



UNIVERSITY OF AMSTERDAM

UvA-DARE (Digital Academic Repository)

Strategic communication: theory and experiment

de Haan, T.

Publication date
2012

[Link to publication](#)

Citation for published version (APA):

de Haan, T. (2012). *Strategic communication: theory and experiment*. [Thesis, fully internal, Universiteit van Amsterdam].

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Chapter 5

Choice architecture and default bias

This chapter is based on Thomas de Haan and Jona Linde, 2011: "Good nudge lullaby", Working paper.

5.1 Introduction

Ample research shows that human behavior is better characterized by bounded rationality than by the full rationality assumed in traditional economic models. Two important findings are that choices can be influenced by normatively irrelevant features of the choice environment (see e.g. Huber, Payne & Puto, 1982 and Kahneman & Tversky, 2000). and that people can make decisions that go against their own interest. Preference reversal (Lichtenstein & Slovic, 1971) and inconsistent time preferences (Frederick, Loewenstein & O'Donoghue, 2002) are well known examples of this second type of behavior.

Combining these two findings Thaler and Sunstein (2003) and Camerer et al. (2003) propose to use our knowledge of the influence of the choice environment on decisions to promote 'better' decisions without changing incentives. They propose adapting the 'choice architecture' by including 'nudges' which steer boundedly rational people towards better decisions without affecting people who are fully rational. Numerous studies validate the effectiveness of this approach (Thaler and Sunstein, 2003) which its inventors call libertarian or asymmetric paternalism⁸⁷.

Although libertarian paternalism has also received its share of criticism, this has mostly focused on the moral justification to influence others' decisions (e.g. Mitchell (2004) and Sugden (2008)). We do not aim to take a side in this philosophical debate but consider a somewhat more practical, possible concern with libertarian paternalism. We ask whether the most popular nudge, a good default option⁸⁸, affects performance in

⁸⁷In the rest of this paper we use the term libertarian paternalism to refer to this approach.

⁸⁸Thaler and Sunstein (2003): "The most common nudge is a default option that the choice architect

subsequent, similar decision situations.

A good default is such an effective nudge because of the well-known status quo or default bias, the behavioral regularity that people are more likely to choose the option presented as the default or the option they currently possess (e.g. Samuelson and Zeckhauser, 1988). So far the implicit assumption underlying libertarian paternalism is that the default bias is a constant predisposition that can be used to influence decisions. However, if the default bias has an ‘endogenous’ component the choice architecture may not only change current decisions but also influence the choice process. If someone receives a nudge in the form of a good default, the choice heuristic to stick with the default performs well. This may reinforce the use of this heuristic and therefore make a person more likely to choose the default in similar decisions in the future, even if she is no longer being nudged but faces a random, or possibly even a bad, default. This would hurt her performance.

Another reason why the default bias can have an endogenous component is that people may view the default option as advice. Madrin and Shea (2001), for example, find that one cause for higher enrollment rates to a pension plan when enrollment is the default is caused by people taking the default as implicit advice from their employer. If the default indeed turns out to be a good suggestion people’s trust in this advice is likely to increase, thereby strengthening the default bias.

A good default may also affect a person’s decision making process more generally, either positively or negatively. Positively if people learn what a good option looks like because the good default draws their attention to good options. Negatively for two possible reasons: people who face a good default may learn to trust the default but not how to make a good decision themselves and they may become spoiled by the good default and therefore unwilling to put in effort now that the default is no longer helpful. In a very recent experiment, Caplin and Martin (2011) find behavior that can be interpreted as becoming spoiled. Their participants put less effort in a choice task when provided with a relatively helpful nudge. They do not, however, look at what happens to performance if the nudge would disappear again.

To test whether a good default affects choice and performance in subsequent decisions we developed an experimental task with an unequivocal best choice which is nevertheless hard to find. Participants face this task for 50 rounds. In the first 25 rounds participants in the “nudge” treatment receive a nudge in the form of a good default. Participants in the control group, on the other hand, receive a random default. In the second 25 periods both groups receive a random default. Any difference in performance between participants in

believes is a good choice for the decision maker.” For the use of a default as a nudge see for example Madrin and Shea (2001), Johnson and Goldstein (2003) and Benartzi and Thaler (2004).

the nudge and control treatments in these second 25 periods reveals the effect of a nudge on subsequent behavior. In section 5.2 we describe the experimental design in detail.

The effect of a nudge on subsequent decisions is not only of academic interest. In real life decisions, a good default may be followed by a worse default for various possible reasons. A first possible reason is commercial interests. Many purchases require several separate decisions (e.g. buying a car, a computer, or a plane ticket). Companies may try to lure consumers into a false sense of security by providing good defaults for the first decisions, but malicious defaults later on (e.g. first recommending economy class and direct routes but later also expensive flight insurance).

A second reason why good defaults may be followed by less helpful ones is that for some decisions, ‘good’ defaults are easier to provide than for others. This happens in one of the most prominent examples of libertarian paternalism, a default enrollment in pension plans (Madrian & Shea, 2001). Saving something, and therefore participating, is probably optimal for the large majority of employees, but there is far more heterogeneity in how much too safe. Setting a ‘good’ default savings rate is therefore far more difficult (Choi et al., 2003). As a result the default savings rate is probably a less helpful default than the default enrollment choice. Good defaults for the decisions how to invest the money saved are probably even harder to provide.

A third reason for following good with random defaults could be legal limitations. Courts may view libertarian paternalism as unwarranted government intervention. In their book “Nudge” Thaler and Sunstein (2008) discuss a program implemented in Maine (USA) to provide Medicare users with a good default health care program. Legal challenges have contributed to the failure of this project to spread to other states. Similar legal challenges may cause libertarian paternalism programs that are already in use to be discontinued. If that happens a person used to helpful defaults may face suboptimal defaults in the future.

For these reasons we believe it is important, both in order to expand our insight in decision heuristics and to design optimal libertarian paternalism policies, to explore the effect of a good default on subsequent decisions. The rest of this paper is structured as follows: section 5.2 elaborates on our experimental design, section 5.3 presents the results of the experiment and section 5.4 concludes.

5.2 Design

The experiment was computerized with php/mysql and conducted at the CREED laboratory of the University of Amsterdam. At the beginning of the experiment participants read the instructions on the computer at their own pace. They then received a summary

Figure 5.1: An example of the choice task as presented to the participants

Choice Task 48						
TIME: 39 BONUS: 20						
Choices	Weight=6	Weight=5	Weight=4	Weight=3	Weight=2	Price
<input type="radio"/> Option 1	12	8	10	4	3	102
<input type="radio"/> Option 2	1	13	11	28	19	138
<input type="radio"/> Option 3	4	9	29	39	13	122
<input checked="" type="radio"/> Option 4	5	20	49	7	13	271
<input type="radio"/> Option 5	28	5	13	21	12	109
<input type="radio"/> Option 6	42	22	6	4	3	348

Make choice

Remarks: Option 4 is the default option in this example.

of the instructions on paper. After reading the instructions, participants had to correctly answer some questions to test their understanding of the instructions.

All participants in the experiment performed the same set of 50 multi-attribute choice tasks. The difference in treatments consisted only of a difference in the nature of the default in the first half of the experiment. Performance in the second half, when all participants face the same task and the same default reveals the effect of being nudged on subsequent decisions. Below we first discuss the choice task and then the difference in the default between the two treatments.

5.2.1 Task

Each round participants chose one option from a list of six. The information on which to base this choice was presented in the form of a table. Each option consisted of a number of points in 6 categories, each with a different weight. The weights were 6, 5, 4, 3, 2 and -1. The category with a weight of -1 was presented as the price of an option. These categories and their weights, but not the points, were the same for each choice task. An option generated an amount of credits equal to the sum of the points in each category multiplied by the weight of that category. The tasks were randomly generated under the conditions that each option generated between 70 and 230 points and that the best option generated at least 10 points more than the second best option. An example of a task is shown in figure 5.1 below.⁸⁹

On top of the credits generated by the chosen option participants received a bonus, starting at 20 credits and decreasing by 1 credit every two seconds the participant used to make a decision. The maximal time a participant could take to choose was 40 seconds. 20 credits is a small amount compared to the gains that could be made by making a better

⁸⁹This task can be seen as a choice between different products, each with a different price and different qualities. The category weights represent the relative importance of different types of characteristics and the points the quality of the product in that characteristic. In that sense the problem is similar to many everyday decision problems, from buying a phone to choosing a medical insurance or an investment plan. This type of task has previously been used by Kalayci and Potters (2011).

decision.⁹⁰ The diminishing bonus can be seen as a small cost put on spending an extra two seconds on the task. After a participant made her decision she had to wait till the time for this round expired before moving on to the next round. In addition there was a 5 second waiting time between rounds.

Without the bonus participants who have already decided could have waited until the full 40 seconds were over without any costs. With the bonus, as soon as a participant has decided, she would want to enter her choice to save on the bonus. We implement this bonus for two reasons. Firstly, in order to have a measure of search effort in the form of time spend on the task. Secondly, to ensure that we know when participants actively choose an option and when they were forced into a decision because time ran out.

Participants performed this task for 50 rounds. All participants faced the same 50 tasks but there were 6 different orders in which the tasks were presented. The order was counterbalanced between treatments.⁹¹ At the end of the experiment one round was randomly selected.⁹² The number of credits earned in that round determined the participant's earnings. Each credit was worth 10 eurocents.

5.2.2 Default and treatment

One of the six options in the table was given a different color and was preselected when participants were presented with the task. This option was the default option. In the instructions it was labeled the recommended option without further specifying why it was recommended. If participants did not select one of the other options, the default automatically became the choice of the participant for that round if time ran out before the participant entered a choice.⁹³ If a participant chose an option different from the default option, a smaller version of the table was shown, containing only the chosen and the default option. They were then asked if they wanted to stay with their original choice or switch to the default.⁹⁴

Participants were randomly assigned to one of two treatments. Both treatments were identical, except for one aspect. In the 'control' treatment, the default option was determined randomly for each task. In the 'nudge' treatment, the default option for the tasks a

⁹⁰For example the difference between the best and the second best option was always at least 10 credits and the difference between the best and the worst option was on average 137 credits.

⁹¹The counterbalancing procedure also ensured that each group as a whole faced the same tasks in the first and the second half of the experiment. Due to a small software error two participants had to be excluded, one in the control treatment and one in the nudge treatment. This affected the counterbalancing slightly, as these two participants had different orders. Leaving out two random participants with these orders in the other treatments does not materially affect our results.

⁹²The same round was selected for all participants in a session but because of the different task orders that was a different task for different participants.

⁹³This happened only 161 times out of 4400 i.e. in 3.66% of all decisions.

⁹⁴People switched a total of 99 times out of 2577, i.e. in 3.8% of all initial non-default choices.

participant faced in the first 25 rounds of the experiment was the option with the highest value. The tasks faced by participants in the ‘nudge’ treatment in the rounds 26 till 50 had the same random defaults as in the control treatments.

5.3 Results

A total of 88 participants participated in the experiment, half of them assigned to the control treatment and half of them to the nudge treatment. Of the participants, 49 were males and 56 were economics students. The experiment lasted on average one hour and the participants earned on average 19.31 euro. To examine the effect of a nudge on subsequent decisions we examine the performance and behavior of participants in the second half of the experiment where participants in both treatments face the same choice tasks and the same (random) defaults. To evaluate choice performance, we take the value of the chosen option as a performance measure. We first test whether receiving a nudge in the form of a good default influences the quality of decisions and then examine possible mechanisms. All reported tests are two-sided and, unless otherwise specified, performed at the individual level.

5.3.1 Treatment effect

The main question our experiment tries to answer is the effect of having received a nudge on performance when that nudge has disappeared. Table 5.1 answers this question. We find that in the second half of the experiment, when all participants faced the same random defaults, participants in the control treatment chose an option worth 5.72 points more on average than participants in the nudge treatment.⁹⁵ A Wilcoxon rank-sum test shows that this difference is marginally significant ($p=0.081$).

The regression in table 5.2 below confirms the main treatment effect. This regression controls for several demographic variables, high school math level and grade as a proxy of skill and time used in the first round as a proxy of effort. As time spent can be influenced by the treatment we take the time spent during the first round as an exogenous measure of effort.⁹⁶ Controlling for these variables in the regression, the treatment effect becomes significant at a 5% level.⁹⁷

⁹⁵For comparison the average value of an option for a person who always chooses the best option is 71.21 points higher than those of a person who chooses randomly.

⁹⁶Time spent in the first round strongly correlates with time spent in later rounds (Spearman correlation coefficient is 0.5265 and $p<0.001$) Using time spent in the entire experiment or only the second half yields the same qualitative results. Time used in the first round might be influenced by the treatment but a Wilcoxon rank-sum test shows that this is not the case (p -value 0.5804).

⁹⁷As can be seen in table 5.2, we find an unhypothesized, but significant gender effect in the regression. Looking at the average earnings in the second half, separately for men and women, a Wilcoxon rank-sum

Table 5.1: Average value of the chosen options in the experiment

	Control treatment	Nudge treatment	p-value of a Wilcoxon rank-sum test
first half	173.0 (14.1)	191.5 (11.6)	0.000
second half	174.8 (15.1)	169.1 (17.5)	0.081
p-value of a Wilcoxon signed rank test	0.398	0.000	

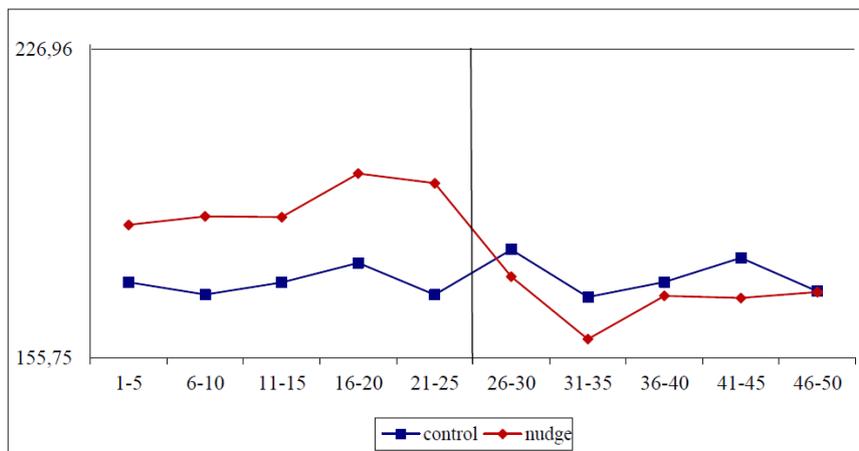
Remarks: Standard deviations at the individual level between brackets.

Table 5.2: Treatment effect when controlling for demographic variables and proxies for effort and skill.

Dependent variable: average value of the chosen option in the second half of the experiment	Coefficient	p-value*
constant	173.08 (13.87)	0.000
nudge treatment	-6.58 (2.89)	0.025
time used in first round	0.52 (0.14)	0.000
Male	7.97 (3.20)	0.015
Age	0.02 (0.02)	0.315
studies economics	0.35 (4.02)	0.930
Dutch	-5.74 (3.50)	0.105
math grade	0.23 (0.87)	0.795
math level	1.15 (4.25)	0.788

Remarks: Round and task fixed effects are included in addition to the controls listed in the table. Standard errors are between brackets.*Standard errors used to calculate the p-value were clustered at the individual level.

Figure 5.2: Average value of the chosen option aggregated over 5 rounds split between the nudge and control treatments



Remarks: The bottom of the graph corresponds to the expected value over the all 50 tasks for a person who chooses randomly (155.75), the line at the top to the average value of the chosen option for a person who always chooses the best option (226.96).

As figure 5.2 shows the treatment effect persists throughout the second half of the experiment. In almost every round average choice earnings were higher for the control than for the treatment group and this difference does not show a tendency to decline.

While we are mainly interested in what happens when the default is no longer optimal, we would expect the nudge to be helpful in the first half of the experiment. Table 5.1 confirms that providing people with a good default indeed helps them to make better decisions. As the first row of table 5.1 shows nudged participants chose an option worth 18.52 points per round more in the first half of the experiment. This is a substantial and highly significant difference, according to a Wilcoxon rank-sum test.

5.3.2 Default bias

Given that a treatment effect exists we will take a closer look at the behavior of participants to explore possible causes. Our first hypothesis is that participants in the nudge treatment exhibit a stronger default bias than participants in the control treatment. As table 5.3 shows this was indeed the case. In the second half of the experiment nudged participants were 11.6 percentage points more likely to pick the default than participants from the control group even though they faced the exact same default.

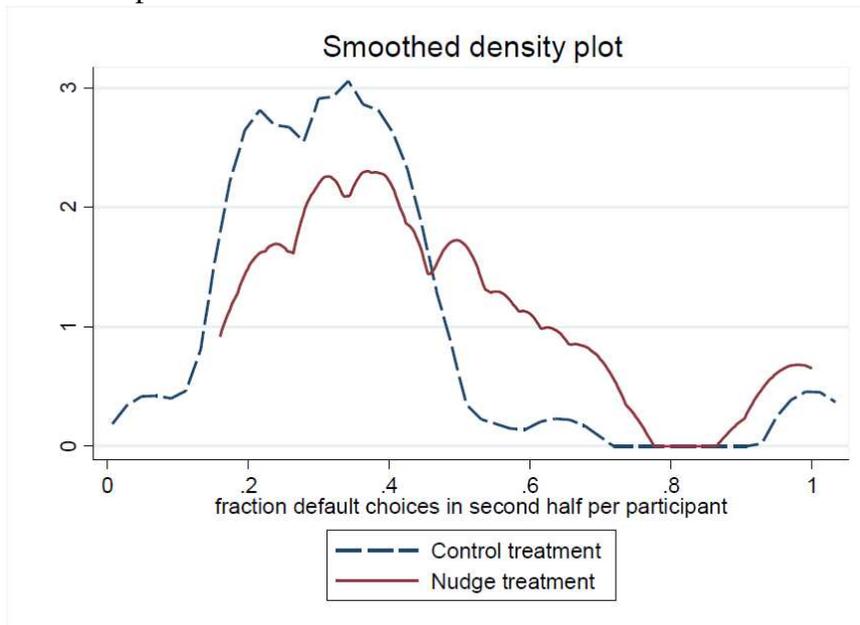
Furthermore, as figure 5.4 shows, the stronger default bias for nudged participants persists throughout the second half of the experiment. Notice also the sharp decline of default choices for nudged participants after round 25. Nudged participants do not simply

test shows there is a significant treatment effect for men ($p=0.028$) but not for women ($p=0.942$). However, including an interaction term between gender and treatment in the regression of table 5.2 shows that the treatment effect is not significantly different for men and women ($p=0.444$).

Table 5.3: Percentage of default choices per round in the experiment.

	Control treatment	Nudge treatment	p-value of a Wilcoxon rank-sum test
first half	30.0%	65.3%	0.0000
second half	33.9%	45.5%	0.0039
p-value of a Wilcoxon signed rank test	0.0511	0.0000	

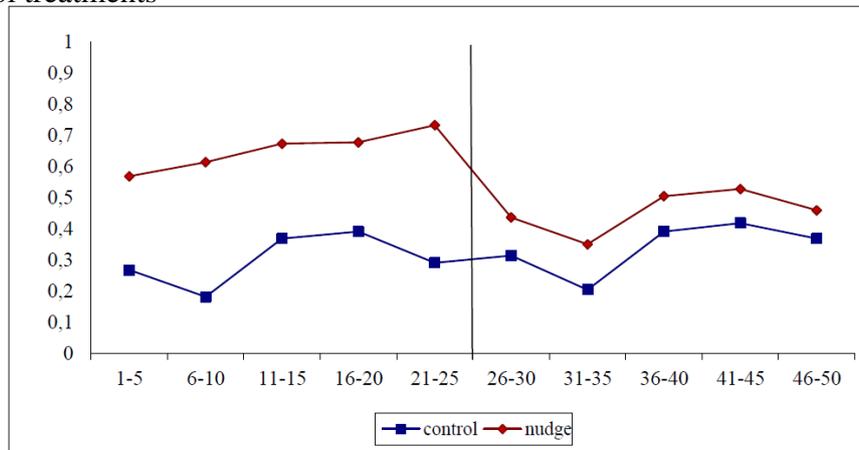
Figure 5.3: Smoothed density of participant's tendency to rely on the default option in the second half of the experiment for both treatments



continue to choose the default at the same rate as they did in the first half. This suggests that although nudged participants choose the default more often, they do realize that the default option is not as good as it was in the first 25 rounds. Figure 5.3 supports this story at the individual level. The greater average reliance on the default is not caused by a small number of nudged participants who have learned to blindly follow the default. Rather the whole distribution of default choices is shifted to the right suggesting that having faced a nudge makes many participants somewhat more likely to choose the default rather than some participants rely heavily on the default while leaving others unaffected.

The difference in the likelihood of default choices between treatments in the first half provides further evidence that reinforcement of the default bias is a cause of the treatment effect. Even in the control treatment participants chose the default more often than the 16.7% expected without a default bias (Wilcoxon rank-sum p-value < 0.001), but nudged participants were significantly more likely to do so (Wilcoxon rank-sum p=0.000). This in itself does not provide evidence for a greater trust in the default in the first half of the nudge treatment. Participants may also have chosen the default more often in the

Figure 5.4: Average proportion of default choices over 5 rounds split between the nudge and control treatments



nudge treatment because it was the best option which they would have chosen anyway, regardless of it being the default or not. However we find that in the first half participants in the nudge treatment were also more likely to choose the default than participants in the control treatment were to choose either the default or the best option (65% vs. 57%, Wilcoxon rank-sum $p=0.0025$). We therefore conclude that nudged participants came to trust the default throughout the first half of the experiment.

In fact as figure 5.4 suggests the tendency to choose the default increases for nudged participants during the first 25 periods⁹⁸. A logit regression with individual and task fixed effects confirms this positive trend ($p<0.001$). There is no significant trend for control group participants in the first half or for either type of participant in the second half of the experiment. Interestingly, while participants in the nudge treatment seem to learn that the default is good, participants in the control treatment do not learn that the default is usually bad. In fact table 5.3 shows that, if anything, they are somewhat more likely to choose the default in the second half of the experiment than in the first half.

5.3.3 Quality of decisions

The stronger default bias evidently has a negative impact on the nudged participants' performance in the second half of the experiment. However the good default may also impact the behavior and performance of nudged participants in other ways. As proposed in the introduction this can be either positively, through learning what a good option looks like, or negatively, because of a lack of motivation, or because nudged participants have not learned how to choose for themselves. Moreover one could ask whether the difference in performance between the two treatments is due to nudged participants having

⁹⁸From an average of 57% default choices in the first five rounds to 73% default choices in rounds 21-25.

Table 5.4: Linear^a Regression examining the likelihood of choosing the default option in the second half of the experiment.

Dependent variable: choice equals the default in the second half	Coefficient	p-value
constant	0.618	0.000
treatment (1=nudge treatment)	0.113	0.008
$\pi(\text{best}) - \pi(\text{default})$	-0.037	0.000
treatment x $\pi(\text{best}) - \pi(\text{default})$ ^b	0.00	0.678

Remarks: Each choice is used as an independent observation, standard errors have been adjusted by treating each participant's choices as a cluster.

^aA probit regression gives the same results, but we chose to report a linear regression as Ai and Norton (2003) have shown that standard statistical packages often miscalculate interaction effects in non-linear regressions.

^bThe interaction term was normalized to prevent multi-collinearity issues with the dummy variable for treatment.

an increased 'default bias', or due to poorer decision making in general, so e.g. also taking account the non-default choices.

To answer this question, we will first look at whether subjects are sensitive to the quality of the default. Table 5.4 reports a regression that indeed shows that subjects are sensitive to the difference in value between the best option and the default option ($\pi_{Best} - \pi_{Default}$). The default option is less often chosen if this difference becomes large. Again we see a significant treatment effect. However, the interaction effect of treatment and $\pi_{Best} - \pi_{Default}$ is zero, indicating that nudged participants make more default choices overall, but are not less sensitive to the value difference between the best and the default option.

Another aspect of the performance is the quality of the non-default choices. If we look at the average non-default choice earnings per participant in the second half, we see a slight difference between the treatment averages (179.63 for the control treatment versus 175.64 for the nudge treatment). However performing a Wilcoxon rank sum test gives no significant difference ($p = 0.1340$).

To answer the question of the effect of the treatment on the quality of participant's decisions in a more general way, we will estimate a modification of the noisy response model as developed by McFadden (1973) (and for example used by McKelvey and Palfrey, 1995), where individuals have both a "choice precision parameter" λ and a "default bias parameter" β . The probability that an individual in one task chooses option i is denoted as follows in this model.

$$P(\text{option } i \text{ is chosen}) = \frac{e^{\lambda \cdot (\pi_i + \delta_i \cdot \beta)}}{\sum_{j=1}^6 e^{\lambda \cdot (\pi_j + \delta_j \cdot \beta)}}$$

Where π_i is the amount of earnings option i would provide and δ_i is a dummy which

Table 5.5: Estimated values of the parameters from the noisy response model

	Control treatment	Nudge treatment	p-value of a Wilcoxon rank-sum test
Average estimation of λ	1.85 (0.78)	1.70 (0.51)	0.2460
Average estimation of β	64.81 (44.10)	91.48 (58.69)	0.0217

Remark: Standard deviations are between brackets.

takes value 1 if option i is the default option⁹⁹. We assume $\lambda > 0$. This model allows for a default bias but also takes the noisiness of the decisions into account, in a sense it separates the default bias from other possible choice irrationality. A positive value of β would mean that an individual is more likely to pick the default option compared to a case where $\beta = 0$. The higher the value of λ , the higher the probability an individual chooses a high value¹⁰⁰ option and therefore λ can be seen as a 'precision parameter'.

We estimated both parameters per individual using a maximum likelihood estimation procedure based on the data of the last 25 periods. The average values of the λ and β parameters are shown in table 5.5 below.¹⁰¹

We see from the table that the estimates of λ are statistically similar in both treatments. This again confirms the picture that emerged from the regression in table 5.4 and the non-default choice earnings, that apart from the observed increase in the default bias, the decision quality did not differ between the two treatments. If for example participants from the nudge treatment would have consistently chosen lower value options given a non-default choice, a significantly lower average level of λ would have been estimated for this treatment. Given that the average levels of λ are similar between treatments, we can also directly compare the average estimated values of β . We see that in both the control treatment and the nudge treatment, a substantial default-bias exists. As expected we do find a significant treatment effect here as the default bias parameter is on average larger for the nudged participants.¹⁰²

The preceding paragraphs show that, except for the larger default bias, the decisions of nudged participants appear to be just as good as those of control group participants.

⁹⁹The model can be derived from assuming the following utility function: $U(\text{choosing option } i) = \pi_i + \delta_i \cdot \beta + \varepsilon_i$, where ε_i is drawn from the extreme value distribution with parameter λ (see McFadden, 1973, for the derivation).

¹⁰⁰In the case of a large default bias β , this is of course not per se true, as here it might be that due to λ and β being large, an individual will choose a sub optimal default with a high probability. However, a high value of λ will still increase the probability of the best option being chosen if the payoff of that option is larger than the sum of the default option payoff and the default bias.

¹⁰¹For the estimation we used the formula $\frac{e^{\lambda \cdot \frac{(\pi_i + \delta_i \cdot \beta)}{100}}}{\sum_{j=1}^6 e^{\lambda \cdot \frac{(\pi_j + \delta_j \cdot \beta)}{100}}}$ to prevent numerical problems.

¹⁰²Due to a few very large estimates of β , caused by a few participants following the default almost all the time, the standard deviation of β is rather large.

Table 5.6: Regression: Interaction between effort and the treatment effect

Dependent variable: average value of the chosen option in the second half of the experiment	Coefficient	p-value*
constant	173.49	0.000
nudge treatment	-6.59 (2.92)	0.027
time used in first round	0.50 (0.21)	0.018
treatment x time used in the first rounds ^a	0.03 (0.29)	0.911
male	8.01 (3.19)	0.014
age	0.02 (0.02)	0.341
studies economics	0.27 (3.98)	0.947
dutch	-5.78 (3.59)	0.111
math grade	0.25 (0.85)	0.774
math level	1.09 (4.22)	0.797

Remarks: Standard errors used to calculate the p-values are clustered at the individual level. Round and task fixed effects are included in addition to the controls listed in the table. Standard errors are between brackets. *Standard errors used to calculate the p-value were clustered at the individual level.^aThe interaction term was normalized to prevent multi-collinearity issues with the dummy variable for treatment.

One other possible behavioral difference that could be caused by the treatments is that nudged participants put less effort/time into the decisions in the second half compared to the control group. We therefore study the time participants took to make decisions during the second half of the experiment. We find that participants in the nudge treatment spend on average slightly more time in the second half of the experiment than participants in the control treatment: 21.7 versus 20.9 seconds. This difference is however far from significant (Wilcoxon rank-sum test $p=0.81$). Nudged participants clearly do not appear to be less willing to put in effort.

Effort might however interact with the treatment effect in another way. If participants who put more effort into the task are more likely to make their own choice and ignore the default, we would expect the treatment to affect them less. If that is the case the treatment effect would be significantly smaller for this group. However, table 5.6 shows that this is not the case. The interaction effect of the treatment with our measure of effort, time taken in the first round, has no significant effect on earnings in the second half of the experiment.¹⁰³

5.4 Conclusion

In this study we showed that providing people with a nudge in the form of a good default can affect the choices and performance in subsequent choices. We find that subjects who have been provided with good defaults before, show a systematically larger default bias than non-nudged participants. Due to this, also the performance of the nudged subjects

¹⁰³As time taken in the second half can be affected by the treatment we take time taken in the first round as a measure of participants willingness to spend time on the task.

is lower. Analysis of the data shows that the larger default bias is the cause of the lower performance and that there were no treatment differences on other choice aspects that influence performance, like the average value of participants' non-default choices. The default-bias shows to be an endogenous phenomenon and the effect we find proves large enough to significantly influence performance.

While this is a single experiment and further studies should assess the robustness of this phenomenon, we do believe this result provides a note of caution to policymakers attempting to improve decisions using a nudge. When implementing a policy it is important to consider possible changes to the policy in the future and the effect a policy has on people's general attitude toward other choices. Of course we certainly do not argue that policy makers should never engage in libertarian paternalism. In fact our experiment showed that the nudge we provided helped participants to make better decisions. Realizing the effects on subsequent decisions may however be an extra element to consider while designing public policies and could also provide a ground to regulate certain business practices.

Stepping away from possible policy implications our experiment provides interesting behavioral insights. Most prominently we observe that the default bias is, at least partly, endogenous. People for whom choosing the default pays off develop a stronger default bias. Furthermore this effect is independent of a person's tendency to put effort into making her choice. Although more effort did improve the performance, those who put in more effort do not 'suffer' less from the negative effect of having previously faced good defaults.

In conclusion our results show that being nudged can affect subsequent decisions and that the choice environment faced influences the extent to which people rely on certain choice heuristics. We hope that these findings induce further research into the wider effects of choice-architecture.

5.5 Appendix: Instructions

General instructions

Welcome to this experiment on decision-making. Please read the following instructions carefully. When everyone has finished reading the instructions and before the experiment starts, you will receive a handout with a summary of the instructions. During the experiment you will be asked to make a number of decisions. Your decisions will determine your earnings. Your earnings will be privately paid to you in cash at the end of the experiment. The experiment will consist of 50 rounds. In each round you will face a choice task with which you can earn credits. The choice task will be explained on the next page. At the end of the experiment one of the rounds will be randomly selected. Your

payoff will be determined by the amount of credits you earned in that round. For each 10 credits you earn, you will receive 1 euro.

Instructions: choice task

In each round you have to select one option from a list. An example is shown below. Each option generates a number of credits equal to the sum of the points in each column multiplied by the weight of that column minus the price of the option: Number of credits generated by an option = 6*points in column 1 + 5*points in column 2 + 4*points in column 3 + 3*points in column 4 + 2*points in column 5 - price.

Choice Task 48						
TIME: 39 BONUS: 20						
Choices	Weight=6	Weight=5	Weight=4	Weight=3	Weight=2	Price
<input type="radio"/> Option 1	12	8	10	4	3	102
<input type="radio"/> Option 2	1	13	11	28	19	138
<input type="radio"/> Option 3	4	9	29	39	13	122
<input checked="" type="radio"/> Option 4	5	20	49	7	13	271
<input type="radio"/> Option 5	28	5	13	21	12	109
<input type="radio"/> Option 6	42	22	6	4	3	348

So for example, in the list above, option 4 will generate: $20*6+0*5+30*4+20*3+10*2-170=150$ points. You can think of this choice task as choosing which product to buy from a set of similar products. The products have different characteristics and prices which determine the value of each product.

Instructions: making a choice

Choosing an option

- You can select an option by clicking on the radio button next to the option.
- The selected option only becomes your choice if you press the 'Make choice' button or when the time limit runs out.

Time limit

- For each choice task you will have a maximum of 40 seconds to decide.
- If you make a choice before the time expires you will earn a bonus.
- This bonus starts at 20 credits but decreases with 1 credit every 2 seconds until you have made your choice.
- The remaining time and the bonus are depicted at the top of the choice task screen.

- After making your choice you will go to the waiting screen for the remaining seconds and 5 extra seconds.

Default option

- At the start of the round, one of the options will be the default option. This option is selected at the start of the round and the option row is shown in green.
- If you do not select any of the other options the default option will automatically become your choice when the time expires or when you press the 'Make choice' button.
- If you make your choice before the time has expired and you do not choose the default option you will be asked whether you want to stick with your original choice or whether you would prefer the default option. The time keeps running until you choose either your original choice or the default option. If the time runs out you stick with your original choice.
- The default option can be seen as a recommendation to buy a certain product.

We would like to remind you that at the end of the experiment one round is randomly selected. Your earnings will be determined by the amount of credits you earned in that round. Remember, for each 10 credits you earn, you will receive 1 euro.