Pre-analysis Plan
Sunk Costs and Exercising

Raphael Epperson* Wolfgang Habla† Matthias Wibral‡

This document provides supplementary information to the pre-registration "Sunk Costs and Exercising" at the AEA RCT Registry.

1 Hypotheses

Our first hypothesis addresses the existence of sunk-cost effects in the domain of exercising and health behavior. In particular, we hypothesize that individuals who receive a discount on their membership fee (and thereby face lower sunk costs) attend the gym less often. Since we expect the reduction in exercising due to lower sunk costs to be most pronounced at the beginning of the intervention and for individuals who receive the discount early, we focus on the behavior of individuals in the treatment Early Discount and in the control group for the months May and June 2023.

Hypothesis 1 *Individuals who receive a 50% discount on their membership fee in May and June 2023 attend the gym less often during this period than individuals in the control group.*

Our two following hypotheses deal with the behavioral implications of the intertemporal distribution of sunk costs. Individuals from treatments Early Discount and Late Discount both experience a reduction in the overall sunk costs associated with their contract. Although the overall reduction in sunk costs is the same, we conjecture that, in the early months, individuals from treatment Early Discount show a stronger behavioral response than individuals from treatment Late Discount because their reduction in sunk costs materializes early on (in the form of the implemented discount).

Hypothesis 2 *Individuals who receive a 50% discount on their membership fee in May and June 2023 attend the gym less often during this period than individuals who receive a 50% discount on their membership fee in November and December 2023.*

Furthermore, we expect this difference to persist in the months following the implementation of the early discount (but before the implementation of the late discount).

*University of Innsbruck.
†Baden-Wuerttemberg Cooperative State University (DHBW).
‡Maastricht University.

Hypothesis 3 *Individuals who receive a 50% discount on their membership fee in May and June 2023 attend the gym less often during the period from July to October 2023 than individuals who receive a 50% discount on their membership fee in November and December 2023.*

2 Data

For each of our participants, we receive data on (i) the date and time of all check-ins at any gym location of the fitness chain during the period from April 1, 2023, until April 30, 2024, (ii) their contract details, like duration or starting date, (iii) individual characteristics, including gender and year of birth, (iv) whether they have opened the first (and second) email, and (v) any follow-up contract in April 2024 (if the initial contract period was 12 months).

3 Empirical Analysis

3.1 Primary

For testing hypotheses 1 and 2, we estimate the following equation based on an OLS regression:

$$y_{i,May-June} = \alpha + \beta_{early} T_{i,early} + \beta_{late} T_{i,late} + \gamma\, y_{i,April} + \boldsymbol{\theta} \boldsymbol{S_i} + \epsilon_i, \tag{1}$$

where $y_{i,May-June}$ refers to the average number of visits per week of individual i during the period from May 1, 2023, until June 30, 2023. $T_{i,early}$ and $T_{i,late}$ are indicator variables for being assigned to the treatment Early Discount or Late Discount, respectively. $y_{i,April}$ is the average number of visits per week during the period from April 1 until April 25, 2023 (i.e., the day before the first intervention email is sent out). S_i are dummy variables representing each of the ten strata used in the randomization.[1] We use heteroskedasticity robust standard errors. To test Hypothesis 1, we conduct a one-sided test based on the normal distribution with $\beta_{early} \geq 0$ as the null hypothesis. To test Hypothesis 2, we conduct a one-sided test based on the normal distribution with $\beta_{early} \geq \beta_{late}$ as the null hypothesis.

For testing Hypothesis 3, we run a similar regression as specified in equation (1). The only difference is that we use the average number of visits per week during the period

[1] The randomization is stratified by prices (with five levels: €19.90–19.95, €24.95, €29.90–29.95, €34.95, and €39.95) and sex (with two levels: male or no male).

from July 1, 2023, until October 24, 2023, as the outcome variable.[2] We then conduct a one-sided test based on the normal distribution with the null hypothesis $\beta_{early} \geq \beta_{late}$.

3.2 Secondary

We are planning the following secondary analyses to shed more light on the effect of the intervention as well as underlying mechanisms. First, we additionally compare the frequency of exercising in May and June 2023 between the treatment Late Discount and the control group.

Second, we estimate differences between the experimental groups for the period of the late discounts (November–December 2023) and the observation period thereafter. In this context, we will also investigate to which extent the presence of a reminder email at the end of October plays a role.

Third, we check whether the discounts have any effect on the probability of staying a member, i.e., whether the initial contract is prolonged or followed by a new contract. We focus on individuals who have signed a 12-month contract in March 2023 (implying a contract period from April 1, 2023, until March 31, 2024) and use the indicator of having a valid contract in April 2024 as an outcome variable.

As a robustness check, we repeat the main analysis specified in Section 2.1 focusing only on individuals who have been identified as opening the first email. The data on whether the email has been opened will be provided by the field partner. Failing to open the email decreases the probability that treated individuals are aware of their discount, which would prevent a reaction to the reduction in sunk costs.[3] Since the subject line of the email is constant across experimental groups, focusing on individuals who open the email should not be compromised by different opening rates.

4 Power Analysis

4.1 Summary

As specified in the pre-registration, our planned sample consists of 3,071 individuals and we already know their distribution across the different strata used for the randomization.

[2]At the end of October, randomly selected subsamples from the treatment Late Discount and the control group receive an additional email (see Experimental Design Details of the pre-registration). The considered period for testing Hypothesis 3 is chosen such that it ends before this additional email is sent out.

[3]However, even if individuals are not identified as opening the email, they might be aware of the discount because they have seen the reduced payment on their bank account or it has just not been captured that they actually opened the email.

We conduct several simulations to assess the statistical power of our empirical analysis under this sample size and randomization procedure.

For this purpose, we have received data on the number of visits in March 2023 of all members who had an active contract at the fitness chain from March 1 until March 31, 2023. The data are supplemented by information on the contract type and individual characteristics. The average number of weekly visits across all individuals in the data set amounts to 0.83 with a standard deviation of 1.15. Since these data include individuals who have been members for a long time and individuals with short-term contracts of only one month, we focus on almost $5,000$ individuals whose contract has a duration of 12 or 24 months and started on February 1, 2023.[4] The average number of weekly visits in this selected sample is 1.27 with a standard deviation of 1.24.[5]

A crucial factor for assessing the power is that some individuals in treatments Early Discount and Late Discount might not be aware of their discount, e.g., because they never open the first email and do not carefully check their bank account. These individuals cannot react to the reduction in sunk costs. Given the reported opening rates of previous emails sent by the field partner to its members, we expect the email opening rate to be around 50%.

Based on the simulations described in detail in Section 4.2, we are confident that we have at least 80% power to identify a sunk-cost effect in Hypothesis 1 at a 5% significance level if 50% of the members in the treatment Early Discount are aware of their discount and as a result reduce their weekly visits during the discount period by 0.25 (or less if this would imply a reduction below zero).[6] This scenario implies at most one visit less per month for those who are aware of the discount and an average reduction of approximately 0.4 visits per month across all individuals in the treatment Early Discount compared to the control group. Any scenario where either the share of individuals who are aware of the discount or the assumed effect on their behavior increases is of course even higher powered. If instead one of the parameters decreases compared to the presented scenario, the change in power depends on the extent to which the other parameter counteracts this decrease. For example, a weaker decrease of 0.125 in the weekly visits for individuals who are aware of the discount might be identified with roughly 80% power if the share of these

[4]Hence, March 2023 is their second month of the contract period. This mirrors the situation for most of our participants in the field experiment: They start their contract on April 1, 2023, and the first month for which we analyze the effect of the intervention on their visits per week is May 2023, their second month of the contract period.

[5]Selecting other individuals with recent contracts gives comparable statistics. For example, focusing on individuals with a 12-month or 24-month contract that has started on January 1, 2023, gives an average number of weekly visits in March 2023 of 1.13 with a standard deviation of 1.28.

[6]Those individuals who would go less than 0.25 times per week in the absence of the treatment are assumed to reduce their weekly visits to zero when being treated. Hence, their weekly visits even decrease by less than 0.25.

individuals reaches about 80% instead of 50%.

The power to identify a more pronounced sunk-cost effect for the treatment Early Discount than for the treatment Late Discount in the first two months after the intervention has started (Hypothesis 2) is also around 80% in the previously described scenario, when assuming that the sunk-cost effect of the treatment Late Discount is close to zero. The power for Hypothesis 3 is similar if we assume that these sunk-cost effects persist during the first six months (i.e., the period before the late discount is implemented). If instead, there is a small sunk-cost effect for individuals in the treatment Late Discount in the first six months, the sunk-cost effect for individuals in treatment Early Discount needs to become stronger to sustain a power of at least 80%.

4.2 Details

In the first simulation approach, we use the observed weekly visits in the data for March 2023 as representative distribution of the weekly visits for the two different periods considered in the primary empirical analysis (May–June 2023 and July–October 2023). Each simulation round proceeds according to the following steps:

1. Draw a random sample of 3,071 observations with replacement from the data for March 2023 (taking into account the distribution of strata in the true sample of our field experiment).
2. Randomize the observations into the three experimental groups according to the pre-registered randomization procedure.
3. Use the data from step 1 to simulate baseline levels ($y_{i,April}$).
4. Use the data from step 1 and the treatment assignment to calculate the outcome variable ($y_{i,May-June}$ or $y_{i,July-October}$) based on (i) an assumed share of individuals in the treatment group who will be aware of the discount (s_{aware}) and (ii) assumed treatment effects for those who are aware of the discount (τ_{early} and τ_{late}).
5. Conduct the pre-registered analyses and store whether the null hypothesis of a given test has been rejected.

These steps are repeated 1,000 times to calculate the power based on the empirical rejection rate across the simulations. To simulate the baseline levels ($y_{i,April}$), we follow three different approaches. First, we simply assume that the baseline level is equal to the initially drawn weekly visits plus a random variable drawn from a discrete uniform distribution. This uniform distribution either ranges from -1.5 to 1.5 with steps of 0.25 (Setting A) or from -2 to 2 with the same step size (Setting B). Note that a change in the weekly visits by 2 implies that the individual visited the gym 8 times less or more often per month.

Second, we use data from the field experiment in Habla and Muller (2021) to derive an empirical distribution of the change in weekly visits between two months. In particular, we focus on the weeks of February 2017 vs. March 2017 (Setting C) or December 2016 vs. June 2017 (Setting D). The latter is based on the available months with the largest distance and should thus serve as a conservative measure. We construct the baseline level by randomly drawing a value from the derived distribution and subtracting it from the number of weekly visits drawn in step 1.

Finally, we also include a worst-case scenario where the baseline level has no predictive power at all (Setting E). This scenario is implemented by randomly drawing a baseline level from the distribution of weekly visits analogously to, but independently from, step 1 of a simulation round.

Figures 1 to 5 show the simulation results for different scenarios either varying the share of individuals in the treatment Early Discount and Late Discount that are aware of their discount (s_{aware}) or varying how these individuals are affected by the discounts (τ_{early} and τ_{late}). Note that in all steps, we account for the fact that individuals can never reduce their weekly visits below zero. Furthermore, it is crucial to understand that τ_{early} and τ_{late} refer to the effect on those who are aware of the discount (given they are not limited in reducing their weekly visits due to low attendance in the absence of treatment). The average effect of being assigned to treatment Early Discount or Late Discount is thus substantially smaller because the share of individuals who are not aware of the discount will not react to it and some who are aware of the discount might not be able to reduce their weekly visits by that amount (but just by whatever is left to reach zero visits per week).

The figures also include the results from a second simulation approach (Setting F). There, we use the distribution of weekly visits in the provided data from the field partner for March 2023 and the changes in weekly visits across time in Habla and Muller (2021) to simulate data for each month in a seven-month period. This allows us to take into account additional fluctuation across months that might feed into the outcome variables defined in Section 3.1. As a starting point, we derive the empirical distribution of changes in average weekly visits from the second month (January 2017) in the data of Habla and Muller (2021) to each of the other months (1st to 7th month, i.e., December 2016–June 2017). We then use these empirical distributions in our simulation where the data from the field partner for March 2023 are used as the distribution of non-treated weekly visits in the second month (the first month after the intervention has started). The simulation steps look as follows:

1. Draw a random sample of 3,071 observations with replacement from the data for March 2023 (taking into account the distribution of strata in the true sample of our field experiment).
2. Randomize the observations into the three experimental groups according to the pre-registered randomization procedure.
3. Use the data from step 1 and the empirical distribution for changes in weekly visits between the 1st and 2nd month to simulate baseline levels.
4. Use the data from step 1 and the treatment assignment to calculate the weekly visits for the 2nd month based on (i) an assumed share of individuals in the treatment group who will be aware of the discount (s_{aware}) and (ii) assumed treatment effects for those who are aware of the discount (τ_{early} and τ_{late}).
5. Use the calculated visits from step 4 to simulate the weekly visits for the 3rd to 7th month by applying the empirical distribution of changes between the 2nd month and the 3rd to the 7th month.
6. Calculate the outcome variable for the different analyses, i.e., the average weekly visits for months 2 to 3 (outcome variable for hypotheses 2 and 3) and the average weekly visits for months 4 to 7 (outcome variable for Hypothesis 3).
7. Conduct the pre-registered analyses and store whether the null hypothesis of a given test has been rejected.

These steps are again repeated 1,000 times to calculate the power based on the empirical rejection rate across the simulations.

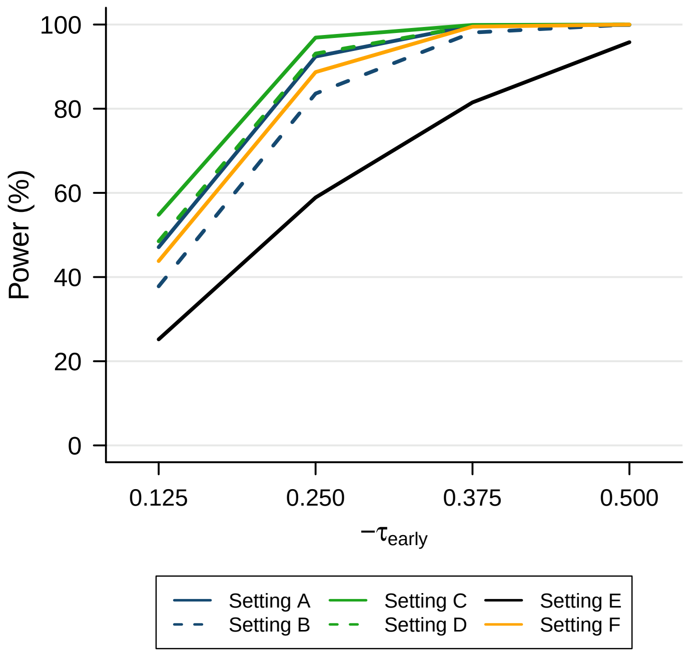

Figure 1: Power Simulations ($s_{share} = 0.5$) – Hypothesis 1

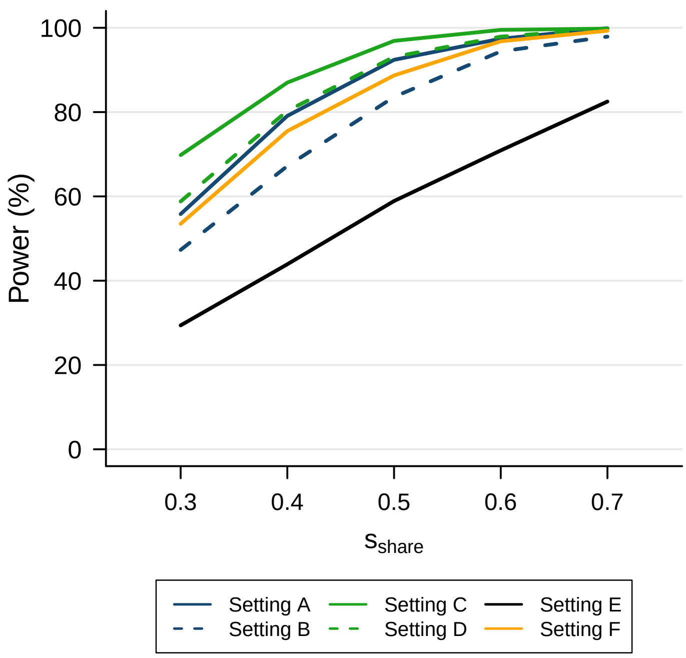

Figure 2: Power Simulations ($\tau_{early} = -0.25$) – Hypothesis 1

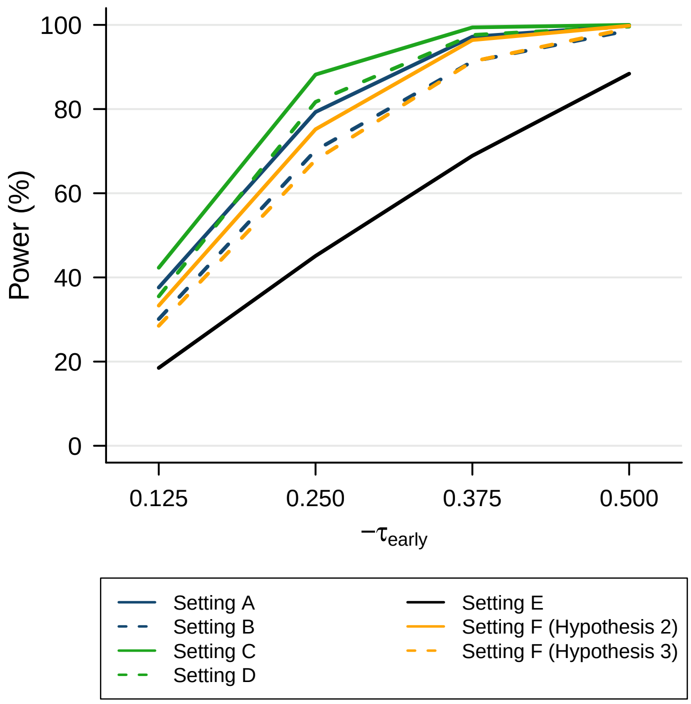

Figure 3: Power Simulations ($s_{share} = 0.5$, $\tau_{late} = 0$) – Hypotheses 2 and 3

Notes: In settings A to E, the power for Hypothesis 2 is exactly the same as for Hypothesis 3.

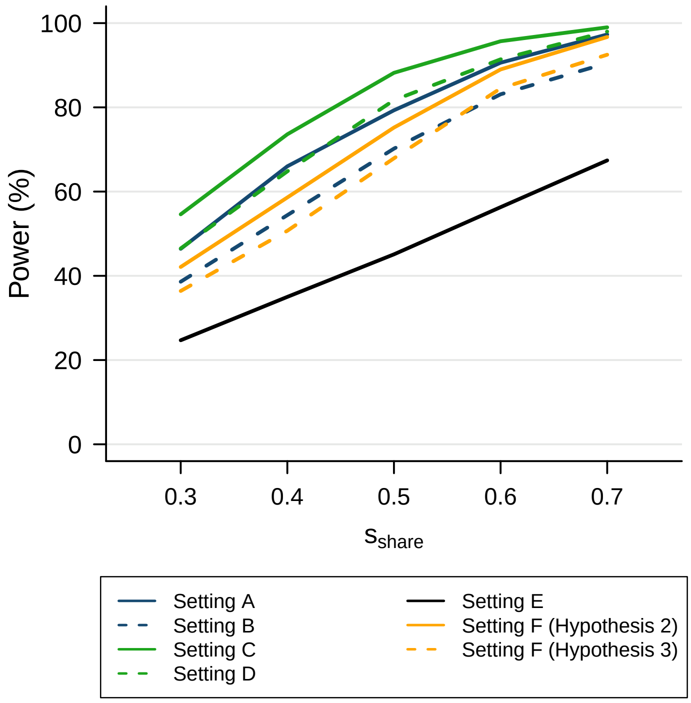

Figure 4: Power Simulations ($\tau_{early} = -0.25$, $\tau_{late} = 0$) – Hypotheses 2 and 3

Notes: In settings A to E, the power for Hypothesis 2 is exactly the same as for Hypothesis 3.

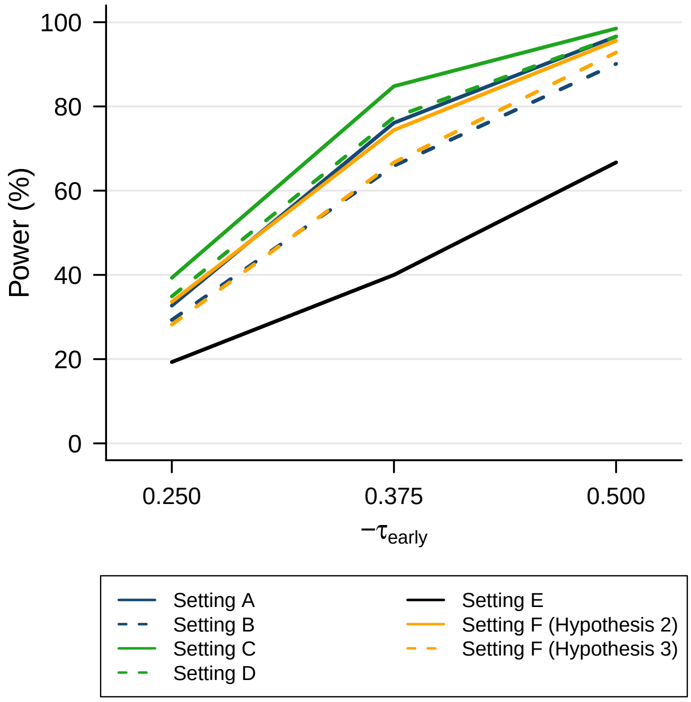

Figure 5: Power Simulations ($s_{share} = 0.5$, $\tau_{late} = -0.125$) – Hypotheses 2 and 3

Notes: In settings A to E, the power for Hypothesis 2 is exactly the same as for Hypothesis 3.

References

Habla, Wolfgang, and Paul Muller. 2021. "Experimental Evidence of Limited Attention at the Gym." *Experimental Economics* 24 (4): 1156–1184.