Nudge lullaby*

Thomas de Haan

CREED, Amsterdam School of Economics, University of Amsterdam

Roetersstraat 11 1018 WB Amsterdam, The Netherlands

thomashaan08@gmail.com

and

Jona Linde

CREED, Amsterdam School of Economics, University of Amsterdam

October 2011

Abstract
Libertarian paternalism (Thaler and Sunstein 2003), using peoples own heuristics and biases to steer them to better choices while preserving choice freedom, has shown to be an effective and popular policy paradigm. It has also sparked many philosophical debates (e.g. Mitchell (2005) and Sugden (2008)) on the role of the government in consumer choice issues. In this paper we address a more practical possible side effect of libertarian paternalism in an experimental study. We have participants perform a number of difficult multi-attribute choice tasks where each option chosen yields a particular payoff. We examine whether participants who are provided with good defaults in the first half of the experiment perform differently compared to a control group, when they perform the task with random defaults in the second half of the experiment. We indeed find that subjects who first were nudged perform worse than the control group. Interestingly this effect holds even if we control for a proxy for effort. We find a significant treatment effect for men but not for women. Attitudes toward defaults thus seem not to be independent of previous experience. Our result has possible implications for government and non-government consumer choice policies.

*We would like to thank Joep Sonnemans, Theo Offerman, other colleagues from the Center for Research in Experimental Economics and political Decision making (CREED) and participants of the CESS-NYU/CREED meeting, September 2011 in New York, and the International ESA conference, July 2011 in Chicago for their helpful comments. Financial support from the University of Amsterdam Research Priority Area in Behavioral Economics is gratefully acknowledged.

1 Introduction

There is ample evidence that people have limited cognitive resources and that as a result they often rely on heuristics when making decisions (see e.g. Kahneman & Tversky, 2000). As a consequence their decisions can deviate from rational choice and 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 kind of behavior.

Knowledge of behavioral biases can be used to influence choices. Thaler and Sunstein (2003) and Camerer et al. (2003) propose to use our knowledge of biases to change the environment in which the choice is made, the "choice architecture", in an attempt to promote "better" decisions without changing incentives. They call their approach libertarian or asymmetric paternalism.[1] A feature of the choice architecture specifically designed to improve decisions is called a "nudge".

Numerous studies validate the effectiveness of this approach (Thaler and Sunstein, 2003). A problem with libertarian paternalism is that although inconsistent behavior shows that people sometimes act against their own interest it does not tell us what decision is optimal. For example, we know that the presence of irrelevant alternatives can influence which product people choose to buy (Huber, Payne & Puto, 1982), but this does not tell us which product they would have bought in the absence of behavioral biases. As a result libertarian paternalism has received its share of criticism focusing on the justification for influencing choices in a certain direction (e.g. Mitchell (2005) and Sugden (2008)).

We do not aim to take a side in this philosophical debate but consider a somewhat different, one might say more practical, possible concern with libertarian paternalism. We ask whether the most popular nudge, a good default option[2], affects performance in subsequent, similar decision situations.

Behavioral Economics research has amply illustrated that people are more likely to choose the option presented as the default or the option they currently possess (e.g.

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

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

Samuelson and Zeckhauser, 1988). This tendency is known as the status quo bias and commonly explained by loss aversion (Kahneman, Knetsch & Thaler, 1991). So far the assumption underlying libertarian paternalism is that the status quo is a constant behavioral tendency that can be used to influence decisions. However the choice architecture may not only change decisions but also more deeply influence the choice process. A behavioral tendency like the status quo bias might have an 'endogenous' component.

If people receive a nudge in the form of a good default, the choice heuristic to stick with the default option performs well. This may reinforce the use of this heuristic. This reinforcement may make the 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.

A nudge can also hurt later decisions because a person who faces good defaults becomes 'spoiled'. If the time comes that the defaults are no longer good options, the consumer might to some degree not have learned how to choose for herself. She may even feel resistance towards putting effort into a task that used to be easy, just 'pick the default', even if she realizes that the nature of the default has changed. 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.

A good default may however also improve performance in subsequent similar decisions without a nudge. A good default draws people's attention to good options which may teach a consumer what a good option looks like. This knowledge may help her make better decisions in the future, even when she is no longer nudged.

To test whether a good default affects 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 the nudge and control treatments in these second 25 periods reveals the effect of a nudge on subsequent behavior. In section 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 three possible reasons[3]. A first 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 that for some decisions, ‘good’ defaults are easier to provide than others. This happens in one of the most prominent examples of libertarian paternalism, a default enrollment in pension plans (Madrin & 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.

A third reason for following good with random defaults is 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 important, both in order to expand our insight in the use of decision heuristics and to design optimal libertarian paternalism policies, to explore the effect of a good default on subsequent decisions. Section 2 elaborates on our

[3] We do not claim to provided an exhaustive list, but we do believe that these three reason are probably the most prevelant.

experimental design. Section 3 presents the results of the experiment. Section 4 concludes.

2 Design

The experiment was computerized with php/mysql and conducted at the CREED laboratory of the University of Amsterdam. A total of 88 participants participated in the experiment, half of them assigned in the control treatment and half of them to the nudge treatment. At the beginning of the experiment participants read the instructions on the computer at their own pace. They then received a summary 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.

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 value. The weight values 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 point 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 1 below.[4]

---

4 This 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

Choice Task 48						
TIME: 39 BONUS: 20						
Choices	Weight=6	Weight=5	Weight=4	Weight=3	Weight=2	Price
Option 1	12	8	10	4	3	102
Option 2	1	13	11	28	19	138
Option 3	4	9	29	39	13	122
Option 4	5	20	49	7	13	271
Option 5	28	5	13	21	12	109
Option 6	42	22	6	4	3	348
Make choice						

Figure 1. An example of the choice task as presented to the participants. Option 4 is the default option in this example.

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 decision.[5] The bonus can be seen as an opportunity cost of spending time 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

---

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).

[5] For example he difference between the best and the second best option was always at least 10 credits.

counterbalanced between treatments[6]. At the end of the experiment one round was randomly selected[7]. The number of credits earned in that round determined the participant's earnings. Each credit was worth 10 eurocents.

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 the time limit ran out before the participant entered a choice.[8] 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.[9]

The experimental design consisted 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 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.

[6] The 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.

[7] The 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.

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

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

3 Results

Because it was in the participants' best interest to choose the option that would yield maximum earnings, we use the value of the chosen options as a performance measure.[10] We will report the main treatment effect in the first subsection and explore the different possible explanations in the second subsection. All tests reported are two-sided and, unless otherwise specified, performed at the individual level.

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 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. [11] A Wilcoxon rank-sum test shows that this difference is marginally significant (p=0.081).

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 1 indeed confirms that providing people with a good default helps them make better decisions. As the first row of table 1 shows nudged participants chose an option worth 18.52 points per round more in the first half of the experiment. That is a substantial and highly significant difference, according to a Wilcoxon rank-sum test.

[10]We leave the bonus payment the subjects received for the speed of their decision out of the analysis for now as this is not the point of focus. Adding the bonus does not significantly change the results. In fact participants in both treatments spend an average of 19.4 seconds per task.

[11] For 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.

	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	

Table 1: Average value of the chosen option split over treatment and first and second half of the experiment (standard deviations at the individual level between brackets).

The regression in table 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 10 rounds as a proxy of effort. As time spend can be influenced by the treatment we take the time spend during the first ten rounds as an exogenous measure of effort.[12] Controlling for these variables in the regression, the treatment effect becomes significant at a 5% level.

Dependent variable: average value of the chosen option in the second half of the experiment	Coefficient	p-value
constant	147.82	0.000
treatment (1=nudge treatment)	-6.62	0.047
time used in first 10 rounds	0.97	0.000
male	11.54	0.001
age	0.01	0.848
studies economics	1.75	0.664
Dutch	-5.56	0.213
math grade	0.20	0.860
math level	3.51	0.393
adjusted r-squared	0.22	

Table 2: Treatment effect when controlling for demographic variables and proxies for effort and skill.

Result 1: Participants who had previously been nudged chose worse options than participants who had not previously been nudged.

[12] Time spend in the first ten rounds strongly correlates with time spend in later rounds (Spearman correlation coefficient is 0.6346 and p-value) Using time spend in the entire experiment or only the second half yields the same qualitative results

In the first few rounds after the default changes from the best to a random option participants might not have noticed that the default is no longer helpful. If that would have been the case, the treatment effect might be an artifact of the first few rounds of the second half. However, as figure 2 shows, the treatment effect persists throughout the second half of the experiment. In almost every round average earnings were higher for the control than for the treatment group and this difference does not show a tendency to decline.

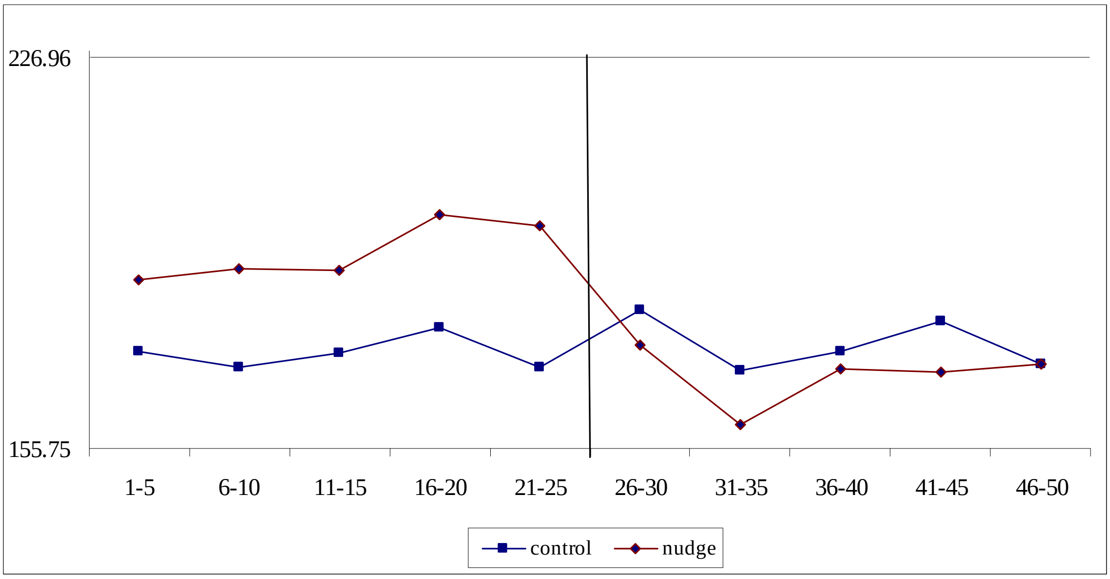

Figure 2: Average value of the chosen option aggregated over 5 rounds split between the nudge and control treatments. The bottom of the graph corresponds to the expected value over the all 50 tasks for a person who chooses randomly (155.75), the line at the top to the average value of the chosen option for a person who always chooses the best option (226.96).

3.2 Causes

As mentioned in the introduction there are two possible causes for the treatment effect. Firstly participants who have received a helpful nudge may have come to rely on default more even though it was no longer helpful. Secondly participants in the nudge treatment may have become spoiled and therefore unwilling to put in effort.

Reliance on the nudge

If placing greater trust in the default is indeed the cause of the treatment effect we should find that also in the second half of the experiment participants in the nudge treatment selected the default option more often than participants in the control treatment. As table 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. Figure 3 also illustrates this effect.

	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	

Table 3: Percentage of default choices per round split over treatment and first and second half of the experiment.

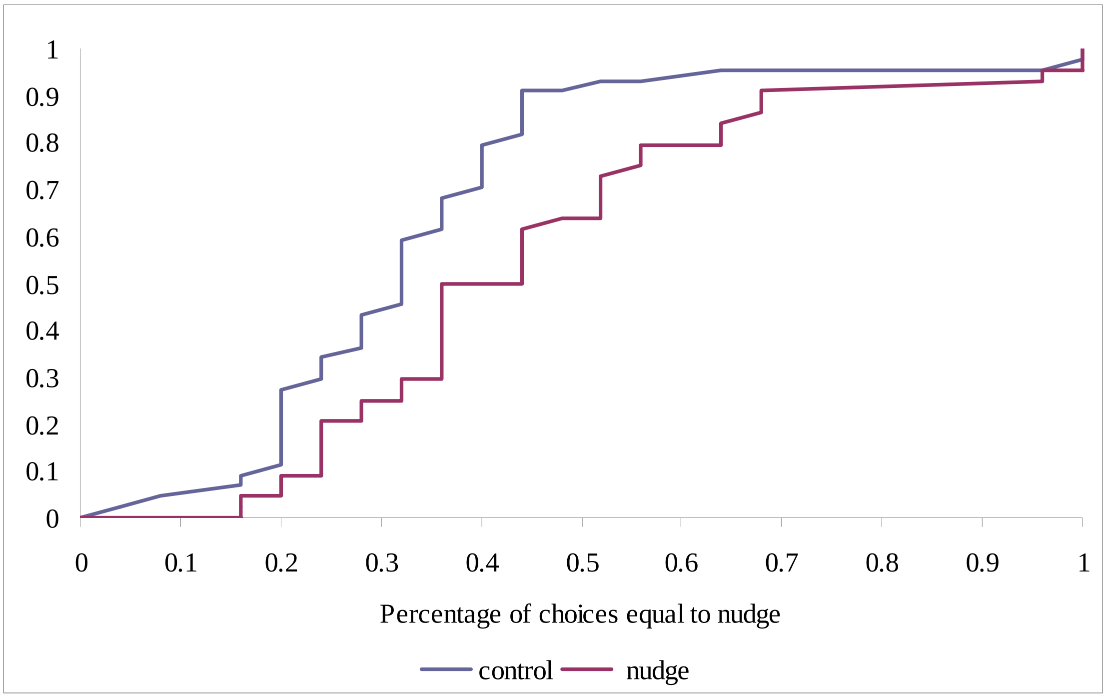

Figure 3: Cumulative distribution of participant's tendency to rely on the default option in the second half of the experiment for both treatments.

Result 2: Participants who have previously been nudged choose the default option more often than participants who have not.

The likelihood of default choices in the first half provides further evidence that participants in the nudge treatment indeed came to rely more on the default than the participants in the control treatment. In both treatments participants chose the default more often than the expected random percentage of 16.7% (rank-sum p-values <0.001), but nudged participants were significantly more likely to do so (rank-sum p=0.000). This in itself does not provide evidence for a greater trust in the default in the first half of nudge treatment. Participants may also have chosen the default more often in the nudge treatment because it was the best option which they could 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%%, rank-sum p=0.0025). We therefore conclude that participants in the nudge treatment came to trust the default more than participants in the control treatment in the first half.

Although this shows that nudged participants were more likely to choose the default, it does not necessarily explain their worse performance in the second half. They could, for example, have chosen the default more often, but only when the default provides relatively good earnings. Conversely their worse performance may also be due not only to choosing the default more often, but also to choosing the default when it is a relatively bad option. The logistic regression in table 4 below shows that neither effect is present. Participants in the nudge treatment are neither more nor less sensitive to the relative value of the default option, they just follow the nudge more often.

Result 3: Participants who have previously been nudged are no less sensitive to the difference in value between the default and the best option.

Dependent variable: choice equals the default in the second half	Coefficient	p-value
constant	0.373	0.000
treatment (1=nudge treatment)	0.295	0.008
difference in value between the default option and the best option	-0.012	0.000
interaction between treatment dummy and the difference in value between the default option and the best option[a]	0.00	0.285

Table 4: Probit regression examining the effect of the treatment, the difference in value between the default option and the best option and their interaction on the likelihood of choosing the default option in the second half of the experiment. Each choice is used as an independent observation, standard errors have been adjusted by treating each participant's choices as a cluster.

[a] The interaction term was normalized to prevent multi-colinearity issues with the dummy variable for treatment.

Effort

We now turn to the second possible explanation of the treatment effect: nudged participants putting in less effort. In opposition to this hypothesis 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.

Result 4: Participants who have previously been nudged do not put less effort into the task than participants who have not.

Effort might however interact with the treatment effect in another way. We saw that nudged participants came to trust the default more and therefore perform worse. It seems that this is less likely to occur for those participants who put more effort into the task as these would be more likely to realize that the default no longer provides guidance. We therefore explore interaction effect between our proxy for effort, time used in the first ten rounds and a treatment dummy (table 5).[13] This regression shows that there is no significant interaction, although the effort proxy on itself does have a significant positive effect. Participants identified as putting in more effort are no less likely to perform worse due to having been nudged.

Dependent variable: average value of the chosen option in the second half	coefficient	p-value
Constant	156.99	0.000
treatment (1=nudge treatment)	-6.87	0.043
time used in the first 10 rounds	0.85	0.007
Interaction treatment and time used in the first 10 rounds[a]	-0.24	0.585

Table 5: Interaction between effort and the treatment effect.

[a] The interaction term was normalized to prevent multi-colinearity issues with the dummy variable for treatment.

Result 5: Participants who (initially) put more effort into the task are no less likely to perform worse due to having been nudged.

[13] We take time taken in the first ten rounds because this time in the second half could be affected by the treatment. Table 2 shows that putting more effort into the task indeed improves performance.

Gender effect

As we can see from the regression in table 2, there is a significant effect of gender on the earnings in the second half. Although we did not hypothesize any gender effect, we feel compelled to analyze the effect further given the size and significance of the effect in the regression.

When we split the results by gender an interesting pattern emerges. Table 6 below shows again the average earnings per round for the first and second 25 rounds, but now separately for men and women. For the first half of the experiment, the picture for men and women is similar. Again, having a good default helps very much. But in the second half of the experiment, we see no difference for females but only a significant treatment difference for men. The reason for this difference between treatment effects appears to be that in the control treatment men improve performance in the second half compared to the first half while women in the control treatment perform worse in the second half.

	Control treatment	Nudge treatment	p-value Wilcoxon rank sum test
first half, females	173.7 (14.2)	190.0 (10.8)	0.000
second half, females	168.3 (15.1)	166.6 (17.9)	0.942
p-value Wilcoxon signed rank test	0.046	0.002	
first half, males	172.4 (14.3)	192.3 (12.5)	0.000
second half, males	179.7 (13.3)	170.9 (17.7)	0.028
p-value Wilcoxon signed rank test	0.007	0.000	

Table 6: Average earnings per round, split between treatments, first and second half of the experiment and males and females.

Result 6: Having been nudged leads to a significantly worse performance for males but not for females.

Conclusion

In this study we showed that providing people with a nudge in the form of a good default can affect performance in subsequent decisions. Participants in our experiment who faced a good default in the past came to rely on the default more than other participants, leading to worse decisions when the default was no longer helpful. We therefore conclude that nudging people not only changes choices, but also seems to affect the choice process, possibly leading to worse decisions when circumstances change.

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 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 direct policy implications our experiment provides interesting behavioral insights. We find that, although putting more effort into the task did improve the performance, our main treatment effect still persists even if we control for effort and on top of that we see that those who put in more effort do not 'suffer' less from the negative effect of having previously faced good defaults. Furthermore, the fact that we observe a significant treatment effect for males but not for females is an intriguing finding.

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.

References

Ariely, D. and Wertenbroch, K. (2002). Procrastination, deadlines, and performance: Self-control by precommitment. *Psychological Science*, 13, 219-224.

Camerer, C., Issacharoff, S., Loewenstein, G., O'Donoghue, T. and Rabin, M. (2003). Regulation for conservatives: Behavioral economics and the case for "asymmetric paternalism." *University of Pennsylvania Law Review*, 151, 1211-1254.

Caplin, A. and Martin, D. (2011). A Testable Theory of Imperfect Perception. Working paper.

Choi, J.J., Laibson, D., Madrian, B.C. and Metrick, A. (2003). Optimal defaults. *American Economic Review*, 93, 180-185.

Frederick, S., Loewenstein, G. and O'donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature*, 40, 351-401.

Huber, J., Payne, J.W. and Puto, C. (1982). Adding Asymmetrically Dominated Alternatives: Violations of Regularity and the Similarity Hypothesis. *Journal of Consumer Research*, 9, 90-98.

Kahneman, D., Knetsch, J. L. and Thaler, R. H. (1991). Anomalies: The endowment effect, loss aversion, and status quo bias. *The Journal of Economic Perspectives*, 5, 193-206.

Kahneman, D. and Tversky, A. (Eds.). (2000). Choices, values and frames. New York: Cambridge University Press.

Kalaycı, K. and Potters, J. (2011). Buyer confusion and market prices. *International Journal of Industrial Organization*, 29, 14-22

Lichtenstein, S. and Slovic, P. (1971). Reversals of Preference between Bids and Choices in Gambling Decisions. *Journal of Experimental Psychology*, 89, 46-55.

Madrian, B.C., and Shea, D.F. (2001). The power of suggestion: Inertia in 401(k) participation and savings behavior. *Quarterly Journal of Economics*, 116, 1149–87

Mitchell, G. (2004). Libertarian paternalism is an oxymoron. *Northwestern University Law Review*, 99, 1245-1277.

Odean, T. (1998). Are Investors Reluctant to Realize Their Losses? *The Journal of Finance*, 53, 1775-1798.

Samuelson, W. and Zeckhauser, R.J. (1988). Status Quo Bias in Decision Making. *Journal of Risk and Uncertainty*, 1, 7-59.

Sunstein, C. R., & Thaler, R. H. (2003). Libertarian paternalism is not an oxymoron. University of Chicago Law Review, 70, 1159-1202.

Sugden, R. (2008). Why incoherent preferences do not justify paternalism. *Constitutional Political Economy*, 19, 226-248.

Thaler, R. H. and Sunstein, C. R. (2003). Libertarian paternalism. *American Economic Review*, 93, 175-179.

Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*, Conneticut (USA), Yale University Press.