first post-incentive trial in prior studies, but has more impact on the net effect because the decision is locked-in for the remainder of Round 3. This locked-in reduction in

math tasks in Round 3 counters the incentive-based increase in Round 2, and as a result the incentives in the no-break reward condition yielded no net increase in math tasks across the study ($M_{control} = 4.70$ vs. $M_{reward, nobreak} = 4.15, t(120) < 1$).

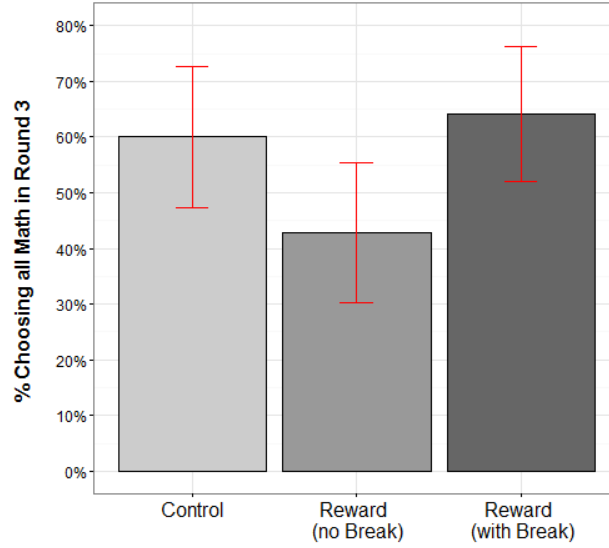


Figure 6.5: Percentage of participants locking-in math tasks for the post-reward period in Study 5. The vertical lines are 95% CIs.

However, participants in the with-break incentive condition, who instead faced the same choice after a 90 second post-incentive break, did not lock-in a reduction and, in fact, chose very similarly to the control condition participants, with 64% opting for the math task in Round 3 ($\chi^2(1) = 0.22, p = .6$). The effort-balancing account predicts that adding a brief break immediately after the rewards ended will help people restore their sense of balance between work and leisure after exerting more effort for the incentive. Consistent with this prediction, people who had a break after the incentive round were willing to commit to doing the math task in Round 3 as much as people in the control condition and significantly more than those in the no-break rewards condition ($\chi^2(1) = 5.7, p = .02$). In fact, simply adding a break resulted in a marginally higher net effect of incentives, compared to rewards without a break ($M_{reward, withbreak} = 5.52$ vs. $M_{reward, nobreak} = 4.15, t(130) = 1.9, p = .054$). Therefore, a brief break immediately after the end of the reward period successfully precluded

a lock-in of the momentary reluctance to do more of the math task after the incentive ended.

Discussion

In everyday life, people often make decisions that are effectively “locked in” (e.g., enrolling in an automatic savings plan, buying a subscription or renewing a trial gym membership), because of the cost or inconvenience of revisiting and changing the decision. Promotions and incentives are often used to increase awareness and generate initial experience, before having people make the “sticky” lock-in decision. Our findings suggest that this strategy may not work well, because of the temporary reduction in interest after the incentive ends. Our results suggest that if the lock-in decision is instead structured to occur after a break, the likelihood of choosing long-term engagement is likely to be higher. This practical implication for how to structure the decision environment to maximize the effectiveness of temporary incentives arises directly from our effort-balancing account, and is not predicted by any of the prior accounts.

CHAPTER 7

ROBUSTNESS OF EFFORT-BALANCING ACCOUNT

In Chapter 6, we examined the key predictions of the Effort-Balancing account. In this chapter, we report two studies that examine the robustness of the proposed psychological account of post-reward behavior. In the studies reported so far, the math task was framed as important and potentially beneficial, in part to highlight the self-control trade-off between goals with more immediate and delayed benefits. This raises the possibility, however, that the framing made participants feel obligated to work on the math task, rather than watch the videos, even after the reward ended. Could this have resulted in a short-lived reduction in engagement? In Study 6 we investigate this potential concern.

When participants in the reward-group were incentivized for the target task, two things happened during the reward-period that were different between the reward group and the no-reward, control group. First, the math task was associated with rewards for the reward-group, but not for the control group. Second, the reward-group participants responded to incentives and exerted more effort than the control-group participants when the rewards were available. Our proposed Effort-Balancing account suggests that the post-reward behavior is driven by a desire to take a ‘break’ after exerting extra-effort, and *not* by the mere association of rewards with a particular task as is predicted by prior theories. Therefore, our account predicts that keeping the effort fixed between the control and reward group, might attenuate the post-reward momentary reduction in task engagement. We examine this in Study 7.

An important question is whether the desire to take a ‘break’ from effort after incentives end, extends to all effortful tasks or if this desire is specific to the target task (and potentially to tasks very similar to the target task). In Study 8 we investigate this question by changing the type of effortful task people encounter after the incentives end.

7.1 Study 6: Investigating Signaling and Desirability

Method

Adult participants were recruited from Amazon MTurk to complete an online survey. A target of 300 participants were requested, yielding 291 surveys. Records with duplicate IP addresses, or who reported having technical problems with viewing the videos or working on the math task, or who failed the basic attention check were removed prior to analysis, yielding 219 valid completes. The proportion of participants included in this sample who reached until the end of Round 1, but then dropped-out part way through was 2.3%.

Participants were randomly assigned to one of four conditions, in a 2 (Control, Reward) x 2 (Math Important, Both Important) between-subjects design. The two replication conditions (Math-Important control and reward) were similar to Study 1. The other two conditions (Both-Important control and reward) were the same, except that participants were told that their data, both from doing math and from the video task was important in the study. Participants were also told that, since the survey was being administered to many people, it was completely up to them to choose what they wanted to do. This framing was designed to remove any signal to the participants that were expected to do the math tasks, and to encourage participants to choose what they truly wanted to do in each round. As a result, if participants' sense of obligation to do the math task had arrested the post-reward reduction in engagement in our studies, we would observe a stronger (and potentially persistent) decrease in engagement in the Both-Important condition.

Results

A manipulation check, collected at the end of the study, confirmed that participants in the Both-Important control condition expressed more agreement that the videos and math task were equally important (on a 9 point scale) than in the Math-Important control condition ($M_{control,both} = 4.43, SD = 1.61$ vs. $M_{control,math} = 3.79, SD = 1.75$; $t(110) = 1.98, p =$

.05).

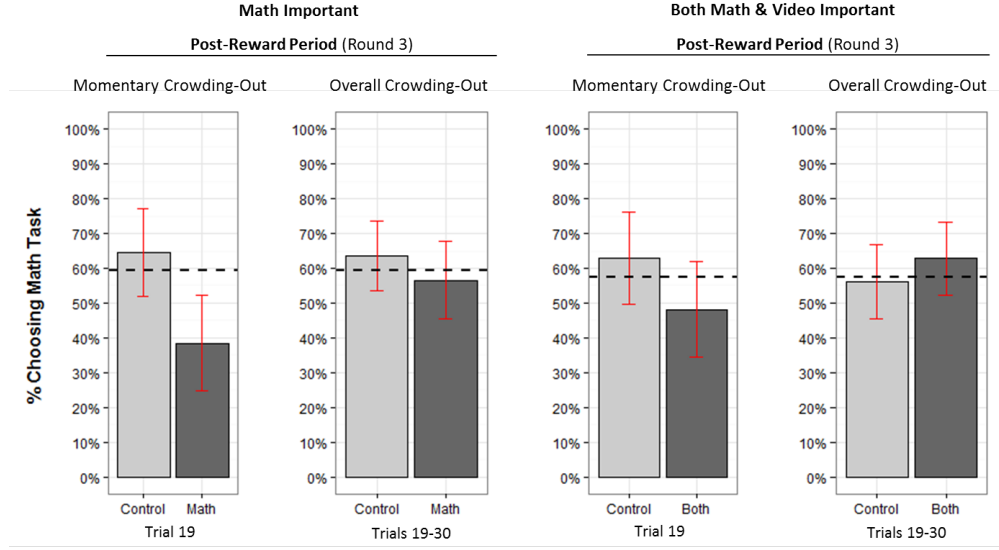


Figure 7.1: The post-reward results in the replication condition and when both math and video were said to be equally important in Study 6. Dotted lines represent the average effort level across the conditions in the baseline (i.e., pre-incentive) period and the vertical bars are 95% CIs.

Since this study included two differently-framed control conditions, we compared each reward condition to the corresponding control condition. We replicated the momentary reduction in engagement behavior when using the same instructions, in the Math-Important incentive and control conditions. Fewer people chose the math task in the first trial of Round 3 in the reward condition after the incentives had ended, compared with the same trial in the control condition (38% vs. 64%; $\beta = -0.18, t = -2.13, p = .04$; Figure 7.1). There was no longer-term reduction in engagement due to incentives, relative to control (57% vs. 63%; $\beta = -0.03, t < 1$). These results were further confirmed in the hierarchical regression models ($\beta_{MOMENTARY} = -4.69, z = -4.00, p < .001$; $\beta_{POST} = -0.04, z = -0.07, p > .250$).

Likewise, when we instead told participants that both the task options (math and video) are equally important, we again replicate the findings in the Both-Importance incentive and control conditions. Fewer people chose the math task in the first trial of Round 3 in the incentive condition after the incentives had ended, compared to in the control condition (48% vs. 63%; $\beta = -0.18, t = -2.42, p = .02$). There was no longer-term reduction in engagement

due to incentives, relative to control (63% vs. 57%; $\beta = 0.03, t < 1$). These results were further confirmed in the hierarchical regression ($\beta_{MOMENTARY} = -2.52, z = -3.33, p < .001; \beta_{POST} = +0.25, z = +0.51, p > .250$).

A hierarchical regression model further confirmed that there was no difference in the extent of momentary reduction in engagement between the two task-framing conditions ($\beta_{MOMENTARY\ interaction} = +1.47, z = +1.08, p > .250$). Likewise, there was no difference in the extent of overall post-reward reduction between the two conditions ($\beta_{POST\ interaction} = +0.27, z = +0.38, p > .250$).

Discussion

The results suggest that the momentary nature of post-reward reduction in engagement cannot be explained by the experimental instructions inducing a feeling of obligation to do math tasks among the participants. Furthermore, the findings are also inconsistent with a self-signaling account, in which participants might continue with the more challenging math tasks (after a short break) to feel good about themselves.

In the next study we separate the effect of incentives from additional effort expended by the reward group during Round 2, on Round 3 behavior. Prior theories predict that the mere association of rewards with a task would have detrimental post-reward consequence, whereas our Effort-Balancing account predicts that mere association of rewards with a task is not harmful. If people invest extra effort under modest rewards, the desire to offset the effort expended with leisure might manifest as a post-reward decrease in engagement. Therefore, our account predicts that if effort is held constant between a group that gets incentives and a group that does not, both the groups will behave in a similar fashion in the immediate and longer-term post-reward period.

7.2 Study 7: Isolating Effort from Rewards

Method

Adult participants were recruited from Amazon MTurk to complete an online survey. A target of 550 participants were requested, yielding 532 surveys. Records with duplicate IP addresses, or who reported having technical problems with viewing the videos or working on the math task, or who failed the basic attention check were removed prior to analysis, yielding 437 valid completes. The proportion of participants included in this sample who reached until the end of Round 1, but then dropped-out part way through was 4.3%.

Participants were randomly assigned to a reward and control group like in previous studies. Before the start of Round 2, participants in both groups were told that the computer will choose whether they will do a math task or a watch a video in Round 2. As a consequence, in this study, participants did not get to *choose* whether they would want to try a math task or watch a video in each trial in Round 2. This made it possible for us to manipulate the effort participants in each of the experimental groups exerted during the round. Specifically, participants in both the control and the reward group were assigned to the same level of effort, and therefore it was possible to isolate the effect of merely associating rewards with an intrinsically motivating task.

In order to further test the prediction of our Effort-Balancing account, we varied in the effort level in the study. Participants in both the control and the reward group were required to either exert low effort (computer randomly choose 6 out of 10 math tasks), or both groups were required to exert high effort (computer randomly choose 9 out of 10 math tasks) in Round 2. The low and high effort level were chosen based on the average effort exerted by the reward and the control group (respectively) during Round 2 in our internal meta-analysis. Thus, in effect, this study employed a 2 (Control, Reward) x 2 (Low Required Effort, High Required Effort) between-subjects design. Our Effort-Balancing account predicted a main effect of effort on immediate post-reward behavior in both the experimental groups, but no

separate effect of incentives on the reward group in post-reward behavior.

Results

When the control and reward groups were required to exert low effort in Round 2, both groups choose to do the math task at the same level immediately after Round 2 (control: 57%; reward: 51%; $\beta = -0.05, t = -0.94, p = .35$; Figure 7.2). Likewise, even when both groups were required to exert high effort, they choose to do the math task at the same level immediately after Round 2 (control: 42%; reward: 41%; $\beta = -0.003, t = -0.06, p = .95$). The difference in immediate post-reward effort level between reward and control were not different in the high and low effort level conditions ($\beta_{interaction} = 0.05, t = .63, p = .53$). Therefore, when reward and control group were required to exogenously exert the same level of effort, associating a task with incentives, did not change the behavior of the reward group with respect to a contemporaneous control in the immediate post-reward period.

Comparing the immediate post-reward engagement across the experimental groups with the average engagement in Round 1 (i.e., baseline shown with a dotted line in Figure 7.2), paired t-tests reveal a significant reduction in the choice of math tasks in both the effort level conditions (low effort: $t(210) = -3.8, p < .001$; high effort: $t(220) = -7.4, p < .001$). These results indicate a momentary reduction in engagement for both reward and control in the immediate post-reward period, when these groups were required to exert similar levels of effort. Furthermore, there was a main effect of effort on immediate post-reward behavior. Across the experimental groups, participants in the high-effort condition choose to do the math task significantly less than participants in the low-effort condition, immediately after the end of Round 2 (High: 42% vs. Low: 54%, $t(435) = 2.5, p = .01$).

After the immediate post-reward period, effort in both the experimental groups settled at the same level in the longer-term for both the effort level conditions (low effort: 59% vs. 58%, $\beta = 0.02, t = 0.46, p = .65$; high effort: 52% vs. 56%, $\beta = -0.03, t = -0.86, p = .39$). Across the experimental conditions, the longer-term post reward level




Figure 7.2: The post-reward results in low required effort and the high required effort conditions in Study 7. Dotted lines represent the average effort level across the conditions in the baseline (i.e., pre-incentive) period and the vertical bars are 95% CIs.

was not different between the two effort level conditions (High: 54% vs. Low: 58%, $t(435) = 1.1, p = .3$).

The results were confirmed in the hierarchical regression models. Compared to contemporaneous control, there was no momentary reduction in engagement when required effort was low ($\beta_{MOMENTARY} = -1.05, z = -1.61, p = .11$), as well as when it was high ($\beta_{MOMENTARY} = 0.38, z = 0.79, p = .43$). However, across the experimental conditions, the momentary reduction in engagement was greater when the effort required was higher than when it was lower ($\beta_{MOMENTARY\ high\ vs\ low} = -1.01, z = -2.58, p = .009$).

The choice of math tasks between in the control and the reward group eventually returned to the same level both in the low effort condition ($\beta_{POST} = 0.13, z = 0.35, p > .250$) as well as in the high effort condition ($\beta_{POST} = -0.34, z = -1.06, p > .250$). Across the experimental conditions, there was no difference in the eventual post-reward effort level between the two different effort conditions ($\beta_{POST\ high\ vs\ low} = -0.31, z = -1.29, p = .20$).

Discussion

Prior theories based on intrinsic motivation suggest that the mere association of extrinsic motivators with an intrinsically motivating task would reduce intrinsic motivation. This would manifest as a reduction in engagement in the post-reward period. In contrast, the proposed Effort-Balancing account suggests a momentary post-reward reduction in effort only when people exert additional effort during the reward period and, as a consequence, do not attend to other goals like pursuing leisure. The results of Study 7 support our prediction. Participants in the reward group showed the same level of immediate post-reward engagement as the control group that was also required to exert the same level of effort. However, across experimental groups, participants who were required to exert more effort showed a greater reduction in immediate post-reward engagement compared to those who were required to exert less effort.

Prior theories based also predict a salutary effect competence feedback on intrinsic motivation. In all our studies so far, only participants in the reward group received feedback about their performance at the end of Round 2. Could this have mitigated the post-reward reduction in engagement, making it only momentary in our experiments? If this was the case then in this particular study, the reward group having received competence feedback (and having exerted the same level of effort as the control group), should have returned to the initial baseline level (i.e., average Round 1 level) after Round 2, and in effect, should have shown an increase in immediate post-reward engagement relative to the control group. The absence of this empirical finding suggests that it was easy for people to figure out themselves if they got the math task correct, and additional feedback was potentially redundant and, therefore, not useful in our studies.

In this study, control group participants exerted the same level of effort as the reward group participants but without any rewards. It seems possible that participants in the control group would be in a state of greater perceived imbalance and therefore should show stronger reduction in engagement immediately after Round 2 compared to the reward group.

However, we did not find this in the study. We suspect that since nothing was mentioned about extra monetary compensation for correctly solving the math task to these participants, they were not even thinking about effort-reward balancing. Had we mentioned to this group that they are not going to earn anything for attempting and getting the math tasks right, they might have showed a greater decrease in post-reward engagement immediately after the round with required effort was over.

In the next study we examine the scope and extent of effort-balancing and attempt to refine our proposed account. Specifically, if effort-balancing extends to all effortful tasks, it would imply people taking a ‘break’ from any effortful task they encounter after a temporary reward ends. In contrast, if effort balancing extends to only the same effortful task that was incentivized (or very similar such tasks), it would imply people taking a ‘break’ from only those kinds of effortful tasks after temporary rewards end. In order to test this, we incentivize a reward group for math tasks, but introduce choices between watching videos and doing verbal spelling tasks after the end of Round 2.

7.3 Study 8: Scope of Effort-Balancing

Method

Adult participants were recruited from Amazon MTurk to complete an online survey. A target of 250 participants were requested, yielding 240 surveys. Records with duplicate IP addresses, or who reported having technical problems with viewing the videos or working on the math task, or who failed the basic attention check were removed prior to analysis, yielding 197 valid completes. The proportion of participants included in this sample who reached until the end of Round 1, but then dropped-out part way through was 4.6%.

At the beginning the study, participants were told that they would be asked to do various tasks in the study like math tasks or spelling tasks. A sample of each task-type was provided to the participants. Round 1 and Round 2 of the study were similar to our experi-

mental paradigm discussion earlier. Participants were randomly assigned to three between-subject conditions: control (no-reward), replication (5 cents for every correct math answer in Round 2), and high-reward (50 cents for every correct math answer in Round 2 with similar manipulation as in Study 2 to make reward magnitude salient). At the end of Round 2, participants in all experimental groups were told that they would make choices between doing spelling tasks and watching videos in the next round. Of course, like in previous studies participants were not told if the next ensuing round is the final round of the study.

A pre-test ensured that our study population perceived both math and spelling tasks as equally effortful, compared to the videos ($M_{math} = 5.08$ vs. $M_{video} = 2.48, paired t(24) = 4.5, p < .001$; $M_{spelling} = 4.83$ vs. $M_{video} = 2.29, paired t(23) = 5.4, p < .001$ 1=low; 9=high). The key question of interest in this study was how would reward group participants behave in Round 3, (compared to control) after they were incentivized for math tasks and then made to encounter a similar effort-leisure trade-off but involving a different effortful task type.

Results

The random group of participants who were paid 5 cents or 50 cents for every correct math task invested significantly more effort compared to the control (5 cents: 86% vs. 63%; $t(130) = 3.9, p < .001$; 50 cents: 93% vs. 63%; $t(130) = 5.6, p < .001$), although as in Study 3, there was no difference between the effort exerted by the 5 cents group and the 50 cents group (86% vs. 93%; $t(130) = 1.5, p = .1$) when incentives were available.

Immediately after the rewards ended, and participants encountered a different type of effortful (e.g., a verbal spelling task) task. As seen in Figure 7.3, participants' effort in both the reward conditions were at the level of the contemporaneous control (5 cents: 63% vs. 65%; $\beta = -0.004, t = -0.05, p = .95$; 50 cents: 70% vs. 65%; $\beta = 0.02, t = 0.80, p = .42$) suggesting no momentary reduction in engagement immediately after the rewards ended. These results suggest that people think of effort-balancing quite narrowly with respect to




Figure 7.3: Choice of the effortful task in Study 8. Round 3 employed an effortful spelling task for all the experimental groups. Dotted lines represent the average effort level across the conditions in the baseline (i.e., pre-incentive) period and the vertical bars are 95% CIs.

the task being incentivized (or, potentially similar types of tasks) and not as trade-offs between leisure and any types of effortful tasks. The results were confirmed in the hierarchical regression models. Compared to contemporaneous control, there was no momentary reduction in engagement in the 5 cents condition ($\beta_{MOMENTARY} = -0.10, z = -1.23, p = .11$), as well as in the 50 cents condition ($\beta_{MOMENTARY} = 0.82, z = 1.06, p = .28$).

Finally, as found in previous studies, all the experimental groups settled at the same level of engagement in the longer run (Control: 46%, 5 cents: 41%, and 15 cents: 49%; 5 cents vs. control: $\beta = -0.03, t = -0.55, p = .58$; 50 cents vs. control: $\beta = 0.01, t = 0.76, p = .4$). Importantly, the group that received a high relative reward did not show a longer-term increase in engagement as found in Study 3. This further highlights the importance of incentives affecting task-specific evaluative associations, which can then be a positive driver of post-reward engagement.

Discussion

The results of Study 8 help to refine our proposed Effort-Balancing account. When people are temporarily given modest incentives for tasks that require immediate effort, people might

seek a ‘break’ from that specific task (or from potentially very similar type of tasks), but this desire does not extend to all types of tasks that require immediate effort. This finding has important theoretical and practical import.

Theoretically, it implies that decision-makers are thinking narrowly about ‘effort’ when they are trying to balance between effort and leisure. Practically, it further implies that temporary incentives for interesting tasks might be less harmful than was perceived based on the predictions from prior theories. Not only do people disengage from the incentivized task only momentarily (vs. persistently), such disengagements do not indiscriminately extend to all effortful tasks with delayed benefits. For example, a cafe manager serving a captive audience and using temporary promotions for , say, salads might worry about a temporary decrease in salad sales after the promotion ends, but might not need to worry about a decrease in transactions provided the cafe has other types of food offerings.

The findings in the high reward condition is also significant. If participants in our study wanted to merely reciprocate to our relatively generous high rewards, we should have seen a post-reward increase in engagement with the spelling task as well. But, the absence of such a finding indicates that the way contingent rewards color associative evaluations of a specific task is important. Therefore, the salutary effect of a ‘rewarding’ experience during the incentive period might spill-over only to the specific target task (or potentially to similar types of tasks). This would result in a new, higher post-reward engagement level as we found in Study 3 reported earlier.

CHAPTER 8

GENERAL DISCUSSION

Effect of Temporary Incentives on Post-Reward Behavior

Temporary incentives can be very effective in motivating beneficial behavior while people are being rewarded. However, a large and influential literature has warned against the use of such policies because of the possibility that external incentives will undermine intrinsic motivation, resulting in a persistent reduction in engagement after the rewards end (Deci et al., 1999). However, these conclusions with adults were based on observing a single initial behavior, immediately after the incentive ended. Noting this shared limitation of the extant studies with adults, the authors pointed out the need for "studies of interesting [e.g., intrinsically motivating] behaviors that examine repeated administration of rewards over time, have appropriate no-reward control groups, and use reasonable sample sizes" (Deci et al., 1999, p. 650). In this project, we have attempted to fill this crucial and long-standing gap in the literature.

Using a novel experimental framework to track choice-by-choice engagement in intrinsically motivating tasks, we do find a momentary reduction in task engagement immediately after the task incentive ends, consistent with prior findings. However, consistent with recent field studies of incentives, we do not find the predicted long-term negative effects after the incentive ends. Instead, we find that the immediate post-incentive reduction is brief, people then return to or even exceed baseline levels of effort, and as a result, there is a strong positive net effect of incentives. Thus, our results reconcile the findings of prior lab studies, which measured initial post-incentive engagement, and field studies, which measured long-term engagement. Across the studies, our findings cannot be explained by incentives reducing people's intrinsic motivation (Studies 1 internal meta-analysis, 3 & 4), rewards undermining people's autonomy (Study 2, Study 7), fatigue, reference points, or variety seeking (Study 3), and desire to confirm to some expected behavior (Study 6).

In particular, our results are incompatible with existing accounts of post-incentive intrinsic motivation. A temporary reduction in engagement contradicts the theoretical premise that intrinsic motivation in the target task is reduced by the incentive. A reduction in intrinsic motivation for the target task would have instead resulted in a persistent decrease in engagement after the incentive ended, since there was no change in information in the decision environment. As noted in our internal meta-analysis, self-reported measures of task interest at the end of our studies were not lower in the incentive conditions, as would be expected if intrinsic motivation had been eliminated. Instead, self-reported task interest was actually higher in the incentive conditions, reflecting a reinforcing effect of incentives. Notably, this lack of evidence for a decline in intrinsic motivation is not unique to our data. As reported in the meta-analysis (Deci et al., 1999) self-reported measures also did not demonstrate a significant decrease in task-interest when using performance-contingent incentives. The findings of Study 7 further support this assertion by showing that mere association of rewards with an intrinsically motivating task does not have detrimental post-reward consequence.

An Effort-Balancing Account of Dynamic Post-Reward Behavior

In this paper, we outlined an Effort-Balancing account as a new way of thinking about the effects of incentives on post-reward effort over time. We propose that increased task effort, induced by incentives, can result in a sense of imbalance between effort and leisure, resulting in a desire to take a ‘break’ from the incentivized task when the incentive period ends. As a result, post-reward task engagement can decrease below the baseline level momentarily, but will return to the baseline level when balance has been restored. When people are given a non-effortful break after the incentive ends (Studies 2 and 5), their need for a break will be satisfied and the post-reward reduction in engagement can be arrested. The idea is similar to recent research suggesting that depletion-like regulatory failures can represent a motivated switching between labor and leisure to strike a balance between two goals, without necessarily reflecting exhaustion of limited resources (Inzlicht et al., 2014; Inzlicht

and Schmeichel, 2012).

Our account makes testable predictions about how characteristics of the available incentive will affect post-incentive engagement. The larger the incentive, holding effort constant, the less people will need to balance out the effort after the incentive ends, and the less people will feel justified in taking a break. Consistent with this prediction, we found that a relatively high incentive eliminated the post-reward momentary reduction in engagement, generating an overall post-incentive increase in engagement relative to control (Study 3). Likewise, when a leisure task was incentivized, we predicted that the relatively low required effort during the incentive period would result in little post-reward reduction in engagement, which was confirmed in Study 4. Notably, prior accounts based on intrinsic motivation would have made the opposite prediction in both cases, stronger and more persistent post-reward reductions in task engagement after larger incentives or after incentivizing a more intrinsically motivating leisure task.

In our account, the reduction in engagement is brief. Building on prior research about positive reinforcement (Anderson et al., 2016; De Houwer et al., 2001; Brooks and Bouton, 1993; Razran, 1954), our account further suggests that people’s longer term engagement will either return to baseline (if the rewards are small) or will be higher if the incentive is larger. We find evidence of this in our internal meta-analysis, as well as more specifically in Study 3, where we vary the magnitude of the incentive. This prediction of our account is also consistent with a comparison across recent field studies. Those studies that provided larger incentives tended to find a positive longer-term increase in engagement (Kane et al., 2004; Volpp et al., 2008; Cawley and Price, 2009; Charness and Gneezy, 2009; Halpern et al., 2015), whereas studies in similar domains with smaller incentives (John et al., 2011; Volpp et al., 2006) tended to find no such positive longer-term behavior (see Appendix G for a summary of past field studies that that looked at post-reward behavior). The importance of positive task-specific associations as a driver of post-reward behavior was further validated in Study 8 where we found changing the post-reward task type does not yield a higher eventual

baseline level.

Implications for Future Research

Our initial findings and the proposed effort-balancing interpretation of these findings relies on data from a single paradigm. Future research should test the generality of our findings, including in the field and in settings not involving a potentially relevant participation fee, as well as over longer periods of time.

It could also be useful to test whether these findings extend to other types of behavior which may be incentivized. Incentives are often used to encourage people to make purchases, for example, and it is not clear if the momentary reduction in engagement after a financial incentive would extend to situations where the effort itself takes the form of spending money. It would also be useful to study the effect of temporary incentives on pro-social behavior. Past research has suggested that people might be less responsive to incentives for such activities (Dickinson, 1989; Benabou and Tirole, 2003; Ariely et al., 2009), and that introducing monetary rewards can undermine the effectiveness of pro-social incentives (Dubé et al., 2015; Heyman and Ariely, 2004; Yang et al., 2014). However, this research has focused on behavior while the incentive was available, and it would be useful to investigate the effects of incentives on dynamic longer-term post-reward pro-social behavior.

Our findings also suggest that we need a better understanding of the contextual factors that can bolster or undermine the effectiveness of incentives post-reward. We found that higher incentives and less effortful tasks yield more post-incentive engagement. Future research may identify low-cost psychological interventions to generate a rewarding experience similar to incentivized effort. For example, highlighting the fun aspects of an effortful but beneficial task or highlighting the enjoyable aspects of temporarily working in groups may make an experience feel more rewarding, resulting in less initial post-incentive reduction and longer-term positive effects. Conversely, building on our finding that even momentary reductions can have substantial negative effects if the initial decisions are locked-in, future

research may identify interventions, beyond providing a break, to counter the initial lock-in effect.

Re-opening the Door to Incentives

Incentives are a cornerstone of economics, and featured prominently in foundational theories developed in the early days of psychology as well. Rewards were considered to be a powerful reinforcer of desirable behavior (Skinner, 1953), an important determinant of motivational force in expectancy-valence theory (Vroom, 1964; Fishbein, 1967), and a means of creating and strengthening expectations of personal efficacy (Bandura, 1977). Despite the reinforcing nature of rewards, incentives have come to be viewed as counter-productive and either non-psychological or subject to psychological backlash. Our results suggest that these concerns may have been over-stated, and that the ways in which people respond to incentives over time involves different psychological mechanisms than previously thought. Our key insight in this project is that post-reward engagement reduction may often have more to do with people wanting a break after investing effort in their work, than with having their enthusiasm for work “smothered” by incentives.

Taking a dynamic perspective on the effects of incentives in this research allows us to uncover the momentary nature of post-reward reductions in engagement, and raises new questions about how incentives and motivation interact over time. What are the psychological factors that extend or inhibit our observed momentary reduction in engagement? What kinds of psychological interventions could be leveraged to make incentives more effective and motivating in the long-term? We believe that our findings should not only re-open the door to testing the use of temporary incentives in consequential domains like health and education, but also to a new re-investigation of this most fundamental psychological driver of human motivation.

REFERENCES

- Ainslie, G. (1975), ‘Specious reward: a behavioral theory of impulsiveness and impulse control.’, *Psychological bulletin* **82**(4), 463.
- Anderson, B. A., Kuwabara, H., Wong, D. F., Gean, E. G., Rahmim, A., Brašić, J. R., George, N., Frolov, B., Courtney, S. M. and Yantis, S. (2016), ‘The role of dopamine in value-based attentional orienting’, *Current Biology* **26**(4), 550–555.
- Ariely, D., Bracha, A. and Meier, S. (2009), ‘Doing good or doing well? image motivation and monetary incentives in behaving prosocially’, *The American Economic Review* pp. 544–555.
- Bandura, A. (1977), ‘Self-efficacy: toward a unifying theory of behavioral change.’, *Psychological review* **84**(2), 191.
- Bazerman, M. H., Tenbrunsel, A. E. and Wade-Benzoni, K. (1998), ‘Negotiating with yourself and losing: Making decisions with competing internal preferences’, *Academy of Management Review* **23**(2), 225–241.
- Bem, D. J. (1972), ‘Self-perception theory’, *Advances in experimental social psychology* **6**, 1–62.
- Benabou, R. and Tirole, J. (2003), ‘Intrinsic and extrinsic motivation’, *The Review of Economic Studies* **70**(3), 489–520.
- Brooks, D. C. and Bouton, M. E. (1993), ‘A retrieval cue for extinction attenuates spontaneous recovery.’, *Journal of Experimental Psychology: Animal Behavior Processes* **19**(1), 77.
- Brown, J. S. (1948), ‘Gradients of approach and avoidance responses and their relation to level of motivation.’, *Journal of comparative and physiological psychology* **41**(6), 450.
- Calder, B. J. and Staw, B. M. (1975), ‘Self-perception of intrinsic and extrinsic motivation.’, *Journal of personality and social psychology* **31**(4), 599.
- Cawley, J. and Price, J. A. (2009), Outcomes in a program that offers financial rewards for weight loss, Technical report, National Bureau of Economic Research.
- Charness, G. and Gneezy, U. (2009), ‘Incentives to exercise’, *Econometrica* **77**(3), 909–931.
- De Houwer, J., Thomas, S. and Baeyens, F. (2001), ‘Association learning of likes and dislikes: A review of 25 years of research on human evaluative conditioning.’, *Psychological bulletin* **127**(6), 853.
- Deci, E. L. (1971), ‘Effects of externally mediated rewards on intrinsic motivation.’, *Journal of personality and Social Psychology* **18**(1), 105.
- Deci, E. L., Koestner, R. and 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., Koestner, R. and Ryan, R. M. (2001), ‘Extrinsic rewards and intrinsic motivation in education: Reconsidered once again’, *Review of educational research* **71**(1), 1–27.
- Deci, E. L. and Ryan, R. M. (1985), ‘The general causality orientations scale: Self-determination in personality’, *Journal of research in personality* **19**(2), 109–134.
- DelVecchio, D., Henard, D. H. and Freling, T. H. (2006), ‘The effect of sales promotion on post-promotion brand preference: A meta-analysis’, *Journal of retailing* **82**(3), 203–213.
- Dhar, R. and Simonson, I. (1999), ‘Making complementary choices in consumption episodes: Highlighting versus balancing’, *Journal of Marketing Research* pp. 29–44.
- Dickinson, A. M. (1989), ‘The detrimental effects of extrinsic reinforcement on intrinsic motivation’, *The Behavior Analyst* **12**(1), 1–15.
- Dubé, J.-P., Luo, X. and Fang, Z. (2015), Self-signaling and prosocial behavior: a cause marketing mobile field experiment, Technical report, National Bureau of Economic Research.
- Esteves-Sorensen, C., Macera, R. and Broce, R. (2013), Do monetary incentives crowd out intrinsic motivation? a field test in the workplace, Technical report, mimeo.
- Fehr, E. and Falk, A. (2002), ‘Psychological foundations of incentives’, *European Economic Review* **46**(4), 687–724.
- Fishbach, A. and Dhar, R. (2005), ‘Goals as excuses or guides: The liberating effect of perceived goal progress on choice’, *Journal of Consumer Research* **32**(3), 370–377.
- Fishbein, M. E. (1967), ‘Readings in attitude theory and measurement.’.
- Förster, J., Liberman, N. and Friedman, R. S. (2007), ‘Seven principles of goal activation: A systematic approach to distinguishing goal priming from priming of non-goal constructs’, *Personality and Social Psychology Review* **11**(3), 211–233.
- Förster, J., Liberman, N. and Higgins, E. T. (2005), ‘Accessibility from active and fulfilled goals’, *Journal of Experimental Social Psychology* **41**(3), 220–239.
- Frey, B. S. and Oberholzer-Gee, F. (1997), ‘The cost of price incentives: An empirical analysis of motivation crowding-out’, *The American economic review* **87**(4), 746–755.
- Gilbert, D. T., Pinel, E. C., Wilson, T. D., Blumberg, S. J. and Wheatley, T. P. (1998), ‘Immune neglect: a source of durability bias in affective forecasting.’, *Journal of personality and social psychology* **75**(3), 617.
- Gneezy, U., Meier, S. and Rey-Biel, P. (2011), ‘When and why incentives (don’t) work to modify behavior’, *The Journal of Economic Perspectives* pp. 191–209.

- Gneezy, U. and Rustichini, A. (2000), ‘Fine is a price, a’, *J. Legal Stud.* **29**, 1.
- Greene, D. and Lepper, M. R. (1978), *The Hidden Costs of Reward: New Perspectives of the Psychology of Human Motivation*, Hillsdale, NJ: L. Erlbaum Associates; New York.
- Halpern, S. D., French, B., Small, D. S., Saulsgiver, K., Harhay, M. O., Audrain-McGovern, J., Loewenstein, G., Brennan, T. A., Asch, D. A. and Volpp, K. G. (2015), ‘Randomized trial of four financial-incentive programs for smoking cessation’, *New England Journal of Medicine* **372**(22), 2108–2117.
- Heyman, J. and Ariely, D. (2004), ‘Effort for payment a tale of two markets’, *Psychological science* **15**(11), 787–793.
- Inzlicht, M. and Schmeichel, B. J. (2012), ‘What is ego depletion? toward a mechanistic revision of the resource model of self-control’, *Perspectives on Psychological Science* **7**(5), 450–463.
- Inzlicht, M., Schmeichel, B. J. and Macrae, C. N. (2014), ‘Why self-control seems (but may not be) limited’, *Trends in cognitive sciences* **18**(3), 127–133.
- Jackson, C. K. (2010), ‘A little now for a lot later a look at a texas advanced placement incentive program’, *Journal of Human Resources* **45**(3), 591–639.
- Jhang, J. H. and Lynch, J. G. (2015), ‘Pardon the interruption: Goal proximity, perceived spare time, and impatience’, *Journal of Consumer Research* **41**(5), 1267–1283.
- John, L. K., Loewenstein, G., Troxel, A. B., Norton, L., Fassbender, J. E. and Volpp, K. G. (2011), ‘Financial incentives for extended weight loss: a randomized, controlled trial’, *Journal of general internal medicine* **26**(6), 621–626.
- Kalyanaram, G. and Winer, R. S. (1995), ‘Empirical generalizations from reference price research’, *Marketing science* **14**(3-supplement), G161–G169.
- Kane, R. L., Johnson, P. E., Town, R. J. and Butler, M. (2004), ‘A structured review of the effect of economic incentives on consumers’ preventive behavior’, *American journal of preventive medicine* **27**(4), 327–352.
- Kelley, H. H. (1973), ‘The processes of causal attribution.’, *American psychologist* **28**(2), 107.
- Khan, U. and Dhar, R. (2006), ‘Licensing effect in consumer choice’, *Journal of Marketing Research* **43**(2), 259–266.
- Kivetz, R. (2003), ‘The effects of effort and intrinsic motivation on risky choice’, *Marketing Science* **22**(4), 477–502.
- Kivetz, R., Urminsky, O. and Zheng, Y. (2006), ‘The goal-gradient hypothesis resurrected: Purchase acceleration, illusionary goal progress, and customer retention’, *Journal of Marketing Research* **43**(1), 39–58.

- Kivetz, R. and Zheng, Y. (2006), ‘Determinants of justification and self-control.’, *Journal of Experimental Psychology: General* **135**(4), 572.
- Kohn, A. (1999), *Punished by rewards: The trouble with gold stars, incentive plans, A’s, praise, and other bribes*, Houghton Mifflin Harcourt.
- Kool, W. and Botvinick, M. (2014), ‘A labor/leisure tradeoff in cognitive control.’, *Journal of Experimental Psychology: General* **143**(1), 131.
- Kruglanski, A. W., Alon, S. and Lewis, T. (1972), ‘Retrospective misattribution and task enjoyment’, *Journal of Experimental Social Psychology* **8**(6), 493–501.
- Kruglanski, A. W., Friedman, I. and Zeevi, G. (1971), ‘The effects of extrinsic incentive on some qualitative aspects of task performance¹’, *Journal of Personality* **39**(4), 606–617.
- Kruglanski, A. W., Shah, J. Y., Fishbach, A., Friedman, R., Chun, W. Y. and Sleeth-Keppler, D. (2002), ‘A theory of goal systems’, *Advances in experimental social psychology* **34**, 331–378.
- Lacetera, N., Macis, M. and Slonim, R. (2013), ‘Economic rewards to motivate blood donations’, *Science* **340**(6135), 927–928.
- Lepper, M. R. and 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, M. R. and Greene, D. (1978), ‘Overjustification research and beyond’, *The hidden costs of reward. Hillsdale, NJ: Erlbaum*.
- Lepper, M. R., Greene, D. and 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, M. R., Henderlong, J. and 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).’.
- Lerner, J. S. and Tetlock, P. E. (1999), ‘Accounting for the effects of accountability.’, *Psychological bulletin* **125**(2), 255.
- Loewenstein, G., Brennan, T. and Volpp, K. G. (2007), ‘Asymmetric paternalism to improve health behaviors’, *Jama* **298**(20), 2415–2417.
- Mazar, N., Amir, O. and Ariely, D. (2008), ‘The dishonesty of honest people: A theory of self-concept maintenance’, *Journal of marketing research* **45**(6), 633–644.
- Mela, C. F., Jedidi, K. and Bowman, D. (1998), ‘The long-term impact of promotions on consumer stockpiling behavior’, *Journal of Marketing research* pp. 250–262.

- Milkman, K. L., Rogers, T. and Bazerman, M. H. (2008), ‘Harnessing our inner angels and demons: What we have learned about want/should conflicts and how that knowledge can help us reduce short-sighted decision making’, *Perspectives on Psychological Science* **3**(4), 324–338.
- Neal, D. T., Wood, W. and Quinn, J. M. (2006), ‘Habits repeat performance’, *Current Directions in Psychological Science* **15**(4), 198–202.
- Nisbett, R. E. and Valins, S. (1971), *Perceiving the causes of one’s own behavior*, General Learning Press.
- Oyer, P. (1998), ‘Fiscal year ends and nonlinear incentive contracts: The effect on business seasonality’, *Quarterly Journal of Economics* pp. 149–185.
- Pink, D. H. (2011), *Drive: The surprising truth about what motivates us*, Penguin.
- Prendergast, C. (1999), ‘The provision of incentives in firms’, *Journal of economic literature* pp. 7–63.
- Raudenbush, S. W. and Bryk, A. S. (2002), *Hierarchical linear models: Applications and data analysis methods*, Vol. 1, Sage.
- Razran, G. (1954), ‘The conditioned evocation of attitudes (cognitive conditioning?).’, *Journal of Experimental Psychology* **48**(4), 278.
- Ryan, R. M. and Deci, E. L. (2000), ‘Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being.’, *American psychologist* **55**(1), 68.
- Schkade, D. A. and Kahneman, D. (1998), ‘Does living in california make people happy? a focusing illusion in judgments of life satisfaction’, *Psychological Science* **9**(5), 340–346.
- Shafir, E., Simonson, I. and Tversky, A. (1993), ‘Reason-based choice’, *Cognition* **49**(1-2), 11–36.
- Shah, J. Y. and Kruglanski, A. W. (2002), ‘Priming against your will: How accessible alternatives affect goal pursuit’, *Journal of Experimental Social Psychology* **38**(4), 368–383.
- Shu, S. B. and Gneezy, A. (2010), ‘Procrastination of enjoyable experiences’, *Journal of Marketing Research* **47**(5), 933–944.
- Skinner, B. F. (1953), *Science and human behavior*, Simon and Schuster.
- Tang, S.-H. and Hall, V. C. (1995), ‘The overjustification effect: A meta-analysis’, *Applied Cognitive Psychology* **9**(5), 365–404.
- Urminsky, O. and Goswami, I. (2015), ‘Impatient to achieve or impatient to receive: How the goal gradient effect underlies time discounting’, *Working Paper, University of Chicago*.

- Urminsky, O. and Kivetz, R. (2011), ‘Scope insensitivity and the mere token effect’, *Journal of Marketing Research* **48**(2), 282–295.
- Urminsky, O. and Zauberaman, G. (2016), *Blackwell Handbook of Judgment and Decision Making*, Wiley-Blackwell, chapter The Psychology of Intertemporal Preferences.
- Vohs, K. D., Baumeister, R. F., Schmeichel, B. J., Twenge, J. M., Nelson, N. M. and Tice, D. M. (2014), ‘Making choices impairs subsequent self-control: a limited-resource account of decision making, self-regulation, and active initiative.’.
- Volpp, K. G., John, L. K., Troxel, A. B., Norton, L., Fassbender, J. and Loewenstein, G. (2008), ‘Financial incentive-based approaches for weight loss: a randomized trial’, *Jama* **300**(22), 2631–2637.
- Volpp, K. G., Levy, A. G., Asch, D. A., Berlin, J. A., Murphy, J. J., Gomez, A., Sox, H., Zhu, J. and Lerman, C. (2006), ‘A randomized controlled trial of financial incentives for smoking cessation’, *Cancer Epidemiology Biomarkers & Prevention* **15**(1), 12–18.
- Vroom, V. H. (1964), ‘Work and motivation. new york: John willey & sons’, *Inc. Vroom Work and Motivation* 1964 .
- White, R. W. (1959), ‘Motivation reconsidered: the concept of competence.’, *Psychological review* **66**(5), 297.
- Yang, A. X., Hsee, C. K. and Urminsky, O. (2014), ‘Eager to help yet reluctant to give: How pro-social effort and pro-social choices diverge’, *Available at SSRN 2733331* .

APPENDIX A

STUDY STIMULI, PRE-TESTS, AND DATA CLEANING PROTOCOL

Study Stimuli

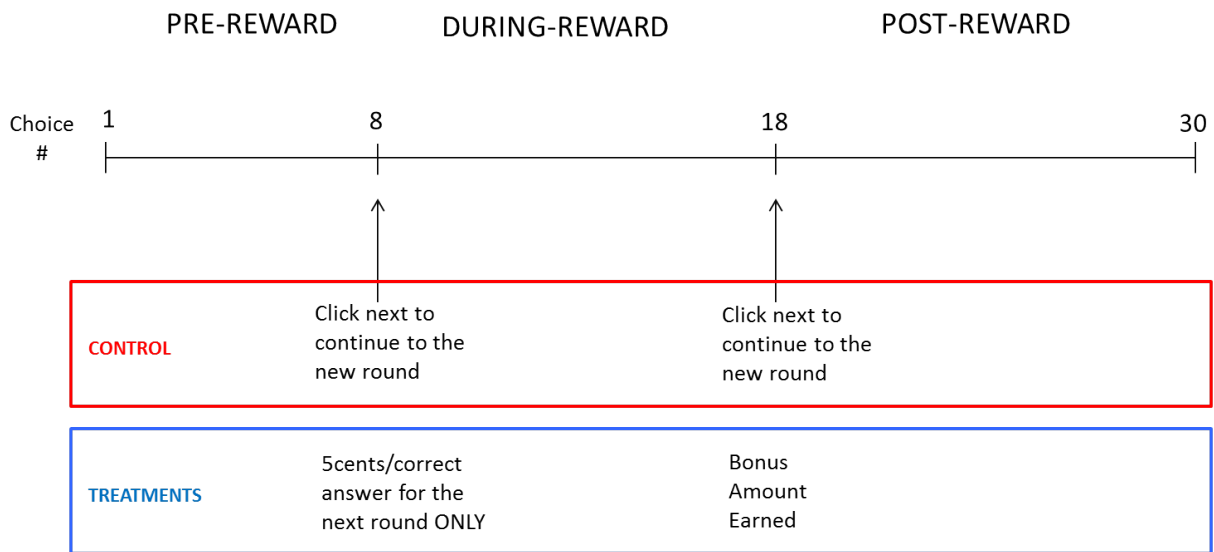


Figure A.1: The exhibit shows how the series of 30 repeated choices (or trials) were organized in the studies. The trials were divided into three rounds for all the experimental groups, with 8, 10, and 12 trials in the pre-, during-, and post-reward rounds respectively. The reward group(s) learnt about the total rewards earned before the start of the third round. However, all payments were made at the very end of the study, and not after the immediate end of the second round.

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.

A typical cognitive task in this survey looks like this.

The task will require **searching and selecting two numbers in a grid such that they add up to 10**. You can select a number by clicking on the box containing the number. An example is shown below.

8.63	4.38	2.68
5.72	1.67	7.38
7.32	3.62	1.29
7.02	5.17	1.62

Note:

1. **All** cognitive tasks in this study have an unique pair of answer.
2. You should **select ONLY two numbers in a grid**, and no more, as shown above.

You will get **30 seconds to complete each such tasks** after which the survey will automatically advance to the next screen. If you are done with the task before the time limit, you can click NEXT to proceed.

Alternatively, you can choose to watch a video clip for the same duration, and click NEXT when you are done viewing.

You will get **30 seconds to complete each such tasks** after which the survey will automatically advance to the next screen. If you are done with the task before the time limit, you can click NEXT to proceed.

Alternatively, you can choose to watch a video clip for the same duration, and click NEXT when you are done viewing.

Figure A.2: Four-stage instructions (shown as separated by dotted lines) provided to all participants before the start of all the studies. After the instructions, all participants sampled a task and a video and were asked if everything worked well. The experiment was aborted for those who reported having a problem.

Please indicate your choice of a topic below.

- ☐ Your favorite genre of Music
- ☐ Your idea of favorite vacation

Please indicate your choice of a topic below.

- ☐ Why are social media sites so important in today's popular culture?
- ☐ Why are tattoos so popular in today's society?

Please indicate your choice below about which you would like to give your opinion.

- ☐ Should kids be given smartphones?
- ☐ Should the minimum age for teenagers to get a driver's license be increased?

Please choose the product category for which you would like to do the matching between logos and brand names.

- ☐ Automobile
- ☐ Sports Good

Please choose the product category for which you would like to do the matching between logos and brand names.

- ☐ Oil and Gas Companies
- ☐ Luxury and Accessories

Please choose the product category for which you would like to do the matching between logos and brand names.

- ☐ Technology Companies
- ☐ Banks and Financial Institutions

Figure A.3: The exhibit shows the types of tasks used for unrelated activities break in Study 2. Participants did all the three tasks of a particular type (Writing: top panel or Brand-name Matching: bottom panel), and each one of these tasks had a time limit of 30 seconds. The exhibit shows the choice condition. In the no-choice condition, one of the options was randomly pre-selected for the participant.

BONUS INFORMATION: PLEASE READ CAREFULLY.

In a previous version of this survey we were able to pay as much as 5 cents for every correct answer, but in this version **we are unable to pay more.**

You will get 1 cent for every cognitive task that you answer correctly.

BONUS INFORMATION: PLEASE READ CAREFULLY.

In a previous version of this survey we were able to pay only 5 cents for every correct answer, but in this version **we are able to pay a LOT more.**

You will get 50 cents for every cognitive task that you answer correctly.

Figure A.4: The exhibit shows the manipulations used to make the perceived reward magnitude salient in Study 3. The top panel shows the low-reward condition, and the bottom panel shows the high-reward condition.

PLEASE READ THE INSTRUCTIONS CAREFULLY.

In this survey you will be asked to do a task. The task is to **solving cognitive math tasks**. WE WILL USE YOUR RESPONSES TO DESIGN EXPERIMENTAL STIMULI FOR A SPATIAL REASONING STUDY.

Since doing the task can be tiring, you will also have an option of a different task, **evaluating videos of television advertisements**, so that you can take a break.

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

PLEASE READ THE INSTRUCTIONS CAREFULLY.

In this survey you will be asked to do a task. The task is to **evaluate videos of television advertisements**. WE WILL USE YOUR RESPONSES TO DESIGN EXPERIMENTAL STIMULI FOR AN ATTENTION AND PERCEPTION STUDY.

Since doing the 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 in each round.

Figure A.5: The exhibit shows the two types of framing used for the math (work: top panel) and the video (leisure: bottom panel) task in Study 4.

Please watch the video and rate it to indicate how much you liked it.



Your rating ★ ★ ★ ★ ★

Figure A.6: The exhibit shows a typical video task used in Study 4. Unlike other studies (Studies 1-3), in this study the participants in the target-task = video condition were asked to rate the video in order to get their rewards.

PLEASE READ CAREFULLY.

In the next round instead of making repeated choices between doing a cognitive task and watching a video, **you will need to choose now what you want to do next during the ENTIRE round.**

If you choose Cognitive Tasks, you will be presented with ONLY Cognitive Tasks in the next round.

If you choose Videos, you will be presented with ONLY Videos in the next round.

Please indicate what would you like to do during the ENTIRE NEXT ROUND.

☐ Do only Cognitive Tasks

☐ Watch only Videos

Figure A.7: The exhibit shows the instructions (two stages, shown separated by a dotted line) given to participants at the end of Round 2 in Study 5. Participants were asked to choose which task they would like to do during the entire ensuing period of Round 3.

PLEASE READ THE INSTRUCTIONS CAREFULLY.

In this survey you will be given a series of **choices between doing two important tasks.** These two tasks are solving math problems and evaluating television advertisements.

Both the tasks are equally important to the researchers, who need people's feedback on both of the tasks.

The cognitive math task will be used to design experimental stimuli for a *reasoning study*.

The TV ad watching task will be used to design experimental stimuli for an *attention study*.

Since we will be collecting data from lots of people, it is completely up to you to choose which task you want to do in each round.

Figure A.8: The exhibit shows the instructions given to participants at the beginning of Study 6. Participants were told that both tasks are equally important to the researchers.

This concludes the first round of the survey.

We will now move to the next round.

In the next round, the computer will choose at each trial whether you will be doing a cognitive task or will be watching a video.

Click NEXT to begin the next round.

This concludes the first round of the survey.

We will now move to the next round.

In the next round, the computer will choose at each trial whether you will be doing a cognitive task or will be watching a video.

Also, you can earn bonus in the next round, over your base payment.

Click NEXT to read more details.

Figure A.9: The exhibit shows the instructions given to participants at the beginning of Round 2 (reward-round for treatment group) in Study 7. Participants in both the control (top panel) and the reward group (bottom panel) were told that the computer will choose the task (math or video) in each trial of the next round. Participants returned to making their own choices after the end of Round 2.

PLEASE READ THE INSTRUCTIONS CAREFULLY.

In this survey you will be asked to do **various tasks**. There will be either **math tasks** or **spelling tasks**.

Instead of doing these tasks you will also have an option of **watching videos of interesting television advertisements** collected from across the world.

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

Please type in the correct spelling of each of the words.

If you think the word is spelled correctly, just retype it. If you think it is spelled incorrectly, please type in the correct spelling.

- 1) harrass
- 2) hiatus
- 3) hirearcy
- 4) homogenity
- 5) heterogenity
- 6) hustle
- 7) hyperbole

Correct spelling:

- 1)
- 2)
- 3)
- 4)
- 5)
- 6)
- 7)

Figure A.10: Top panel shows the instruction given to participants at the very beginning of Study 8. This was followed by a sample of each task-type. Participants in both the control and the reward group made multiple choices between doing a spelling task or watching a video in Round 3. The bottom panel shows an actual example of a spelling task used in Round 3.

Pre-test Results

A pretest ($N = 47$) was done to examine how people felt about the math task (i.e., the target task) and the video task (i.e., the alternative task). A random sample of participants was chosen from the same population and they judged the two tasks on several attributes. Participants judged the math task as relatively more work-like compared to the video ($M_{math} = 6.87, SD = 2.19$ vs. $M_{video} = 2.49, SD = 1.96; t(46) = 9.77, p < .001$), on 9-point scales, but considered the video more leisure-like compared to the math ($M_{math} = 3.89, SD = 2.63$ vs. $M_{video} = 6.78, SD = 2.37; t(46) = 6.24, p < .001$). The math task was also judged as relatively more effortful ($M_{math} = 5.59, SD = 2.44$ vs. $M_{video} = 1.95, SD = 1.52; t(46) = 9.05, p < .001$) and less entertaining compared to the video ($M_{math} = 4.64, SD = 2.32$ vs. $M_{video} = 6.43, SD = 2.10; t(46) = 4.26, p < .001$).

Participants felt that the math task had more long-term benefits whereas the videos had higher immediate benefits ($M_{math} = 5.38, SD = 1.97$ vs. $M_{video} = 3.97, SD = 1.65; t(46) = 3.65, p < .001$). Participants also felt that more justification (on a 1 = Less to 9 = More scale) was needed for choosing to watch the video task over doing the math tasks, than for the opposite choice ($M_{choose math} = 2.70, SD = 2.28$ vs. $M_{choose video} = 4.72, SD = 2.97; t(46) = 3.53, p < .001$).

Most importantly, both the tasks satisfied the pre-condition required of tasks that can be deemed appropriate for testing theories of intrinsic motivation (Deci et al., 1999). Both task had a rating of higher than the mid-point on a scale measuring how “interesting and enjoyable” the task is (1=Low, 9=High): math task ($M_{math} = 6.02, SD = 2.19$) and the video task ($M_{video} = 6.64, SD = 2.08$).

Data Cleaning Protocol

Every study started with an initial sampling of both types of tasks (math, video) after which participants were asked if they faced any technical problems. If a problem was reported, the study was aborted and data from these participants were discarded from further analysis.

Our experimental paradigm was specifically designed to capture dynamic changes in behavior over time, and could distinguish between temporary and permanent disengagements. A temporary decrease in motivation to do the target task would be reflected in the participant choosing to watch the video for a few trials before choosing to do the math task again. A more persistent decrease in motivation to do the target activity could be reflected in two ways. Participants could “quit” within the study, by repeatedly deciding to only watch the videos for the remaining duration. Alternatively, participants could quit by ending the study part way through and not completing the remaining trials. We tracked all dropouts, and included participants who dropped-out of the study after completing the pre-reward baseline period coding their participation as zero for the target task. The reward for the math task to the treatment group was announced at the end of the pre-reward baseline period, and therefore this analysis strategy ensured that we included anyone whose behavior could have been impacted by the incentives, whether they finished the study or not.

Participant’s data containing duplicate IP addresses were removed prior to analysis. Finally, an attention check question was administered at the end of the survey, and data from participants who reached till the end of the survey but failed the attention check were discarded prior to analysis. Participants who quitted part way through and therefore did not answer the attention check question were given the benefit of doubt and were included in the analysis.

APPENDIX B

HIERARCHICAL REGRESSION SPECIFICATIONS

Model for Momentary post-reward engagement reduction

We capture total momentary post-reward engagement reduction using a functional form assumption about how effort returns to baseline over time in the post-reward period. Assuming a non-linear return of effort (i.e., likelihood of choosing the math task) to baseline over time (the number of periods t since the incentive ended), we parameterize momentary post-reward engagement reduction (*MOMENTARY*) as:

$$MOMENTARY_t = \frac{1}{t} \quad (\text{B.1})$$

Using this parameterization¹, the probability of individual i choosing to do the math task in post-reward (Round 3) during trial t can be written as:

$$Pr(Y_{ti} = 1) = \phi(\beta_{0i} + \beta_{Mi} * MOMENTARY_t) \quad (\text{B.2})$$

In our model we set ϕ to the logit link function and β_{0i} is a person-specific intercept and β_{Mi} is a person-specific momentary post-reward engagement reduction behavior. In the hierarchical regression the parameters in Equation B.2 are a function of time-invariant individual-level covariates, to account for the repeated observations per person.

$$\beta_{0i} = \beta_{00} + \beta_{01} * C_i + \beta_{02} * \mathbf{X}_i + u_{0i} \quad (\text{B.3})$$

The person-specific baseline parameter β_{0i} is a function of the experimental condition C_i that individual i has been randomly assigned to, the total number of math task choices by individual i in the pre-incentive Round 1 \mathbf{X}_i , as well as the population baseline β_{00} and

1. We report robustness analysis with other parameterizations in Appendix D

time-invariant person-specific error term u_{0i} .

$$\beta_{Mi} = \beta_{10} + \beta_{11} * C_i + u_{1i} \quad (\text{B.4})$$

The person-specific momentary post-reward engagement reduction behavior β_{Mi} is estimated as a function of experimental condition C_i , as well as the baseline β_{10} and the individual-specific error term u_{1i} . The random effects for the intercept and the slope for every individual i , u_{0i} , u_{1i} , are assumed to follow a bi-variate normal distribution with zero-mean, variances τ_{00} , τ_{11} and common co-variance τ_{01} . This error structure accounts for the potentially correlated repeated-measures for each individual. Combining equations B.2, B.3, and B.4 yields an “intercepts and slopes-as-outcomes” model (Raudenbush and Bryk, 2002).

The expected proportion of math tasks chosen in each trial t of Round 3 is:

$$\begin{aligned} Pr(Y_{ti} = 1) = & \phi(\beta_{00} + \beta_{10} * MOMENTARY_t + \beta_{01} * C_i + \\ & \beta_{02} * \mathbf{X}_i + \beta_{MOMENTARY} * C_i * MOMENTARY_t) \end{aligned} \quad (\text{B.5})$$

The coefficient $\beta_{MOMENTARY}$ in equation B.5, which is β_{11} from equation B.4 renamed for ease of exposition, tests for a difference in the extent of momentary post-reward engagement reduction in the experimental condition ($C_i = 1$), compared to the corresponding time periods in the control condition ($C_i = 0$). A significant and negative $\beta_{MOMENTARY}$ generally indicates momentary post-reward engagement reduction after the incentive ended, compared to the corresponding trials in the control condition, controlling for individual differences in baseline effort \mathbf{X}_i . An important exception to such an interpretation arises when there is an overall increase in effort of the reward group relative to the control group in the post-reward period. In this case a significant negative $\beta_{MOMENTARY}$ might indicate an immediate decrease in the effort of the reward group relative to its longer-run steady-state level, but not a momentary post-reward engagement reduction relative to the contemporaneous

control group.

In a similar vein, an estimate of $\beta_{MOMENTARY}$ that is not statistically distinguishable from zero represents a consistent level of effort throughout Round 3, with two very different potential interpretations. A non-significant $\beta_{MOMENTARY}$ could indicate that no post-reward reduction in engagement has occurred or it could represent a consistent overall increase or decrease in engagement in the reward group in Round 3. Hence, it will be important to also estimate the overall effects of incentives on choices in Round 3, in addition to momentary reduction in engagement behavior and interpret these parameters jointly. Next, we describe the tests we use to estimate the overall effects of incentives.

Difference-in-Difference Model for Overall Effects

We use a hierarchical non-linear difference-in-difference model to estimate differences in the overall probability of choosing the math task between two experimental conditions $C_i = 0$ or $C_i = 1$ and between two experimental rounds $R_t = 0$ or $R_t = 1$. The general specification can be written as follows:

$$Pr(Y_{ti} = 1) = \phi(\beta_0 + \beta_1 * R_t + \beta_2 * C_i + \beta_3 * R_t * C_i) \quad (\text{B.6})$$

The interpretation of the key coefficient β_3 depends on how the rounds (R_t) and conditions (C_i) are coded. To estimate the effect of incentives on during-reward performance β_{REWARD} we compare during-incentive Round 2 ($R_t = 1$) to baseline Round 1 ($R_t = 0$) and exclude Round 3 data. To estimate the overall post-reward engagement level β_{POST} we compare post-reward Round 3 ($R_t = 1$) to baseline Round 1 ($R_t = 0$), and exclude Round 2. The net effect of incentives (e.g., during-reward and post-reward behavior) β_{NET} is estimated by comparing the combined during-reward Round 2 and post-reward Round 3 trials ($R_t = 1$) to baseline Round 1 ($R_t = 0$).

APPENDIX C

STUDIES USED IN INTERNAL META-ANALYSIS

Table C.1: List of Studies included in the internal meta-analysis. Not all studies are included in the manuscript as indicated in the Remarks column.

Studies	Remarks	Sample Sizes	β	SE	z	p
1	Study 1 in paper	C = 39; R = 38	-2.6	1	-2.6	.009***
2	Study 2 in paper	C = 41; R = 46	-2.5	0.9	-2.5	.01**
3	Study 3 in paper	C = 68; R = 56	-3.5	1.1	-3.1	.002**
4	Study 6 in paper	C = 59; R = 52	-4.7	1.2	-3.8	<.001**
5	Replication of momentary reduction in engagement with notice about temporary nature of rewards at the end of Round 2	C = 31; R = 33	-2.2	0.9	-2.3	0.02
6	Replication of momentary reduction in engagement (different project)	C = 72; R = 74	-1.6	0.9	-1.7	.09*
7	Replication of Study with autonomy vs non-autonomy breaks	C = 42; R = 41	-2.3	1.6	-1.4	.15
8	Replication of study with different reward magnitude	C = 35; R = 36	-1.5	1.2	-1.3	.2
9	Study with and without a pre-reward round (different project)	C = 56; R = 96	-2.6	1.3	-2	.048**
10	Replication of Study with incentives for Math vs Leisure	C = 87; R = 96	-3.2	0.7	-4.6	<.001***

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Note: Study 4 in the main paper was not included in the meta-analysis because it used a different math task. The results will not change (in fact, will become stronger) if that study is also included.

APPENDIX D

ROBUSTNESS CHECKS

Different Parametric Specifications

Table D.1: Estimates for Momentary Reduction in Engagement $\beta_{MOMENTARY}$ in meta-analysis with various specifications, SEs are in parenthesis

	$M_t = \frac{1}{t}$	$M_t = \frac{1}{t^2}$	$M_t = \frac{1}{\sqrt{t}}$	$M_t = \frac{1}{t}$ without dropouts
Constant	-5.1 (0.3)***	-4.8 (0.3)***	-5.1 (0.3)***	-4.4 (0.3)***
M_t	1.5 (0.3)***	1.5 (0.3)***	1.5 (0.3)***	1.2 (0.2)***
Condition = Reward	2.2 (0.3)***	0.9 (0.2)***	2.2 (0.3)***	1.3 (0.2)***
Total Attempts in Round 1	1.0 (0.05)***	1.1 (0.05)***	1.0 (0.05)***	1.0 (0.04)***
M_t * Condition	-3.8 (0.4)***	-2.8 (0.4)***	-3.7 (0.4)***	-2.9 (0.3)***
Dropouts Included	Yes	Yes	Yes	No
N	1098	1098	1098	1046
BIC	9821	9899	9881	9526

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Note: For the sake of parsimony and simplicity, we model $MOMENTARY_t$ (shown as M_t) as per specification 1 in this paper which also has the lowest BIC.

Columns 1-3 compare different functional form specifications for the $MOMENTARY_t$ parameter (shown as M_t), and column 4 excludes the dropouts. Dropouts are participants who dropped-out of the study after completing Round 1. The parameter t indicates the post-reward trial or choice number.

Non-parametric Specification

We used the following flexible specification for predicting probability of choosing the math task in the post-reward round by individual i at trial t :

$$Pr(Y_{ti} = 1) = \phi(T * C_i + \mathbf{X}_i + S) \quad (\text{D.1})$$

Here,

ϕ = Logit link function

T = Cubic splines ($df = 3$) for post-reward trial t , $t \in [1 \dots 12]$ i.e., 12 post-reward choices

C_i = Experimental condition to which individual i is randomly assigned

\mathbf{X}_i = Total number of math attempts by individual i in the pre-incentive round. This is a measure of individual level differences in ability or interest in the target-task of the experiment

S = Study Fixed Effect

The model specification did not use any assumption about how effort (e.g., attempting the math task) returns to baseline after incentives are stopped. Using flexible cubic-splines a piece-wise third order polynomial is used to fit individual post-reward behavior after the incentive ends.

Also, instead of using distributional assumptions to draw statistical inferences, we predicted the probability of attempting the math task for each of the post-reward trials in both the control and the treatment group. Using a 1000-sample bootstrap test, we examined if this difference did not contain zero to infer statistical difference in behavior between these two experimental groups.

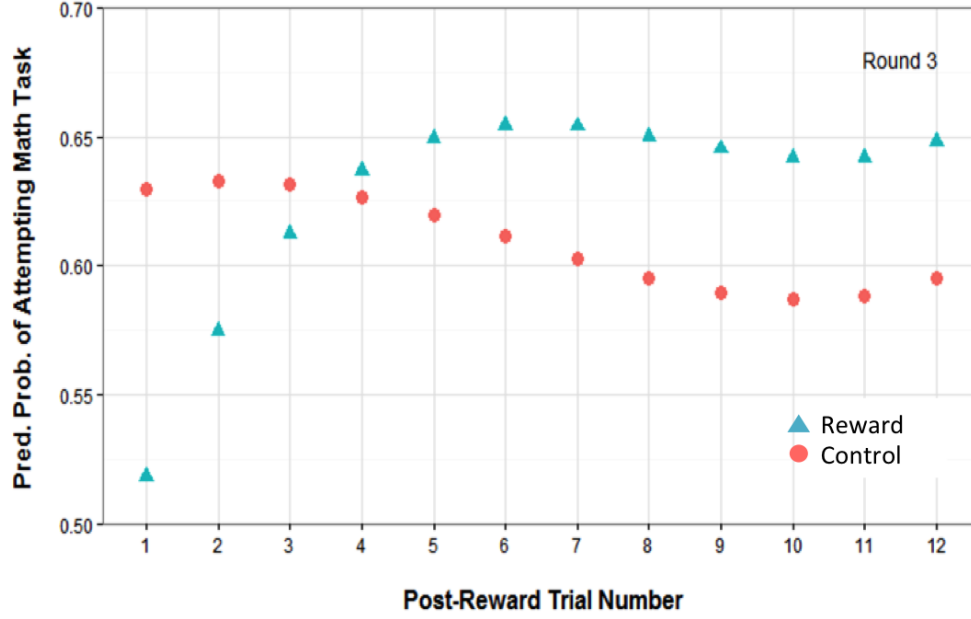


Figure D.1: Figure shows the predicted probability of choosing the math task in the reward group (blue triangular dot) and the control group (red circular dot) in Round 3 based on a typical run of the non-parametric model.

Post-reward Choice #	Bootstrapped 95% CI	
Choice 1	-19.6%	-10.6%
Choice 2	-12.6%	-4.8%
Choice 3	-7.8%	0.0%
Choice 4	-4.3%	3.6%
Choice 5	-1.9%	6.1%
Choice 6	-0.2%	7.7%
Choice 7	0.9%	8.7%
Choice 8	1.4%	9.3%
Choice 9	1.7%	9.6%
Choice 10	1.8%	9.7%
Choice 11	2.0%	9.8%
Choice 12	1.7%	10.7%

Table D.2: The table shows the bootstrapped 95% CI from the non-parametric model for the difference in the predicted probability of choosing the math task between the reward and the control condition.

APPENDIX E

HETEROGENEITY IN POST REWARD BEHAVIOR

Heterogeneity by Initial Post-reward Behavior




Figure E.1: The figure shows reward and the control groups further divided into two sub-groups based on initial post-reward behavior (initial post reward behavior lower or higher than individual-level pre-reward average). For example, ‘Reward - Initially Low’ are reward-group participants who showed an initial dip in engagement (compared to their own Round 1 baseline) at the end of Round 2.

As shown in Figure E.1, the reward group that shows initial post-reward reduction in engagement eventually settles at a level higher than the corresponding control group. This suggests that the aggregate behavior was *not* driven by two types of reward group participants - one that showed a persistent post-reward reduction in engagement as predicted by prior theories (relative to corresponding control) and another that showed a persistent post-reward increase in engagement (relative to corresponding control). As shown in the table below, the final post-reward behavior does not differ between the control and the reward as a function on their initial post-reward behavior.

Table E.1: Regression results showing that longer-term post-reward behavior (i.e., return to baseline) not different between control and reward group as a function of their initial post-reward behavior

	DV = Final Post-reward Behavior (Normalized)
Condition = Reward	0.121 (-0.014)***
Initial Reduction in Engagement (Normalized)	0.627 (-0.03)***
Condition = Reward x Initial Reduction in Engagement	0.016 (-0.039)
Constant	-0.048 (-0.01)***
N	1098
R-Squared	0.503

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$ (SEs are in parenthesis)

Notes: Normalized Initial Reduction in Engagement means Post-reward Initial Behavior *minus* Pre-reward Baseline

Normalized Final Post-reward Behavior (DV) means Final Post-reward Behavior *minus* Pre-reward Baseline

Heterogeneity by Initial Motivation to do the Math Task




Figure E.2: Post-reward behavior of participants with high initial task interest using the internal meta-analysis data. The figure shows a significant post-reward decrease in engagement after incentives ended.




Figure E.3: Post-reward behavior of participants who low initial task interest using the internal meta-analysis data. The figure shows a significant post-reward decrease in engagement after incentives ended.

As shown in figures E.2 and E.3, the result of temporary incentives on post-reward behavior is very similar for the group of participants who exerted more versus less effort in the pre-reward round. This difference in initial effort represents difference in initial intrinsic motivation because in the pre-reward period participants did not know about any impending rewards. Both groups show an increase in effort when the rewards are available, followed by a decrease in the choice of the math task in the immediate post-reward period (low intrinsic motivation: $\beta_{MOMENTARY} = -2.11, z = -2.71, p = .007$; high intrinsic motivation: $\beta_{MOMENTARY} = -3.22, z = -8.84, p < .001$), and initial interest level does not moderate the momentary reduction in engagement behavior ($\beta_{MOMENTARY\ interaction} = 0.15, z < 1$).

APPENDIX F

EFFECT ON ACCURACY AND NET OUTCOME




Figure F.1: Raw data showing the percentage of participants accurately solving the math task, conditional on choosing to attempt it. Incentives did not affect accuracy in our experiments.

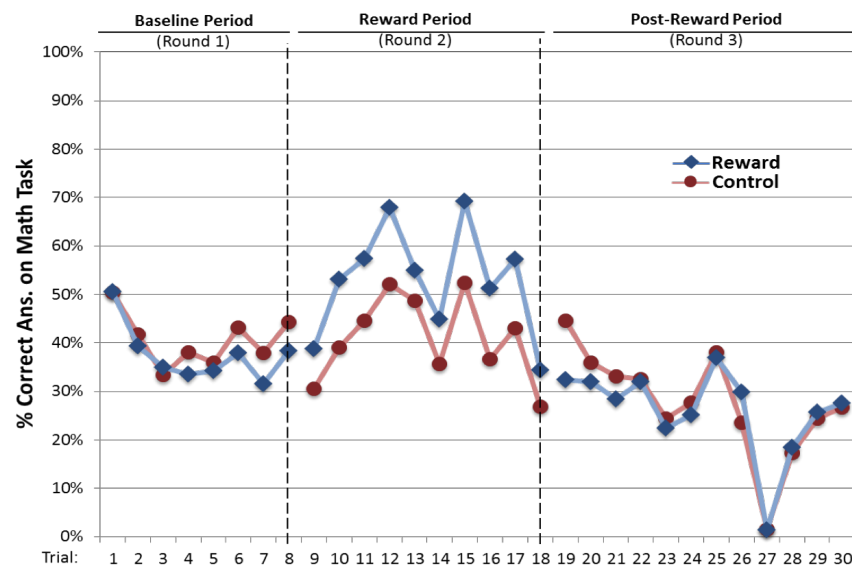


Figure F.2: Raw data showing proportion of correct answers for every trial in each round. Incentives had a net positive effect on the total number of correct answers, driven by the reward period, despite a significant post-reward decrease after incentives ended.

We primarily focused on effort (e.g., choosing to do the target task) as the key variable of interest, consistent with the approach used in the intrinsic motivation literature because effort is a behavioral outcome variable that represents a person’s motivation level. Our flexible experimental paradigm also allowed us to also examine the effect of temporary incentives on accuracy (probability of correctly answering the math task after deciding to attempt it) and net outcome (total number of correct answers). The incentive could have resulted in people choosing the math task without being able to answer correctly, resulting in a decrease in accuracy compared to the control condition. However, we did not observe any such effects and temporary incentives did not affect accuracy at all. Therefore, as shown in figures F.1 and F.2, the same conclusions hold for effort and net outcome - a significant positive effect of incentives on the total number of math problems solved correctly ($\beta_{NET} = +1.25, z = +5.21, p < .001$).

APPENDIX G

FIELD STUDIES USING TEMPORARY CONTINGENT INCENTIVES

Table G.1: List of past field studies that have used temporary incentives and looked at post-reward behavior. Only studies that (a) have measured actual post-reward behavior (and not self-report), (b) have studied adults (including a Study with 11th and 12th graders), and (c) where the target task is not a pro-social activity (e.g., contribution to charity, blood donation etc.) are included.

	Domain	Target Group	Incentive size / type	Time-point(s) when post-reward behavior was measured	Finding(s)
Garbarino & Slonim, 2005	Education	University students at a private school	\$10 for passing a test	Immediate post-reward behavior (e.g, the subsequent test)	On Average, 2 fewer questions attempted in incentive group vs control ($p < .05$)
Volpp et al., 2006	Smoking Cessation	Smokers at Philadelphia Veteran Affairs Medical Center	\$100 for quitting to smoke	6 months after incentive to quit	Quit rates in incentive group (6.5%) not different from control (4.6%, $p > .2$)
Jackson, 2007	Evaluation of Advanced Placement Incentive Program (APIP) in Education	11th and 12th grade students (and teachers) in Texas schools serving underprivileged populations	Between \$100 and \$500 for getting a score of 3 or more in each eligible test subject	Future test scores and college graduation	13% increase in number of students scoring about 1100/24 on SAT/ACT ($p < .05$) and 5% increase in students matriculating in college ($p < .10$)
Volpp et al., 2008	Warfarin Adherence	Warfarin patients at the UPenn Anticoagulation Management Center	Lottery with daily expected value of \$5(Study 1) or \$3(Study 2)	Not-reported	Regulation of anticoagulation levels changed from 35% (pre) to 42% (post, w/S, Study 1; ns) and from 65% (pre) to 60% (post, w/S, Study 2; ns)
Volpp et al., 2008	Weight Loss	Healthy adults age 30-70 with a BMI of 30-40 from the Philadelphia VA Medical Center	Lottery incentive (expected value \$3/day), or deposit contract with matching incentives (max \$8.4/day)	7 months after end of intervention	Both in Lottery ($\Delta \approx -9 lbs.$, $p = .01$) and in deposit contract ($\Delta \approx -6 lbs.$, $p = .03$) participants weighed less than the beginning of the study
Angrist et al., 2009	Education	Entering first-year undergraduates at a primarily commuter school	Merit scholarship or merit scholarship with support service	1 year after end of intervention	0.28 percentage points ($p < .01$) increase in grade points for women; longer-term effect on men non-significant

Table G.1: (continued)

	Domain	Target Group	Incentive size / type	Time-point(s) when post-reward behavior was measured	Finding(s)
Cawley & Price, 2009	Weight Loss	Employees from employer that has contract with Company X	Various quarterly monetary rewards or lottery plus refundable bonds	1 year after end of intervention	No significant difference with quarterly rewards w.r.t baseline. For lottery + bonds 3.6 lbs. ($p < .05$) loss w.r.t baseline
Charness & Gneezy, 2009	Gym Attendance	University of Chicago undergrad students	Low Reward: \$25 to attend gym once in week. High Reward: \$100 to attend gym 8 times in 4 weeks	One attendance measure per week, 7 weeks after intervention	Higher post-intervention gym attendance in high-reward vs control (0.67 visits/week) and vs. low-reward group (0.58 visits/week)
Acland & Levy, 2010	Gym Attendance	Self-reported non-regular gym attenders	\$25 to attend gym once in week and then \$100 to attend gym 2 times every week for 4 weeks	One attendance measure per week, 5 weeks post-treatment and following weeks into next semester	Higher post-intervention gym attendance in reward vs control (0.256 visits/week)
John et. al, 2010	Weight Loss	Patients at the Philadelphia Veterans Affairs Medical Center with BMIs of 30-40, age 30-70	Deposit contracts in which participants put \$0-3 daily of their own money at risk (matched 1:1)	Weigh in 36 weeks after end of intervention	No longer-term difference in weight loss between treatment (1.2 lbs.) and control (0.27 lbs; $p = .76$)
Kimmel et. al, 2012	Warfarin Adherence	Warfarin patients at the UPenn Anticoagulation Management Center	Lottery with daily expected value of \$3	6-months after end of intervention	No difference on anticoagulation levels between reward (23%) and control (25.9%; ns)
Royer et. al, 2012	Gym Attendance	Employees from Midwest Fortune 500 company	\$10 for visiting company gym (up to 3x per wk.) over 4 wks., free membership, and \$20 for new members; w/ or w/o self-funded commitment contracts	Gym use via login records 5-13 weeks and 14-52 weeks after end of incentives	Significantly higher post-intervention gym attendance in incentives vs control (0.11 visits/week; $p < .05$) in weeks 5-13 after incentives end. The results are directionally positive but ns in weeks 14-52 after incentives end
Bareket-Bojmel et. al 2014	Work productivity	Technicians at a global high-tech semiconductor company working at a fabrication plant in Israel.	\$25, family pizza meal voucher, verbal reward, or own choice if performance level exceeded base productivity.	Productivity on the first day, second day, and third day after rewards were stopped.	Removal of the cash bonus significantly reduced performance by 13.2% relative to base productivity on day 1. However, with verbal praise productivity was 4.2% higher than baseline on day 1.

Table G.1: (continued)

	Domain	Target Group	Incentive size / type	Time-point(s) when post-reward behavior was measured	Finding(s)
Sen et. al, 2014	Adherence to medical regimen among diabetes patients	Patients of a Primary Care Medical Home at UPenn	Lottery incentive with expected daily value of \$2.80 or \$1.40 for daily monitoring	Every month for three months after end of incentives	After three months, adherence rate was 62 % in low, 35 % in high ($p = .015$) and 27 % in control ($p = .002$ vs. low incentives).
Halpern et. al, 2015	Smoking Cessation	CVS Caremark employees	\$800 or refundable deposit of \$150 with an opportunity to win \$650 in rewards	6 months and 12 months after end of incentives	Compared to the abstinence rate in control (6.0%), abstinence with individual rewards was significantly higher after 6 months (15%, $p < .01$) and after 12 months (7.5%, $p < .05$).
Huffman & Bog-nanno, 2015	Work Productivity	Workers hired to register people for a company database during street festival	Hourly wage (\$18) with a per sign-up monetary bonus (\$5/sign-up) for 1 hour	Every hour for three hours after incentive ended	Immediately after incentives ended Reward group recruited more than Control (7% higher)
Mochan et. al, 2015	Purchase of healthy grocery items	Households participating in Points-based Healthy Food program	Forfeiting Healthy Food discount for failing to increase healthy food items purchased by 5% for the month.	Supermarket shopping data in following 6 months after intervention ended	0.49 percentage-points ($p < .1$) average increase in healthy items purchased in the first three months, and 0.79 percentage-points ($p < .01$) average increase in healthy items purchased in the next three months
Wang et. al, 2016	Number of hotel nights	Loyalty program members at a major international hotel chain	Bonus points during an 8-month period	Hotel stay in 8 months after the intervention	On Average, compared to the control group, the treatment group stayed one-night more in the post-reward period ($p < .01$)