Australian Government

Department of the Prime Minister and Cabinet

BETA

Easy as pie: encouraging greater survey completion

June 2018

Who?

Who are we?

We are the Behavioural Economics Team of the Australian Government, or BETA. We are the Australian Government's first central unit applying behavioural economics to improve public policy, programs and processes.

We use behavioural economics, science and psychology to improve policy outcomes. Our mission is to advance the wellbeing of Australians through the application and rigorous evaluation of behavioural insights to public policy and administration.

What is behavioural economics?

Economics has traditionally assumed people always make decisions in their best interests. Behavioural economics challenges this view by providing a more realistic model of human behaviour. It recognises we are systematically biased (for example, we tend to satisfy our present self rather than planning for the future) and can make decisions that conflict with our own interests.

What are behavioural insights and how are they useful for policy design?

Behavioural insights apply behavioural economics concepts to the real world by drawing on empirically-tested results. These new tools can inform the design of government interventions to improve the welfare of citizens.

Rather than expect citizens to be optimal decision makers, drawing on behavioural insights ensures policy makers will design policies that go with the grain of human behaviour. For example, citizens may struggle to make choices in their own best interests, such as saving more money. Policy makers can apply behavioural insights that preserve freedom, but encourage a different choice – by helping citizens to set a plan to save regularly.

Contents

Executive summary

Including behaviourally-informed progress updates is a simple and cost effective way to increase survey completion rates.

The Community Grants Hub, implemented by the Department of Social Services (DSS), supports the administration and delivery of over 21,000 community grants across government. These grants, worth around $7.8 billion, provide valuable support services to individuals, families and their communities across Australia.

To enable the collection of meaningful data about service delivery outcomes from organisations receiving community grants, DSS introduced a new reporting platform called the Data Exchange. As part of this platform, DSS piloted a client survey to enable clients to self-report on why they sought assistance and the outcomes of accessing the service. We partnered with DSS to develop and test ways to encourage greater survey completion rates.

There are many reasons why people may not complete a survey in its entirety. People may not understand why the survey is important, may not anticipate how long it will take, or simply think the required effort is just not worth it.

To address these challenges, we incorporated small behaviourally-informed progress updates into the Data Exchange survey. Updates were presented in a form of a gradually completing pie-chart, accompanied by short messages with references to social norms and highlighting how completing the survey would benefit the individual and their community.

We ran a trial to test whether our re-designed survey worked. We found our changes had a significant impact, increasing survey completion by 6.8 percentage points, from 84.7 per cent to 91.5 per cent.

The outcomes of this trial show how simple low-cost changes to survey design may increase survey completion. With many government and private organisations using surveys to better understand and service their clients, the significance of this finding extends beyond future redesigns of the Data Exchange survey and should be considered by other organisations looking to survey their clients.

Why?

Policy context

The Community Grants Hub, implemented by DSS, delivers grant administration services on behalf of Australian Government departments, agencies and organisations who deliver community grants to individuals and the community sector.

These funds allow organisations to focus on what they do best: help people and communities facing difficult circumstances. Services offered by these organisations assist families, children and older people, enhance the wellbeing of people with high needs, assist those who need help with care, and support a diverse and harmonious society (DSS, 2017).

The Data Exchange reporting platform revamps the way DSS collects and uses administrative data, shifting the focus from outputs to more meaningful information about service delivery outcomes. The Data Exchange client survey provides a complementary data source on outcomes achieved by clients. The survey enquires about nine broad areas of a client's life and provides DSS funded organisations with key client insights. This information helps improve services being delivered and ultimately leads to better outcomes for clients, families and their communities.

The problem

In recent years there's been a decrease in the number of people who undertake and complete surveys (Australian Bureau of Statistics, 2010). This decrease is of concern as government and private organisations regularly use surveys to better understand their clients and gauge the effectiveness of their services.

High Data Exchange survey completion rates are important for DSS to ensure results reflect representative views and opinions of all the clients of a particular cohort (Fogliani, 1999).

BETA partnered with DSS' Community Grants Hub to explore ways of encouraging more people to complete the Data Exchange survey.

What?

A number of behavioural factors are relevant when completing a survey. We tested whether incorporating small behaviourally-informed progress updates into the Data Exchange survey would increase the completion rate.

Behavioural Analysis

A number of behavioural biases are at play when completing or choosing to complete a task such as a voluntary survey. These biases can prevent people from finishing the task or choosing to participate at all. People may not understand why completion is important, may not think it is relevant to them, may not anticipate how long it will take, put it off to complete later, or simply think the required effort is just not worth it. We outline biases underpinning these assumptions in Box 1.

Box 1 – Behavioural biases impacting survey completion rates

Present bias is when people overweight the immediate implications of their decisions compared with the impacts they have in the future. There is an immediate cost incurred in terms of time and energy spent, particularly when the benefits are uncertain and occur in the future.

Procrastination can occur because people are short-sighted and often put off decisions or behaviours, even those in their best interests.

Uncertainty bias may stop people committing to an action or a decision when they are unfamiliar with processes, expectations or systems. In the context of the Data Exchange survey, being asked to complete an unexpected survey may cause people to disengage due to uncertainty around the new task.

Simple behavioural insights techniques have been shown to be effective in addressing these biases. There is strong evidence to suggest surveys which are personalised, anticipated, concise, and/or use social norming messages are more likely to be completed than their non-modified counterparts (Edwards, et al., 2009; Nulty, 2008).

Box 2 summarises key behavioural techniques found to be effective in encouraging greater survey participation and completion rates.

Box 2 - Behavioural science techniques to increase survey completion rates

Social norms are behavioural expectations or rules within a group of people signalling appropriate behaviour. Social norms can have a significant effect on people's behaviour and influence their actions. A number of studies have shown people are more likely to complete a survey once presented with a social norming message.[1]

Salience is the act of making the state or quality of something stand out relative to its neighbours. Highlighting the benefits of completing a survey may mitigate **present bias** and encourage more people to continue with the survey.

Altruism is the desire to help others without gaining anything in return. More people may be inclined to complete the survey if they understand their responses improve program service delivery in their communities.

Pseudo sets theory suggests arbitrarily grouping items or tasks together as part of an apparent 'set' of items or tasks motivates people to reach perceived completion points: to finish a pseudo-set. This provides a form of goal-setting to encourage people to reach a target (Barasz, et al., 2017). Adapted to provide **progress updates**, this approach may be used to encourage more people to complete a greater portion of the survey.

BETA worked with DSS to incorporate these behavioural insights into the design of the Data Exchange survey and then tested the response.

---

[1] See Edwards and others (2009) for a systematic review of methods to increase response rate to postal and electronic questionnaires.

The study

We ran a trial to see whether providing behaviourally-informed progress updates would encourage more people to complete the Data Exchange survey. We also intended to test whether inserting a pre-commitment page at the end of the first survey would lead to more people participating in a second, follow-up survey.

The study was designed as a randomised controlled trial (RCT), which involved splitting people into three groups, with two groups completing our redesigned surveys (either with the behaviourally-informed progress updates or the pre-commitment page) and the control group completing the standard survey. For more information on how RCTs work, see Box 3.

We intended participants would be randomly allocated into the three different groups but the actual allocation mechanism was not strictly random. We nonetheless consider the allocation was as good as random: see the Limitations section for further discussion.
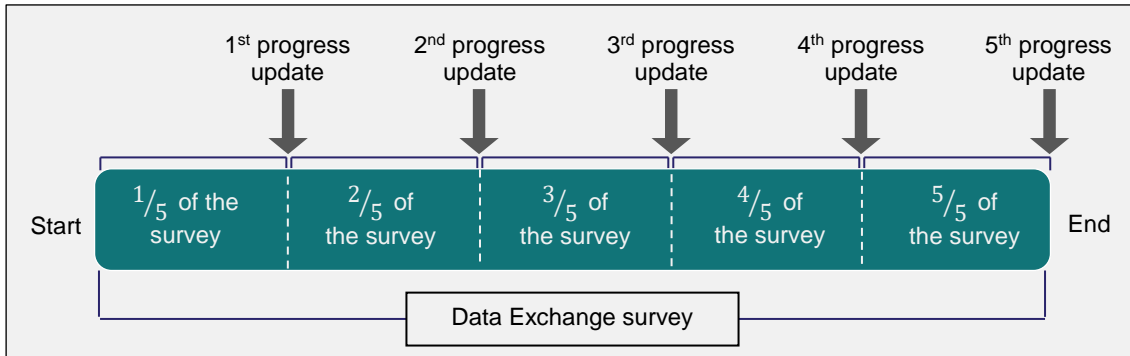
Box 3: What is a randomised controlled trial?

Well-designed RCTs provide the best empirical method for determining a policy's quantifiable impacts. In this respect, RCTs are considered the 'gold standard' for impact evaluation. RCTs work by randomly separating people into two or more groups, in a manner similar to flipping a coin. People in a 'treatment' group receive an intervention (new policy) while people in the 'control' group receive the business as usual experience. On average, the difference in outcomes between people in a treatment group and in the control group reflects the causal impact of the new policy.

Unfortunately, not enough people were offered the opportunity to complete the second survey so we were unable to evaluate the impact of the pre-commitment page. We anticipated this possibility in our pre-registration of this trial and pre-specified, if this were to eventuate, we would combine the pre-commitment page group with the control group. As such, this report focuses on the impact of inserting behaviourally-informed progress updates into the survey, and hereafter the 'control group' refers to the combination of the original control group and the pre-commitment group. The details of the pre-commitment trial are set out in Appendix C.
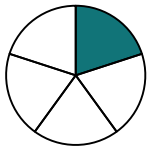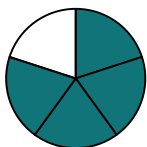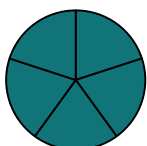
For our behaviourally-informed survey, we inserted five extra pages into the survey. One of the five pages appeared every time a person completed another fifth of the survey (Figure 1). These pages contained a written and graphical update on their progress.

Figure 1: Sequence of progress updates

To enhance the effectiveness of inserted updates, we incorporated social norming messaging, appealed to altruistic motivations, and made the benefit of completing the survey more salient. Based on promising results of appealing to people's innate desire to see things through to completion, we designed the graphical update in the form of a gradually completing pie-chart (see discussion on 'pseudo-sets' in Box 2). Table 1 contains the progress updates used in our study.

Table 1: Content of progress updates

Progress update	Message displayed	Behavioural insight	Graphic displayed
First progress update	*You just finished the first section. Thank you for joining <u>thousands of other contributors</u> by telling us your story*	**Social norm** was used to signal appropriate behaviour	
Second progress update	*Thank you, you're <u>nearly half way through the survey</u>*	Advising of the progress through the survey **reduced uncertainty** around its length	
Third progress update	*You have finished more than half of the survey. Your input helps us <u>improve our services to you and those in similar situations</u>*	Benefits of completing the survey were made **salient**, while highlighting benefits to the community appealed to **altruistic motivations**	
Fourth progress update	*Another section complete, <u>last section coming up</u>*	A further progress update to **reduce uncertainty** around the length of the survey	
Fifth progress update	*Thanks, you've now completed the survey. This information will be grouped together and used to help improve the services provided to you and your community*		

People involved in the trial were clients of various DSS service providers and were enrolled into the trial over a seven-month period from 13 March to 10 October 2017. These clients were assigned to the control group (standard survey, n=366) or treatment group (survey with behaviourally-informed progress updates, n=172).

In this trial, we expected people would be more likely to complete a survey if they received periodic progress updates. Accordingly, our primary outcome variable was the survey completion rate.
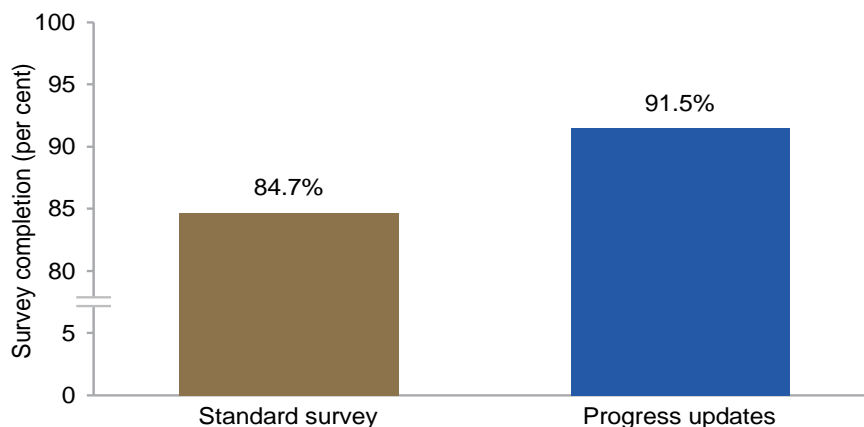
More detail on the trial design can be found in Appendix A.

Results

People who received behaviourally-informed progress updates were more likely to finish the survey.

Receiving progress updates – in the form of a gradually completing pie-chart with accompanying text – increased survey completion by 6.8 percentage points, from 84.7 per cent to 91.5 percent (p=0.02, Figure 2).[2] In other words, these progress updates reduced survey non-completion by about 44 per cent[3].

Figure 2: Progress updates increase survey completion rates

Displaying behaviourally-informed progress updates led to a 6.8 percentage point increase in the survey completion rate (p=0.02, n=512).

We also conducted exploratory analysis looking at whether the impact of these progress updates varied among different groups of people. We found no strong evidence to suggest progress updates work differently for any demographic group (such as a person's location, age, gender, income, or employment status).[4]

---

[2] We are aware there is a lively academic debate about the merits of testing for 'statistical significance', the appropriateness of conventional thresholds such as $p<0.05$ (or any thresholds at all), and even the use of p-values generally. See the 'The American Statistical Association Statement on Statistical Significance and P-Values' (Wasserstein and Lazar 2016).

[3] Standard survey non-completion was 15.3 per cent, introducing behaviourally-informed progress updates reduced non-completion to 8.5 per cent, a reduction of approximately 44 per cent.

[4] See 'subgroup analysis' in Appendix B for further discussion.

Limitations

RCTs work by randomly allocating people to different groups. The way people were allocated into groups in this study was not strictly random. Instead, the allocation was ordered: the first person in an organisation to take a survey was allocated to the control group, the second person was allocated to treatment one, the third person was allocated to treatment two (which we subsequently combined with the control group), the fourth was allocated to control again, and so on. Some organisations had multiple outlets, and clients were assigned in order across all outlets of a given organisation. The ordering procedure operated independently from one organisation to another and was not reset each day.

Given ordered assignment operated across many outlets within many organisations, there is no reason to suspect individuals arriving at certain times of day, at certain organisations or outlets, or with certain characteristics would always be assigned to the same group. There is also no way the staff in participating organisations could know which group assignment was due next, so there was no opportunity to ensure individuals deemed likely to complete the survey were assigned to treatment. Consequently, we consider the allocation was as good as random.

We removed a number of observations prior to our final analysis because three organisations deviated from the agreed procedure for survey distribution by assisting their clients to undertake the survey. They did this for understandable reasons, such as their clients having poor English or IT literacy skills. This additional assistance had the potential to distort the results of the trial because we cannot be sure participants saw the treatment. Consequently, although we did not anticipate this issue in our pre-analysis plan, we excluded 26 observations from those organisations from our analysis. We report the results for the full sample in Appendix B.

It is possible the trial population is not representative of the broader population of service recipients because organisations could exercise discretion over which clients received the survey. Even so, we think it plausible for these progress updates to have a similar impact on a wide range of people and so should generalise well to different populations.

Discussion and conclusions

Including behaviourally-informed progress updates increased the survey completion rate. This finding may be useful for government and private organisations seeking to survey their clients.

It is vitally important for organisations, government or otherwise, to understand the people who use their services and to gauge the impact of these services. One of the more common ways of collecting this information is by asking clients to complete surveys. Unfortunately, with survey participation and completion rates falling in recent years, it is becoming increasingly challenging to encourage people to participate.

Based on the latest behavioural science literature, we modified the Data Exchange survey to increase the number of people completing it. We included periodic written progress updates and presented people with an opportunity to complete an arbitrary 'set' in the form of a pie-chart.

In this trial, providing written and graphical progress updates increased survey completion by 6.8 percentage points, from 84.7 per cent to 91.5 per cent. Our sample included a diverse range of individuals, male and female, young and old, employed and unemployed, across all states and territories of Australia. We found no strong evidence the effect of progress updates differed between these groups. Given similar strong effects have also been demonstrated in another trial in a very different context (Barasz et al., 2017), our results suggest these simple tweaks could be effective across different populations and settings.

This trial showed providing behaviourally-informed progress updates can be an effective, cheap, and low risk improvement to standard survey design. Encouraging more people to share their story through surveys enables government to improve its programs and services, leading to better outcomes for Australians and their communities.

Appendices

Appendix A: Technical details

Trial overview

We ran a field trial over a seven-month period from 13 March to 10 October 2017. Our unit of study was the individual client taking an online survey. The survey was administered by 39 organisations through 70 outlets across all states and territories of Australia. These organisations provide a number of services across eight broad categories (see Table A.1).

Clients receiving services in these organisations were asked to complete two surveys by staff: the first survey was offered at the time these services commenced, and the second follow-up survey was offered once these services concluded. Both treatments (behaviourally-informed progress updates and the pre-commitment page) appeared in the first survey only and were presented to two separate treatment groups. Clients could complete surveys on a computer terminal, tablet device or smart phone. The survey platform automatically allocated individuals using a pseudo-random algorithm upon survey initiation. Researchers at BETA were blind to treatment assignment. Trial participants could see the treatment when administered but were unlikely to know they were involved in a trial.

The Data Exchange pilot survey received ethical approval from the Australian Institute of Family Studies' Human Research Ethics Committee (application number 16/08). BETA's trial was included in this application.

Table A.1: Participating organisations by broad program category and type

Organisation broad program category	Program activity
Housing assistance or homelessness support	Assistance with care and housing
	Reconnect (housing assistance for young people)
Carer supports	Care relationships and carer support
	National carer counselling programme
	Mental health respite: carer support
Children and families	Children and parent support services
	Children's contact services
	Family and relationship services
	Family and relationship services – specialised family violence
	Communities for children – facilitating partners
Financial wellbeing and capability	Commonwealth financial counselling and financial capability
	Financial counselling, capability and resilience IM hubs
	Financial crisis and material aid – emergency relief
	Problem gambling
Family law services	Family dispute resolution
	Family law counselling
	Family relationship centre centres
	Parenting orders programme
	Regional family dispute resolution
Mental health	Personal helpers and mentors
Home support and care	Community and home support
Settlement services	Settlement grants

Pre-registration and analysis plan

Prior to the trial launch, we pre-registered this trial on both the American Economic Association RCT Registry (reference AEARCTR-0002110) and the BETA website. This pre-registration includes details for our proposed analysis including our hypotheses, outcome variables, power calculations and experimental design.

We subsequently prepared a more detailed analysis plan in January 2018. We finalised this plan after the conclusion of the trial and after we had undertaken some analysis with preliminary data. This analysis plan largely elaborates on details already set out in the pre-registration. Where relevant, we refer to 'the pre-registration' to refer to anything that was pre-specified and 'the analysis plan' for any details that we added later.

As discussed in the Limitations section, we made one major unanticipated change to our analysis: we excluded 26 observations because they received assistance to undertake the survey from their service provider. We removed these observations because the assistance to these participants had the potential to distort the trial results. To be transparent, we also report estimates when these observations are included: see Appendix B.

Decision to pool the control group and the second treatment group

We specified in our pre-registration that we might combine participants assigned to the pre-commitment survey group into the control group. For reasons discussed in Appendix C, our power to detect an effect between the control and pre-commitment group was likely very low so we decided to pool them. This is a valid comparison because the pre-commitment intervention was delivered at the end of the survey after the pseudo-set treatment had been delivered and the relevant outcomes measured.
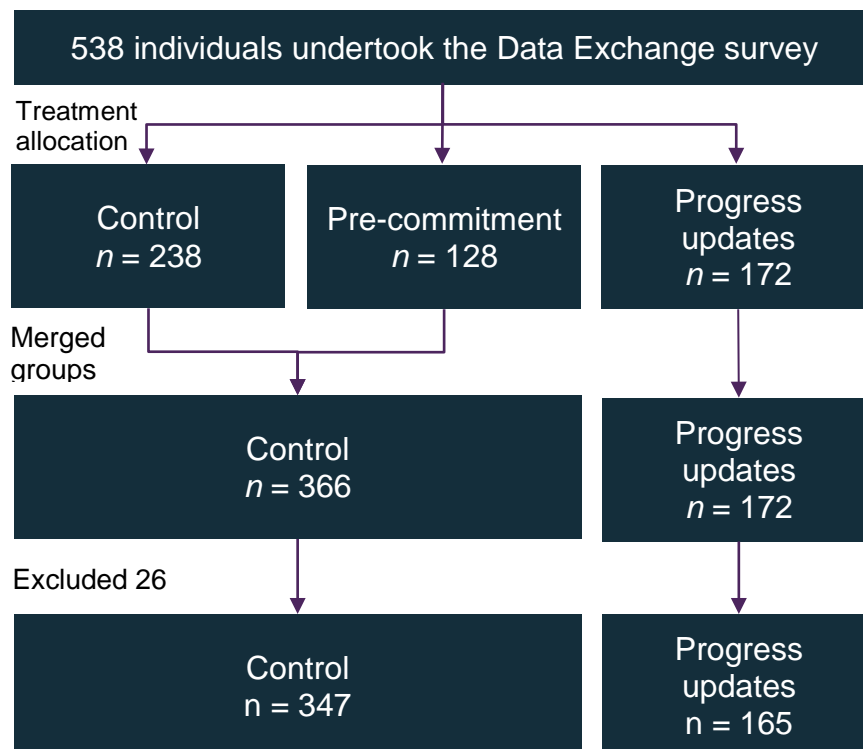
Power calculations and sample size

Prior to the trial commencing, we calculated that a total sample of 423 (three groups, 141 units per group) would provide 80% power at a 5% significance level to detect a large change in survey completion from 85% to 95%. See the pre-registration for further details.

Allocation mechanism

We intended participants would be randomly allocated to treatment or control however the allocation mechanism used an ordered algorithm. As discussed in the Limitations section, this algorithm allocated individuals to each arm in order of their arrival. The algorithm operated across all outlets of a given organisation but independently for each organisation and was not reset each day.

Figure A.1 – Trial flowchart

Balance and missingness checks

Additional variables were collected as part of the survey after treatment assignment. Some variables were missing records and we found evidence these missing records were related to treatment assignment. Specifically, 10% of individuals were missing data in the control group compared to 3% in the treatment group ($p = 0.023$). Consequently, we only tested for balance using organisation type (for which we had no missing records) and main language spoken at home (missing 11 records).

To test for balance, we regressed a binary treatment indicator on these two variables – organisation type and main language spoken at home – and calculated a robust Wald statistic with the null hypothesis that all slope coefficients were zero. We found no evidence of imbalance larger than expected from chance (F = 0.14 on 9 and 502, $p = 0.998$).

Covariate selection

Prior to analysis, we selected covariates likely to predict outcomes and used pre-trial data to confirm that selection. In our analysis plan, we specified main language spoken at home, highest level of education obtained, and health status as candidates to adjust our primary analysis.

However, data for education and health status were collected late in the survey and we found evidence that missing data on these covariates was related to treatment status (see 'Balance and missingness checks' section above). Therefore, we did not adjust for them. Only 11 records were missing for the 'main language spoken at home' variable so we adjusted for this in our main model. Comparisons between unadjusted and adjusted models show little difference (Table B.5, Appendix B).

Outcome variables

Our pre-registered primary outcomes were:

1) proportion of subjects who completed the initial survey,

2) average percentage of survey questions completed, and

3) proportion of subjects who commenced the follow-up survey.

Due to our decision to pool the control group and pre-commitment group we treated Outcome 3 as a secondary outcome. Because of data quality issues we did not assess Outcome 2 at all.

Method of analysis

The principal analysis of the effect of the intervention was an adjusted comparison of our primary outcome across the treatment and control groups. This estimate, confidence intervals and p-values were derived from a linear regression model with the following specification:

$$y_i = \alpha + \tau T_i + \beta x_i + \gamma x_i T_i + \epsilon_i$$

where T_i is an indicator for treatment and x_i is a vector of two mean-centred variables, one a covariate indicating if English was an individual's main language spoken at home and the other a missingness indicator for this covariate. These variables were interacted with the treatment indicator as per Lin (2013).

We calculated robust (HC2) standard errors for all linear models, including robustness checks, balance and missingness checks and secondary analyses. Because our primary outcome was binary, we also ran robustness checks using equivalent logistic regression specifications. After calculating average marginal effects, we found little difference between model types (Table B.3, Appendix B).

Appendix B: Key statistical tables

This appendix presents the full statistical tables which underlie the results section. It also includes robustness checks and additional analyses not included in the main body of the report.

Main results

Table B.1 – Main result

	N	Completed survey % (no.)	Percentage point difference	95% confidence interval	p-value
Control	347	84.7% (294)			
Progress updates	165	91.5% (151)	6.8	1.1 to 12.5	0.02

Note: This table shows the total number of individuals in the control and treatment groups as well as the number and percentage of individuals completing the survey in each group. The percentage point difference, 95% confidence interval and p-value are from an adjusted linear regression model with robust standard errors. We excluded 26 individuals from this analysis (see 'Robustness checks and missing data' section below).

Table B.2 - Participant characteristics and balance

		Control	Treatment
Survey completion	*n*	347	165
	Completed	294 (84.7%)	151 (91.5%)
Primary language spoken at home	*n*	340	161
	English	282 (82.9%)	134 (83.2%)
Age	*N*	301	152
	mean ± SD	45.2 ± 17.2	46.2 ± 19.1
Gender	*N*	301	152
	Female	198 (65.8%)	102 (67.1%)
Income	*n*	298	154
	Less than $60,000	161 (54.0%)	84 (54.5%)
	$60,000 or more	100 (33.6%)	56 (36.4%)
	Not answered	37 (12.4%)	14 (9.1%)
Employment	*n*	298	154

		Control	**Treatment**
	Employed	162 (54.4%)	87 (56.5%)
	Unemployed	121 (40.6%)	63 (40.9%)
	Not answered	15 (5.0%)	4 (2.6%)
Self-assessed health status	*n*	309	158
	Excellent / Very Good	121 (39.2%)	63 (39.9%)
	Good	90 (29.1%)	54 (34.2%)
	Fair / Poor	66 (21.4%)	25 (15.8%)
	Not answered	32 (10.4%)	16 (10.1%)
Completed Year 12	*n*	298	154
	Year 12	178 (59.7%)	97 (63.0%)
	No Year 12	101 (33.9%)	50 (32.5%)
	Not answered	19 (6.4%)	7 (4.5%)
Organisation type	*n*	347	165
	Family law services	125 (36.0%)	59 (35.8%)
	Children and families	94 (27.1%)	45 (27.3%)
	Home support and care	43 (12.4%)	23 (13.9%)
	Financial wellbeing and capability	33 (9.5%)	15 (9.1%)
	Carer supports	21 (6.1%)	7 (4.2%)
	Settlement services	17 (4.9%)	9 (5.5%)
	Housing or homelessness support	7 (2.0%)	4 (2.4%)
	Mental health	7 (2.0%)	3 (1.8%)

Note: Percentages for some characteristics do not add to 100% due to rounding error. Some individuals are missing data on some characteristics due to unsubmitted responses or filter questions, which attempted to present relevant questions based on responses to previous questions. These missing records were not included in the denominator for percentage calculations.

Robustness checks and missing data

We report two robustness checks. First, we compared OLS regression (our primary method of analysis) with an equivalent logistic regression specification. Second, we ran the regressions with covariate adjustment (our primary method of analysis) and without adjustment. After calculating average marginal effects for the logistic regression, we found little difference between model types (Table B.3).

We excluded 26 observations from our main analysis. For transparency, we report the characteristics of the excluded units and the results of the primary analysis without exclusions (Table B.4 and Table B.5).

Table B.3 – Primary analysis treatment effects – robustness check

Model	Treatment effect (percentage points)	95% confidence interval	df	p-value
Unadjusted OLS	6.8	1.1 to 12.5	510	0.020
Adjusted OLS	6.8	1.1 to 12.5	506	0.020
Unadjusted logit	6.8	1.1 to 12.5	510	0.019
Adjusted logit	6.8	1.1 to 12.5	506	0.019

Note: Estimates and 95% confidence intervals are in percentage points. Logistic regression results are average marginal effects. For OLS, robust standard errors (HC2) were used and for logistic regression standard errors were calculated using the delta method.

Table B.4 - Excluded units

	Control	Treatment	Total
Total *n* (no exclusions)	366	172	538
Excluded no. (%)	19 (5.2%)	7 (4.1%)	26 (4.8%)

Table B.5 – Primary analysis without exclusions

	N	Completed survey % (no.)	Percentage point difference	95% Confidence interval	p-value
Control	366	85.5% (313)			
Progress updates	172	90.1% (155)	4.5	-1.3 to 10.3	0.125

Note: This table shows the total number of individuals (without excluding individuals who started the survey in an organisation that deviated from the trial protocols) in the control and treatment groups as well as the number and percentage of individuals completing the survey in each group. The percentage point difference, 95% confidence interval and p-value are from an adjusted linear regression model with robust standard errors.

Subgroup analyses

We performed exploratory subgroup analyses to determine if progress updates worked better among different groups of people.

We found weak evidence to suggest progress updates work better among those with higher incomes (a six percentage point improvement compared to those with lower incomes). The result must be interpreted with caution because our experiment was not powered for subgroup analyses and we had no pre-specified hypothesis that progress updates would be more effective for those with higher incomes. The p-value for this comparison ($p = 0.03$) should also be interpreted in the context of a large number of comparisons. There was no evidence of other subgroup effects.

Table B.6 – Subgroup analyses

Subgroup	Level	Treat - Control difference (95% CI)	p-value	Difference across levels (95% CI)	p-value
Main language spoken at home	Other	0.4 (-17.9 to 18.8)	0.96		
	English	7.1 (1.2 to 13)	0.02	6.7 (-12.4 to 25.7)	0.49
Age	< 40	1.5 (-2.4 to 5.4)	0.45		
	40+	1.9 (-0.2 to 4.0)	0.08	0.4 (-4.0 to 4.8)	0.87
Sex	Male	1.9 (-0.8 to 4.6)	0.16		
	Female	1.5 (-1.4 to 4.5)	0.30	-0.4 (-4.4 to 3.6)	0.84
Income	< $60,000	-3.0 (-7.1 to 1.2)	0.17		
	$60,000+	3.0 (-0.4 to 6.4)	0.08	6.0 (0.6 to 11.3)	0.03
Education	No Year 12	-4.0 (-11.3 to 3.2)	0.28		
	Has Year 12	1.1 (-0.4 to 2.7)	0.16	5.1 (-2.2 to 12.5)	0.17
Employment	Unemployed	-0.8 (-4.3 to 2.8)	0.67		
	Employed	-0.4 (-4.3 to 3.4)	0.81	-0.3 (-5.5 to 4.9)	0.90
Self-reported health status	Excellent/ very good	0.1 (-5.3 to 5.6)	0.96		
	Good	-3.0 (-11.3 to 5.4)	0.48	-3.1 (-13 to 6.8)	0.54
	Fair/poor	6.1 (0.2 to 11.9)	0.04	5.9 (-2 to 13.9)	0.14

Note: Subgroup analyses were performed using linear regression with robust standard errors. The difference across levels was tested by interacting an indicator for subgroup membership with an indicator for treatment.

Appendix C: Pre-commitment trial

Our trial was designed with a second treatment group, in which individuals who completed the first survey saw a message asking them to pre-commit to completing a second follow-up survey after the conclusion of the relevant services.

The sample size for the pre-commitment group was somewhat smaller than planned and, of this group, very few individuals were offered the opportunity to commence the follow-up survey. As a result, our power to detect an effect between the control and pre-commitment group was likely very low. Consequently, we decided to combine the pre-commitment treatment group into the control group and treat the results from the pre-commitment message as a secondary (or exploratory) analysis.

Across the control and pre-commitment groups, only 14 people commenced the follow-up survey. This was 2.5% of the pre-commitment group compared to 4.9% of the control group (Table C.1).

Although it would appear that seeing a pre-commitment message reduced follow-up survey initiation, the evidence for this is weak. The p-value for the pre-commitment/control comparison was 0.28 and our 95% confidence interval contained values consistent with both reductions and improvements in survey initiation.

Therefore, we cannot conclude whether a pre-commitment device causes more or fewer people to start a follow-up survey.

Table C.1 – Follow-up survey commencement – pre-commitment trial

	n	Commenced follow-up survey % (no.)	Percentage point difference	95% Confidence interval	p-value
Control	226	4.9% (11)			
Pre-commitment	121	2.5% (3)	-2.2	-6.2 to 1.8	0.28

Note: The percentage point difference, 95% confidence interval and p-value are from an adjusted linear regression model with robust standard errors.

References

Australian Bureau of Statistics, 2010. *Empirical Models for Survey Cost, Response Rate and Bias Using Paradata,* s.l.: s.n.

Barasz, K., John, L. K., Keenan, E. & Norton, M., 2017. Pseudo-set framing. *Journal of Experimental Psychology,* 146(10), pp. 1460-1477.

DSS, 2017. *Department of Social Services Annual Report 2016-17,* Canberra, Australia: DSS.

DSS, 2018. *Grants.* [Online]
Available at: https://www.dss.gov.au/grants/grants-funding

Edwards, P. J. et al., 2009. *Methods to increase response to postal and electronic questionnaires,* s.l.: Cochrane.

Fogliani, M., 1999. Low Response Rates and Their Effects on Survey Results. *Methodology Advisory Committee paper.*

Lin, W., 2013. Agnostic notes on regression adjustments to experimental data: Reexamining Freedman's critique. *The Annals of Applied Statistics,* 7(1), pp. 295-318.

Nulty, D. D., 2008. *The adequacy of response rates to online and paper surveys: what can be done?.* s.l.:s.n.

Australian Government

BETA

Behavioural Economics Team of the Australian Government