Pre-analysis plan Workplace giving 3

Note this pre-analysis plan was edited for layout prior to final publication, the content has not changed.

Policy problem

Workplace giving (WPG), also referred to as payroll giving, provides employees with an automated way to donate to their charity of choice. In terms of this trial, when we refer to WPG we are referring to regular giving, that is giving an amount each pay period. (While it is also possible to make one of donations via the WPG system, we do not include such donations as WPG for this purposes of this study).

Participation in WPG is low with the national average at 4.7 per cent. This is similar to the rate at our corporate partner for this trial.

Trial aim

This trial aims to look at the impact of present bias on participation in WPG. Specifically we will test the effect of encouraging staff at our corporate partner to sign up now to start giving at a point in the future.

Outcomes

The primary outcome for this trial will be WPG sign-up. This will be derived from a binary outcome variable where registering to donate for WPG during the trial period (for any amount over $0) will count as 1 and otherwise 0.

We will assess the average amount (in dollars) given by staff who donate through WPG as a secondary outcome.

We will also seek to collect data on whether people followed through on their registration and actually became a regular workplace giver. We will aim to collect data from late February/early March 2021, that indicates whether staff that registered are still giving after multiple pay periods. This should offer an indication of whether they will continue as longer-term workplace givers.

Interventions

All staff members at our corporate partner enrolled in this trial will receive an email encouraging them to sign up to WPG, plus a second email 11 days later. The trial will be a two-arm design, with two treatment groups (but no control –past trials have shown that

people rarely spontaneously sign-up for WPG, it tends to be as a result of outside events, such as a WPG drive). There will be 2 key data collection points (see Figure 1 below).

Treatment 1 – Group A – sequenced give later

1st email - Give now: Staff will receive an email that encourages them to sign up to start donating through WPG now.

2nd email - Give later: Staff will receive an email that acknowledges that now may not be a good time for them to start giving, but maybe they would like to register now to start giving in the New Year?

Treatment 2 – Group B – Up front give later

1st email - Give later: Staff will receive an email that encourages them to sign up now, but with donations not starting until the New Year (ie 2-3 months in the future).

2nd email – simple reminder: Staff will receive a simple reminder email about registering now to start WPG in the New Year.

Data collection points

The first data collection point will occur immediately before the second emails are sent on day 11. (That is, everyone who signs up between the first email being sent and the second email being sent will be seen as attributable to the first emails.) The second data collection point will be on day 16. This will mark the end of the end of the trial period. We will not be collecting signup data past this point.

Figure 1. Overview of trial flow

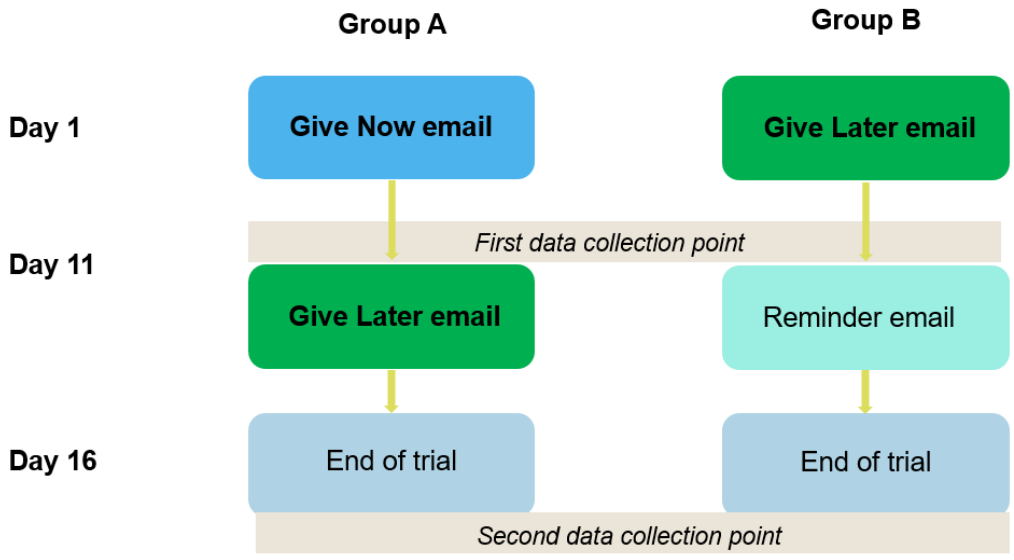

Hypotheses

H1. Give Later sign-up > Give Now sign-up (up to first data collection point)

H2. Give Later actually started giving > Give Now actually started giving (using follow-up data from New Year, but only those that sign-up before the first data collection point)

H3. Sequenced Give Later (ie Group A) ≠ Up front Give Later (ie Group B) (second data collection point)

H4. (secondary outcome) give later mean amount given > give now mean amount given (up to first data collection point)

Note that we do not propose to adjust for multiple hypothesis testing. We expect that the outcomes for H1 and H2 will be highly correlated.

Sample selection

All staff working at our corporate partner in the week of 12 October 2020 are eligible for inclusion, besides current workplace givers and staff involved in the design and implementation of the trial. This gives a sample size of 981 individuals.

Power calculations

The sample for the RCT will be 981 individual staff. We performed power calculations that indicated that at an alpha of 5% we will have 80% power to detect a standardised effect of 0.159. If we assume that post-intervention group A has a giving rate of 1% (based on results of our previous WPG trial, which also involved an email intervention) after the intervention, this would be equivalent to a 3.2% or a 3.2pp increase for group B.

Randomisation

Randomisation will be at the level of individual staff members. Participants will be assigned to either Group A or Group B using complete random assignment. Assignment will be balanced (that is, an equal number in each treatment group) to the extent that final participant numbers allow.

Randomisation will be implemented via an R script using the ‘complete_ra’ command from the ‘randomizr’ package . We will set a seed in order to ensure the reproducibility of the randomisation process. The randomisation process and code will be verified by another BETA staff member not directly involved in the project.

Trial threats

Blinding

Individuals enrolled in the trial will be aware of the interventions they are assigned to, but unaware that they are involved in a trial. BETA will perform the randomisation and analysis using a de identified dataset and will not be delivering the intervention.

Spillovers

There is a chance that individuals enrolled in the trial will compare emails with those in other treatment groups, however, we do not expect this will be widespread.

Attrition/missing data

From the point in time when the initial randomization data is gathered, through to the end of the trial period, some staff at our corporate partner will likely exit their jobs, or take leave, as a

part of the normal operations of the organisation. These individuals may not receive the intervention, and will be lost-to-follow up. We may not be aware of which staff members fall into this category. However, due to the timing of the trial (ie, mid-year, not during school holidays, weak labor market), we expect that only a small proportion of staff will be impacted. This attrition should be independent of treatment assignment. These individuals will be included in the final analysis as “non-givers”.

Post-trial survey

We may perform a survey on individuals enrolled in the trial after data collection ends. We will treat this survey as exploratory.

Analysis

The principal analysis of the effect of the intervention will be intent-to-treat and will consist of a covariate-adjusted comparison of our primary outcome for our two arms (corresponding to H1, see 4. Hypotheses). This estimate, confidence intervals and p values will be derived from a linear regression model with the following specification:

$$y = \alpha + \tau A + \delta x + \gamma x A + \epsilon$$

The coefficient on A is the main effect of the ‘give later’ message, and x is a vector of mean centred covariates indicating individuals wage. These variables will be interacted with treatment indicators.

Our hypothesis test for H1 will be one sided (see ‘Section 4. Hypotheses’). See section 13 for a discussion on the interpretation of resulting p-values.

We will run a similar analysis for H2, comparing the eventual giving outcomes for the two groups.

The analysis will also be similar for H3, where we compare the impacts of the sequenced give later message vs the upfront give later message. In this case A will be the effect of the upfront give later, while the vector x remains the same. The hypothesis test will be two-sided for this analysis.

For our secondary analysis, we will again use the same regression in order to estimate the effect of the treatments on average amount donated (H4).

We will calculate robust (HC2) standard errors for all linear models. Because our primary outcome is binary, we will run a robustness check using an equivalent logistic regression specification. We will calculate and report average marginal effects from this model.

Interpretation

We will make use of p-values to aid in the interpretation of our results. However, we will avoid taking a ‘bright line’ approach, in which a threshold (in this case the conventional 0.05) is used to determine a meaningful finding. Instead, we will consider the p-value together with prior evidence, effect size, outcome variability and design limitations in order to assess the strength of a finding.

There is some small risk associated with acting upon a false positive result for this intervention – delaying the start of the giving will likely mean lower overall donations if there is no increase in giving prevalence (depending on whether there is a large enough increase in average giving amounts). Therefore, we will be cautious in interpreting the p value.

We expect that the delayed giving intervention will have a positive effect. This is based on the substantial evidence that present bias impacts decision making. Due to the nature of the intervention (delayed giving), we have little interest in differentiating between no significant impact and a negative impact. That is, a result of no significant increase in giving prevalence is a poor result because the delay in giving will lower overall donations compared to the control. Therefore, we will make use of a one-tailed significance test.

Reporting

We will report the n, group means or proportions for all treatment groups on our primary and secondary outcomes. We will also report average treatment effects, 95% CIs and p-values for all comparisons and hypothesis tests that we run.

Pre-analysis plan commitments

- No trial data have been collected/no analysis has been undertaken prior to the completion of this pre-analysis plan.
- We will be transparent about, and provide justification for, any deviations (additions or omissions) from this plan.