WARWICK
ECONOMICS

To the Depths of the Sunk Cost: Experiments Revisiting the Elusive Effect

George Beknazar-Yuzbashev, Sota Ichiba & Mateusz Stalinski

January 2025 No: 1544

Warwick Economics Research Papers

ISSN 2059-4283 (online)
ISSN 0083-7350 (print)

To the Depths of the Sunk Cost: Experiments Revisiting the Elusive Effect*

George Beknazar-Yuzbashev[†] **Sota Ichiba**[‡]
Mateusz Stalinski[§]

December 2024

Abstract

Despite being often discussed both in practice and academic circles, the sunk cost effect remains empirically elusive. Our model based on reference point dependence suggests that the traditional way of testing it—by assigning discounts—may not produce the desired effect. Motivated by this, we evaluate it across the gain-loss divide in two pre-registered experiments. In an online study, we randomize the price (low, medium, or high) of a ticket to enter a real-effort task and observe its effect on play time. Despite varying the sunk cost by $2 for a 14-minute task and the sample size of N=1,806, we detect only a small effect (0.09 SD or 1.1 minutes). We further explore the economic applications of the effect in a field experiment on YouTube with N=11,328 videos in which we randomize whether the time until a pre-video ad becomes skippable is shortened (0 s), default (5 s), or extended (10 s). The intervention has an overall insignificant effect on video engagement. This is driven by a sizable negative effect on the extensive margin, a channel which is not present in the online study. Specifically, more users leave before the video starts in the extended treatment (5.2 pp. or 28% more relative to the shortened treatment). Taking the results of both studies together, we offer a cautionary tale that applying even the most intuitive behavioral effects in policy settings can prove challenging.

JEL Classification: C91, C93, D04, D11, D12, D90, M31, M37

Keywords: sunk cost, sunk cost fallacy, loss aversion, mental accounting, regret, digital platforms

*We would like to thank Leonardo Bursztyn, Patricio Dalton, Mark Dean, Alex Imas, Boris van Leeuwen, John List, Devin Pope, Jan Potters, Avner Strulov-Shlain, George Wu, and numerous seminar participants for helpful comments and suggestions. This study was approved by the University of Chicago Institutional Review Board (IRB21-0450 for the online study, IRB22-0073 for the field study) and the Columbia University Institutional Review Board (AAAT6659 for the online study, AAAT9887 for the field study), and registered in the American Economic Association Registry for randomized control trials under trial numbers AEARCTR-0009017 (online study) and AEARCTR-0010157 (field study).

[†]Columbia University gb2683@columbia.edu.

[‡]Tilburg University s.ichiba@tilburguniversity.edu.

[§]University of Warwick and CAGE mateusz.stalinski@warwick.ac.uk.

1 Introduction

We revisit the gaps left by the previous inconclusive attempts to harness the psychology of the sunk cost effect.[1] Naturally, due to its policy-relevant implication that the usage of durable goods may increase by assigning a higher upfront cost, it came into focus of economists in settings ranging from development to education. Despite wide acceptance of the sunk cost effect in psychology (e.g. Olivola, 2018; Krefeld-Schwalb et al., 2024), several influential field experiments investigating its power to improve welfare report a null effect (Ashraf et al., 2010; Cohen and Dupas, 2010; Ketel et al., 2016). Our paper evaluates this paradox through the lens of a reference point dependence model and two pre-registered experiments, designed to test for the effect in settings where the theory predicts its existence.

An important feature of the prior field experiments is that while they attempt to investigate the effect of a higher price on product usage,[2] the existing designs rely on discounts—which measure the effect of a *lower* price. We re-examine the economic intuition offered by Thaler (1980), which suggests that looking at the discount side, relative to the regular price, may be seen as a gain, not producing the desired effect.[3] In addition to the lack of testing for the sunk cost effect in the loss region, much of the existing literature suffers from identification and power limitations. The former is exemplified by Arkes and Blumer (1985), whose positive results are undermined by the price signaling confound—a higher price could indicate better quality. The latter issue, with typical studies having a sample size in the range of 100-500, poses a problem for interpreting the null effects in the context of stark discrepancies regarding the existence of the sunk cost effect reported in the literature.

We overcome all of these issues by investigating both the gain and the loss regions, and avoiding the price signaling confound. This approach allows us to comprehensively evaluate the question of the existence and magnitude of the sunk cost effect. In our first experiment, we test the sunk cost effect and its gain-loss asymmetry in a high-powered online study

[1] The titular sunk costs are costs that cannot be recovered (Friedman et al., 2007). Since there is no action that can undo the previous expense, the neoclassical model predicts that the costs have no impact on future decision making. The failure of this principle is known as the sunk cost effect. There are multiple synonyms for the sunk cost effect in the literature such as sunk cost fallacy and sunk cost bias.

[2] Ashraf et al. (2010) highlight this important issue in the abstract of their paper: "The controversy over how much to charge for health products in the developing world rests, in part, on whether higher prices can increase use, either by targeting distribution to high-use households (a screening effect), or by stimulating use psychologically through a sunk-cost effect."

[3] In Section 2, we outline a model based on reference point dependence and mental accounting, as well as a regret-based model—following Eyster (2002) and Eyster et al. (2021). We explain that our reasoning is consistent with both approaches.

(N=1,806). Second, we conduct a field experiment on YouTube (N=11,328), varying sunk cost in the form of advertisement time, thus applying the novel gain-loss design to study the effect in a policy-relevant setting.

In the online experiment, we randomize the price of a ticket to enter a real-effort task (counting zeros in a sequence) and observe its effect on the time spent on the task (play time). The price ($0.5, $1.5, or $2.5) is realized after the purchase of the ticket and is taken out of the balance acquired by the participant during an earlier task. A critical feature of our design is that the distribution of the entry price is fully disclosed prior to the purchase, allowing us to eliminate the price signaling channel. Furthermore, the participants are informed about the average earnings, which crystallizes the interpretation of the gain and loss regions. In particular, the payoff of the counting game is constructed such that the high price constitutes a clear loss while the low price leads to a gain (see Section 3.4 for details).

Our intervention induced sizable variation in the sunk cost—a $2 difference or 133% of the participants' outside option in a task lasting an average of 13.8 minutes. Moreover, we rely on a sample size of $1,806$, which makes us ex-ante powered to detect small effects. Despite these features, we report only a small magnitude effect of increasing the sunk cost on play time (0.09 SD, or 1.1 minutes, p=0.054).[4] The result is robust to controlling for demographics (p=0.044) as well as restricting the sample to individuals who provided correct answers to the manipulation checks, including the treatment status (p=0.027). The specification combining these two elements also gives a significant result (p=0.023). Interestingly, the observed effect size lies outside of what was detectable in the prior studies. In Section 3.6, we provide evidence against alternative explanations, such as the intervention's impact on efficiency, accuracy, and possibility of income targeting. Finally, we do not find sufficient evidence to reject the symmetry of the sunk cost effect between the gain and loss regions, which is not surprising given the small overall effect.

The field experiment on YouTube retains the main feature of the online study—clear gain and loss regions. Using a within-subjects design, we randomize whether the time until a pre-video ad becomes skippable is shortened (0 s), default (5 s), or extended (10 s). Since the time spent on an ad cannot be recovered, it constitutes a form of sunk cost—and the sunk cost effect implies that user's engagement time with the video increases with the time

[4] Following unequivocal theoretical predictions, in our pre-registration, we specified the sunk cost hypothesis as one-sided. Thus, when evaluating the existence of the sunk cost effect, though not its asymmetry, we provide one-sided p-values.

spent on the ads. Our intervention was administered through a browser extension that the participants previously installed as a part of another study targeting a broader set of platforms.[5] There are several strengths of our approach. First, our intervention is completely seamless and does not change user experience beyond the difference in the minimum ad duration.[6] Second, prior to the experiment, all relevant ads were skippable after 5 seconds, providing a universal reference point. This is crucial in generating predictions based on the model. Lastly, the participants were unlikely to attribute the intervention to the extension due to the obfuscated roll-out of the feature.

The intervention successfully induces significant variation in sunk cost. In particular, participants in the extended condition spend 6.8 seconds or 35% ($p<0.001$) more time watching pre-video ads in comparison to the shortened condition. This difference is meaningful in light of the literature showing high sensitivity of online users to time costs even as small as 2 seconds (e.g., Nah, 2004). We report no overall effect of the intervention on video engagement time. However, this masks an important counterforce. Focusing on the extensive margin decision to watch the video, we find that increasing the sunk cost causes a marked increase in the likelihood of leaving the video page before the end of the ad segment (i.e., before the video starts). In particular, the proportion of people leaving the page during the ad is 5.2 pp. lower (28%, $p<0.001$) in the shortened condition compared to the extended condition.

Lastly, we report the effect of the intervention on the likelihood of clicking the ad, which is our secondary outcome. We find limited evidence that the intervention alters the chance of ad clicking. We report that the likelihood of clicking the ad in the default condition (5 s) is higher than in the shortened condition (0 s) by 0.3 pp. ($p=0.09$). There is no further gain in moving from the default length to the extended length.

Taking the online and field studies together, we offer limited support for the existence of the sunk cost effect. Applying a theoretically well-founded and novel design accounting for the gain-loss divide and the price signaling channel, we demonstrate the sunk cost effect of

[5] The participants were originally recruited to install the extension for an unrelated project by Beknazar-Yuzbashev et al. (2022). They were not specifically informed that we would conduct an intervention manipulating the minimum ad duration on YouTube. Instead, during the initial recruitment, the extension functionality was outlined in high-level terms i.e., the extension can alter page content on three social media platforms (Facebook, Twitter, and YouTube).

[6] Given that YouTube ordinarily shows a counter displaying the number of seconds until skipping becomes possible, in the extended treatment (10 s), we modified it to show 10 seconds in the beginning and count down to 0. As a result, user experience was as seamless as in the other two treatment conditions.

0.09 SD in the online study. The effect is not present in the field experiment. This is likely due to the extensive margin effect, which is negligible in the online study.[7] Our study offers important implications for policy. On one hand, a careful application of designs leveraging losses may enable harnessing the sunk cost effect. However, the overall magnitudes are small and might be dominated by other considerations, such as the extensive margin effects, in many settings like YouTube advertisement. More broadly, our paper serves as a cautionary tale for the naive application of psychological insights to influence economic decision making at scale.

Our paper contributes to the rich, yet often inconsistent, literature on the sunk cost effect contrasting the upfront cost and the subsequent usage of a durable good. The previous field experimental attempts yield null results across a variety of settings. In particular, developmental studies by Ashraf et al. (2010) and Cohen and Dupas (2010) report no sunk cost effect in the context of water purification and mosquito nets, respectively. Similarly, Ketel et al. (2016) finds a null result in an educational setting looking at extracurricular tutorial attendance. On the other hand, Hidalgo et al. (2013), whose study focuses on schools in poor areas of Ecuador, reports higher attendance in schools where students had to pay for a school uniform. All of the above studies only leverage discounts—i.e., are in the *gain* region. This approach originates in Arkes and Blumer (1985), who randomly assigned discounts (\$7, \$2, \$0) to buyers of seasonal theater tickets and find that discounts reduce play attendance. The experimental studies are complemented by research relying on observational data (Ho et al., 2018).

Our paper maintains the key elements of the prior field attempts—which involve a random discount and usage as the outcome variable—while offering improvements in several critical aspects. First, we argue using both mental accounting and regret-based models that the discount design may be ineffective. Second, we employ novel interventions enabling us to explore the sunk cost effect in the loss region. This approach, combined with a large sample size in the online study, allows us to detect effects that remained outside the scope of the previous experiments. Third, our designs eliminate the price signaling channel, which is present in the field attempts.[8] In the online study, we do so by disclosing the distribution of

[7] See Section 4.5 for details of why the extensive margin effect is not present in the online study.

[8] The issue was acknowledged by Ashraf et al. (2010), who asked hypothetical questions about the product quality. Ketel et al. (2016) is the only prior field experiment which addresses the issue at the design stage—the size of the discount is determined by randomly picking a closed envelope. They report an overall null result with a sample size of N=327. We improve on their study by offering a much higher sample size and

the price prior to the purchase decision.

Within the literature on laboratory experiments on sunk cost, there is a sizable body of work on the escalation of commitment to risky investments, first explored by Staw (1976), and subsequently by Martens and Orzen (2021). In the experiment, the participants were asked to make hypothetical two-stage investment decisions, where the viability of the second stage investment was determined by the first-stage outcome. The main finding is that the individuals were more likely to pursue the second-stage investment after receiving a *negative* outcome in the first stage. An important distinction between the literature contrasting the upfront cost and the subsequent usage of a durable good and the studies on the escalation of commitment is the lack of explicit variation in the sunk cost: the cost is only indirectly incurred through a reduced probability of success. Negrini et al. (2022) show that this difference may be pivotal—their results suggest the reverse sunk cost effect, where a higher upfront cost leads to a lower commitment.

Our paper offers a complementary finding to those studies by testing the sunk cost effect in the cost-usage paradigm. This makes our experiments more comparable to the prior field attempts and applicable to assessing the sunk cost effect as a policy instrument. We achieve the latter objective especially through our field experiment, set in a consequential setting of digital advertising.

Our work also adds to the literature on sunk cost outside of both the cost-usage and the escalation of commitment paradigms (Augenblick, 2016; Camerer and Weber, 1999; Friedman et al., 2007; Offerman and Potters, 2006; Phillips et al., 1991; Ronayne et al., 2021; Staw and Hoang, 1995; Weigel, 2018).

Lastly, to our knowledge, this constitutes one of a few economic experiments to leverage a browser extension to study behavior on digital platforms (e.g., Levy, 2021; Beknazar-Yuzbashev and Stalinski, 2022; Aridor et al., 2023), an area of major interest given their welfare impacts (e.g., Allcott et al., 2020) and the lack of transparency. Furthermore, our study supplements an even smaller set of papers which directly alter user experience through a browser extension (Aridor and Goldberg, 2024; Beknazar-Yuzbashev et al., 2022).

The paper proceeds as follows. Section 2 introduces two models to conceptualize the sunk cost effect and the asymmetry between gain and loss regions. Section 3 discusses the experimental design and the results of the online experiment. This is followed by Section 4, which summarizes the field study. Section 5 concludes.

testing for the effect across the gain-loss divide.

2 Theoretical Framework

In this section, we argue that two popular accounts of the sunk cost effect—mental accounting and regret—lead to a prediction where the effect size depends on whether the sunk cost intervention is coded as a gain or a loss. Specifically, we show that experiments relying solely on discounts, like the previous field attempts, may not be able to detect an effect. This underscores the importance of employing the loss intervention in understanding the potency of the effect.

2.1 Setup and Solution

Consider an agent who first decides whether to purchase a good with price p that offers benefits conditional on usage (e.g., gym membership, cars, mosquito nets), then chooses how much time t to spend using it, with the cumulative benefit denoted as $f(t)$ and the cost $c(t)$.[9] Without loss of generality, the utility associated with not purchasing the good is normalized to 0.

Conditional on purchasing, the decision problem of a broad bracketing agent is

$$\max_t f(t) - p - c(t).$$

Assuming $f' > 0$, $f'' \leq 0$, $c' > 0$, $c'' \geq 0$, the optimal solution t_{OPT} is characterized by the following first order condition:

$$f'(t_{OPT}) = c'(t_{OPT}),$$

and hence it is independent of the sunk cost p.

As Thaler (1980) and Thaler (1999) argue, people often put the upfront cost in the same mental account as usage, and chase to "recoup a loss". To formalize the idea, we suppose that an agent with mental accounting (MA) solves the following problem:

$$\max_t v(f(t) - p|R) - c(t), \tag{1}$$

[9] This setup can be mapped into most of the experiments related to the sunk cost effect. In the prototypical experimental framework used by Arkes and Blumer (1985), $f(t)$ can be interpreted as the utility of attending the theater t times and $c(t)$ is the opportunity cost of attending. Similarly, in the case of Ashraf et al. (2010) the payoff $f(t)$ corresponds to the utility from treating water t times, and $c(t)$ is the cost, including future repurchase and hassle of usage.

where

$$v(x|R) = \begin{cases} (x-R) & x \geq R \\ \lambda(x-R) & x < R \end{cases}$$

is a utility function with loss aversion used in Kőszegi and Rabin (2006), and R is a pre-specified reference point.[10] The first order conditions for the agent's problem are:

$$f'(t_{MA}) = c'(t_{MA}) \text{ if } f(t_{MA}) - p > R \tag{2}$$

$$t_{MA} = f^{-1}(R+p) \text{ if } \exists\, \alpha \in [1, \lambda]\ s.t.\ \alpha f'(t^R) = c'(t^R) \tag{3}$$

$$\lambda f'(t_{MA}) = c'(t_{MA}) \text{ if } f(t_{MA}) - p < R. \tag{4}$$

Note that $f^{-1}(R+p)$ is the time that places the agent exactly at the reference point R.

Using the first order conditions above, the optimal choice of t as a function of p is given by:

$$t_{MA}(p) = \begin{cases} \underline{t} & p \leq f(\underline{t}) + R \\ f^{-1}(R+p) & f(\underline{t}) + R < p < f(\bar{t}) + R \\ \bar{t} & p \geq f(\bar{t}) + R \end{cases},$$

where $\underline{t}$ satisfies $f'(\underline{t}) = c'(\underline{t})$ and $\bar{t}$ satisfies $\lambda f'(\bar{t}) = c'(\bar{t})$. If the price is very low, the agent's gain-loss utility $v(\cdot)$ is in the gain region. This renders equation (2) as the relevant first order condition. When p is moderately high, equation (3) becomes the first order condition to satisfy, where the agent just reaches the reference point R. Equation (4) is satisfied if p puts the agent's gain-loss utility in the loss region. Figure 1 plots $t_{MA}(p)$ when $f(t)$ is linear.

2.2 Interpretation

The optimal time varies with price only when the agent is located exactly at the reference point, or the price change is large enough to put the agent in a different utility region. Price changes within regions FOC (2) or (4) (i.e., such that $p \leq f(\underline{t}) + R$ or $p \geq f(\bar{t}) + R$) do not trigger any change in the optimal behavior. This leads to two phenomena. First, the model predicts the traditional sunk cost effect, where a high price in FOC (4) region leads to a higher t than a low price in FOC (2). Second, the model produces what we call the asymmetry of the sunk cost effect, where local price changes within FOC (2) or (4) region do not produce an effect. It is active only when the price change crosses the boundaries

[10] In Section 3.4 and Section 4.4, we provide a discussion of how the reference point may be formed in our online study and YouTube study, respectively.

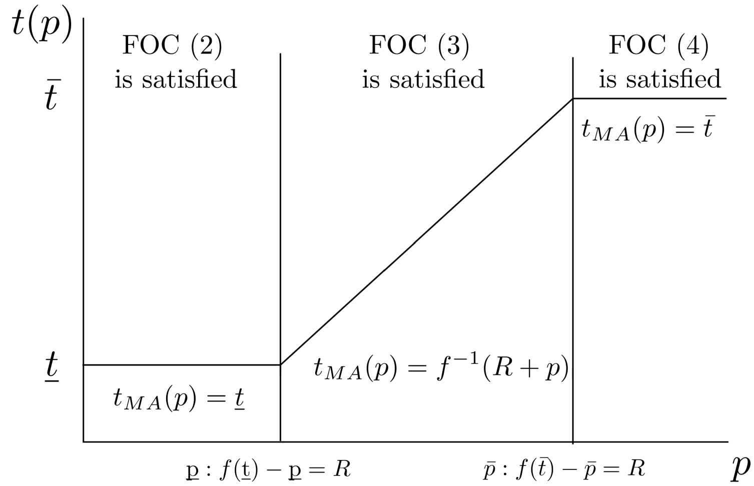

FIGURE 1: THE EFFECT OF SUNK COST FOR AN AGENT WITH MENTAL ACCOUNTING

Note: This figure depicts how an agent with mental accounting preferences adjusts their engagement with a durable good (vertical axis) to different sunk costs (horizontal axis).

separating the FOCs. To understand the asymmetry argument better, consider the reaction to changes in price by an agent facing some price p_0. Suppose that their optimal usage time is t_0 such that $f(t_0) - p_0$ is just above R (i.e. FOC (2) is satisfied). As p decreases, the left hand side of $f(t_0) - p > R$ is even higher, which keeps t_0 optimal. However, as p increases, there exists $p_1 > p_0$ such that $f(t_0) - p_1 < R$. Under p_1, the new usage time t_1 increases to satisfy $f(t_1) - p_1 = R$, or $\lambda f'(t_1) = c'(t_1)$. Intuitively, the agent needs to satisfy another "constraint" of meeting the utility target R. If they are exceeding the target already, lowering the target will not change their action (the constraint remains slack), but if the target increases enough, the usage time goes up so that they get closer to the new target level. The asymmetry argument can be reversed to show that increasing the price does not change t if $f(t_0)$ puts the utility below the reference point R.

To summarize, the mental accounting model predicts that in many cases, only one direction of sunk cost manipulation creates a sizable sunk cost effect—in cases where reaching the utility target is "easy", only an increase in sunk cost will produce the effect, but the same increase is not effective if the utility target is not easily reachable.

2.3 Alternative Specification: Regret

A class of models that augment the utility function with a regret utility can also generate the asymmetric sunk cost effect. One such specification by Eyster et al. (2021) for our decision problem is as follows:

$$\max_t f(t) - p - c(t) - \rho[\max_{\theta \in \Theta} \theta f(t) - p - c(t) - \max\{\theta f(t) - c(t) - p, 0\}], \tag{5}$$

where $\theta \in \Theta \subseteq \mathbb{R}$, $\{1\} \subseteq \Theta$ is the decision maker's "rationales" that only influences the regret utility.[11] To see how this generates the asymmetry, first note that the regret function always remains 0 if the price p is lowered—the ex-ante optimal action of purchasing is still ex-post optimal in that case. Thus, the optimal t remains at the rational level. However, if Θ is bounded from above with $\sup \Theta = \bar{\theta}$, there exists $\bar{p} = \bar{\theta} f(t^*) - c(t^*)$ such that the regret function is strictly negative at the rational optimal level. This causes an upward pressure for the optimal t to reduce regret, creating the sunk cost effect. Crucially, the model predicts the effect only in the loss region, as otherwise the ex-post optimal action would be the same.

3 Online Experiment

3.1 Overview

The goal of the online experiment is to explore the existence and asymmetry of the sunk cost effect in a setting that closely matches the theoretical model. We induce the utility (i.e., $f(t)$) via a real-effort task, which gives a monetary payoff as a function of time. Our intervention for the upfront cost involves randomizing the price of entering the task, where a low price (\$0.5), a medium price (\$1.5), or a high price (\$2.5) are equally likely. The payoff and the prices are calibrated such that the high price clearly constitutes a loss, whereas the low price results in a gain.

3.2 Sample

We recruited a sample of 1,898 US-based adult participants on Prolific, a platform commonly used for recruiting individuals to research studies (e.g., Bursztyn et al., 2023; Oprea, 2024), in March 2022. To avoid potential skewness of the age distribution in our sample, we targeted two age groups separately: 18-34 (one-third) and 35+ (two-thirds). We additionally enforced balance on gender. We report that 1,818 individuals were assigned one of the three treatment groups—606 participants paid \$0.5 for the ticket to enter the real effort task, 606 people paid

[11] Eyster (2002) studies a similar regret function to this, which in our case corresponds to $\Theta = \{1\}$.

$1.5, and 606 paid $2.5. Table 1 demonstrates balance on covariates by treatment group, indicating no significant differences across the conditions. Out of the 1,818 individuals who paid for the ticket, 1,806 people successfully completed the study (99.2%). This matches the planned sample size specified in the pre-registration—approximately 1,800 individuals (600 per treatment arm). The study participants spent, on average, 24 minutes engaging with our Qualtrics survey and earned $7.20, including a fixed fee of $3.25.

Table 1: Balance Table (Online Experiment)

	0.5 (N=606)		1.5 (N=606)		2.5 (N=606)	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Accessed from Mobile	0.16	0.37	0.18	0.38	0.19	0.39
Age	42.46	13.47	41.90	13.15	41.74	13.27
Bachelor or Above	0.59	0.49	0.59	0.49	0.60	0.49
High Income	0.43	0.50	0.39	0.49	0.43	0.50
White	0.80	0.40	0.82	0.38	0.82	0.39
Hispanic	0.06	0.23	0.07	0.25	0.06	0.24
West	0.19	0.40	0.20	0.40	0.19	0.39
Midwest	0.21	0.41	0.21	0.41	0.21	0.40
South	0.36	0.48	0.36	0.48	0.36	0.48
Northeast	0.23	0.42	0.23	0.42	0.24	0.43

Note: This table compares characteristics of participants assigned to different treatment groups. The table presents means and standard deviations. “Accessed from Mobile” takes a value of one if the participant accesses our study from a mobile device (e.g., Android and iOS devices). “Age” variable is the participants’ age at the time of the survey (March 2022). “Male”, “White”, “Hispanic” are dummy variables that take a value of one if the participant identifies themselves as such. “Bachelor or Above” takes a value of one if the participant has at least a bachelor degree. “High Income” takes a value of one if the participant’s reported annual household income in 2021 is above $75,000. The four regional variables take a value of one if the participant reports living in a state corresponding to the region.

3.3 Experiment Design

Preliminaries Individuals recruited for the study played two games, with the experimental variation applied to the price of a ticket to enter the second one. In the very beginning, we displayed a “warning” to the participants outlining the payoff risks associated with the survey and offered them an opportunity to leave the study.[12] In particular, we informed them that while the advertised average bonus is $4, it can be as low as $0.25. The purpose of the screen is to induce a sense of agency and responsibility, which is an important driver of regret (Zeelenberg et al., 1998; Ordónez et al., 2000). After collecting demographics, we directed everyone to the first game, which was designed to give participants an endowment required for the later stage. We chose a real effort task to minimize the likelihood that they

[12] In the online appendix, we provide the wording of instructions and questions used in our survey.

perceive the endowment as a windfall gain. In each round, participants needed to memorize a sequence of buttons that light up and then press them in the same order. The game lasted 120 seconds, with the participants earning either \$2.75 or \$1.00, depending on their performance in the game. We set a low bar for the high reward (only two participants did not qualify) to maximize the subject pool for the experimental intervention.[13]

Counting Game After learning how much they earned from the first game, the participants proceeded to the main part of the study—the counting game—which required payment for entry. They were provided with information on the rules of the game. Figure 2 depicts an example of the counting game screen. The objective of the game was to earn points by counting the number of zeros in the sequence presented. Points were given only for rounds in which participants answered correctly (entered the correct number of zeros in the text box). The participants were told that the number of points per round decreases over time. Lastly, they were also informed that they are allowed to stop playing at any point by clicking "Finish" on the game screen. If they played for more than 60 minutes, the game automatically terminated when they clicked to proceed to the next round. Upon leaving, the game points were translated into the bonus payment according to the following formula:

$$\text{Bonus} = \frac{\text{Points}}{1000} - \text{Price}.$$

Treatments After the discussion of the rules of the counting game, the participants were directed to a "purchasing screen", where they paid for the ticket to enter. There were three possible prices ($p_L = \$0.5$, $p_M = \$1.5$, $p_H = \$2.5$), each equally likely to be selected. The distribution was fully disclosed to preclude the possibility that the randomized price could be a meaningful signal, e.g., about the quality or the value of the game. Another vital element of this screen was the information about the average bonus that participants get from the study (\$4), which helped in forming the reference point. After the price had been randomized, we informed everyone of the cost of the ticket and their remaining balance. Immediately afterwards, we included two manipulation checks—questions asking the participants about the price they ended up paying and the average bonus for the study. The game followed right after the checks. The primary outcome variable is the time an individual spent playing

[13] Note that the two participants who did not qualify for the high reward were not assigned a treatment group.

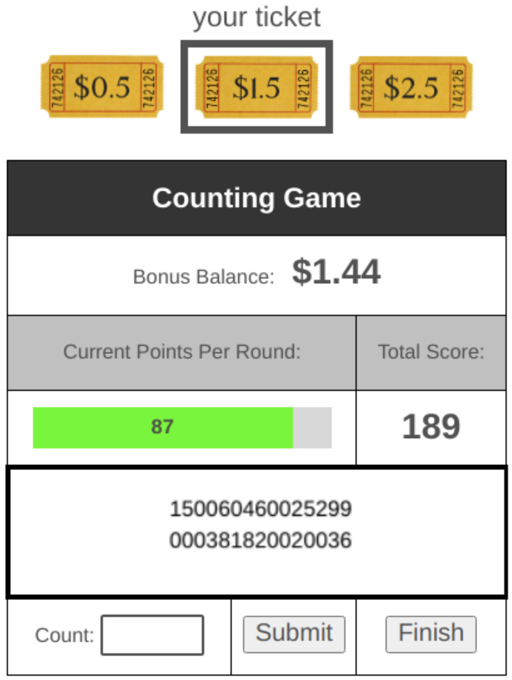

FIGURE 2: GAME SCREEN

Note: This figure depicts the main task screen (counting game) for the online experiment, where participants earn money by counting the number of zeros in a sequence of numbers. The entry price (sunk cost) for the task is shown at the top of the screen. "Bonus Balance" row keeps track of the payoff from the whole experiment ($2.75 - entry cost + payoff from the task). The payoff per round and the total payoff (in an experimental currency with an exchange rate of 1000 to $1) are shown in the next row.

the counting game ("play time").

3.4 Hypotheses

In this section, we make predictions about how the optimal play time $t(p)$ depends on the price to enter the counting game. Assuming that individuals convert monetary payoffs linearly to utility, their decision problem about play time is isomorphic to the model outlined in Section 2. In particular, $f(t)$ is the cumulative payoff function for the counting game, $c(t)$ is the cost of performing the task up to time t (inclusive of the opportunity cost), and R is the reference point. Figure 3 depicts the payoff function $f(t)$ for the average participant in the experiment.

Under the mental accounting model (1), if $f(t(p_M)) - p_M > R$, it follows from the argument in Section 2 that we have $t(p_M) = t(p_L)$ for any $p_L < p_M$, and $t(p_M) < t(p_H)$ as long as p_H is high enough to make the optimal solution satisfy first order condition (2) or (3) i.e. $f(t(p_M)) - p_H < R$. On the other hand, if $f(t(p_M)) - p_M < R$, we expect $t(p_L) < t(p_M) = t(p_H)$. This prediction gives a formal foundation to the asymmetry hypothesis. We

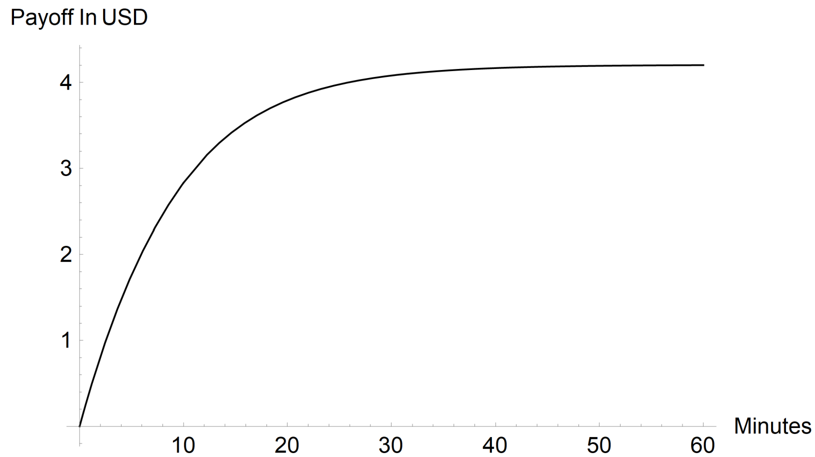

FIGURE 3: CUMULATIVE PAYOFF FUNCTION

Note: The figure depicts the cumulative payoff for the counting game. To construct this graph we assumed the average efficiency of 4.47 rounds per minute (this quantity is based on the average efficiency in the pilot experiment, which we used to calibrate the main study).

design the study with an explicit goal to narrow down potential values of R. Specifically, the participants were told at the point of recruitment and reminded before the counting game that the average payoff for the whole study is \$4. Given that they earned \$2.75 prior to the counting game, we assume that their reference point is $R = \$1.25$. Therefore, our model states that the direction of the asymmetry depends on whether $f(t(p_M)) - p_M$ is above or below \$1.25. To ensure the directionality of the effect, the cumulative payoff curve is calibrated such that reaching $R = \$1.25$ requires high effort (19.2 minutes based on the pilot study) when the upfront cost is \$2.5. The time effort is much higher than the average time spent, making this a loss intervention (see Section 3.5). At the same time, reaching the same R requires below-the-average effort when the upfront cost is either \$1.5 or \$0.5. Consequently, the model predicts $t(p_L) = t(p_M) < t(p_H)$.

We pre-registered two empirical hypotheses that encompass the theoretical predictions. First, we test the existence of the sunk cost effect—specified in Hypothesis 1. Here, we use the largest variation in sunk cost present in our experimental design. Given the clarity of the theoretical predictions, in our pre-registration, we specified the sunk cost hypothesis as one-sided. Thus, when evaluating the existence of the sunk cost effect, we provide one-sided p-values.

Hypothesis 1 *(The Sunk Cost Effect) The average play time in the high price group is greater than in the low price group:* $E[t(p_H)] > E[t(p_L)]$.

Second, we test the asymmetry hypothesis that is based on our modeling framework. We choose this particular form of Hypothesis 2, to ensure that the null hypothesis of the effect symmetry can be rejected using the evidence provided by the experimental data.

Hypothesis 2 *(The Asymmetry Hypothesis) The difference in average play time between the high price group and the medium price group is not equal to that between the medium price group and the low price group:* $E[t(p_H)] - E[(t(p_M)] \neq E[(t(p_M)] - E[t(p_L)]$.

3.5 Results

Table 2 summarizes the results of the online experiment. Overall, we find support for the existence of the sunk cost effect (Hypothesis 1), although its magnitude is small. Specifically, the participants whose entry price was \$2.5 played 1.1 minutes (or 0.09 SD) longer than those charged \$0.5 (Column 1, p=0.053). The result is robust to including demographic controls[14] (Column 3, p=0.044) as well as restricting the sample to individuals who provided correct answers to the manipulation checks[15] (Column 2, p=0.027). The combined specification, including both the restriction on the manipulation checks and the demographic controls, also gives a significant effect (p=0.023). Additionally, the reported effect sizes are similar across specifications (1)-(4), ranging from 1.1 to 1.3 minutes.

Additionally, we do not find evidence supporting the asymmetry hypothesis (Hypothesis 2). The Wald test for the inequality of the effect below and above the medium price of \$1.5 fails to reject the symmetry with a p-value of 0.93. Table 2 shows the result across our specifications. The finding is not surprising given the small magnitude of the overall sunk cost effect.

3.6 Discussion and Robustness

3.6.1 Intervention Strength

Given the small magnitude of the detected sunk cost effect, it is crucial to underscore the intervention strength. At the time of our experiment, Prolific guaranteed its respondents

[14] Demographic controls include the variables reported in Table 1, as well as dummy variables for the regions of the US.

[15] We asked participants about their assigned price and the average total payoff which we expressly told them about. Out of the 1,806 participants, 99.0% correctly identified their treatment group and 97.8% recalled the average bonus. Since the attention checks took place after the treatment assignment, we test for differential attrition. The F-test gives a p-value of 0.418.

Table 2: Existence and Asymmetry of the Sunk Cost Effect

	Play Time (Minutes)			
	(1)	(2)	(3)	(4)
Price = 1.5	0.584	0.659	0.652	0.711
	(0.656)	(0.664)	(0.655)	(0.663)
Price = 2.5	1.063*	1.285**	1.123**	1.329**
	(0.656)	(0.666)	(0.654)	(0.664)
Attention		X		X
Controls			X	X
Mean	13.706	13.849	13.706	13.849
Wald P-value	0.93	0.98	0.87	0.93
N	1,806	1,749	1,806	1,749

Note: This table reports the results on the existence and asymmetry of the sunk cost effect. Column 1 shows a regression of game play time (in minutes) on a dummy variable equal to one if a participant was assigned the medium price ($1.5) and a dummy variable equal to one if a participant was assigned the high price ($2.5). Column 2 additionally includes controls. Column 3 restricts the sample to those participants who correctly answered manipulation checks about the price they were assigned and the average bonus payment in the study. Column 4 extends the previous specification by including controls. Robust standard errors are parenthesized. Following Hypothesis 1, we rely on one-sided p-values when reporting the significance of results in the row "Price=$2.5". Additionally, we report p-values based on a Wald test against a null hypothesis outlined in Hypothesis 2 in a row "Wald P-value". *p<0.1; **p<0.05; ***p<0.01.

an hourly wage of at least $6.5, which, taking into account the average play time of 13.7 minutes, translates into an outside option worth $1.5. Therefore, our sunk cost intervention between the low price and the high price constituted 133% of their outside option. Moreover, previous experimental literature shows that loss aversion—a crucial theoretical driver of the sunk cost effect—is significant within the stake sizes we consider (Bleichrodt and L'haridon, 2023; Brown et al., 2024). Additionally, taking into account the sample size, we were ex-ante powered to detect the sunk cost effect as small as 0.1 SD. This is important as the size of the sunk cost effect that we detect in our study is such that the previous literature that reported null results would not have been able to detect.

Lastly, our experiment is explicitly designed with a goal of enabling mental accounting in a way that is in line with the theoretical model. In particular, we repeatedly displayed information about the average earning, and the entry price was saliently featured in the counting game screen (see Figure 2).

3.6.2 Effect on Efficiency and Accuracy

Our pre-registered outcome of interest is play time, which closely mirrors the theoretical models. However, since the payoff from the counting game also depends on how efficient participants are in performing the task, we test if the efficiency, defined as the number of rounds that a participant correctly solves per minute, differs across the conditions. We find that the difference is negligible—individuals in the \$2.5 group solved 4.33 rounds per minute, as opposed to 4.35 rounds in the \$0.5 group (p=0.79). Similarly, we find no differences in the accuracy, defined as the proportion of correctly solved rounds, between the conditions (p=0.53). Table A1 and Table A2 present the results.

3.6.3 Income Targeting

Our experimental design features monetary payoff as a proxy of utility. Therefore, one might argue that participants may work until a certain income target is met (Camerer et al., 1997). Using our modeling framework, this can be expressed as a mental accounting agent having the following utility:

$$v(x|R) = \begin{cases} 0 & \text{if } x < R \\ \bar{v} & \text{if } x \geq R \end{cases}$$

for a sufficiently high $\bar{v}$. What sets our model apart from the income targeting is the implication on the total earnings—under the income targeting model, the agent's total earning $f(t_{IT}) - p$ is fixed regardless of the initial upfront cost p. We reject this hypothesis, as the difference in total earnings between \$0.5 and \$2.5 groups is \$1.96 (p<0.01). Table A3 demonstrates the result.

3.6.4 Capped Play Time

As discussed in Section 3.3, if participants play the counting game for more than 60 minutes, the game would terminate when they click to proceed to the next round. In our sample, there are seven individuals whose play time exceeds 60 minutes, likely due to the delay in proceeding to the next round after the 60-minute mark. We verify that our results are robust to capping play time of these individuals at 60 minutes. Table A4 replicates Table 2 with the capped play time. All of the results remain unchanged.

4 Field Experiment

4.1 Overview

We aim to parallel the structure of the online study in the context of video engagement behavior on YouTube, using ad duration as a form of sunk cost. Ads preceding videos on YouTube can be one of two types—skippable or unskippable. As the name suggests, users have to watch unskippable ads in full. On the other hand, skippable ads allow users to skip the ad and move to the video by clicking "Skip" button, which universally appears after 5 seconds (as of October 2022).[16] Our study relies on a browser extension which targets skippable ads, and has the capacity to seamlessly vary the time until a user is allowed to skip such an ad (henceforth—"minimum ad duration" or MAD). The existence of the pre-treatment level of MAD provides a perfect testbed for the theory, which requires a clear expectation of the "normal" price.

4.2 Sample Description

Participants in our field experiment were recruited from the pool of existing users (as of October 2, 2022) of a browser extension called *Social Media Research*, which broadly recorded users' browsing behavior on YouTube (and other social media platforms) and could alter the structure of YouTube pages, including the minimum ad duration. The extension operated on desktop versions of Chromium-based browsers and Firefox and was listed both on Chrome Web Store and Firefox Browser Add-ons. The users originally installed the extension via Twitter ads for the purpose of an unrelated study by Beknazar-Yuzbashev et al. (2022). We enabled the sunk cost intervention on October 2, 2022, with the experiment lasting four weeks, as pre-registered. Please note that the original consent form was structured such that both studies are fully covered.

Within this period, the extension recorded 410 users active on YouTube who saw at least one video with a skippable ad.[17] Table A5 demonstrates summary statistics outlining the demographics of the users, which were collected during an onboarding survey. To mitigate the concern of selection into the study based on preferences for ad length and video engagement, the original recruitment materials stated the study purpose only in very general terms, as "improving user experience on social media as well as understanding user preferences and decision making".

The extension recorded a total of 11,486 unique video-user pairs with a skippable ad at the beginning of the video.[18] This is consistent with our pre-registration, where we

[16] https://support.google.com/youtube/answer/2467968?hl=en, accessed: 2023-04-04.

[17] Users are identified based on unique user IDs assigned at the time of extension installation. See Beknazar-Yuzbashev et al. (2022) for details.

[18] An ad is considered to occur at the beginning of the video if (1) it starts within the first 10 seconds since the page load and it is the first ad of the video, or (2) it starts within 120 seconds of the page load and it is the second ad of the video. At the time of the experiment, the maximum number of ads at the beginning of the video was two. The outlined approach is needed as the extension relies on timestamps of events, such as

highlighted that we will focus on cases where the intervened skippable ad occurs at the front of the video. Our final sample for analysis is comprised of 11,328 observations (98.6%), after discarding observations which correspond to a live stream. Live streams are a special category of YouTube videos and their length is hard to interpret. In Section 4.5.4, we show robustness of the main results to their inclusion. Table 3 reports balance by treatment group for the final sample of videos. We demonstrate that the sample is well-balanced on observable characteristics of the videos, such as video length, ad length, and topic categories.

4.3 Experiment Design

4.3.1 Intervention

We opt for a within-subjects design in our field experiment—the minimum ad duration (MAD) for the first skippable ad is randomized to be one of $\{0, 5, 10\}$ seconds, with each realization equally likely. In the case of the *shortened* treatment (MAD=0), the skip button appears immediately (Figure B1). The *default* treatment (MAD=5) allows the user to skip an ad after 5 seconds (Figure B2)—here, user experience does not deviate from the standard YouTube experience (outside the study). Lastly, in the *extended* treatment (MAD=10), the user has to wait for 10 seconds (Figure B3) before being able to skip the ad. Given that YouTube ordinarily shows a counter displaying the number of seconds until skipping becomes possible, we modified it to show 10 seconds in the beginning and count down to 0. As a result, user experience in the MAD=10 group was as seamless as in the other two treatment conditions. To minimize the concern that a participant might refresh the page until they get a draw of MAD=0, we enforce the same treatment assignment if the user refreshes the page or transitions to another video within 8 seconds of the page load. In Section 4.5.4, we show robustness of the main results to excluding such multi-video observations.

Importantly, the participants were unlikely to associate the intervention with the extension. They were recruited for a separate study that improves user experience on social media by filtering out toxic content several months prior (Beknazar-Yuzbashev et al., 2022). Relying on the broad permissions for modifying experience on Facebook, Twitter, and YouTube, we unexpectedly enabled the minimum ad duration intervention. As a result, it is plausible that the participants attributed the ad intervention to YouTube piloting new ad settings and thought that the original extension intervention is ongoing.

page load or ad load, and these may occur with a small delay.

Table 3: Balance Table (YouTube Study)

	Low (N=3705)		Medium (N=3803)		High (N=3820)	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Video Length (Minutes)	21.95	74.02	21.81	75.13	21.20	60.65
Video Length Known	1.00	0.03	1.00	0.02	1.00	0.04
Treated Ad Length (Minutes)	0.99	1.91	1.07	3.59	1.02	2.61
Ad Length Known	0.91	0.29	0.90	0.31	0.89	0.31
Category Known	0.83	0.37	0.85	0.36	0.85	0.36
Category: Music	0.20	0.40	0.20	0.40	0.21	0.41
Category: Entertainment	0.12	0.33	0.12	0.33	0.12	0.33
Category: Politics	0.05	0.22	0.05	0.23	0.06	0.23

Note: This table compares characteristics of video assigned to different treatment groups. The table presents means and standard deviations. Video length is retrieved from YouTube API, or the recording from our browser extension when the former is not available. Ad lengths come from the reading of the browser extension, and the video category variable is retrieved from YouTube API. "Video Length Known" takes a value of one if at least one of the video length measures is available for the video, and "Ad Length Known" is correspondingly defined. "Category: X" variables take a value of one if the registered video category includes the word "X".

4.3.2 Outcomes

Pre-Registered Outcomes We pre-registered one primary outcome: video engagement time. Below, we provide details on how it was measured by the browser extension. We pre-registered an additional secondary outcome—the likelihood of clicking an ad. Lastly, our pre-registration specifies that we will look at heterogeneity of the treatment effects by the length of the video, presence of long music videos, and duration of exposure to the sunk cost intervention. Any additional analysis reported in the paper is exploratory and intended to test the robustness of the results and interpret the effect sizes.

Time Measurement In order to conceptualize the way in which the extension measures user's interactions with videos, consider three timestamps associated with each user-video pair: *InTime*, *OutTime* and *VideoTime*. We describe each of them in turn (see Figure 4 for a visualization). First, *InTime* is the best estimate of the loading time of the video page (the timestamp of the first event for that video). Second, *VideoTime* refers to the loading time of the main video after the segment of initial ads.[19] Third, *OutTime* corresponds to the the first indication that the user is not watching the video. For example, it is triggered

[19] In order to identify the initial ad segment, it is required that no more than 10 second of cumulative non-ad time occurs till the end of the segment. The non-zero allowance is needed as the extension relies on timestamps of events, such as video load or ad load, and these may occur with a delay. The 10-second rule does not apply to cases where engagement occurs after the last ad to account for potential short-lived engagement with the video.

when the user switches to another page (both within and outside of YouTube), minimizes the tab, or closes the browser. As pre-registered, any user activity after their first transition out of the video was not considered even if they came back to it later.

With these three timestamps, we define the sunk time to be the duration between *InTime* and *VideoTime*, and the engagement time to be the duration between *VideoTime* and *OutTime*. We cap the sunk time at the total recorded ad length, and the engagement time at the recorded video length.[20]

Specification To estimate the effect of the experimental intervention on engagement, our primary outcome, we use the following pre-registered specification:

$$Y_{ij} = \alpha_i + \sum_{k \in \{5,10\}} \beta_k I(D_{ij} = k) + \epsilon_{ij}, \tag{6}$$

where Y_{ij} is the video engagement for user i watching their j-th video in our sample, D_{ij} is the minimum ad duration associated with that video, α_i is the user fixed effect and ϵ_{ij} is an iid random error. Moreover, we use the same specification to report the effect of the intervention on the likelihood of clicking the ad—our secondary outcome.

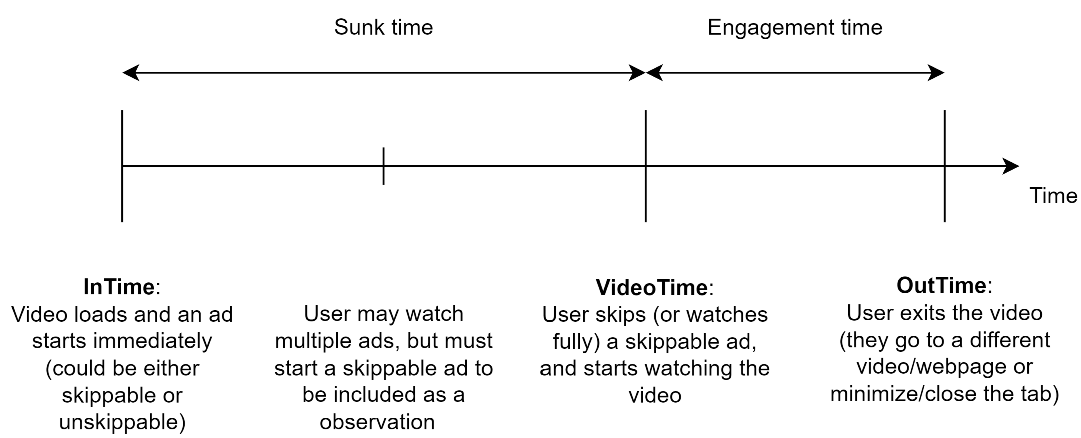

FIGURE 4: SUNK TIME AND ENGAGEMENT TIME

Note: This is a graphical representation of how sunk time and engagement time are measured in the YouTube experiment.

[20]We do not cap in the rare cases where the corresponding length is unavailable.

4.4 Hypotheses

Since the minimum ad duration is universally set by YouTube to 5 seconds, we posit that people's reference point R in the mental accounting model is based on the minimum ad duration of 5 seconds. Consequently, our predictions for the field study mirror those for the online experiment.

Prediction 1 *The average engagement time in the extended minimum ad duration group is greater than in the shortened minimum ad duration group:* $E[t(MAD = 10)] > E[t(MAD = 0)]$.

Prediction 2 *The difference in the average engagement time between the extended minimum ad duration group and the normal minimum ad duration group is different from that between the normal minimum ad duration group and the shortened minimum ad duration group:* $E[t(MAD = 10)] - E[t(MAD = 5)] \neq E[t(MAD = 5)] - E[t(MAD = 0)]$.

Under the main specification (6), the model predictions can be tested by the following empirical hypotheses:

Hypothesis 3 *The average engagement time in the extended minimum ad duration group is greater than in the shortened minimum ad duration group:* $\beta_{10} > 0$.[21]

Hypothesis 4 *The difference in the average engagement time between the extended minimum ad duration group and the normal minimum ad duration group is different from that between the normal minimum ad duration group and the shortened minimum ad duration group:* $\beta_{10} \neq 2\beta_5$.

Heterogeneity Analysis In addition to the main hypotheses, we pre-registered our intention to explore heterogeneity of the effect across two dimensions: video length and data collection period. For video length, we categorize observations by video length in the following way—shorter than 3 minutes, 6 minutes, 10 minutes, and 20 minutes. The heterogeneity analysis by video length is motivated twofold. First, there is suggestive evidence from the previous literature that the effect of sunk cost wanes over time (Arkes and Blumer, 1985; Ho et al., 2018). Second, we expect that the longest videos will add a lot of noise to the

[21] Due to the clarity of the theoretical predictions, we pre-registered using one-sided p-values to test this hypothesis.

estimation. Additionally, we look at the subsample of videos where we exclude music videos, which might be intended to be played in the background. We identify those videos as ones that are more than 10 minutes long and are in the "music" category according to YouTube API data. Finally, we test whether the duration of exposure to factors affecting the sunk cost changes the strength of the effect. In particular, we compare the effect size in the first two weeks and the last two weeks.

4.5 Results

The results of the main specification estimating the treatment effect of the minimum ad duration (MAD) intervention are summarized in Table 4.

4.5.1 Effects on Sunk Time

We begin by reporting on the effects of MAD intervention on the sunk time. We find that the intervention meaningfully affects the sunk time—those with a longer minimum ad duration indeed spend more time watching ads. Specifically, we find that having the extended treatment (MAD=10) increases the average sunk time by 6.8 seconds ($p<0.001$) relative to the shortened treatment (MAD=0). The induced difference between the default group (MAD=5) and the shortened group (MAD=0) is equal to 2.5 seconds ($p=0.091$). These results are summarized in Column 1 of Table 4.

The relatively low statistical significance between MAD=5 and MAD=0 is largely driven by outliers. Excluding the top one percentile of sunk time considerably reduces the standard error and yields a difference in the sunk time of 3.0 seconds ($p<0.001$) between MAD=5 and MAD=0. The corresponding difference in sunk time between MAD=10 and MAD=0 groups equals 6.2 seconds ($p<0.001$). Table A14 provides the full set of results using this sample.

4.5.2 Effects on Engagement Time

Primary Outcome Results In Column 2 of Table 4, we document the results on video engagement time. Our point estimates indicate that participants receiving the extended treatment (MAD=10) engaged with the videos for 1.9 seconds less, on average, than those with the shortened treatment (MAD=0). As the sign of the point estimate is opposite to the hypothesized effect, we cannot reject Hypothesis 3. Moreover, we do not find support for

Table 4: YouTube Study: Main Results

	Sunk Time	Engagement Time	Engagement = 0	Ad Click
	(1)	(2)	(3)	(4)
MAD = 5	2.531*	−4.921	0.033***	0.003*
	(1.500)	(9.106)	(0.009)	(0.002)
MAD = 10	6.829***	−1.850	0.052***	0.002
	(1.586)	(9.808)	(0.009)	(0.002)
Sample	Full	Full	Full	Full
Individual FE	X	X	X	X
Mean	22.87	228.64	0.217	0.005
Wald P-value	NA	0.623	NA	NA
N	11,328	11,328	11,328	11,328

Note: This table summarizes the effect of the minimum ad duration intervention on various outcomes of interest using the main regression specification that includes dummy variables for each treatment status (rows MAD=5 and MAD=10) and individual fixed effects. Column 1 reports the effect on sunk time. The sunk time is defined as the interval between the best estimate of the loading time of the video page and the end of the initial ad segment. Column 2 reports the effect on video engagement time. The engagement time is the interval between the end of the initial ad segment and the first indication that the user is not watching the video. Details are provided in Section 4.3.2. Column 3 reports the effect on no engagement behavior, defined as having an engagement time of 0. Column 4 reports the effect on ad click rate. All specifications use the full sample. For the result on engagement time, we also report p-values based on a Wald test against a null hypothesis outlined in Hypothesis 4 in a row "Wald P-value". Robust standard errors are parenthesized. *p<0.1; **p<0.05; ***p<0.01.

Hypothesis 4—the effects in the gain and the loss region do not significantly differ (p=0.62, Wald test).

Discussion Outline Below, we explore factors that might have contributed to the overall null result. First, we analyze the importance of the extensive margin of the engagement decision—whether to watch the video at all. Second, we discuss whether the intervention strength was sufficient to rule out meaningful effect sizes, especially for shorter videos. To that end, we leverage the pre-registered angle of heterogeneity by video length. Lastly, we discuss the engagement results in the field study in light of the results of the online experiment (Section 3.5).

Extensive Margin of Engagement Column 3 of Table 4 documents a significant increase in the likelihood of users giving up on the video (i.e., leaving before it starts) when faced with a longer sunk time. Specifically, in the extended treatment (MAD=10), the proportion of users not starting the video was higher than in the shortened treatment (MAD=0) by 5.2 pp. (p<0.001), a rise of 28%. We also find a significant effect when comparing the default treatment (MAD=5) to the shortened treatment (MAD=0). The likelihood of no engagement was higher in the former group by 3.3 pp. (p<0.001), an increase of 18%. The extensive margin analysis suggests that "frustration" with the upfront cost may play a role in the users' engagement decisions, which goes against finding an overall sunk cost effect. The size and significance of the extensive margin effect is consistent across all video length categories (Table A10).

Intervention Strength The manipulation in the sunk cost is done through manipulating the minimum ad duration, which ranged from 0 to 10 seconds. First, we document that our intervention is strong in relative terms. Specifically, the extended treatment translates into a 35% increase in the observed upfront cost. Second, one may worry about individuals being insensitive to small changes in time cost in absolute terms. However, previous literature indicates that online users are extremely sensitive to time cost, with 1 second being the limit for user's flow of thought to stay uninterrupted, 2 seconds being the limit for user's willingness to wait for a simple information retrieval task, and 10 seconds sufficient to disengage them completely (Nah, 2004). The significant effects reported on the extensive margin decisions using our sample further corroborate this point.

Effect Size and Heterogeneity by Video Length We continue our discussion of the effect size by looking at its heterogeneity with respect to the video length. First, Table A8 indicates that our intervention successfully induced exogenous variation in sunk time regardless of the video length. Table A9 demonstrates that the overall null effect of the intervention is stable across the categories. As expected, our estimates are less noisy for shorter videos. This allows us to rule out sunk cost effect of meaningful magnitude. For example, our analysis for videos shorter than 10 minutes (the median length is 9.8 minutes) rules out sunk cost effects grater than 2.4 seconds between MAD=10 and MAD=0. The induced variation in the sunk time for that subsample is 5.9 seconds, hence we rule out effects larger than 40% of the cost difference. The corresponding analysis for videos shorter than 6 minutes rule out effects larger than 68% of the cost difference. Table A7 offers additional insights into effect sizes that we can rule out. The table provides treatment effects for an alternative engagement metric—the proportion of the video watched. Here, we rule out sunk cost effect greater than 0.4 pp. of the share of the video watched. Simultaneously, we report a 15.7 pp. intervention-induced variation in the proportion of the ads watched.

Comparison to the Online Study We now contrast the results to those obtained in the online experiment, where we find a small yet significant sunk cost effect. Our analysis reveals that the significant negative extensive margin effect in the field study is a likely driving force behind the difference. Note that the channel is negligible in the online study—participants face a substantial loss of not getting the fixed fee of $3.25 by leaving the study without finishing the counting game. Overall, settings where sunk cost effect can be employed as a policy intervention vary by the likely potency of the extensive margin effect, with the risk higher for settings where people have a strong prior about the default cost, or are aware of the associated temporal or cross-sectional variation. Taking the results of the two studies together, we conclude that (1) the magnitude of the sunk cost effect is small even in the absence of the extensive margin effect, (2) the extensive margin effect likely dominates any sunk cost effect in policy-relevant settings, and (3) asymmetry of the sunk cost effect, while theoretically well-founded, is not the first order concern in light of the overall magnitude of the effect.

4.5.3 Ad Clicks

We now proceed to the analysis of the impact of the experimental intervention on the likelihood of clicking the ad. The results are reported in Column 4 of Table 4. We find that users who experienced the default treatment (MAD=5) are more likely to click the ad in comparison to the shortened (MAD=0) treatment by 103% (p=0.09). The click rate in the MAD=10 group is higher than in the MAD=0 group by 83%, but the result is statistically insignificant (p=0.117). In other words, if skipping the ad is immediately an option, users are less likely to click. However, prolonging the MAD from 5 to 10 seconds has no additional impact on the outcome. This muted effect is consistent with the counterfactual simulation by Chiong et al. (2024). Ad clicks are one of the key metrics relevant for YouTube's ad revenue. We demonstrate that when attempting to leverage temporal sunk cost effect to increase engagement on a digital platform, the impact on behavior while waiting (including clicks, search for information etc.) should be carefully taken into account, as it may be significantly affected in ways that matter for business profitability.

4.5.4 Robustness

We perform a number of additional robustness checks. First, we report that our results are consistent with the two-stage least square specification, where the sunk time is instrumented by the treatment status (Table A6). Furthermore, we show that our results are unchanged when excluding observations where the user refreshed the page or switched to another video within 8 seconds, which triggered assignment of the same treatment status again (Table A12). We also demonstrate that inclusion of live streams in the sample does not alter our main results (Table A13).

Lastly, we report additional heterogeneity results specified in our pre-registration. First, Column 6 in Tables A8, A9, and A10 show that excluding long music videos, which are likely played in the background, does not change the results. We also do not find evidence that treatment effects on video engagement are heterogeneous by intervention duration (the last two weeks vs. the first two weeks). The results are reported in Table A11.

5 Conclusion

We report two experiments that probe the potential of the sunk cost effect to be applied in practice, exploring its magnitude and economic significance. Our approach offers multiple methodological improvements over the prior studies measuring the effect of upfront cost on the subsequent usage of a good or service. Specifically, our intervention enables us to investigate the sunk cost effect in the loss region and allows us to eliminate the price signaling channel. The former is particularly important in light of theoretical predictions, offered by models based on reference point dependence and regret, indicating the gain-loss divide as a potential reason for the conflicting results provided by the previous attempts.

The results of our online study point towards a small magnitude sunk cost effect that may not be robust across settings. Despite a novel approach to testing for the sunk cost effect in the field—with a clear gain and loss region relative to the 5-second reference wait time—we do not report a significant result. This corroborates the previous field experimental findings that the sunk cost effect might be difficult to leverage in policy settings. Our paper explores the impact on the extensive margin as one of the contributing factors. In particular, we observe that a higher sunk time results in a greater share of people leaving before the video starts. These type of considerations can be dominant in some, mostly field settings, while remaining insignificant in others (e.g., our online study). We hope that future academic work will further illuminate the interaction between the intensive and extensive margin effects of varying sunk costs. In particular, we encourage further exploration in digital contexts, where platforms have more flexibility in leveraging different forms of sunk cost. More broadly, our paper cautions against leveraging psychological factors to influence economic decision making when the exact mechanism behind the effect is unknown. More rigorous inquiry into the nature of intuitive "biases" such as the sunk cost effect is needed before applications can be considered.

References

Allcott, H., L. Braghieri, S. Eichmeyer, and M. Gentzkow (2020): "The Welfare Effects of Social Media," *American Economic Review*, 110, 629–676.

Aridor, G., Y.-K. Che, and T. Salz (2023): "The Effect of Privacy Regulation on the Data Industry: Empirical Evidence from GDPR," *RAND Journal of Economics*, 54, 695–730.

Aridor, G. and S. Goldberg (2024): "Does Transparency Matter in Opaque Product Markets? Evidence from Privacy," *Working Paper*.

Arkes, H. R. and C. Blumer (1985): "The Psychology of Sunk Cost," *Organizational Behavior and Human Decision processes*, 35, 124–140.

Ashraf, N., J. Berry, and J. M. Shapiro (2010): "Can Higher Prices Stimulate Product use? Evidence from a Field Experiment in Zambia," *American Economic Review*, 100, 2383–2413.

Augenblick, N. (2016): "The Sunk-Cost Fallacy in Penny Auctions," *The Review of Economic Studies*, 83, 58–86.

Beknazar-Yuzbashev, G., R. Jiménez Durán, J. McCrosky, and M. Stalinski (2022): "Toxic Content and User Engagement on Social Media: Evidence from a Field Experiment," *Available at SSRN*.

Beknazar-Yuzbashev, G. and M. Stalinski (2022): "Do Social Media Ads Matter for Political Behavior? A Field Experiment," *Journal of Public Economics*, 214, 104735.

Bleichrodt, H. and O. L'haridon (2023): "Prospect Theory's Loss Aversion Is Robust To Stake Size," *Judgment and Decision Making*, 18, e14.

Brown, A. L., T. Imai, F. M. Vieider, and C. F. Camerer (2024): "Meta-Analysis of Empirical Estimates of Loss Aversion," *Journal Of Economic Literature*, 62, 485–516.

Bursztyn, L., G. Egorov, I. Haaland, A. Rao, and C. Roth (2023): "Justifying Dissent," *The Quarterly Journal of Economics*, 138, 1403–1451.

Camerer, C., L. Babcock, G. Loewenstein, and R. Thaler (1997): "Labor Supply of New York City Cabdrivers: One Day at a Time," *The Quarterly Journal of Economics*, 112, 407–441.

Camerer, C. F. and R. A. Weber (1999): "The Econometrics and Behavioral Economics of Escalation of Commitment: A Re-Examination of Staw and Hoang's NBA Data," *Journal of Economic Behavior & Organization*, 39, 59–82.

Chiong, K. X., M. Shum, R. Webb, and R. Chen (2024): "Combining Choice and Response Time Data: A Drift-Diffusion Model of Mobile Advertisements," *Management Science*, 70, 1238–1257.

Cohen, J. and P. Dupas (2010): "Free Distribution or Cost-Sharing? Evidence from a Randomized Malaria Prevention Experiment," *The Quarterly Journal of Economics*, 1–45.

Eyster, E. (2002): "Rationalizing the Past: A Taste for Consistency," *Nuffield College Mimeograph*.

Eyster, E., S. Li, and S. Ridout (2021): "A Theory of Ex Post Rationalization," *arXiv preprint arXiv:2107.07491*.

Friedman, D., K. Pommerenke, R. Lukose, G. Milam, and B. A. Huberman (2007): "Searching for the Sunk Cost Fallacy," *Experimental Economics*, 10, 79–104.

Hidalgo, D., M. Onofa, H. Oosterbeek, and J. Ponce (2013): "Can Provision of Free School Uniforms Harm Attendance? Evidence from Ecuador," *Journal of Development Economics*, 103, 43–51.

Ho, T.-H., I. P. Png, and S. Reza (2018): "Sunk Cost Fallacy in Driving the World's Costliest Cars," *Management Science*, 64, 1761–1778.

Ketel, N., J. Linde, H. Oosterbeek, and B. van der Klaauw (2016): "Tuition Fees and Sunk-Cost Effects," *The Economic Journal*, 126, 2342–2362.

Kőszegi, B. and M. Rabin (2006): "A Model of Reference-Dependent Preferences," *The Quarterly Journal of Economics*, 121, 1133–1165.

Krefeld-Schwalb, A., E. R. Sugerman, and E. J. Johnson (2024): "Exposing Omitted Moderators: Explaining Why Effect Sizes Differ in The Social Sciences," *Proceedings of The National Academy of Sciences*, 121, e2306281121.

Levy, R. (2021): "Social Media, News Consumption, and Polarization: Evidence from a Field Experiment," *American Economic Review*, 111, 831–70.

Martens, N. and H. Orzen (2021): "Escalating Commitment to a Failing Course of Action—A Re-Examination," *European Economic Review*, 137, 103811.

Nah, F. F.-H. (2004): "A Study on Tolerable Waiting Time: How Long Are Web Users Willing to Wait?" *Behaviour & Information Technology*, 23, 153–163.

Negrini, M., A. Riedl, and M. Wibral (2022): "Sunk Cost in Investment Decisions," *Journal of Economic Behavior & Organization*, 200, 1105–1135.

Offerman, T. and J. Potters (2006): "Does Auctioning of Entry Licences Induce Collusion? An Experimental Study," *The Review Of Economic Studies*, 73, 769–791.

Olivola, C. Y. (2018): "The Interpersonal Sunk-Cost Effect," *Psychological Science*, 29, 1072–1083.

Oprea, R. (2024): "Decisions Under Risk Are Decisions Under Complexity," *American Economic Review*, 114, 3789–3811.

Ordóñez, L. D., T. Connolly, et al. (2000): "Regret and Responsibility: A Reply to Zeelenberg et al.(1998)," *Organizational Behavior and Human Decision Processes*, 81, 132–142.

Phillips, O. R., R. C. Battalio, and C. A. Kogut (1991): "Sunk and Opportunity Costs in Valuation and Bidding," *Southern Economic Journal*, 112–128.

Ronayne, D., D. Sgroi, and A. Tuckwell (2021): "Evaluating the Sunk Cost Effect," *Journal of Economic Behavior & Organization*, 186, 318–327.

Staw, B. M. (1976): "Knee-Deep in the Big Muddy: A Study of Escalating Commitment to a Chosen Course of Action," *Organizational Behavior and Human Performance*, 16, 27–44.

Staw, B. M. and H. Hoang (1995): "Sunk Costs in the NBA: Why Draft Order Affects Playing Time and Survival in Professional Basketball," *Administrative Science Quarterly*, 474–494.

Thaler, R. (1980): "Toward a Positive Theory of Consumer Choice," *Journal of Economic Behavior & Organization*, 1, 39–60.

Thaler, R. H. (1999): "Mental Accounting Matters," *Journal of Behavioral Decision Making*, 12, 183–206.

Weigel, C. (2018): "Don't Stop Now: The Sunk Cost Effect in an Incentivized Lab Experiment," *Available at SSRN 3207806.*

Zeelenberg, M., W. W. van Dijk, and A. S. Manstead (1998): "Reconsidering the Relation Between Regret and Responsibility," *Organizational Behavior and Human Decision Processes*, 74, 254–272.

A Tables

Table A1: Efficiency in the Counting Game

	Round Per Minute			
	(1)	(2)	(3)	(4)
Price = 1.5	−0.063	−0.056	−0.082	−0.068
	(0.087)	(0.087)	(0.084)	(0.084)
Price = 2.5	−0.023	−0.044	−0.053	−0.071
	(0.087)	(0.087)	(0.084)	(0.084)
Attention		X		X
Controls			X	X
Mean	4.325	4.347	4.325	4.347
N	1,806	1,749	1,806	1,749

Note: This table reports the results of the sunk cost treatments on efficiency, defined as the number of rounds completed per minute. Column 1 shows a regression of the number of rounds completed per minute on a dummy variable equal to one if a participant was assigned the medium price ($1.5) and a dummy variable equal to one if a participant was assigned the high price ($2.5). Column 2 additionally includes controls. Column 3 restricts the sample to those participants who correctly answered manipulation checks about the price they were assigned and the average bonus payment in the study. Column 4 extends the previous specification by including controls. Robust standard errors are parenthesized. *p<0.1; **p<0.05; ***p<0.01.

Table A2: Accuracy in the Counting Game

	Proportion of Correct Round			
	(1)	(2)	(3)	(4)
Price = 1.5	−0.008	−0.009	−0.008	−0.008
	(0.009)	(0.009)	(0.009)	(0.009)
Price = 2.5	−0.006	−0.005	−0.006	−0.006
	(0.009)	(0.009)	(0.009)	(0.009)
Attention		X		X
Controls			X	X
Mean	0.914	0.916	0.914	0.916
N	1,781	1,727	1,781	1,727

Note: This table reports the results of the sunk cost treatments on accuracy, defined as the proportion of correctly solved rounds conditional on having completed at least one round. Column 1 shows a regression of the proportion of correctly solved rounds on a dummy variable equal to one if a participant was assigned the medium price ($1.5) and a dummy variable equal to one if a participant was assigned the high price ($2.5). Column 2 additionally includes controls. Column 3 restricts the sample to those participants who correctly answered manipulation checks about the price they were assigned and the average bonus payment in the study. Column 4 extends the previous specification by including controls. Robust standard errors are parenthesized. *p<0.1; **p<0.05; ***p<0.01.

Table A3: Earnings in the Counting Game

	Total Earning			
	(1)	(2)	(3)	(4)
Price = 1.5	−1.032***	−1.034***	−1.039***	−1.039***
	(0.079)	(0.079)	(0.079)	(0.079)
Price = 2.5	−1.966***	−1.963***	−1.977***	−1.973***
	(0.079)	(0.080)	(0.079)	(0.079)
Attention		X		X
Controls			X	X
Mean	3.949	3.978	3.949	3.978
N	1,806	1,749	1,806	1,749

Note: This table reports the results of the sunk cost treatments on the total earnings from the experiment. Column 1 shows a regression of the total earnings on a dummy variable equal to one if a participant was assigned the medium price ($1.5) and a dummy variable equal to one if a participant was assigned the high price ($2.5). Column 2 additionally includes controls. Column 3 restricts the sample to those participants who correctly answered manipulation checks about the price they were assigned and the average bonus payment in the study. Column 4 extends the previous specification by including controls. Robust standard errors are parenthesized. *p<0.1; **p<0.05; ***p<0.01.

Table A4: Online Study: Capped Play Time

	Capped Playtime (Minutes)			
	(1)	(2)	(3)	(4)
Price = 1.5	0.584	0.659	0.652	0.711
	(0.652)	(0.661)	(0.652)	(0.660)
Price = 2.5	1.029*	1.249**	1.091**	1.296**
	(0.653)	(0.662)	(0.651)	(0.660)
Attention		X		X
Controls			X	X
Mean	13.695	13.837	13.695	13.837
Wald P-value	0.9	0.95	0.85	0.91
N	1,806	1,749	1,806	1,749

Note: This table reports the results on the existence and asymmetry of the sunk cost effect. Column 1 shows a regression of game playtime (in minutes, capped at 60 minutes) on a dummy variable equal to one if a participant was assigned the medium price ($1.5) and a dummy variable equal to one if a participant was assigned the high price ($2.5). Column 2 additionally includes controls. Column 3 restricts the sample to those participants who correctly answered manipulation checks about the price they were assigned and the average bonus payment in the study. Column 4 extends the previous specification by including controls. Robust standard errors are parenthesized. Following Hypothesis 1, we compute one-sided p-values in a row "Price=$2.5". Additionally, we report p-values based on a Wald test against a null hypothesis outlined in Hypothesis 2 in a row "Wald P-value". *p<0.1; **p<0.05; ***p<0.01.

Table A5: Characteristics of Users in the YouTube Study

	Mean	Std. Dev.
Number of Users	410	
Available on Qualtrics	0.62	0.49
Age	41.51	16.79
Male	0.44	0.50
White	0.64	0.48
Hispanic	0.13	0.33
West	0.23	0.42
Midwest	0.21	0.41
South	0.31	0.46
Northeast	0.25	0.44

Note: This table provides characteristics of users with a record of seeing at least one skippable ad. The table presents means and standard deviations (parenthesized). "Available on Qualtrics" variable equals one if a user takes the intake survey, which provides all the variables below it. Age is defined as the difference between 2022 and the year of birth. "White" and "Hispanic" variables equal one if the user reports identifying themselves as such. Last four regional variables equal one if the user reports living in a state corresponding to the region.

Table A6: YouTube Study: 2SLS

	Engagement Time	Engagement=0
	(1)	(2)
Sunk Time (Instrumented)	−0.160	0.007***
	(1.437)	(0.002)
Sample	Full	Full
Individual FE	X	X
Mean	228.64	0.22
N	11,328	11,328

Note: The table reports the effect of sunk time on the video engagement. The sunk time is defined as the interval between the best estimate of the loading time of the video page and the end of the initial ad segment. The engagement time is the interval between the end of the initial ad segment and the first indication that the user is not watching the video. Details are provided in Section 4.3.2. Column 1 shows a regression of engagement time (in seconds) on individual fixed effects, and sunk time which is instrumented with the assigned minimum ad duration. Column 2 reports the effect on no engagement behavior, defined as having an engagement time of 0. Robust standard errors are parenthesized. *p<0.1; **p<0.05; ***p<0.01.

Table A7: YouTube Study: Outcomes as Proportions

	Proportion of Sunk Time	Proportion of Engagement Time
	(1)	(2)
MAD = 5	0.088***	−0.011
	(0.007)	(0.008)
MAD = 10	0.157***	−0.012
	(0.008)	(0.008)
Sample	Full	Full
Individual FE	X	X
Mean	0.401	0.352
Wald P-value	NA	0.496
N	11,328	11,221

Note: This table summarizes the effect of the minimum ad duration intervention on various outcomes of interest using the main regression specification that includes dummy variables for each treatment status (rows MAD=5 and MAD=10) and individual fixed effects. The sunk time is defined as the interval between the best estimate of the loading time of the video page and the end of the initial ad segment. The engagement time is the interval between the end of the initial ad segment and the first indication that the user is not watching the video. Details are provided in Section 4.3.2. Column 1 reports the effect on the proportion of sunk time relative to the total ad length. Column 2 reports the effect on the proportion of video engagement time relative to the video length. For the result on engagement time, we also report p-values based on a Wald test against a null hypothesis outlined in Hypothesis 4 in a row "Wald P-value". Robust standard errors are parenthesized. *p<0.1; **p<0.05; ***p<0.01.

Table A8: Heterogeneity of Sunk Time by Video Length

	Sunk Time					
	(1)	(2)	(3)	(4)	(5)	(6)
MAD = 5	4.673	3.949***	3.132**	2.210*	2.531*	3.015**
	(2.889)	(1.438)	(1.370)	(1.262)	(1.500)	(1.392)
MAD = 10	4.360	5.854***	5.874***	6.480***	6.829***	7.422***
	(3.024)	(1.596)	(1.552)	(1.379)	(1.586)	(1.547)
Sample	<3 mins	<6 mins	<10 mins	<20 mins	Full	NoMV
Individual FE	X	X	X	X	X	X
Mean	19.50	19.41	20.77	21.63	22.87	22.88
N	1,361	3,886	5,767	8,734	11,328	10,904

Note: The table reports the effect of the minimum ad duration intervention on the sunk time across different subsamples based on the video length and category. In all columns, we report a regression of sunk time (in seconds) on individual fixed effects, a dummy variable equal to one if the minimum ad duration is 5 seconds (MAD=5) and a dummy variable equal to one if the minimum ad duration is 10 seconds (MAD=10). The sunk time is defined as the interval between the best estimate of the loading time of the video page and the end of the initial ad segment. Details are provided in Section 4.3.2. Column 1 restricts the sample to videos shorter than 3 minutes. Column 2 restricts the sample to videos shorter than 6 minutes. Column 3 restricts the sample to videos shorter than 10 minutes. Column 4 restricts the sample to videos shorter than 20 minutes. Column 5 reports the regression based on the full sample, and Column 6 reports the subsample that excludes videos that are in "Music" category and are more than 10 minutes long. Robust standard errors are parenthesized. *p<0.1; **p<0.05; ***p<0.01.

Table A9: Heterogeneity of Engagement Time by Video Length

	Engagement Time					
	(1)	(2)	(3)	(4)	(5)	(6)
MAD = 5	−2.649	−3.691	−3.000	0.109	−4.921	−7.578
	(3.887)	(3.665)	(4.316)	(5.582)	(9.106)	(9.127)
MAD = 10	1.334	−3.116	−6.509	−4.230	−1.850	0.255
	(3.645)	(3.628)	(4.294)	(5.634)	(9.808)	(9.993)
Sample	<3 mins	<6 mins	<10 mins	<20 mins	Full	NoMV
Individual FE	X	X	X	X	X	X
Mean	57.99	89.25	120.45	175.17	228.64	229.14
Wald P-value	0.337	0.498	0.946	0.643	0.623	0.345
N	1,361	3,886	5,767	8,734	11,328	10,904

Note: The table reports the effect of the minimum ad duration intervention on the engagement time across different subsamples based on the video length and category. In all columns, we report a regression of engagement time (in seconds) on individual fixed effects, a dummy variable equal to one if the minimum ad duration is 5 seconds (MAD=5) and a dummy variable equal to one if the minimum ad duration is 10 seconds (MAD=10). The engagement time is the interval between the end of the initial ad segment and the first indication that the user is not watching the video. Details are provided in Section 4.3.2. Column 1 restricts the sample to videos shorter than 3 minutes. Column 2 restricts the sample to videos shorter than 6 minutes. Column 3 restricts the sample to videos shorter than 10 minutes. Column 4 restricts the sample to videos shorter than 20 minutes. Column 5 reports the regression based on the full sample, and Column 6 reports the subsample that excludes videos that are in "Music" category and are more than 10 minutes long. Robust standard errors are parenthesized. *p<0.1; **p<0.05; ***p<0.01.

Table A10: Heterogeneity of No Engagement Behavior by Video Length

	Engagement Time=0					
	(1)	(2)	(3)	(4)	(5)	(6)
MAD = 5	0.077**	0.050***	0.037***	0.029***	0.033***	0.033***
	(0.031)	(0.015)	(0.013)	(0.010)	(0.009)	(0.009)
MAD = 10	0.057*	0.061***	0.049***	0.045***	0.052***	0.053***
	(0.029)	(0.016)	(0.013)	(0.010)	(0.009)	(0.009)
Sample	<3 mins	<6 mins	<10 mins	<20 mins	Full	NoMV
Individual FE	X	X	X	X	X	X
Mean	0.259	0.214	0.219	0.214	0.217	0.215
N	1,361	3,886	5,767	8,734	11,328	10,904

Note: The table reports the effect of the minimum ad duration intervention on the likelihood of no engagement behavior, defined as having an engagement time of 0, across different subsamples based on the video length and category. In all columns, we report a regression of a dummy variable equal to one if the engagement time is 0 on individual fixed effects, a dummy variable equal to one if the minimum ad duration is 5 seconds (MAD=5) and a dummy variable equal to one if the minimum ad duration is 10 seconds (MAD=10). The engagement time is the interval between the end of the initial ad segment and the first indication that the user is not watching the video. Details are provided in Section 4.3.2. Column 1 restricts the sample to videos shorter than 3 minutes. Column 2 restricts the sample to videos shorter than 6 minutes. Column 3 restricts the sample to videos shorter than 10 minutes. Column 4 restricts the sample to videos shorter than 20 minutes. Column 5 reports the regression based on the full sample, and Column 6 reports the subsample that excludes videos that are in "Music" category and are more than 10 minutes long. Robust standard errors are parenthesized. *p<0.1; **p<0.05; ***p<0.01.

Table A11: YouTube Study: Temporal Stability

	Sunk Time	Engagement Time	Engagement = 0	Ad Click
	(1)	(2)	(3)	(4)
MAD = 5 × Last	0.591	17.146	0.007	0.004
	(2.909)	(18.295)	(0.018)	(0.003)
MAD = 10 × Last	−2.103	−1.249	0.015	0.006**
	(3.077)	(19.601)	(0.018)	(0.003)
Sample	Full	Full	Full	Full
Individual FE	X	X	X	X
Mean	22.87	228.64	0.217	0.005
N	11,328	11,328	11,328	11,328

Note: This table summarizes the effect of the minimum ad duration intervention by treatment period on various outcomes of interest using a regression specification that adds to the main specification a dummy that is one if the video is watched in the last two weeks of the intervention period ("Last"). This dummy is interacted with the dummy variables for each treatment status (rows MAD = 5 × Last and MAD = 10 × Last). Column 1 reports the effect on sunk time. The sunk time is defined as the interval between the best estimate of the loading time of the video page and the end of the initial ad segment. Column 2 reports the effect on video engagement time. The engagement time is the interval between the end of the initial ad segment and the first indication that the user is not watching the video. Details are provided in Section 4.3.2. Column 3 reports the effect on no engagement behavior, defined as having an engagement time of 0. Column 4 reports the effect on ad click rate. All specifications use the full sample. Robust standard errors are parenthesized. *p<0.1; **p<0.05; ***p<0.01.

Table A12: YouTube Study: Excluding Multi-Video Observations

	Sunk Time	Engagement Time	Engagement = 0	Ad Click
	(1)	(2)	(3)	(4)
MAD = 5	2.533*	−4.678	0.032***	0.003*
	(1.504)	(9.128)	(0.009)	(0.002)
MAD = 10	6.606***	−1.682	0.051***	0.002
	(1.572)	(9.839)	(0.009)	(0.002)
Sample	No Multi	No Multi	No Multi	No Multi
Individual FE	X	X	X	X
Mean	22.84	229.16	0.216	0.005
Wald P-value	NA	0.638	NA	NA
N	11,291	11,291	11,291	11,291

Note: This table summarizes the effect of the minimum ad duration intervention on various outcomes of interest using the main regression specification that includes dummy variables for each treatment status (rows MAD=5 and MAD=10) and individual fixed effects. Column 1 reports the effect on sunk time. The sunk time is defined as the interval between the best estimate of the loading time of the video page and the end of the initial ad segment. Column 2 reports the effect on video engagement time. The engagement time is the interval between the end of the initial ad segment and the first indication that the user is not watching the video. Details are provided in Section 4.3.2. Column 3 reports the effect on no engagement behavior, defined as having an engagement time of 0. Column 4 reports the effect on ad click rate. All specification use a subset of the full sample that excludes observations where the same treatment status is enforced by the extension for multiple consecutive videos. For the result on engagement time, we also report p-values based on a Wald test against a null hypothesis outlined in Hypothesis 4 in a row "Wald P-value". Robust standard errors are parenthesized. *p<0.1; **p<0.05; ***p<0.01.

Table A13: YouTube Study: Including Video Streams

	Sunk Time	Engagement Time	Engagement = 0	Ad Click
	(1)	(2)	(3)	(4)
MAD = 5	3.168**	−4.565	0.033***	0.002
	(1.594)	(9.083)	(0.009)	(0.002)
MAD = 10	6.931***	−2.357	0.053***	0.002
	(1.589)	(9.723)	(0.009)	(0.002)
Sample	Include Streams	Include Streams	Include Streams	Include Streams
Individual FE	X	X	X	X
Mean	23.00	227.72	0.216	0.005
Wald P-value	NA	0.677	NA	NA
N	11,486	11,486	11,486	11,486

Note: This table summarizes the effect of the minimum ad duration intervention on various outcomes of interest using the main regression specification that includes dummy variables for each treatment status (rows MAD=5 and MAD=10) and individual fixed effects. Column 1 reports the effect on sunk time. The sunk time is defined as the interval between the best estimate of the loading time of the video page and the end of the initial ad segment. Column 2 reports the effect on video engagement time. The engagement time is the interval between the end of the initial ad segment and the first indication that the user is not watching the video. Details are provided in Section 4.3.2. Column 3 reports the effect on no engagement behavior, defined as having an engagement time of 0. Column 4 reports the effect on ad click rate. All specifications use a sample that adds video streams to the full sample. For the result on engagement time, we also report p-values based on a Wald test against a null hypothesis outlined in Hypothesis 4 in a row "Wald P-value". Robust standard errors are parenthesized. *p<0.1; **p<0.05; ***p<0.01.

Table A14: YouTube Study: Excluding Top One Percentile by Sunk Time

	Sunk Time	Engagement Time	Engagement = 0	Ad Click
	(1)	(2)	(3)	(4)
MAD = 5	2.985***	−5.793	0.033***	0.003
	(0.566)	(9.166)	(0.009)	(0.002)
MAD = 10	6.247***	−2.249	0.052***	0.003
	(0.570)	(9.880)	(0.009)	(0.002)
Sample	ST99	ST99	ST99	ST99
Individual FE	X	X	X	X
Mean	18.11	229.53	0.212	0.005
Wald P-value	NA	0.569	NA	NA
N	11,215	11,215	11,215	11,215

Note: This table summarizes the effect of the minimum ad duration intervention on various outcomes of interest using the main regression specification that includes dummy variables for each treatment status (rows MAD=5 and MAD=10) and individual fixed effects. Column 1 reports the effect on sunk time. The sunk time is defined as the interval between the best estimate of the loading time of the video page and the end of the initial ad segment. Column 2 reports the effect on video engagement time. The engagement time is the interval between the end of the initial ad segment and the first indication that the user is not watching the video. Details are provided in Section 4.3.2. Column 3 reports the effect on no engagement behavior, defined as having an engagement time of 0. Column 4 reports the effect on ad click rate. All specifications use a sample that excludes videos within the top one percentile by sunk time from the full sample. For the result on engagement time, we also report p-values based on a Wald test against a null hypothesis outlined in Hypothesis 4 in a row "Wald P-value". Robust standard errors are parenthesized. *p<0.1; **p<0.05; ***p<0.01.

B Figures

FIGURE B1: AD INTERFACE: MAD=0

Note: This is an example of user experience under MAD=0 treatment, which shows the “Skip Ads” button immediately after ad starts.

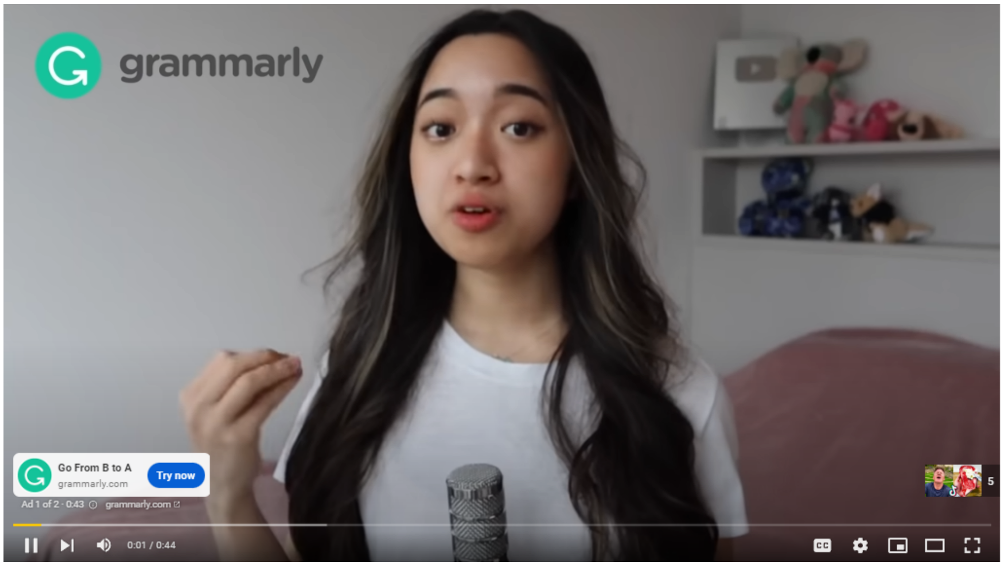

FIGURE B2: AD INTERFACE: MAD=5

Note: This is an example of user experience under MAD=5 treatment, which shows the “Skip Ads” button after 5 seconds of ad starting.

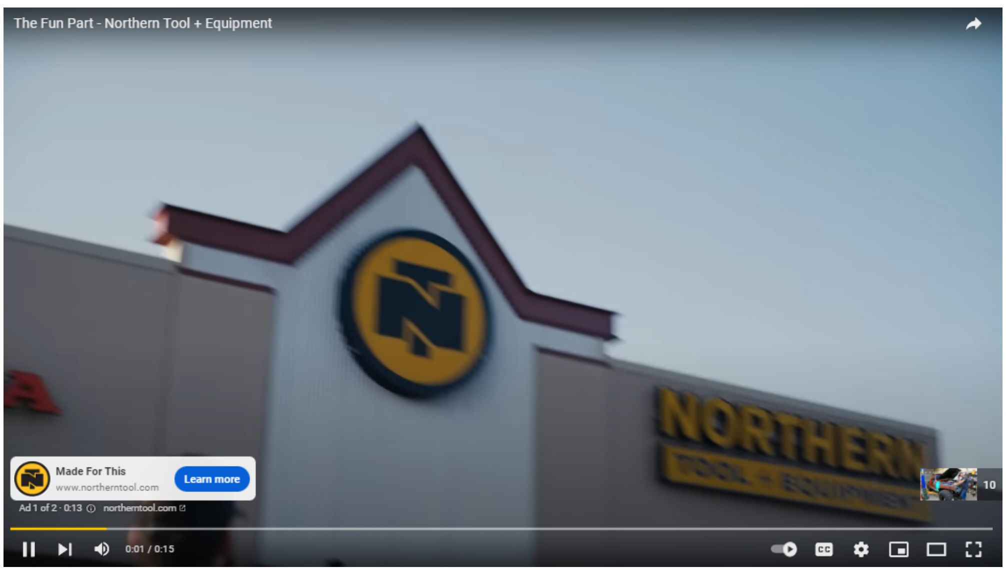

FIGURE B3: AD INTERFACE: MAD=10

Note: This is an example of user experience under MAD=5 treatment, which shows the "Skip Ads" button after 10 seconds of ad starting.

ONLINE APPENDIX

Survey Instructions

Below, we provide survey instructions (printouts of relevant screens) for the online epxeriment described in Section 3 of the paper.

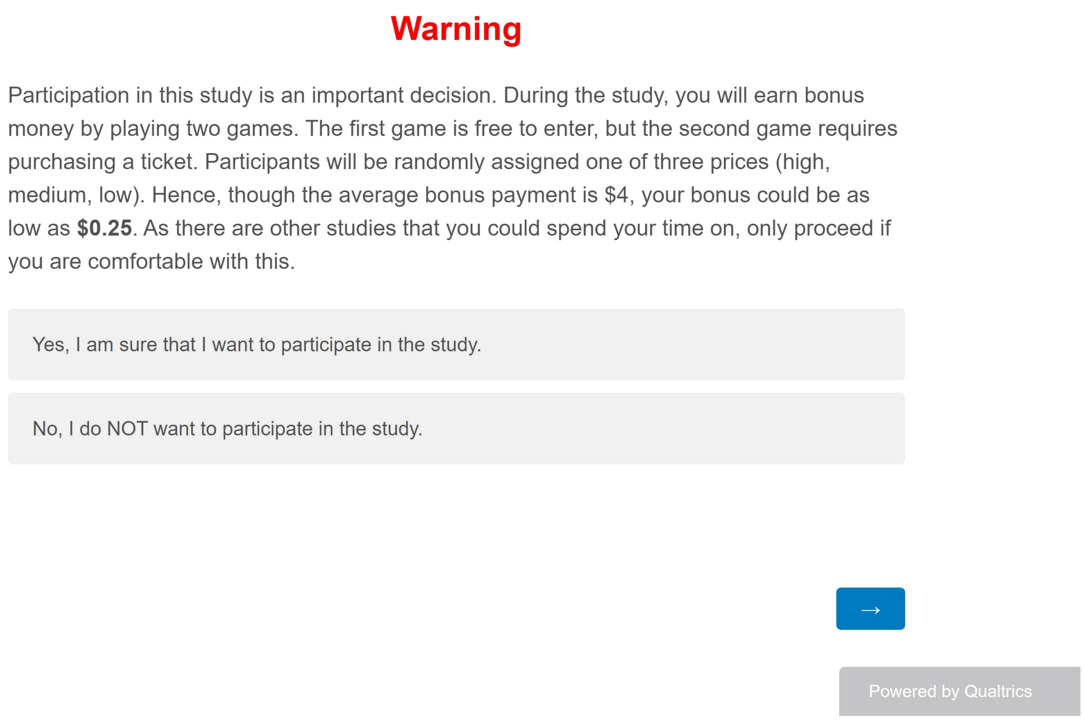

Warning

Participation in this study is an important decision. During the study, you will earn bonus money by playing two games. The first game is free to enter, but the second game requires purchasing a ticket. Participants will be randomly assigned one of three prices (high, medium, low). Hence, though the average bonus payment is $4, your bonus could be as low as **$0.25**. As there are other studies that you could spend your time on, only proceed if you are comfortable with this.

Yes, I am sure that I want to participate in the study.

No, I do NOT want to participate in the study.

→

Powered by Qualtrics

Figure OA1: Survey Instructions: Initial Warning Screen

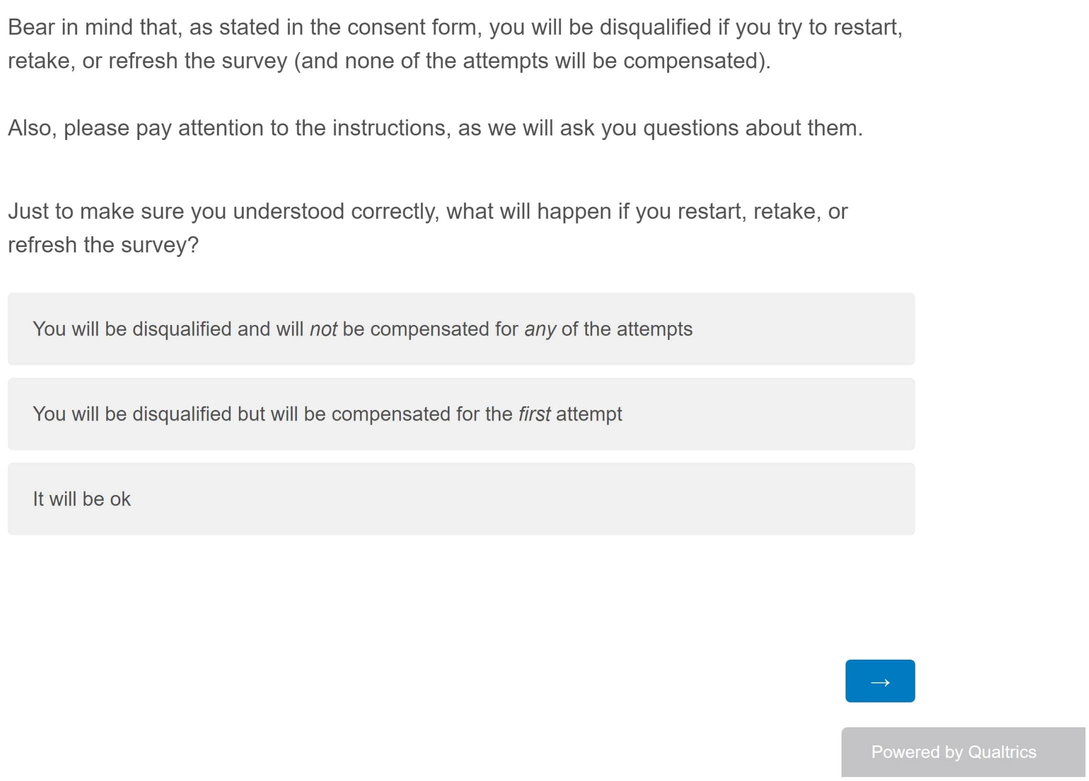
Bear in mind that, as stated in the consent form, you will be disqualified if you try to restart, retake, or refresh the survey (and none of the attempts will be compensated).

Also, please pay attention to the instructions, as we will ask you questions about them.

Just to make sure you understood correctly, what will happen if you restart, retake, or refresh the survey?

You will be disqualified and will *not* be compensated for *any* of the attempts

You will be disqualified but will be compensated for the *first* attempt

It will be ok

→

Powered by Qualtrics

Figure OA2: Survey Instructions: Screen Warning Against Refreshing

In this part of the study, we would like to ask you some basic demographic questions. There are no "right" answers. We would appreciate if you provide honest answers to the questions.

What is your year of birth?

What is your sex?

- Male
- Female

Please select your state of residence.

State:

What was your family's gross household income in 2020 in US dollars?

- Less than $25,000
- $25,000 to $49,999
- $50,000 to $74,999
- $75,000 to $99,999
- $100,000 to $149,000
- $150,000 to $199,999
- $200,000 or more

Figure OA3: Survey Instructions: Demographic Questions 1

Which of the following best describe your race or ethnicity? You can select more than one option.

- African American/Black
- Asian/Asian American
- Caucasian/White
- Native American, Inuit or Aleut
- Native Hawaiian/Pacific Islander
- Other

Are you of Hispanic, Latino, or Spanish origin?

- Yes
- No
- Prefer not to answer

Figure OA4: Survey Instructions: Demographic Questions 2

What is the highest level of education you have completed or the highest degree you have received?

- Less than high school degree
- High school graduate (high school diploma or equivalent including GED)
- Some college but no degree
- Associate degree in college (2-year)
- Bachelor's degree in college (4-year)
- Master's degree
- Doctoral degree (PhD, JD, MD)

→

Powered by Qualtrics

Figure OA5: Survey Instructions: Demographic Questions 3

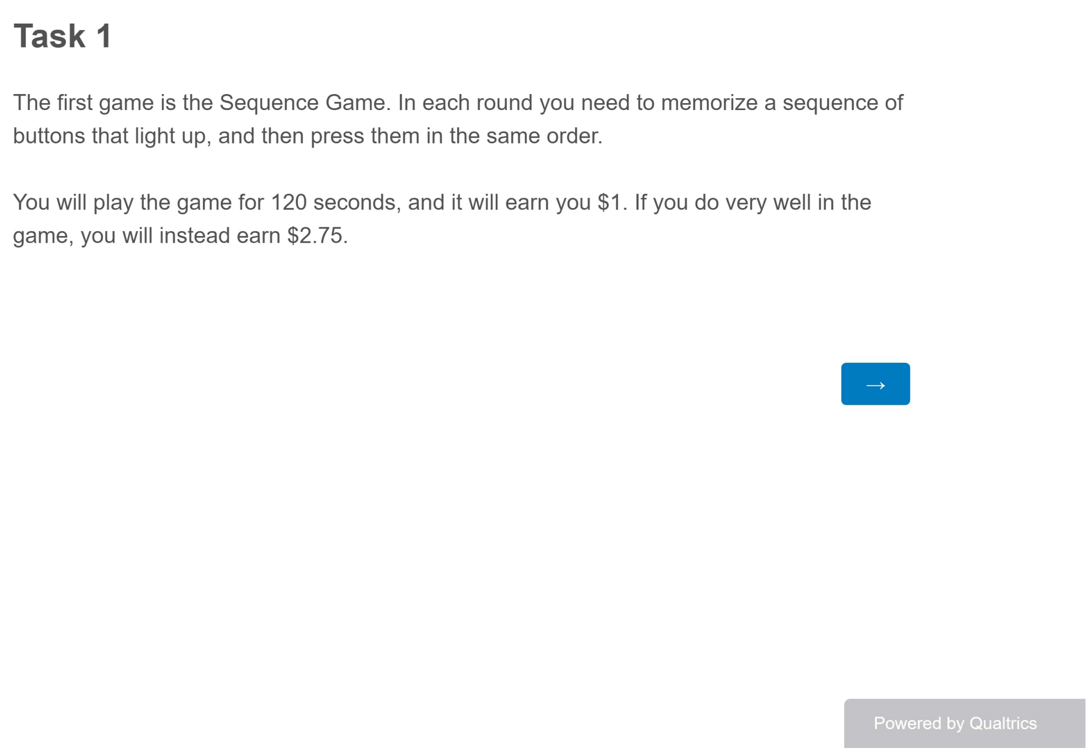

Figure OA6: Survey Instructions: Introduction to the First Task

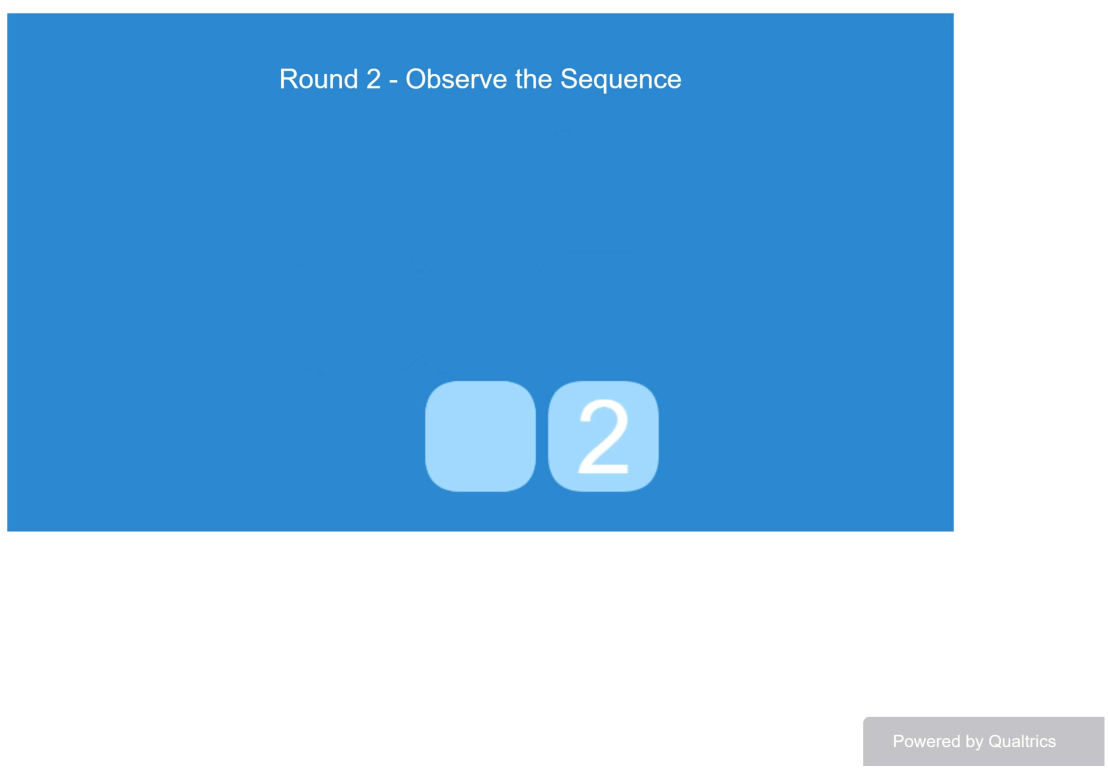

Figure OA7: Survey Instructions: The First Task

Congratulations! We evaluated your performance in the game. Your score was above the threshold for the high bonus. You earned $2.75.

→

Powered by Qualtrics

Figure OA8: Survey Instructions: Result of the First Task

The Counting Game

You are now proceeding to the Counting Game. Playing this game requires purchasing an entry ticket.

Throughout the game players earn points. In each round we will show you a set of 30 digits. Your task will be to count the number of zeros.

If you are correct, you will earn the number of points indicated on the screen. If you make a mistake, you will get *no* points for the round.

Upon completion, you will earn the rewards based on the points you collected, where **1000 points = $1**.

As an example, here is an image of one round.

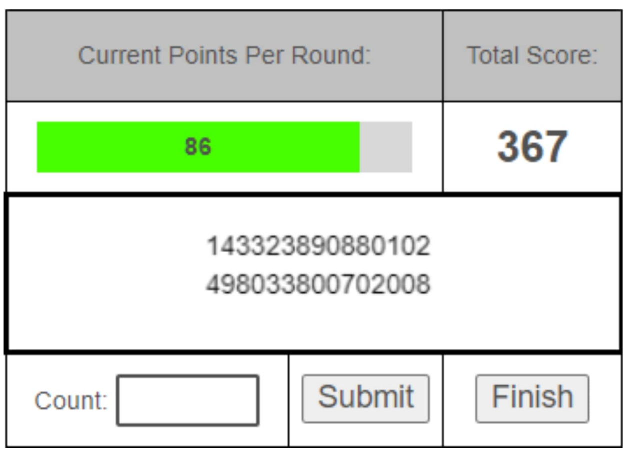

You can see (top-left) how many points you can earn in this round (**this will decrease over time**) and how many points you've accumulated over all rounds (top-right).

You need to count the number of zeros in a sequence in the center; this one has 9. You should write the number in the text box (bottom-left) and click "Submit".

You can play until you decide to finish. Once you decide to do that, you should click the button “Finish” to collect the reward based on the points that you have accumulated.

Figure OA9: Survey Instructions: Introduction to the Main Task 1

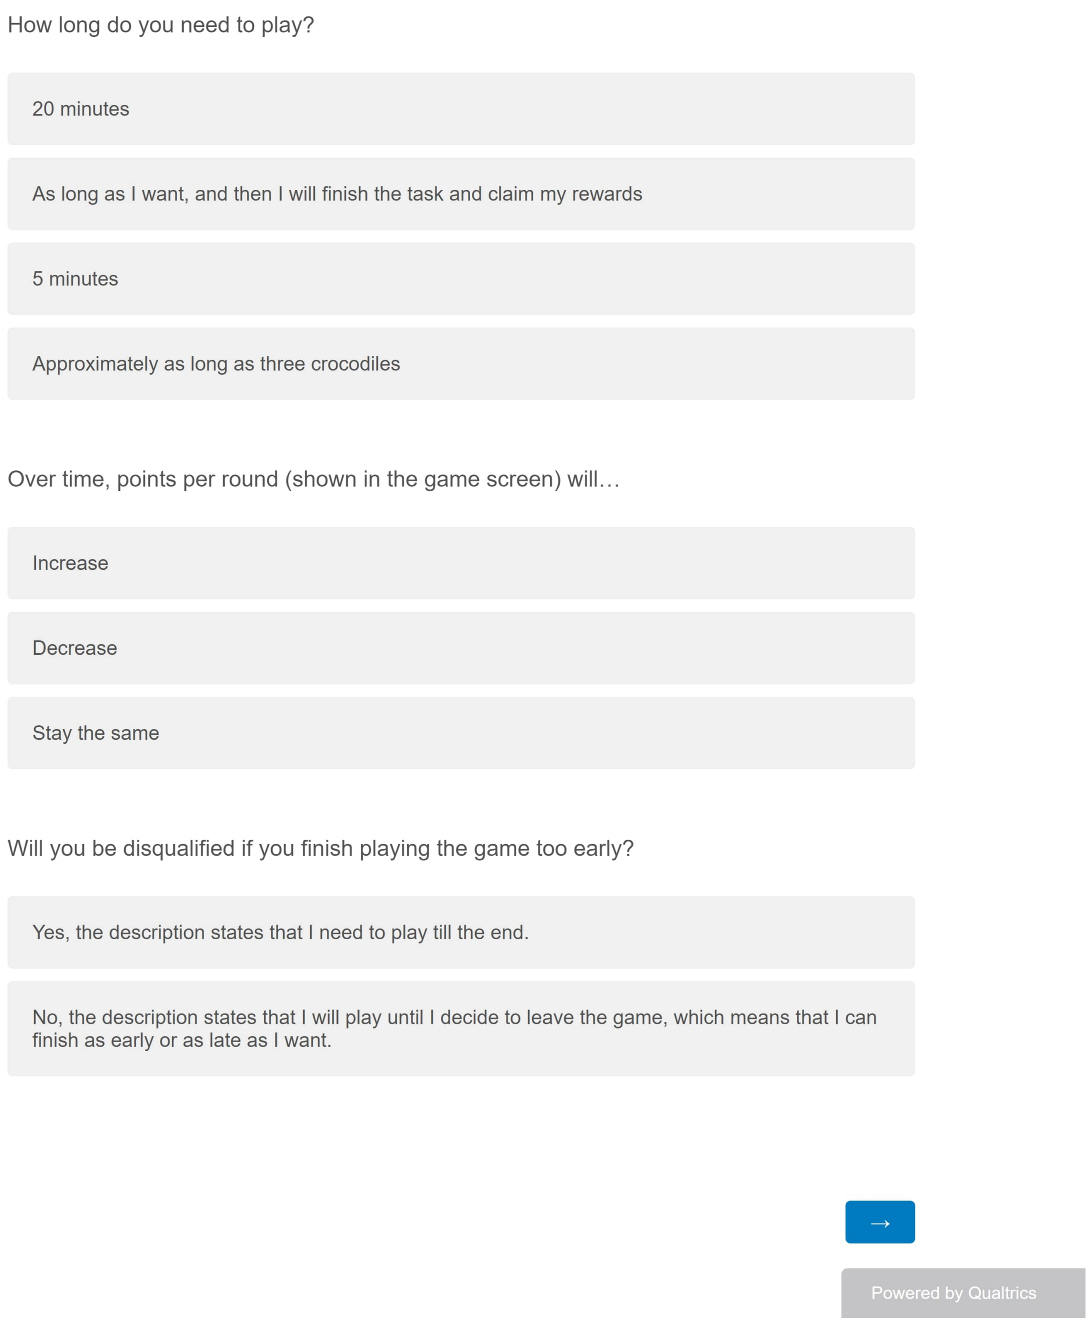

How long do you need to play?

- 20 minutes
- As long as I want, and then I will finish the task and claim my rewards
- 5 minutes
- Approximately as long as three crocodiles

Over time, points per round (shown in the game screen) will…

- Increase
- Decrease
- Stay the same

Will you be disqualified if you finish playing the game too early?

- Yes, the description states that I need to play till the end.
- No, the description states that I will play until I decide to leave the game, which means that I can finish as early or as late as I want.

→

Powered by Qualtrics

Figure OA10: Survey Instructions: Introduction to the Main Task 2

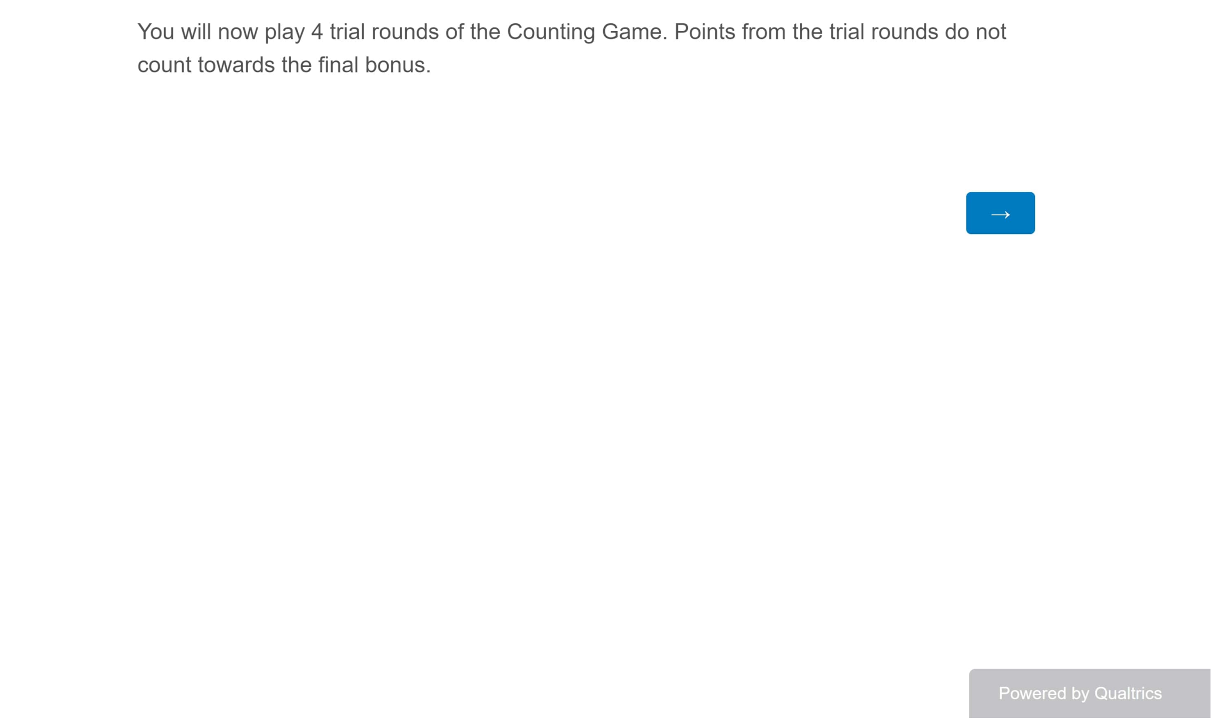

Figure OA11: Survey Instructions: Trial Rounds of the Main Task 1

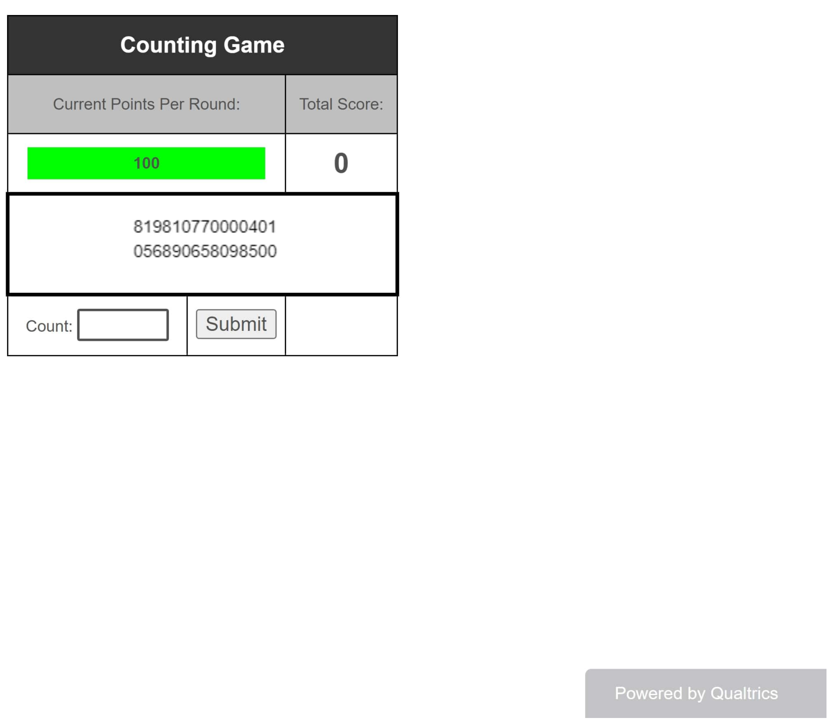

Figure OA12: Survey Instructions: Trial Rounds of the Main Task 2

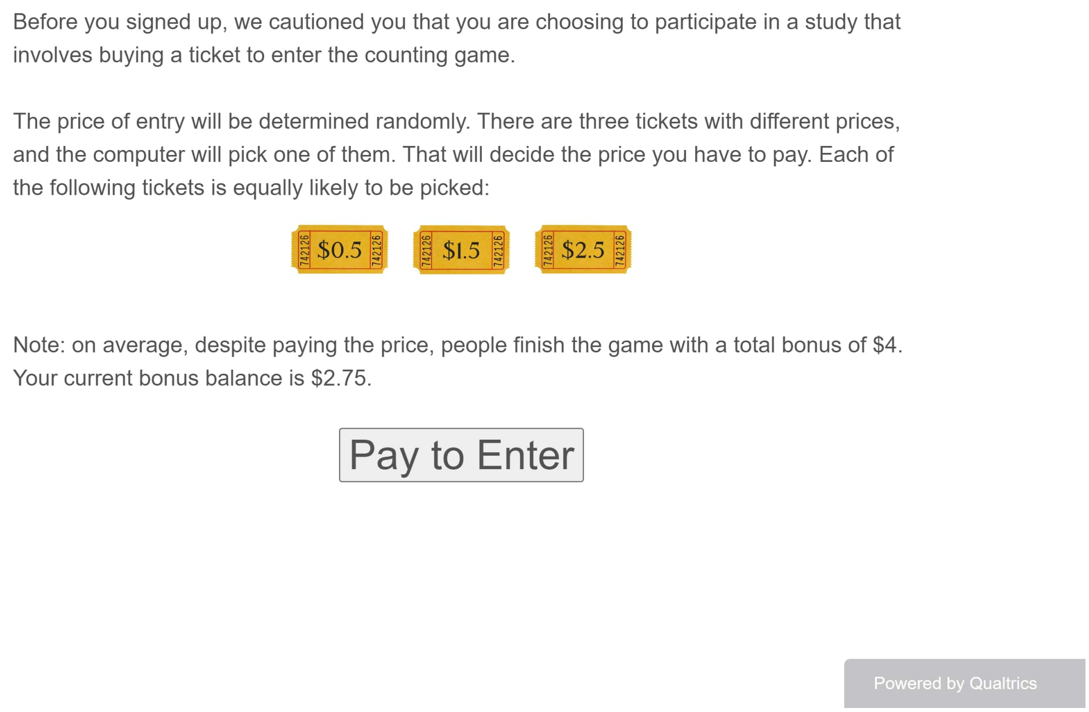

Figure OA13: Survey Instructions: Description of the Tickets to Enter the Main Task

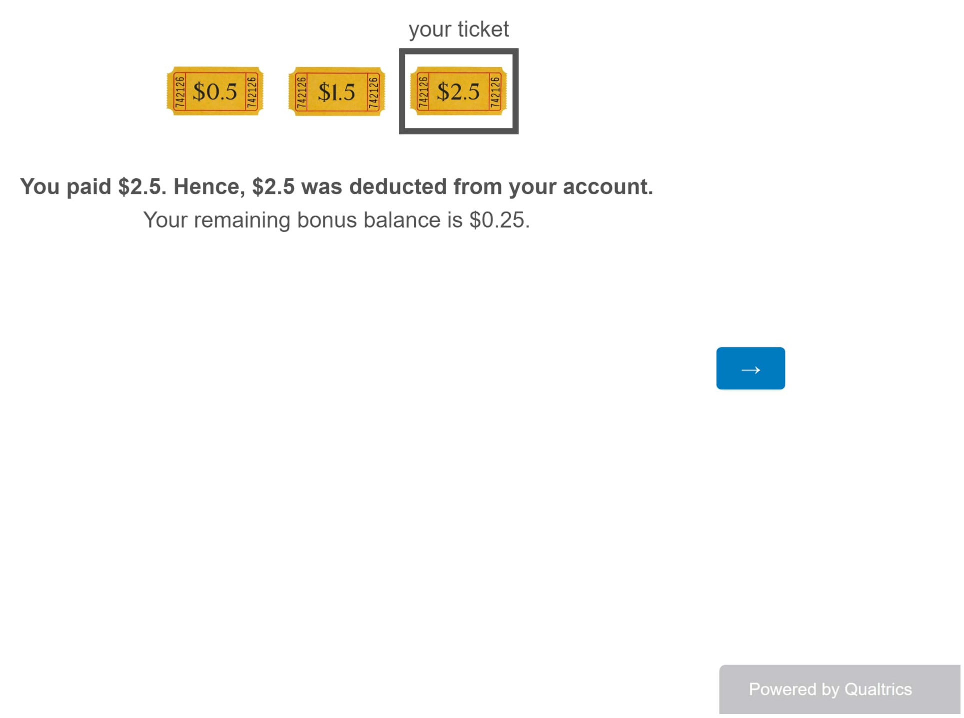

Figure OA14: Survey Instructions: Realization of the Price of the Ticket

Which price did you end up paying?

- $0.5
- $1.5
- $2.5

Upon finishing the game, people on average receive a total bonus of…?

- $3
- $4
- $5

→

Powered by Qualtrics

Figure OA15: Survey Instructions: Attention Checks

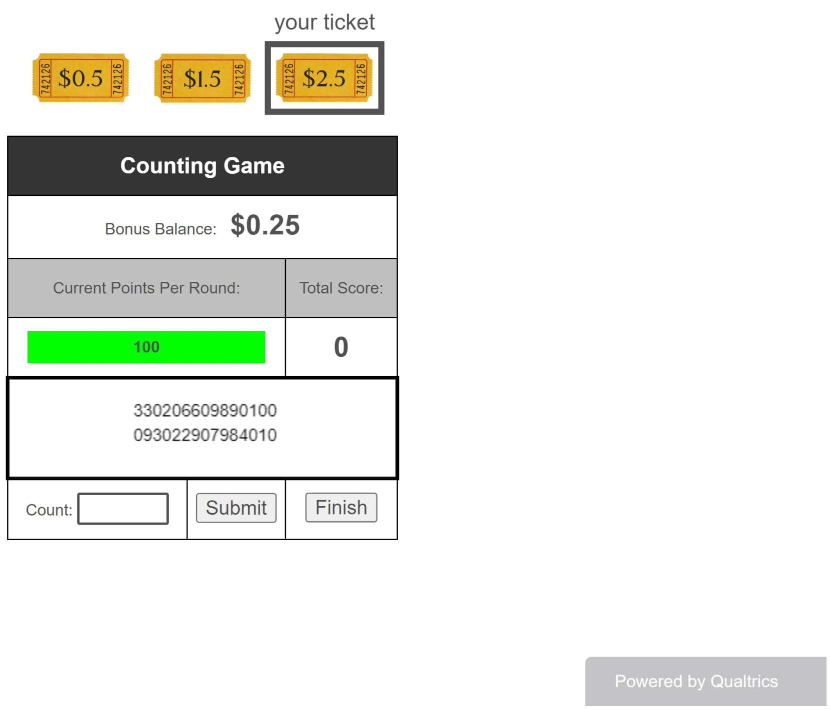

Figure OA16: Survey Instructions: The Main Task