the behaviouralist

February 2020

Using experiments to inform OBIE's Customer Experience Guidelines for VRP and Sweeping

Preliminary Analysis

Executive summary

This report presents preliminary findings from two large-scale (and nationally representative) experiments conducted by The Behaviouralist on behalf of the Open Banking Implementation Entity (OBIE). The objective of the experiments is to inform OBIE's Customer Experience Guidelines (CEGs) related to Variable Recurring Payments (VRP) and Sweeping.[1]

The first experiment examines customer preferences related to Sweeping. More specifically, it measures the extent to which customers prefer Sweeping propositions facilitated by VRP relative to other transfer methods (direct debit, free but infrequent transfers using a debit card, or paid-for but frequent transfers using a debit card). The main findings are that:

1. Around 39% of customers said that they would use the service facilitated by VRP.
2. The other options performed significantly worse than the VRP proposition. For example, those in the direct debit group were 5 percentage points less likely to say that they would like to download the app. The worst performing proposition was the paid-for debit card service.

The second experiment investigates how VRP Sweeping app journeys can be designed in order to maximise customer well-being. This experiment involved: 1) asking users to click through a prototype version of a savings app (called the 'Saving Robot') and set up (hypothetical) VRP Sweeping transfers, 2) randomly varying elements of the app (e.g., the ways in which terms and conditions are displayed), and 3) recording a series of outcomes (e.g., whether users will download the app if it becomes available and if they comprehend their rights as customers). More specifically, the users were asked to click through the following app journey:

1. Proposition screen: shows users how much they could save and asks if they want to set up an automated saving service.
2. Payment setup screen: customers are asked to set up their transfer by, for example, deciding how often money can be transferred, the duration of the VRP, and a threshold balance above which money can be transferred (this is referred to as setting up 'consent parameters').
3. Permission screen: shows customers a summary of the automatic transfer that they are setting up, displays terms and conditions, and asks them to confirm whether they would like to proceed.

[1] VRP refers to propositions where customers give long-lived consent to a Payment Initiation Service Provider (PISP) to initiate ongoing payments of variable amounts on their behalf; Sweeping refers to the automated movement of funds on behalf of a customer between two accounts in their name.

4. ASPSP screen: customers are transferred to their bank app and asked to confirm the set-up.
5. Complete screen: customers are told that they have set up the automatic transfers.
6. Notification screen: customers are notified before/after a payment is being made using the VRP they set up.

The experiment involved varying the content displayed in these screens. Users were first randomly allocated to one of the following three groups (in each case, the content was varied when on the payment setup screen, and their choices were also displayed on the permissions screen):

1. Consent parameter group 1: users are asked to set the maximum number of transactions per week and the maximum amount that can be transferred per week.
2. Consent parameter group 2: users are only asked to set the maximum amount that can be transferred per week.
3. Consent parameter group 3: users are asked to set the maximum amount that can be transferred per week and the maximum amount that can be transferred per payment.

They were then, once again, randomly allocated to one of three groups:

1. T&Cs group 1: users are shown the T&Cs on the permissions screen.
2. T&Cs group 2: users are shown T&Cs as a pop-up while on the permissions screen.
3. T&Cs group 3: users have to click a link on the permissions screen to read the T&Cs.

Next, they were randomly allocated to a group that was asked to 'confirm' their consent when setting up the VRP (when on the ASPSP screen) or to a group where consent was automatically provided using a biometric fingerprint reading. Finally, they were allocated to three 'notification' groups that were shown:

1. Notification group 1: no notification.
2. Notification group 2: pre-transfer notification after the VRP has been set up.
3. Notification group 3: post-transfer notification after the VRP has been set up.

The main results from experiment 2 are that:

1. The app variants performed roughly as well in terms of customer comprehension of their rights if they suffer material damage (i.e., there are no statistically significant differences between groups). Further, we do not see any significant differences in terms of customer comprehension of their rights.
2. The post-transfer notification is most popular when individuals are asked which they would prefer. Further, customers are around 2 percentage points more likely to click

the post-transfer notification than the pre-transfer notification (although this difference is not statistically significant).

3. Those in consent parameter group 1 ('maximum number of transactions per week + maximum amount per week') were significantly less likely to understand the consent parameters (i.e., individuals in this condition were the worst at identifying when the parameters were violated in hypothetical payment scenarios).
4. Most users were positive toward the Saving Robot app. Moreover, those in T&Cs group 2 (pop-up T&Cs) found it more difficult to understand the information conveyed in the app than those in other T&Cs groups. Further, the post-payment notification made users feel like the information in the app was easier to understand and made them less likely to state that the app 'was unnecessarily complicated'.
5. Around 36% said that they would download the app when asked, and around 32% said that they would recommend it to family or friends. However, the pre-emptive notification reduced users' willingness to recommend the app by around 3 percentage points.
6. Users were asked about how easy they thought it would be to cancel a payment and how worried they are that the app would make a transfer without their knowledge. Adding a CTA (i.e., a 'confirm' button) when providing consent in the ASPSP app had a positive effect on a range of customer perceptions, for example making individuals 2 percentage points less likely to say they are worried that the app would make a transfer without their knowledge. The only other variant to have a significant effect on these outcomes was T&Cs group 3 (the 'link' variant), which made people around 4 percentage points less likely to say that it would be easy to move money back from the Saving Robot to their current account.

While the findings are preliminary (we have recruited 6,000 out of 10,000 participants), they are robust. In other words, we have sufficient statistical power to detect differences between groups, and the additional sample will primarily be used to facilitate segment analyses (i.e., examining how effects differ for various sub-groups of the population). Further analyses will be conducted when the full sample has been collected.

Table of contents

1 Introduction

This report presents preliminary findings from two large-scale experiments conducted by The Behaviouralist on behalf of OBIE. The objective of the experiments is to inform OBIE's Customer Experience Guidelines (CEGs) related to Variable Recurring Payments (VRP) and Sweeping.[2]

The first experiment examines customer preferences related to Sweeping. More specifically, it measures the extent to which customers prefer Sweeping propositions facilitated by VRP relative to other transfer methods (e.g., direct debit or having a debit card on file).

The second experiment investigates how VRP Sweeping app journeys can be designed in order to maximise customer well-being. This experiment involved: 1) asking users to click through a prototype version of a savings app and set up (hypothetical) VRP Sweeping transfers, 2) randomly varying elements of the app (e.g., the ways in which terms and conditions are displayed), and 3) recording a series of outcomes (e.g., whether users will download the app if it becomes available and if they comprehend their rights as customers).

Both experiments were conducted online with a nationally representative sample of UK adults. The target sample size for both experiments is 10,000 individuals; this report presents results from analyses conducted with around 6,000 individuals. Fortunately, the full sample of 10,000 individuals is not required in order to conduct analyses comparing aggregate differences between experimental groups (this can be done with a sample of 5,000 individuals).[3] The full sample of 10,000 individuals is only required in order to conduct cohort/segment analyses (e.g., comparing the efficacy of different 'treatments' for younger and older users). The remaining sample will be recruited throughout February 2021.

Randomised experiments are often considered to be the most robust (or the 'gold standard') evaluation methodology, and are used by researchers, policymakers, and product developers alike. Randomising individuals to 'treatment' and 'control' groups ensures that the groups are 'statistically identical' at baseline, meaning that it is possible to interpret differences in outcomes between the groups as the *causal* effect of the treatment. This type of causal interpretation is typically not possible with observational (non-experimental) data, or when conducting qualitative analyses (e.g., focus groups).

[2] VRP refers to propositions where customers give long-lived consent to a Payment Initiation Service Provider (PISP) to initiate ongoing payments of variable amounts on their behalf; Sweeping refers to the automated movement of funds on behalf of a customer between two accounts in their name.
[3] Please see the Methodology section for information about sample size requirements and power calculations.

The experiments replicate real-world decision environments, meaning that they provide generalisable results that are likely to hold in other settings. Further, the analysis is conducted using rich data on user attitudes, customer satisfaction, and incentivised questions that measure customer comprehension of the transfer they set up.

The remainder of the document is structured as follows. Section 2 discusses the experimental methodology. Section 3 concludes by presenting and discussing the results from the experiments.

2 Methodology

2.1 Research questions

The experiments were designed to answer a series of research questions, including:

1. Do customers prefer Sweeping propositions that are facilitated by VRP, direct debit, card-on-file (free, but with infrequent payments), or card-on-file (paid-for service with frequent payments)?
2. Does altering the 'consent parameters' (e.g., the maximum amount that can be transferred per week) associated with a VRP Sweeping proposition influence customer satisfaction? Do customers understand some consent parameters better than others?
3. How should Terms & Conditions (T&Cs) be displayed to customers? Do customers better understand their rights if T&Cs are displayed in-screen as they complete a VRP Sweeping journey, as a pop-up during that journey, or if they can be accessed via a link? Is there a trade-off between customer comprehension of their rights and customer satisfaction?
4. Does including a 'Call to Action' (CTA) to confirm consent when customers are directed from the PISP app to their ASPSP app change customer perceptions about the VRP Sweeping transfer they just set up?
5. What are the effects of including notifications that inform customers: a) prior to a transfer being made using VRP and b) after a transfer has been made using VRP? Do customers prefer one type of notification over the other?

Question 1 is answered in the first experiment. Questions 2–5 are answered in experiment 2.

2.2 Experimental design

Both experiments took place within an online survey that was fielded in January/February 2021. Participants were recruited via Prolific Academic and Panelbase; the sample is nationally representative of the UK adult population. The target sample size is 10,000 (recruitment is not yet complete). The experiments were conducted using a custom-coded Qualtrics survey.

Upon entering the survey, participants answered a series of questions about themselves, their savings habits, and their use of similar apps. Participants were told:

> In this survey, we will ask you a series of questions about a new savings app called Saving Robot. Saving Robot will help users save money by automatically moving money between different bank accounts. This will ensure that users:
>
> 1) avoid paying unnecessary fees, and
> 2) maximise their interest rate earnings.
>
> The Saving Robot will be able to analyse your spending and determine if money should, for example, be moved from your current account to your savings account.
>
> The Saving Robot will be compliant with all relevant UK regulations. Users will need to grant the Saving Robot access to their bank accounts in order to make the automatic transfers. The automatic transfers can be cancelled at any time.
>
> Users will need to provide the app with some restrictions when setting up the automatic transfers (e.g., limiting the amount of money that can be moved during a given month).

They were then randomised into one of four groups for the first experiment. Each group was presented with a description of a hypothetical sweeping service, where we randomly varied the features of the service (e.g., whether it was administered via direct debit or via a debit card). Those in the first group were shown:

> The example below illustrates how the Saving Robot will operate. Please read this text carefully.
>
> When setting up the app, you instruct the Saving Robot to move money from your current account to your savings account if you have more than £2000 in your current account. Also:
>
> - The Saving Robot is only able to make one transfer per month and makes this transfer at the end of every month (the money usually arrives in your savings account within 24-48 hours).
> - These transfers are made using direct debit.
> - The Saving Robot service is free.
>
> For example, if you end the month with £2500 in your current account, the Robot will automatically transfer £500 to your savings account at the end of the month.

Those in the second group were shown:

> The example below illustrates how the Saving Robot will operate. Please read this text carefully.
>
> When setting up the app, you instruct the Saving Robot to move money from your current account to your savings account if you have more than £2000 in your current account. Also:
>
> - The Saving Robot is only able to make one transfer per month and makes this transfer at the end of every month (the money usually arrives in your savings account within 24-48 hours).
> - These transfers are made using a debit card that is linked to your current account.
> - The Saving Robot service is free.
>
> For example, if you end the month with £2500 in your current account, the Robot will automatically transfer £500 to your savings account at the end of the month.

Those in the third group were shown:

> The example below illustrates how the Saving Robot will operate. Please read this text carefully.
>
> When setting up the app, you instruct the Saving Robot to move money from your current account to your savings account if you have more than £2000 in your current account. Also:
>
> - The Saving Robot is able to move money between accounts on a daily basis, and the money usually arrives in your savings account within 24-48 hours.
> - These transfers are made using a debit card that is linked to your current account.
> - The Saving Robot service costs you £5 per month.
>
> For example, if you have £2500 in your current account, the Robot will automatically transfer £500 to your savings account.

Finally, those in the fourth group were shown:

> The example below illustrates how the Saving Robot will operate. Please read this text carefully.
>
> When setting up the app, you instruct the Saving Robot to move money from your current account to your savings account if you have more than £2000 in your current account. Also:
>
> - The Saving Robot is able to move money between accounts on a daily basis, and the money arrives immediately in your savings account.
> - These transfers are made using OBIE.
> - The Saving Robot service is free.
>
> For example, if you have £2500 in your current account, the Robot will automatically transfer £500 to your savings account.

After having read this information, participants were asked: "Imagine that the Saving Robot app is launched next week. How likely is it that you would download and use the Saving Robot app?" and could respond on a 5-point Likert scale (from 'likely' to 'unlikely'). Responses to this question constitute the main outcome variable for the first experiment. Participants also responded to an open-ended follow-up question that asked them to justify their choice. As participants were randomised to different conditions, it is possible to deduce which type of Sweeping proposition they prefer by examining mean differences between groups.

Next, participants were provided with information about the second experiment:

> We will now ask you to click through a prototype of the Saving Robot.
>
> More specifically, we will ask you to click through an example which involves setting up automatic transfers from your current account to your savings account when you have more than £2000 in your current account.
>
> The purpose of the transfers is to maximise your interest rate earnings while minimising the payment of unnecessary fees. The Saving Robot will analyse your data and determine when it is suitable to make transfers.
>
> In this scenario, the automatic transfers will be made using OBIE. The Saving Robot service will be free of cost and the app can make several transfers per month.
>
> We will ask you a series of questions about the app once you have clicked through the prototype.
>
> You will be able to earn more money if you answer these questions correctly.
>
> It is therefore crucial that you pay attention when clicking through the app, and that you spend the same amount of time that you would spend when using a similar app in real life.
>
> Not all buttons and features in the app will be clickable, and some information will have been pre-filled. You will mainly be able to go back or forward between screens in the app.
>
> Please click next to begin testing the app.

These instructions were displayed across 3 pages, and participants spent a minimum of 20 seconds reading the instructions. The instructions were intended to explain the task that the participants would have to complete, while also incentivising them to spend approximately as much time walking through the app as they would in a similar real-life situation. (Indeed, when analysing the data, we find that participants spend roughly as much time clicking through the app as they would spend setting up a payment in real life.)

Participants were then randomised into different groups for the second experiment, which involved clicking through the Saving Robot app, and setting up a hypothetical VRP for sweeping. Each group was shown a different variant of the Saving Robot app. The randomisation took place in stages. First, participants were randomly allocated to groups that

were shown one of three possible sets of consent parameters when using the app. Second, participants were randomly allocated to groups that were shown one of three types of terms and conditions. Third, participants were randomly allocated to groups that either were, or were not, asked to click an additional call to action (CTA) in order to set up their payment. Fourth, participants were randomly allocated to groups that were shown a notification after a payment had been made using the service they just set up, a notification before a payment was made, or no notification at all. The randomisation is illustrated below:

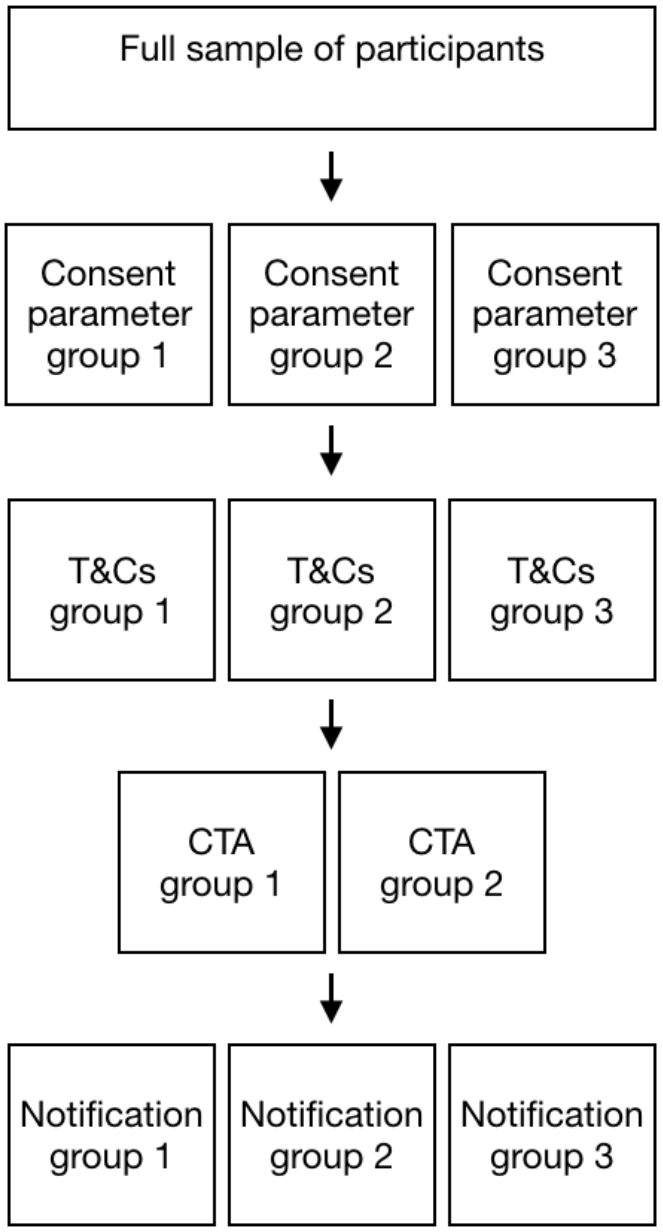

As such, some individuals were, for example, shown the following 'variant' of the app: 1) consent parameter group 1, 2) T&Cs group 2, 3) CTA group 2, 4) notification group 1. While this multi-layered randomisation meant that there were 54 distinct groups (one per permutation, as we have a 3 x 3 x 2 x 3 set-up), it is possible to pool these groups when

conducting the analysis (e.g., it is possible to compare the outcomes for everybody in the first T&Cs group to everybody in the second T&Cs group).

The Saving Robot app consisted of the following steps (the information provided within each step was randomly varied):

1. Proposition screen: shows users how much they could save and asks if they want to set up an automated saving service.
2. Payment setup screen: customers are asked to set up their transfer by, for example, deciding how often money can be transferred, the duration of the VRP, and a threshold balance above which money can be transferred (this is referred to as setting up 'consent parameters').
3. Permissions screen: shows customers a summary of the automatic transfer that they are setting up, displays terms and conditions, and asks them to confirm whether they would like to proceed.
4. ASPSP screen: customers are transferred to their bank app and asked to confirm the set-up.
5. Complete screen: customers are told that they have set up the automatic transfers.
6. Notification screen: customers are notified before/after a payment is being made using the VRP they set up.

The groups mentioned above are defined as follows:

1. Consent parameter group 1: users are asked to set the maximum number of transactions per week and the maximum amount that can be transferred per week (when on the payment setup screen; these choices are then displayed on the permissions page).
2. Consent parameter group 2: users are asked to set the maximum amount that can be transferred per week (when on the payment setup screen; these choices are then displayed on the permissions page).
3. Consent parameter group 3: users are asked to set the maximum amount that can be transferred per week and the maximum amount that can be transferred per payment (when on the payment setup screen; these choices are then displayed on the permissions page).
4. T&Cs group 1: users are shown the T&Cs on the permissions screen.
5. T&Cs group 2: users are shown T&Cs as a pop-up while on the permissions screen.
6. T&Cs group 3: users have to click a link on the permissions screen to read the T&Cs.
7. CTA group 1: consent is automatically given when the users enter the ASPSP (using biometric authentication).
8. CTA group 2: users need to click "confirm" in order to provide consent when they enter the ASPSP.
9. Notification group 1: no notification.

10. Notification group 2: pre-transfer notification after the VRP has been set up.
11. Notification group 3: post-transfer notification after the VRP has been set up.

The screens below provide an example of differences between the app variants. More specifically, they display how the 'payment setup' screen looked like in the app for the three respective consent parameter groups:

Consent group 1

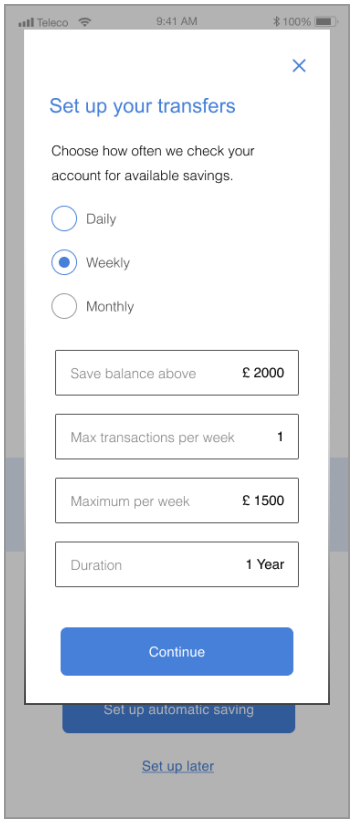

Consent group 2

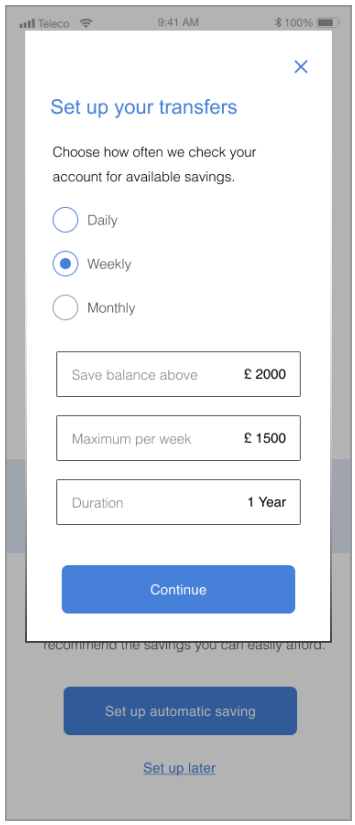

Consent group 3

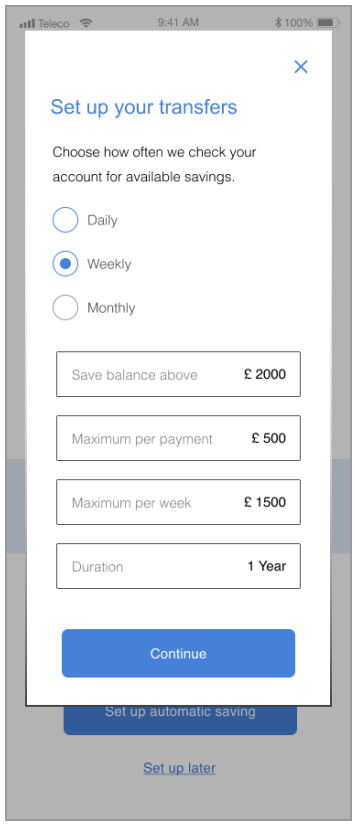

Once participants had completed the VRP sweeping journey, they were asked the following question:

> Imagine that you have a standing direct debit for your rent payments. These payments are made from your current account on the 15th of every month. Your rent is £2200 per month.
>
> In order to make your rent payments, you top up your current account on the 14th of every month, so that you have exactly £2200 in your account.
>
> Now imagine that the Saving Robot detects that you have over £2000 in your account (remember: you set the threshold at £2000) and moves £200 to your savings account on the 14th.
>
> You did not notice the Saving Robot transfer, and the direct debit subsequently fails to make your rent payment on the 15th. As a consequence, you were fined £100 in late payment fees.
>
> Do you think that you would have the right to ask the company responsible for the Saving Robot for compensation (i.e., that they pay for the £100 incurred in late fees) if this happened?
>
> - Yes
> - No
> - Don't know

This question captures the extent to which users understand that they have the right to ask for compensation if the app makes sub-optimal decisions that cause material harm. This question is included as it helps us understand the effects of the terms and condition variants.

They were then asked an incentive-compatible question related to whether they would ask for compensation in different scenarios:

> Now, imagine that you set up the recurring transfer that you were shown when clicking through the app.
>
> You have the right to ask the app company for compensation if the Savings Robot makes a transfer that breaks the terms and conditions that you agreed to.
>
> Please read over the transfers made on your behalf and determine whether you would:
>
> - Request compensation (will cost you £0.05 of your survey earnings per compensation request)
> - Do nothing (no cost)
>
> You will earn £0.10 (minus the cost of asking for compensation) for each request that is successful. These earnings will be added to your regular survey completion payment, and you can never earn less than £0.
>
> 1. The app makes a transfer from your Santander account
> 2. The app makes a transfer from your Standard Bank account
> 3. The app makes a transfer from your Barclays account
> 4. The app makes a transfer to your designated savings account in May 2022
> 5. The app makes a transfer to your designated savings account in March 2022
> 6. The app makes a transfer to your designated savings account in January 2022
> 7. The app makes a transfer that brings your main account balance down to £1000
> 8. The app makes a transfer that brings your main account balance down to £1600
> 9. The app makes a transfer that brings your main account balance down to £2100

This question helps understand the extent to which individuals comprehend the VRP sweeping service that they just set up, and how that depends on the app variant they were shown (e.g., the different T&Cs or consent parameters).

Finally, the survey concluded by asking a similar incentive-compatible question related to the consent parameters, a series of questions about how comfortable or worried they would feel about the service if they used the app in real life, questions about whether they would download and recommend the app, and a series of questions about the usability of the app.

2.3 Sample size calculations

It is crucial to have a sufficiently large sample—and adequate statistical power—when conducting experiments. Statistical power refers to the probability of detecting a difference between treatment conditions, assuming that there truly is a difference. Further, most academic studies aim to achieve a statistical power of at least 80%.

There are a number of factors that influence the statistical power of an experiment. These include:

1. The sample size
2. The average and standard deviation in the comparison group

3. The minimum detectable effect size (i.e., the smallest effect that can be detected with a given level of statistical power)
4. How the sample is distributed between groups (power is typically maximised when the sample is split evenly between conditions)
5. The significance level (typically set to 5%)

Most of the outcomes in this study are binary. This means that statistical power is the lowest if 50% of individuals in the comparison group answered 'yes' to that question (as this maximises the variance in the outcome). In order to conduct conservative power calculations, 50% is therefore used as the average in the control group. Further, the MDE is set to 5 percentage points, and the sample will be split evenly between groups. This means that a sample of 1,565 is required per group in order to obtain a statistical power of 80%.

The target sample size is 10,000, and the analyses in this report is based on approximately 6,000 responses. This means that it is possible to conduct robust statistical comparisons with the existing sample, as there are never more than 3 groups that will be compared to each other (e.g., the three T&Cs groups or the three consent parameter groups). It is, however, not yet possible to conduct segment analyses that examine how app variants influence different groups of the population.

2.4 Analytical approach

The experimental data is analysed using Ordinary Least Squares and Probit regressions. Most analyses in this report use the following type of specification:

$$y_i = b_0 + b_1 * CTA_i + \sum_{j=2}^{3} b_j * T\&C_{ji} + \sum_{j=4}^{5} b_j * Consent_{ji} + \sum_{j=6}^{7} b_j * Notification_{ji} + e_i$$

where y_i refers to the outcome of interest (often a binary variable), b_0 refers to the mean in the control group. In this case, the control group is the condition where individuals are allocated to the app variant with: 1) no CTA, 2) with the first T&Cs variant, 3) the first consent parameter variant, and 4) no notification. The remaining b_j variables represent to the treatment coefficients for the different app variants, CTA_i refers to whether individuals have been allocated to the second CTA group or not, $T\&C_{ji}$ refers to the T&Cs group an individual has been allocated to, $Consent_{ji}$ refers to the consent group that participants have been allocated to, and $Notification_{ji}$ refers to the notification group that participants have been allocated to. Finally, e_i refers to the error term.

The analysis is conducted using the statistical software Stata. Robust standard errors are used throughout the analysis.

3 Results

3.1 Experiment 1

Table 1 presents the results from experiment 1. The outcome of interest is whether participants said they would like to download the app if it became available. As can be seen, around 39% of those who were told that transfers would be made using VRP stated that they would like to download and use the Saving Robot app. The other options performed significantly worse than the VRP proposition. For example, those in the direct debit group were 5 percentage points less likely to say that they would like to download the app. The worst performing proposition was the paid-for debit card service.

Table 1. Consumer preferences

	(1) Are you likely to download the app?
Direct debit	-0.05***
	(0.02)
Debit card (free)	-0.01
	(0.02)
Debit card (not free)	-0.17***
	(0.02)
Constant	0.39***
	(0.01)
Observations	6045
R-squared	0.022

In table 1, the 'constant' represents the share of those who were assigned to the VRP group that would like to download the app. The other rows provide the *difference* between the share that would like to download the app in the VRP group and the three other groups. The stars indicate whether the results are statistically significant or not (one star means significant on a 10%-level, two stars means significant on a 5%-level and three stars means significant on a 1%-level). The estimation was conducted using a Linear Probability Model regression.

3.2 Experiment 2

The analysis of experiment 2 begins by examining the effects of the variants on (1) whether users thought they could ask for compensation if the Saving Robot made a transfer that caused material damage, (2) the share of comprehension questions users answered correctly (e.g., whether a given payment was in violation of the T&Cs they agreed to), and (3) the amount users earned when answering the comprehension questions (they earned more if they got the questions right, and lost money if they got the questions wrong, and earned/lost nothing if they did not choose a response).

The app variants performed roughly as well in terms of customer comprehension of their rights if they suffer material damage (i.e., there are no statistically significant differences between groups). Further, we do not see any significant differences in terms of customer comprehension of their rights.

Moving on to customer preferences regarding notifications, the analysis shows that the post-payment notification is most popular when individuals are asked which they would prefer. Further, customers are around 2 percentage points more likely to click the post-transfer notification than the pre-transfer notification (although this difference is not statistically significant).

The experiment also measures the effects of the variants on whether participants understand the consent parameters they were shown. The analysis reveals that those in the 'maximum transaction per week + maximum amount per week' consent parameter group were significantly less likely to understand the consent parameters (i.e., individuals in this condition were the worst at identifying when the parameters were violated in hypothetical payment scenarios).

Participants were asked questions about whether they felt like the Saving Robot app was easy to use, felt safe, was unnecessarily complicated, and whether the information conveyed in the app is easy to understand. Most users were positive toward the app (50-70% depending on the question). Further, those in the pop-up T&Cs group found it more difficult to understand the information in the app than those in other T&Cs groups. Finally, the post-payment notification made users feel like the information in the app was easier to understand and made them less likely to state that the app 'was unnecessarily complicated'.

Around 36% said that they would download the app when asked, and around 32% said that they would recommend it to family or friends. The only significant effect we find here is that the pre-emptive notification *reduced* users' willingness to recommend the app by around 3 percentage points.

Finally, the analysis examines how the variants affected participants' opinions about the app. Users were, for example, asked about how easy they thought it would be to cancel a payment and how worried they are that the app would make a transfer without their knowledge. The analysis shows that the CTA variant had a positive effect on most questions, for example making individuals 2 percentage points less likely to say they are worried that the app would make a transfer without their knowledge. The only other variant to have a significant effect on these outcomes was the T&C 'link' variant, which made people around 4 percentage points less likely to say that it was easy to move money back from the Saving Robot to their current account.