

# 'Good Nudge Lullaby': Choice Architecture and Default **Bias Reinforcement**

Citation for published version (APA):

de Haan, T., & Linde, J. (2018). 'Good Nudge Lullaby': Choice Architecture and Default Bias Reinforcement. The Economic Journal, 128(610), 1180-1206. https://doi.org/10.1111/ecoj.12440

Document status and date: Published: 01/05/2018

DOI: 10.1111/ecoj.12440

**Document Version:** Publisher's PDF, also known as Version of record

**Document license:** Taverne

#### Please check the document version of this publication:

 A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.

• The final author version and the galley proof are versions of the publication after peer review.

 The final published version features the final layout of the paper including the volume, issue and page numbers.

Link to publication

#### General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these riahts.

Users may download and print one copy of any publication from the public portal for the purpose of private study or research.

- You may not further distribute the material or use it for any profit-making activity or commercial gain
   You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

#### Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.





The Economic Journal, 128 (May), 1180–1206. Doi: 10.1111/ecoj.12440 © 2016 Royal Economic Society. Published by John Wiley & Sons, 9600 Garsington Road, Oxford OX4 2DQ, UK and 350 Main Street, Malden, MA 02148, USA.

# 'GOOD NUDGE LULLABY': CHOICE ARCHITECTURE AND DEFAULT BIAS REINFORCEMENT\*

#### Thomas de Haan and Jona Linde

Because people disproportionally follow defaults, both libertarian paternalists and marketers try to present options they want to promote as the default. However, setting certain defaults and thereby influencing current decisions, may also affect choices in later, similar decisions. In this article, we explore experimentally whether the default bias can be reinforced by providing good defaults. We show that people who faced better defaults in the past are more likely to follow defaults than people who faced random defaults, hurting their later performance. This malleability of the default bias explains certain marketing practices and serves as an insight for libertarian paternalists.

Influencing people's decisions is an important goal of both clever marketers and ambitious policy makers. Because human behaviour is better characterised by bounded rationality than by the full rationality assumed in traditional economic models, 'choice architects', well-meaning or self-interested, can increase the probability that their favoured option is chosen by influencing the environment in which people make their decisions without changing incentives or prices (Thaler and Sunstein, 2003). One prominent tactic of both marketing executives and the so-called libertarian paternalists is to make an endorsed option the default choice, because many people disproportionally follow defaults (Samuelson and Zeckhauser, 1988). In this article, we study experimentally whether this 'default bias' is affected by recent experiences with defaults. Specifically we study whether people who were recently exposed to good defaults are more likely to follow random defaults compared to people who faced random defaults. If the nature of the choice architecture of recent choice situations affects subsequent behaviour this could mean that techniques such as those proposed by libertarian paternalists can have a 'side effect'. It can also explain ways in which defaults are used in marketing.

Marketing executives are generally thought to set the option that yields the highest profit as the default. However, research by Brown and Krishna (2004) and Campbell (2007) suggests people deduce the quality of the default from their beliefs about the default setter and adjust their behaviour accordingly, which reduces the effectiveness of such a policy. Therefore, it may be more effective to first lull consumers into a false sense of security by setting defaults in their best interest, only to follow up with profit

<sup>\*</sup> Corresponding author: Thomas de Haan, Norwegian School of Economics, Helleveien 30, 5045 Bergen, Norway. Email: thomas.dehaan@nhh.no.

We thank Gary Charness (UCSB), Joep Sonnemans, our colleagues from the Center of Research in Experimental Economics and Political Decision making (CREED), the Human Capital research group at the University of Amsterdam, the Economics of Innovative Change group at the Friedrich Schiller Universität and the Max Planck Institute of Economics, Jena and the participants of the 2011 CCC-workshop in Nottingham, the 2011 ESA world conference and the NYU/CREED graduate student conference 2011 for their helpful comments. Financial support from the University of Amsterdam Research Priority Area in Behavioural Economics and from the University of Stirling is gratefully acknowledged. Jona Linde's research is supported by a VIDI grant from the Netherlands Organisation for Scientific Research (NWO, grant nr. 452-11-005).

maximising defaults later on. Sellers can do this when consumers repeatedly purchase from the same supplier and when a purchase requires several separate decisions (e.g. buying a car, a computer, or a plane ticket), because in both cases sellers can set many consecutive defaults.<sup>1</sup> Indeed anecdotal evidence suggests that, when a purchase requires several decisions, initial defaults are often better than the later ones: first you are recommended economy class and direct routes but later buying expensive flight insurance or extra luggage allowance may well be the default option.

Libertarian or asymmetric paternalists (Camerer et al., 2003; Thaler and Sunstein, 2003) do not set defaults to further their own interests but to help people make better decisions.<sup>2</sup> Behavioural economic research shows that people can make decisions that go against their own interest (Hoch and Loewenstein, 1991) and that their choices can be influenced by features of the choice environment that are irrelevant according to standard rational choice theory (Kahneman and Tversky, 2000). Myopic time preferences (see for an overview Frederick et al., 2002) and framing (Tversky and Kahneman, 1981) are well known examples of the first and the second type of behaviour respectively. Libertarian or asymmetric paternalists therefore advocate designing the choice environment to promote 'better' decisions without changing incentives. They propose to introduce 'nudges', which steer boundedly rational people towards better decisions without affecting people who are sufficiently rational. Numerous studies validate the effectiveness of this approach (Thaler and Sunstein, 2003). Policy makers recently appear more and more open to the idea of nudging people as evident from the establishment of the, now independent, 'Behavioural Insights Team'<sup>3</sup> by the British government and similar teams in other countries. In a recent report the World Bank (2015) commends the approach as a tool for economic development. Libertarian paternalism has also sparked a lively media debate.<sup>4</sup>

Because the default bias is a common and well established behavioural bias (Samuelson and Zeckhauser, 1988; Madrian and Shea, 2001; Johnson and Goldstein, 2003) setting a good default is a prevalent nudge. However, if facing good defaults reinforces the default bias, this type of nudge may have an unintended side effect if good defaults are followed by inferior ones. One reason why that may happen 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: default enrolment in pension plans (Madrian and Shea, 2001) to prevent under-saving due to time inconsistency. Saving something, and therefore participating, is probably optimal for the large majority of employees but there is far more heterogeneity in how much people should ideally save. Setting a 'good' default savings rate is therefore far more

<sup>&</sup>lt;sup>1</sup> Our experiment mainly focuses on situations where decision makers repeatedly face similar decisions. Levav *et al.* (2010) provide an interesting investigation of settings where consumers have to take a series of decisions to customise a product. They show, among other things, that consumers are more likely to follow the default if earlier decision situations had more options and therefore presumably required more effort.

<sup>&</sup>lt;sup>2</sup> In the remainder of this article, we use the term libertarian paternalism to refer to this approach. <sup>3</sup> http://www.behaviouralinsights.co.uk/, their blog chronicles numerous applications of behavioural paternalism around the world.

<sup>&</sup>lt;sup>4</sup> see e.g. http://www.economist.com/news/britain/21625871-behavioural-economics-changing-regulation-payday-lending-target-nudge-nudge-think-think and http://www.ft.com/intl/cms/s/2/9d7d31a4-aea8-11e3-aaa6-00144feab7de.html

<sup>© 2016</sup> Royal Economic Society.

difficult (Choi *et al.*, 2003). As a result, the level of the default savings rate is probably a less helpful default than the default to enrol and save at least something.

A second reason why libertarian paternalism can lead to good defaults followed by worse ones is legal limitations. Courts may view libertarian paternalism as unwarranted government intervention. In their book *Nudge* Thaler and Sunstein (2008) discuss a programme implemented in Maine (USA) to provide Medicare users with a good default health care programme. Legal challenges have contributed to the failure of this project to spread to other states. Similar legal challenges may cause abandonment of libertarian paternalism programmes already in place. If that happens, a person used to helpful defaults may face suboptimal (e.g. random) defaults in the future. Lobbying efforts by affected companies or a change in the ideological nature of the government can have similar effects.

The idea that facing good defaults increases the default bias assumes that the default bias is not a constant tendency. Caplin and Martin (2013) show that many people indeed do not blindly follow defaults, but consider whether a default is likely to help them or not. In their experiment giving participants explicit information about how often the default equals the best option strongly impacts participants' default bias and the effort they put into making an active decision.<sup>5</sup> Experiments by Brown and Krishna (2004) and Campbell (2007), already referred to above, show that people deduce the quality of the default from their beliefs about the default setter. In this article, we examine a related issue: whether people also change their behaviour based on recent experiences with a default. Given that the default bias is pliable it seems plausible that earlier experiences may also affect the default bias, either through beliefs about the quality of the defaults, as suggested by the findings mentioned in this paragraph, or through some other mechanisms.

Whether facing good defaults indeed reinforces the default bias is likely to depend on what causes the default bias in the first place. There are several explanations as to why people exhibit a default bias. One possibility is that a preselected option simply attracts more attention (Caplin and Martin, 2013). A second possibility is based on loss aversion (Kahneman and Miller, 1986, Kahneman et al., 1991). If the default is treated as a reference point, people will compare other options to the default. Due to loss aversion, dimensions on which the rival option is worse than the default get relatively more weight than dimensions on which the rival option is better. A third possible reason why people may be inclined to choose the default option is that a default option is perceived as advice. Madrian and Shea (2001), for example, find that one cause for higher enrolment rates into a pension plan when enrolment is the default is that people take the default as implicit advice from their employer. A fourth possibility is that choosing the default is a choice heuristic (Tversky and Kahneman, 1974). If making a choice is difficult, people are likely to use simple rules and following the default is one possible rule. This could also be due to people trying to avoid the mental cost of making a difficult decision (Anderson, 2003). Given a limited capacity to process information the last two

MAY

 $<sup>^{5}</sup>$  In their experiment this response is so strong that it may actually negate the positive effect of a good default.

<sup>© 2016</sup> Royal Economic Society.

explanations are not necessarily biases but possibly a boundedly optimal way to make decisions.  $^{6}$ 

All of the possible explanations of the default bias mentioned in the previous paragraph allow for an effect of an earlier helpful default on the default bias although they do not necessarily predict it. A default option may well draw more attention if it is associated with good outcomes. If the default bias is driven by loss aversion, it is important that the agent's reference point coincides with the default. Facing good defaults for some time could possibly ensure this and therefore establish a stronger default bias. The default is also more likely to be seen as advice if it has proven to be good advice in the past. If the default bias is a choice heuristic, past experiences can also affect how often people choose the default because successful heuristics could be reinforced (Anufriev and Hommes, 2012). An individual who faced good defaults is therefore more likely to use this heuristic, even if it is not a good heuristic now. Although in this article, we do not distinguish between these possible explanations of the default bias we do examine whether the default bias is affected by earlier experiences with defaults.

To test whether facing a good default reinforces the default bias we developed an experimental task with an unequivocal best choice which is nevertheless hard to find. The task is designed to resemble the type of task for which a nudge is likely to be useful: difficult and hard to translate into easily understandable terms and difficult to learn.<sup>7</sup> Participants face this task for 50 rounds. In the first 25 rounds participants in the 'good-default' treatments receive a nudge in the form of a good default, while participants in the 'random-default' treatments receive a random default. In the second 25 rounds, both groups receive a random default. Any difference in choices between participants in the good-default and random-default treatments in these second 25 rounds reveals the effect of a good default on subsequent behaviour.

Our findings indeed suggest an increased reliance on the default for participants who faced good defaults in the first 25 rounds. Although good-default participants immediately choose the default at a lower rate in the rounds when the default becomes random, they keep choosing the default about 10–20 percentage points more often than participants from the random treatment. We also identify a significant negative effect on the quality of the decisions for good-default treatment participants in the second half, however the choices in the first 25 rounds are much better for the good-default participants due to the very good defaults they receive. In the next Section, we describe the experimental design in detail. Section 2 presents the results of the experiment and introduces a behavioural model to explain these results. Section 3 discusses an extension to bad defaults and Section 4 concludes.

<sup>&</sup>lt;sup>6</sup> Nevertheless, we refer to people's tendency to follow the default as the default bias.

<sup>&</sup>lt;sup>7</sup> Thaler and Sunstein (2008, pp. 79–88) recommend these criteria to determine whether a nudge is warranted. They also mention choices where benefits and costs are separated in time. For practical reasons this aspect does not play a big role in our experiment, although participants do need to put effort into the task to perform well and do not receive benefits from this until the end of the experiment.

# 1. Design

The experiment was computerised 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.<sup>8</sup> They then received a summary of the instructions on paper. After reading the instructions, participants had to answer some questions correctly to test their understanding of the instructions.

All participants in the experiment performed the same set of 50 multi-attribute choice tasks.<sup>9</sup> We implement a two-by-two design where one dimension is whether participants get feedback on the value of their decision or not and the other dimension is a difference in the nature of the default, either the best option (good-default treatments) or a random option (random-default treatments), in the first half of the experiment. This results in four treatments: no-feedback good-default, no-feedback random-default, feedback good-default and feedback random-default. Decisions in the second half, when all participants face the same task and the same, random, default reveals the effect of being nudged by a good default on subsequent decisions. This effect is the main focus of this article. Comparing the size of this effect between treatments with and without feedback allows us to explore in what kind of setting the nature of the default has a larger effect on the default bias and further explore the cause of any difference between the good-default and random-default treatments. Below we first discuss the choice task, then the difference in the nature of the default in different treatments and, lastly, the feedback participants receive in the different treatments.

#### 1.1. Task

Each round of 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 six 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. Points in a category ranged from 0 to 399 for the price and from 0 to 49 for the other categories.<sup>10</sup> The 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 credits and that the best option generated at least 10 credits more than the second best option. An example of a task is shown in Figure 1.<sup>11</sup>

© 2016 Royal Economic Society.

<sup>&</sup>lt;sup>8</sup> The instructions can be found in online Appendix A.

<sup>&</sup>lt;sup>9</sup> These tasks were presented in different orders, counterbalanced between treatments, see below for details.

<sup>&</sup>lt;sup>10</sup> Participants were not informed about these ranges.

<sup>&</sup>lt;sup>11</sup> 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 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 policy or an investment plan. This type of task has previously been used by Kalayci and Potters (2011).

| Choice Task 48           |              |              |              |              |              |       |
|--------------------------|--------------|--------------|--------------|--------------|--------------|-------|
|                          |              | TIME         | : 39 BONUS   | S: 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 |              |              |              |              |              | 348   |
| Make Choice              |              |              |              |              |              |       |

Fig. 1. Example of the Choice Task as Presented to Participants Notes. Option 4 is the default option in this example. Colour figure can be viewed at wileyonlinelibrary.com.

A participant could select an option by clicking on the button in front of the option and could change his selection as often as he wanted. The selected option was recorded as her choice when she pressed the 'make choice' button. A participant had 40 seconds in which to make a choice. If these 40 seconds ran out the option that was selected at that time was recorded as the participant's choice.<sup>12</sup> After a participant made his decision he had to wait until the time alloted for this round expired before moving on to the next round. In addition, there was a 5-second waiting time between rounds. Participants were allowed to use pen and paper but no calculators or other devices that could calculate for them.

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 diminishing bonus puts a small cost on spending more time on the task.<sup>13</sup> We implemented this bonus to ensure that a participant would enter his (or her) choice as soon as (s)he had decided, which was desirable for two reasons. First, in order to have a measure of search effort in the form of time spent on the task. Second, 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 in which all participants faced the same 50 tasks in six different orders. The order was counterbalanced between treatments.<sup>14</sup> At the end of the experiment, one round was randomly selected.<sup>15</sup> The number of credits earned in that round determined the participant's earnings. Each credit was worth 10 euro cents.

<sup>15</sup> 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.

© 2016 Royal Economic Society.

 $<sup>^{12}</sup>$  This happened 325 times out of 8,700 decisions, i.e. in 3.74% of all decisions.

 $<sup>^{13}</sup>$  Even all 20 credits were a small amount compared to the gains that could be made by making a better decision. For 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.

<sup>&</sup>lt;sup>14</sup> 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 no-feedback good-default treatment and one in the no-feedback random-default 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.

The task we chose for this experiment has four advantages. First, it has an unambiguous outcome measure in the value of the chosen option, while the value of the options is not easily determined by participants. In fact our task is difficult in exactly the way decision situations for which Thaler and Sunstein (2008, p. 79) suggest people need a nudge are difficult: because it is a task where people 'have trouble translating aspects of the situation into terms that they can easily understand' in this case difficulty translating points with a weight into an aggregate value.<sup>16</sup> Second, because each choice takes only 40 seconds participants can perform relatively many tasks in a limited period of time. This allows for a substantial number of tasks to be faced in which participants can become used to a certain kind of default and a subsequent set of tasks to evaluate the effect of the earlier default.<sup>17</sup> Third, because it is a choice between options it naturally allows for the setting of a default. Fourth, because it is a difficult task, a good default can be very useful for participants, which is likely to strengthen the effect of having faced such a default. The combinations of these features results in a task that is likely to reveal a substantial effect of facing good defaults on the default bias in subsequent decision if such an effect exists. The difficulty increases the usefulness of a good default but also depresses performance in the random-default group. For that reason, other types of tasks may well reveal a larger effect of a stronger default bias on performance.

#### 1.2. Default

In every round, one of the six options in the table was given a different colour and was preselected when participants were presented with the task (see option 4 in Figure 1). This option was the default option. In the instructions, it was labelled 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 selected another option.<sup>18</sup> If a participant did not chose 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.<sup>19</sup>

Participants were randomly assigned to a 'random-default' or 'good-default' treatment. Both types of treatment were identical, except for the nature of the default option in the first half of the experiment. In the random-default treatments, the default option was always determined randomly for each task. In the good-default treatments, 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

<sup>&</sup>lt;sup>16</sup> We present evidence about the difficulty of the task at the beginning of Section 2.

<sup>&</sup>lt;sup>17</sup> Thaler and Sunstein suggest that nudges are most important for tasks which are rare. Given the need for repetition we cannot incorporate this feature in our experiment. However, libertarian paternalism has also been advocated, including by Thaler and Sunstein, for decisions which are repeated often, such as food choices in cafeteria.

<sup>&</sup>lt;sup>18</sup> This happened only 87 times out of 8,700 decisions, i.e. in 1% of all decisions.

<sup>&</sup>lt;sup>19</sup> People switched a total of 246 times out of 6,668, i.e. in only 3.7% of all initial non-default choices.

<sup>© 2016</sup> Royal Economic Society.

the good-default treatments in the rounds 26 to 50 had the same random defaults as in the random-default treatments.  $^{20}$ 

These treatments are selected to maximise the impact of earlier defaults on the default bias in an effort to explore whether such an effect can occur. This is why we have chosen a rather extreme situation where participants face either perfect or purely random defaults. Very good early defaults are most likely to enhance the default bias and random defaults will likely leave a default bias unaffected in the random-default group or even decrease it. For the same reason, we provide participants with very little explicit information on the nature of the default. As a result, they only have their own experience to rely on, which we think increases the effect of that experience. At the same time, we believe our experimental design does provide an appropriate model for the type of situation we intend to study. As in our experiment, real world defaults are also often either as good as possible, for example when they are set by a wellintentioned choice architect or a marketer trying to increase the default bias of customers; or random, for example, if a choice architect does not consciously set a default. The lack of information about the nature of the default is similarly realistic. People can of course try to deduce the nature of the default from the default setter as Brown and Krishna (2004) and Campbell (2007) show but this is not always straightforward. Businesses may set genuinely good defaults in an effort to stimulate consumer satisfaction and repeat purchases (Goldstein et al., 2008). Governments have other goals than helping people make better decisions, as is evident in 'nudges' which aim for societal or environmental benefits rather than better individual decisions (Cabinet Office Behavioural Insights Team, 2011). To model other situations or further examine the way in which different kinds of defaults affect the default bias variations of our design could be applied.

#### 1.3. Feedback

The feedback participants receive during the experiment depends on the treatment. In the no-feedback treatments participants receive no pay-off feedback until the end of the experiment, when they learn the value of the option they chose in the round which is randomly selected to determine their payout. Conversely, in the feedback treatments participants are shown the value, in credits, of the option they chose plus separately the time bonus, immediately after each round. Because there is a small waiting time between rounds, participants who only recorded a choice when the decision time runs out also get a chance to see the value of their chosen option.

Thaler and Sunstein (2008, p. 79) promote nudges especially for decisions 'for which [people] do not get prompt feedback'. Indeed, in many decision situations for which a nudge has been proposed important aspects of the value of a chosen option are unknown until far in the future. For example, the health effects of food choices and future consumption possibilities resulting from savings decisions usually do not become clear until many years later. In other cases where nudges have been proposed

 $<sup>^{20}</sup>$  The random default option was predetermined for each task, so participants who perform the same task also have the same default option, except participants in the good-default treatments in the first half of the experiment for whom the default option is the option with the highest value.

<sup>© 2016</sup> Royal Economic Society.

there is some feedback about the value of a chosen option but this feedback is very noisy, so it takes a lot of experience to learn the expected value of an option. This is the case with, for example, insurance and investment decisions. In decision situations with noisy and/or delayed feedback a decision maker faces many similar decisions before the true value of a chosen option becomes clear. To model the effect of a good default on subsequent decisions in such situations it therefore seems most appropriate not to provide participants with feedback about the quality of their decisions, as we do in the no-feedback treatments.

No-feedback is representative of most situations in which people are nudged but it is less representative of many situations where marketers set defaults, particularly for situations in which you repeatedly purchase from the same supplier.<sup>21</sup> If you bought the flight suggested by a search engine, a car with the standard options, or a computer with default settings, you will soon find out what the pay-off of your choice is. Our treatments with feedback more closely resemble such situations.

Besides being representative of different types of situations in which people face defaults in real life, comparing the effect of a good default on subsequent decisions in situations with and without feedback may also reveal something about the reasons why earlier defaults affect decisions. Assuming that having received a good defaults increases the default bias, a first thing that might be revealed is whether the increased default bias is caused by either a conscious decision to trust the default rather than rely on your own decisions, or a less conscious increased attraction to the default option. A second thing that comparing treatments with and without feedback can reveal is the interaction between the nature of the default and the quality of decisions. In feedback treatments participants can evaluate the quality of their decisions.<sup>22</sup> It is possible that this reduces reliance on the default, perhaps only in the random-default treatment, in which case it could lead to a stronger effect of having been nudged.

#### 2. Results

A total of 174 participants participated in the experiment, 44 in both treatments with feedback and 43 in both treatments without feedback.<sup>23</sup> Eighty eight participants were male and 113 were economics or business students. On average the sessions lasted about an hour and the average earnings were 19.25 euro.

The results from the random-default treatments show that the task fulfils the criteria listed in subsection 1.1: difficult but not impossible and with limited opportunities to

<sup>&</sup>lt;sup>21</sup> For purchase decisions which require several separate consecutive decisions (e.g. buying a plane ticket, a computer, or a car) and therefore allow for consecutive defaults, it does hold that you do not receive feedback about the quality of your decision before facing later defaults. The same is true for defaults for commercial products which have the same characteristics as the products for which nudges have been proposed, for example, insurance.

<sup>&</sup>lt;sup>22</sup> Because participants only receive feedback about the value of the chosen option and not about the value of the other options they cannot be sure whether they made a good or a bad decision but the value of an option, especially relative to the value of the chosen option in other rounds, does give an indication of the quality of a decision.

 $<sup>^{1}</sup>$   $^{23}$  We aimed to recruit 50 participants per treatment, which we believe is a reasonable number for an individual decision making experiment. Due to slightly lower than expected attendance the number of actual participants is slightly lower.

<sup>© 2016</sup> Royal Economic Society.

learn how to select a good option. In the random-default treatments participants on average chose an option with a value of approximately 174, with or without feedback. That means they managed to realise around 25% of the difference in expected value between picking a random option (155.75) and always picking the best option (226.96). Without feedback, we find no learning effect but with feedback there is some evidence for learning. A regression using data from the random-default treatment with feedback of the value of the chosen option on the number of the round, with task fixed effects and standard errors clustered at the individual level reveals that the round has significant positive effect (p < 0.0001).<sup>24</sup>

To examine the effect of a good default on subsequent decisions we examine the behaviour of participants in the second half of the experiment where participants in good-default and random-default treatments face the same choice tasks and the same (random) defaults. We first test whether receiving a nudge in the form of a good default increases the default bias. After showing that this is indeed the case, we explore whether experiencing good defaults affects the quality of subsequent decisions, either through an increased default bias, or for other reasons. Lastly, we present a behavioural model which captures the elements of our data that cannot be explained by existing models. All reported non-parametric tests are two-sided and performed at the individual level. Regressions use decisions as the unit of observation and cluster standard errors by individual.

#### 2.1. Default Bias

Our main hypothesis is that participants who have faced a good default in the past exhibit a stronger default bias. As Table 1 shows this was indeed the case for both the treatments with and without feedback. In the second half of the experiment participants in the good-default treatments were on average 11.6 and 18.2 percentage points more likely to pick the default than participants in the random-default group in the treatments without and with feedback respectively, even though they faced the exact same defaults. Furthermore, as Figure 2 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 good-default participants after round 25. This decline is especially rapid in the good-default treatment without feedback, from 70% in round 25 to 36% in round 26 (Wilcoxon p = 0.0014) and from 73% in rounds 21–25 to 44% in rounds 26–30 (Wilcoxon p < 0.0001). In the gooddefault treatment with feedback the decline takes somewhat longer. The drop between rounds 25 and 26 is only 14 percentage points, from 86% to 72% and not significant at conventional levels (Wilcoxon p = 0.1138). However it does drop to 47% in round 27 (Wilcoxon p = 0.0164 compared to round 25). When comparing rounds 21–25 to 26–30 the drop is very clear, from 82% to 56% (Wilcoxon p < 0.0001). Eventually the decline in default choices for good-default participants is significantly larger with feedback than without: comparing default choices in rounds 21–25 to those in rounds

<sup>&</sup>lt;sup>24</sup> Feedback random-default treatment participants did however start out picking a somewhat worse option than no-feedback random-default treatment participants and actually perform somewhat worse in the first half of the experiment en somewhat better in the second half, see Table 2.



Fig. 2. Average Proportion of Default Choices Per 5 Round Intervals by Treatment Note. Colour figure can be viewed at wileyonlinelibrary.com.

| Percentage of Default Choices in | n Each Half of the Experiment |     |
|----------------------------------|-------------------------------|-----|
| Random default                   | Good default                  | Man |

Table 1

|                          | Random default | Good default | Mann–Whitney p |
|--------------------------|----------------|--------------|----------------|
| First half, no-feedback  | 30.0%          | 65.3%        | < 0.0001       |
| First half, feedback     | 28.3%          | 73.0%        | < 0.0001       |
| Mann–Whitney p           | 0.7944         | 0.0205       |                |
| Second half, no-feedback | 33.9%          | 45.5%        | 0.0039         |
| Second half, feedback    | 26.3%          | 44.5%        | < 0.0001       |
| Mann–Whitney p           | 0.0401         | 0.8781       |                |

41-50 we find a drop from 73.2% to 49.3% without feedback and from 81.9% to 44.9% with feedback (Mann–Whitney p = 0.0225).

Clearly participants in neither of the good-default treatments simply continue to choose the default at the same rate as they did in the first half. This suggests that, although good-default participants choose the default relatively more often than random-default group participants, they do realise that the default option is not as good as it was in the first 25 rounds. Figure 3 supports this story at the individual level. The greater average reliance on the default is not caused by a small number of good-default treatment 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 good-default 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 random-default treatments participants chose the default



Fig. 3. Smoothed Density of Participants' Tendency to Choose the Default Option in the Second Half by Treatment Note. Colour figure can be viewed at wileyonlinelibrary.com.

more often than the 16.7% expected without a default bias (Wilcoxon p < 0.001) but good-default participants were significantly more likely to do so (Mann–Whitney p < 0.001 in both random-default treatments). This in itself does not provide evidence for a stronger default bias in the first half of the good-default treatment. Participants may also have chosen the default more often in the good-default 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 good-default treatments were more likely to choose the default treat than participants in the random-default treatments were to choose either the default or the best option (65% versus 57%, Mann–Whitney p = 0.0025 without feedback and 73% versus 53.8%, Mann–Whitney p < 0.0001 with feedback).

The linear probability regression reported in Table 2 explores the difference between good-default and random-default, and feedback and no-feeback treatments further. This regression controls for task and round fixed effects, several demographic variables, high school mathematics 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 proxy for effort.<sup>25</sup> The

 $<sup>^{25}</sup>$  Time spent in the first round strongly correlates with time spent in later rounds (Spearman correlation coefficient is 0.3172 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 have been influenced by the treatment (good-default *versus* random-default) but a Mann–Whitney test shows that this is not the case (p = 0.5804 for the non-feedback treatments and p = 0.1537 for the feedback treatments).

<sup>© 2016</sup> Royal Economic Society.

| Dependent variable: choice equals the default option, second half only | Coefficient    | p-value |
|--|----------------|---------|
| Good default   | 0.133 (0.044)  | 0.003   |
| Feedback   | -0.076(0.035)  | 0.030   |
| Good-default feedback interaction                                      | 0.056 (0.056)  | 0.319   |
| Time used in first round*  | -0.003(0.002)  | 0.032   |
| Good default $\times$ time used in the first round <sup>*</sup>        | -0.000(0.003)  | 0.910   |
| Male   | -0.034(0.030)  | 0.268   |
| Age  | 0.001 (0.000)  | 0.029   |
| Studies economics  | 0.027(0.032)   | 0.409   |
| Dutch  | 0.074(0.036)   | 0.040   |
| Maths grade  | -0.002(0.007)  | 0.814   |
| Maths level  | -0.022 (0.043) | 0.614   |

 Table 2

 Linear Probability Regression: Interaction Between Effort and the Treatment Effect

*Notes.* Standard errors used to calculate the p-values are clustered at the individual level. Task and round fixed effects are included in addition to the controls listed in the Table. We also included a dummy variable for 11 participants who did not state their age and/or gender, so their choices could be included in the regression. The coefficients for gender and age are relevant to the other participants who did state their age and/or gender. Standard errors are in parenthesis. \*The variable 'time used in the first round' was normalised to have mean zero before interacting it with the treatment dummy to facilitate interpretation of the coefficient on the uninteracted treatment dummy.

regression confirms that having been exposed to good defaults increases the likelihood of choosing the default. The presence of explicit feedback decreases the likelihood of second half default choices but only significantly so for random-default-treatment participants.<sup>26</sup> The effect of having been nudged with good suggestions may be a little larger if participants receive feedback but this difference is far from significant. The regression further shows that participants who take longer to make their decision choose the default less often but the interaction between time used in the first round and the good-default treatment indication shows that such participants do not exhibit a significantly weaker treatment effect.

Figure 2 also shows that the tendency to choose the default increases for gooddefault participants during the first 25 rounds, from an average of 57% (58%) default choices in the first five rounds, to 73% (82%) default choices in rounds 21–25 for the good-default treatment without (with) feedback. Linear regressions with task fixed effects confirm this positive trend both with and without feedback (p < 0.001). A linear regression using data from good-default treatments with and without feedback including an interaction between feedback and rounds reveals that the trend is somewhat stronger with feedback (p = 0.026). Without feedback there is no significant trend for random-default group participants in the first half. With feedback randomdefault group participants did become significantly less likely to choose the default in the first half of the experiment (p = 0.005). This negative trend is also present in the

 $<sup>^{26}</sup>$  A regression without the good-default-feedback interaction estimates a feedback coefficient of -0.0474444 with a p-value of 0.111, furthermore testing whether the sum of the feedback and good-default-feedback interaction coefficient from the regression of Table 2 is different from zero (using an F-test) gives a p-value of 0.6653).

<sup>© 2016</sup> Royal Economic Society.

feedback-good-default treatment in the second half of the experiment (p = 0.001). No other treatment exhibits a trend in the second half of the experiment. We therefore conclude that good-default participants' default bias was reinforced during the first half of the experiment both with and without the presence of feedback. Feedback participants who receive random defaults do slightly reduce their default bias, also if they have received good defaults before. However, a significant default bias remains and the increase in the default bias caused by having received good defaults does not disappear.

#### 2.2. Performance

Nudging participants with good defaults clearly reinforces their default bias but does this also hurt their performance in this experimental task? To answer this question, we use the value of the chosen option as a performance measure. Table 3 shows that in the second half of the experiment, when all participants faced the same random defaults, participants in the random-default treatments chose a better option than those in the good-default treatments, independent of whether they received feedback or not. In the treatments without feedback, the difference is 5.7 points, or just over 1/3 of a standard deviation; with feedback the difference is 8.8 points, or about 2/3 of a standard deviation. This means that whereas random-default participants managed to realise 28.8% of the difference in expected value between picking a random option and picking the best option without feedback and 30% with feedback, participants in the good-default treatment realised only 18.7% of this difference without feedback and 17.6% with feedback. By this measure good-default participants performed 29.9% worse without feedback and 41.2% worse with feedback. As Table 3 shows this difference is strongly significant with feedback and marginally significant without.<sup>27</sup>

The regression in Table 4, which pools the data from all four treatments, confirms the treatment effect on performance. This regression controls for round and task fixed

| Average Value of the Chosen Options |                          |                        |                |
|-------------------------------------|--------------------------|------------------------|----------------|
|                                     | Random-default treatment | Good-default treatment | Mann–Whitney p |
| First half, no feedback             | 173.0 (14.1)             | 191.5 (11.6)           | 0.0000         |
| First half, feedback                | 170.0 (9.7)              | 195.1 (13.6)           | 0.0000         |
| Mann–Whitney p                      | 0.0739                   | 0.0618                 |                |
| Second half, no feedback            | 174.8 (15.1)             | 169.1 (17.5)           | 0.081          |
| Second half, feedback               | 177.1 (9.0)              | 168.3 (12.4)           | 0.0008         |
| Mann–Whitney p                      | 0.9898                   | 0.3043                 |                |

|         | -     | Гa | ble | 3      |                |
|---------|-------|----|-----|--------|----------------|
| Average | Value | of | the | Chosen | <b>Options</b> |

Notes. Standard deviations reported in parenthesis.

 $^{27}$  Despite signs of hurting performance when the good default disappears in the second half, there is a strong direct effect of the good first half defaults on performance in the first half where good-default participants realise 50.2% and 55.3% of the difference in expected value between picking a random option and always picking the best option without and with feedback respectively. As Table 3 shows this is significantly better than the performance of the random-default treatments.

| Dependent variable: value of the chosen option                  |                |         |
|---|----------------|---------|
| in the second half of the experiment                            | Coefficient    | p-value |
| Good default  | -6.528(3.035)  | 0.033   |
| Feedback  | 3.407 (2.530)  | 0.180   |
| Good-default feedback interaction                               | -2.167(3.815)  | 0.571   |
| Time used in first round*                                       | 0.370(0.129)   | 0.005   |
| Good default $\times$ time used in the first round <sup>*</sup> | -0.022(0.182)  | 0.902   |
| Male  | 7.090 (2.047)  | 0.001   |
| Age   | 0.014 (0.018)  | 0.455   |
| Studies economics or business                                   | 0.018 (2.110)  | 0.993   |
| Dutch   | -3.187 (2.458) | 0.196   |
| Maths grade   | 0.065 (0.497)  | 0.896   |
| Maths level   | 2.164 (2.626)  | 0.411   |

|             | Tab       | le 4      |               |
|-------------|-----------|-----------|---------------|
| Regression: | Treatment | Effect or | n Performance |

*Notes.* Standard errors used to calculate the p-values are clustered at the individual level. Task and round fixed effects are included in addition to the controls listed in the Table. We also included a dummy variable for 11 participants who did not state their age and/or gender, so their choices could be included in the regression. The coefficients for gender and age are relevant to the other participants who did state their age and/or gender. Standard errors are in parenthesis. \*The variable 'time used in the first round' was normalised to have mean zero before interacting it with the treatment dummy to facilitate interpretation of the coefficient on the uninteracted treatment dummy.

effects, several demographic variables, high school mathematics level and grade as a proxy of skill, and time used in the first round as a proxy of effort. Controlling for these variables in the regression, the treatment effect of the default becomes significant at a 5% level in the no-feedback treatment. The effect of the default may be somewhat larger in the feedback treatment but, as with the treatment effect on default choice, the difference is far from significant.<sup>28</sup>

The regression in Table 4 shows that our proxy for effort, time used in the first round, predicts participants' performance in the second half. At the same time, it might be the case that experiencing a good default, or the higher expected earnings it brings, affects the time and/or effort required to make decisions or participants' willingness to spend time and effort on the task and, perhaps, through that channel, performance and/or their inclination to choose the default.<sup>29</sup> This effect could run in both directions. The fact that participants in the good-default treatment got used to helpful defaults and high (expected) earnings could have made them unwilling to expand effort later on, or prevented them from learning how to make a good decision themselves.<sup>30</sup> However, Abeler *et al.* (2011) show that the effect could also run in the opposite direction. In their experiment participants were either paid a predetermined amount, or their earnings from a task. They found that participants

<sup>&</sup>lt;sup>28</sup> As can be seen in Table 4, we find an unhypothesised but significant gender effect. However, including an interaction term between gender and treatment in the regression of Table 4 shows that the treatment effect is not significantly different for men and women (p = 0.578).

<sup>&</sup>lt;sup>29</sup> Caplin and Martin (2013) find that participants who know they face better defaults indeed spend less time to make a decision.

<sup>&</sup>lt;sup>30</sup> Carlin's *et al.* (2013) model shows this reduced incentive to learn can also hurt others when acquired knowledge can spread through the population.

<sup>© 2016</sup> Royal Economic Society.

for whom the predetermined amount was higher spent more effort on the task, presumably to prevent the disappointment from getting a lower amount when the task is paid out. Because only one round is paid out in our experiment these results suggests that higher earnings if a round from the first half is paid out should lead to increased effort in the second half. As good-default participants chose better options in the first half this should increase their effort in the second half. We therefore study the time participants took to make decisions during the second half of the experiment.

Without feedback, participants in the good-default treatment spend on average slightly less time in the second half of the experiment than participants in the random-default treatment: 18.3 *versus* 19.1 seconds. Without feedback time spent was similar with good or random defaults: 14.5 seconds *versus* 14.7 seconds. Both of these differences are far from significant (Mann–Whitney test p = 0.81 for the no-feedback treatments and p = 0.77 for the feedback treatments). Given that we find no treatment effect on our proxy of effort we do not think a difference in effort drives our results. Furthermore, the findings of Abeler *et al.* suggest that more effort by good-default participants is also a likely scenario, which would mean that if there is a treatment effect on effort our results would be a conservative estimate.

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 on the likelihood of default choices would be significantly smaller for this group. However, Tables 2 and 4 show that this is not the case. The interaction effect of the treatment with our measure of effort inclination, time taken in the first round, has no significant effect on either the likelihood to choose the default or on the performance in the second half of the experiment.

#### 2.3. Behavioural Model

If we roughly summarise the choice behaviour of the participants in our experiment, we observe the following main patterns:

- (i) Less than optimal but better than random choices.
- (ii) A clear and persistent tendency to choose the default over other options even in the random-default treatment where the default is random.
- (*iii*) This default bias reduces over the course of the experiment when randomdefault group participants receive feedback but not when there is no feedback.
- (iv) Participants in the good-default treatment come to choose the default more and more often during the first 25 rounds, even when they do not receive explicit feedback regarding the value of the chosen option, which they could use to learn that the default is a good choice.
- (v) This increase in the default bias for good-default participants is stronger with feedback.
- (*vi*) There is a sharp drop in default choices after round 25 for the good-default treatments with and without feedback.

- (vii) This drop is less fast, but eventually larger, with than without feedback.
- (*viii*) Despite this drop good-default participants show a substantially higher default bias in the second half than the random-default group in both the no-feedback and feedback treatments, which is the main result of the experiment.

In this Section, we present a behavioural model that can capture all these patterns in the data: a noisy response model combined with an effect we call 'choice reinforcement' and the possibility for reinforcement learning (Erev and Roth, 1998). Choice reinforcement here means that the attraction of an option increases after one has chosen it and decreases when one has not chosen it.<sup>31</sup> Besides such choice reinforcement participants' choices can also be affected by a more explicit learning process about the relative value of default and non-default options. We model this by reinforcement learning. The main idea of reinforcement learning is that participants' expectations of (only) the values of chosen options are updated, based on the experienced pay-off. In the feedback treatment, participants receive explicit feedback about the value of their chosen option, so reinforcement learning is more plausible in those treatments. For this reason the model only incorporates reinforcement learning in those treatments. In contrast to reinforcement learning, choice reinforcement does not need any specific performance feedback so the model allows for this effect to occur in all treatments.

We want to mention here that although the model only incorporates reinforcement learning in the treatments with feedback, we do not mean to imply that there is potentially no learning happening in the no feedback treatment. The acquaintance with the task could help participants recognise over time how high-valued options look like. Even when participants are not able to check their chosen score via explicit feedback, they might still be able to figure this out by their own estimation. This 'learning without explicit feedback' could possibly resemble the structure of reinforcement learning, as the option a participant chose might very well be the option the participant studied most closely. However, the effect is likely to be (far) weaker without explicit feedback and later in this Section, we present some results which indeed suggest that choice dynamics following a reinforcement-learning pattern are hardly present in the no-feedback treatments. Nevertheless, it is important to consider this when interpreting the parameters of the model because the choice reinforcement part of the model may capture the reinforcement learning that does occur in the no-feedback treatments and as a consequence the model may underestimate the importance of reinforcement learning.

The model works as follows: Each round for each participant the probability that an option i is chosen is given by:

$$P(\text{option } i \text{ is chosen}) = \frac{\alpha_i \times e^{\lambda \times \left[\pi_i + \mathbf{1}_{\text{option } i \text{ is the default}} \times (\pi_d^R - \pi_{\neg d}^R)\right]}}{\sum_{j=1}^6 \alpha_j \times e^{\lambda \times \left[\pi_j + \mathbf{1}_{\text{option } j \text{ is the default}} \times (\pi_d^R - \pi_{\neg d}^R)\right]}}.$$

Here  $\lambda$  is the rationality parameter and  $\pi_i$  is the value of option *i*. Choice reinforcement is incorporated into the model via the  $\alpha$ -parameters.  $\alpha_i$  is the attraction

<sup>&</sup>lt;sup>31</sup> One way to see the reinforcement is that people remember and value their own past evaluation effort that has led to previous choices.

<sup>© 2016</sup> Royal Economic Society.

that an option has for a participant in a given round. In every round and for every participant  $\sum_{i=1}^{6} \alpha_i = 1$ , so that  $\alpha_i$  can be seen as the weight, or attention that an option gets in a participants' choice deliberation. We distinguish only between default and non-default choices. So, for example, if option *i* is the default in one round and  $\alpha_i$  has a value of 1/3, all other  $\alpha_i$ s will have a value of [1-(1/3)]/5 = 2/15. The choice reinforcement of the default in this model is based on participants' past choices in the following way. After a default choice the default attraction becomes:  $\alpha_{d,t+1} = \alpha_{d,t}^{\delta}$ where  $\delta \leq 1$  is the reinforcement parameter. The attraction parameters for the non-default choices in round t + 1 are set at  $\alpha_{j \neq d, t+1} = (1 - \alpha_{d, t+1})/5$ . If a non-default choice *i* is made in round *t*,  $\alpha_{d, t+1} = \alpha_{d, t} \times (1 - \alpha_{j \neq d, t}^{\delta})/(1 - \alpha_{j \neq d, t})$  and (again)  $\alpha_{j \neq d, t+1} = (1 - \alpha_{d, t+1})/5$ . So, in case of a non-default choice, first the attraction of a not chosen default option is decreased as if the attraction of one of the non-default choices was choice-reinforced, which would have resulted in multiplying the attraction of each not-chosen option by  $(1 - \alpha_{j \neq d,l}^{\delta})/(1 - \alpha_{j \neq d,l})$ . In this case only the attraction of the default option is multiplied by  $(1 - \alpha_{j \neq d,l}^{\delta})/(-\alpha_{j \neq d,l})$ . Then subsequently, because we only distinguish between the default option and the non-default options, all nondefault attractions are set equal to  $(1 - \alpha_{d,t+1})/5$  such that again  $\sum_{i=1}^{6} \alpha_i = 1$ . The consequence of this updating mechanism is that after one default choice, the default reinforcement is usually stronger than the deinforcement of the default attraction after one non-default choice. A picture of how the value of  $\alpha_{d,t}$  changes from one round to the next is given in Figure 4.<sup>32</sup> To allow for an inherent default bias the initial attraction of the default,  $\alpha_{d,1}$ , is the third free parameter.

As explained above the model incorporates reinforcement learning in addition to the choice reinforcement, but only in the feedback treatments where participants



Fig. 4. Attraction Parameter Dynamics Note. Colour figure can be viewed at wileyonlinelibrary.com.

1197

 $<sup>^{32}</sup>$  The Figure is based on a value of  $\delta$  equal to 0.9537. This is a value we estimated based on the participant's choice data. See below for more on our estimation procedure.

<sup>© 2016</sup> Royal Economic Society.

[MAY

receive explicit feedback on the value of their chosen option. We use a simple implementation of reinforcement learning, so we can have a single model that would predict participants' choices both for the treatments with and without feedback. The main intuition of this implementation is that via the pay-off feedback  $\pi_d^R$ , the reinforcement learning estimated average value of choosing a default option, and  $\pi_{-d}^R$ , the reinforcement learning estimated average value of choosing a non-default option, are updated. In the first round  $\pi_d^R = \pi_{-d}^R = \pi_1$  so there is no inherent 'conscious' belief that the default option has a higher value.<sup>33</sup> For the treatments with feedback, each round, the values are updated as follows. After a default choice  $\pi_{d,t+1}^R = \gamma \times \pi_{chosen option,t} + (1 - \gamma) \times \pi_{-d,t}^R$  and  $\pi_{d,t+1}^R = \pi_{-d,t}^R$ . The parameter  $\gamma$  reflects the intensity of the reinforcement learning. As stated in the no-feedback treatments,  $\pi_d^R$  and  $\pi_{-d}^R = \pi_d^R - \pi_{-d}^R$  will always be equal to 0 for the no feedback treatments. We end up with one model that, for given parameters  $\lambda$ ,  $\alpha_{d,1}$ ,  $\delta$ ,  $\gamma$  and  $\pi_1$ , makes choice predictions.

Table 5 shows the parameter estimates and log likelihoods of a maximum likelihood estimation of the five parameters on the data of our experiment. We estimate one model on the combined data of all four treatments. In addition to our full model, we present three other models with different restrictions on the estimated parameters. First, because the model we present here is based on the well-known logit choice model developed by McFadden (1973)<sup>34</sup> we compare our model to that model. Second, to separate the effect of introducing a default bias *per se* and allowing the default bias to be reinforced we also present a logit noisy choice model without choice reinforcement. Third, to look at the impact of reinforcement learning we estimate a model without the possibility of reinforcement learning. Again, the model's reinforcement learning only has an influence for the part of the observations coming from the treatments with feedback.

|                      | Maximum Likelihood Estimation |                               |                                 |                                      |
|----------------------|-------------------------------|-------------------------------|---------------------------------|--------------------------------------|
|                      | Full model                    | Standard<br>noisy<br>response | Without choice<br>reinforcement | Without<br>reinforcement<br>learning |
| λ                    | 1.4714 (0.0275)               | 1.6358 (0.0260)               | 1.4370 (0.0270)                 | 1.5773 (0.0287)                      |
| $\hat{\alpha}_{d,1}$ | 0.2958 (0.0086)               | 0.167*                        | 0.4139(0.0085)                  | 0.2891 (0.0079)                      |
| $\hat{\delta}$       | 0.9537 (0.0025)               | 1*                            | 1*                              | 0.9455 (0.0023)                      |
| Ŷ                    | $0.3044 \ (0.0306)$           | 0*                            | $0.1241 \ (0.0135)$             | 0*                                   |
| $\hat{\pi}_1$        | 98.37 (0.0306)                | _*                            | 28.67 (9.71)                    | _*                                   |
| $\log L$             | -11,616                       | -13,367                       | -11,800                         | -11,666                              |

|         | Table 5    |            |
|---------|------------|------------|
| Maximum | Likelihood | Estimation |

*Notes.* Standard errors are in parenthesis. Parameter values denoted with \* are not estimated, but fixed due to model restrictions. For the estimation we divided all option point totals by 100.

<sup>33</sup> Although such initial default bias could exist, this would be the same in the feedback and no-feedback treatments, so we capture it with the  $\alpha_{d,1}$  parameter.

<sup>34</sup> McKelvey and Palfrey (1995), among others, use this model.

© 2016 Royal Economic Society.

1199

All four models estimate a fairly similar rationality parameter,  $\lambda$ . Both the constant and the reinforced default bias models reflect the substantial (initial) default bias observed in both treatments with an estimate of  $\alpha_{d,1}$  that is well above 1/6, but substantially higher if we do not allow for choice reinforcement. The choice reinforcement parameter  $\delta$  is estimated at a level below 1, which indicates positive choice reinforcement. The estimation also gives a positive value for the reinforcement learning intensity parameter  $\gamma$ . Looking at the log-likelihood values of the different models, we see that setting the choice reinforcement parameter  $\delta$  equal to one has a substantial impact on the likelihood score. This impact is larger than that of setting the reinforcement learning parameter  $\gamma$  equal to zero. The estimate of  $\pi_1$  is lower than the average value of a (random) option (156). One explanation for this is that a low value of  $\pi_1$  produces an effect that is similar to what the choice reinforcement is doing. With a low value of  $\pi_1$  participants, in the feedback-good-default treatment, who do not often choose the non-default option and, hence, not often update  $\pi_{\neg d}^R$ , will have a relatively low reinforced expectation of  $(\pi_{\neg d}^R)$  by round 25, compared to participants in the feedback-random-default treatment. This will lead to the model producing our main effect of more default choices for good-default participants, on top of the effect of choice reinforcement. We indeed see that in the model without choice reinforcement (when  $\delta \equiv 1$ ), the estimated value of  $\pi_1$  is even lower, as a low value of  $\pi_1$  has now become the only way for the model to produce the second half default choices effect. We estimate the value of  $\pi_1$  in this model together with the choice reinforcement part that brings along the term  $\alpha_{d,1}$ , the initial default bias. This allows us to estimate one model for both the feedback and no-feedback treatments but makes a direct interpretation of the estimated value of  $\pi_1$  difficult. The reinforcement learning model needs a starting parameter  $\pi_1$  and our main assumption here is that both the initial reinforcement learning valuation of default and no-default are equal and otherwise we allow this to be a free parameter of the model. Importantly, per construction of the model, this reinforcement learning feature of the model based on explicit pay-off feedback for the chosen option, cannot explain our main treatment difference in the second half default choice in the case of no feedback. Our proposed mechanism of choice reinforcement is crucial for the model to produce our main treatment effect both in the feedback and no-feedback treatments. All in all, our model estimations suggest that choice reinforcement can explain a good part of the choice pattern for both the no-feedback and feedback treatments. However the addition of a simple reinforcement learning mechanism does help to capture the differences between the data patterns of the no-feedback and feedback treatments.

The model assumes that reinforcement learning is absent without feedback. If we do allow for reinforcement learning (also) in the no-feedback treatments, by adding a separate  $\gamma_{no-feedback}$  parameter to the model, the maximum likelihood optimisation estimates this parameter to be equal to 0.0072 (with an estimated standard deviation of 0.0032) which is very small compared to the  $\gamma$  estimated for the feedback treatments.<sup>35</sup> This suggests that there is indeed little reinforcement learning going on in the

 $<sup>^{35}</sup>$  Including this parameter to the model also adds very little in terms of explanatory power, yielding a log likelihood of -11,613 compared to -11,616 in the original model.

<sup>© 2016</sup> Royal Economic Society.

no-feedback treatments and validates not including it in the model. Nevertheless, strictly speaking, the difference between the no-feedback and feedback treatments and the effect of reinforcement learning in the model and simulations could best be interpreted as the result of the additional reinforcement learning that occurs as a result of participants receiving explicit pay-off feedback about the chosen option.

Figure 5 shows the percentage of default choices in five round intervals predicted by the model in 1,000 simulations using the estimated parameters compared to the actual choice data. The top graph shows the simulation and the actual data for the no-feedback treatments and the bottom graph shows the simulation and data for the feedback treatments. Note that the simulations for treatments with and without feedback treatments are based on the same parameters. The only difference in the simulations between feedback and no-feedback is that the estimated parameters  $\hat{\gamma}$ 



Fig. 5. Simulations and Data of Default Choices. (a, b) Default Choices Over Different Periods Note. Colour figure can be viewed at wileyonlinelibrary.com.

<sup>© 2016</sup> Royal Economic Society.

and  $\hat{\pi}_1$  have no influence in the simulation of choices without feedback. The simulations capture the qualitative features of the choice behaviour in all four treatments. Most importantly they capture the steady rise of default choices in the first half for good-default participants and the main experimental result that the share of default choices in the last 25 rounds is higher in the good-default treatments than in the random-default treatments. As can be seen, the choice reinforcement model increases the attraction of the default option in the first 25 rounds in the good-default treatment, because it is chosen relatively often as it is a good option. This in turn increases the relative chance that a default option is chosen, also in the second half.<sup>36</sup>

As mentioned, allowing the model to use reinforcement learning in the treatment with feedback, but not in those without, allows the model to capture the qualitative differences between these treatments. First, adding reinforcement learning reduces the share of default choices in the simulation of the feedback random-default treatment. The choice reinforcement element in the model, combined with the relatively high initial default bias  $\alpha_{d,1}$  causes the default option to be chosen 'too often' in the random-default treatments. Given this, the pay-off from a chosen default option will typically be lower than the pay-off from a chosen non-default choice in any round. In a treatment with feedback, the choice model with reinforcement learning will consequently update  $\pi_d^R$  to be lower than  $\pi_{\neg d}^R$ , lowering the likelihood of a default choice. Second, adding feedback and reinforcement learning will cause someone who chooses along the lines of the behavioural model to choose the default more often in the first half of the good-default treatment. As the default option is always the highest,  $\pi_d^R$  will quickly update to a high level, which increases the default choice likelihood, on top of the choice reinforcement effect. Third, in the second half of the experiment, when the value of the default option is random, the high value of  $\pi_d^R$  will initially lead to more default choices with than without feedback explaining the slower drop in default choices. However, the high default choice rate for good-default decision makers with feedback will cause  $\pi_d^R$  to decrease, leading to an additional decrease in the second half default choice likelihood for good-default participants with feedback. In this way, the model can reproduce the slower and in the end larger second half drop in default choices with feedback.

### 3. Extension to Bad Defaults

Our analysis so far has focused on the effect of good defaults on subsequent decisions. Good defaults are the most relevant case because, good defaults are a common 'nudge' and, good defaults are the most plausible default to use when someone, for example,

<sup>&</sup>lt;sup>36</sup> A crucial feature of the model which allows for the drop in the share of default choices for the gooddefault group after round 25 is that a given level of skewness in the  $\alpha$ s only has a real impact on choices when the differences in the actual pay-offs of the options are not too large and  $\lambda$  is not too large. So here attraction differences play a role when the decision maker is 'confused' enough (of course if the  $\alpha$ s are skewed enough to one option this can still determine the choice probabilities for a given large value of  $\lambda$  and  $\pi_i$ s). This feature allows us to replicate the pattern in the data that when a default in some round has a very low pay-off, very few will choose the option, even when someone has very much reinforced the attraction level  $\alpha$  for the default.

<sup>© 2016</sup> Royal Economic Society.

a marketer, wants to induce people to follow their defaults. However, the effect of earlier defaults on subsequent decisions and the theory we used to explain this effect should in principle apply equally to other defaults. A conspicuous case is one with initial bad defaults. Can bad defaults perhaps 'debias' people, reducing their default bias in subsequent decisions?<sup>37</sup> As Figure 6 shows, our model with the estimated parameters (based on the other treatments) predicts that the answer is yes. To explore this possibility we conducted an additional 'bad-defaults' treatment where in the first half of the experiment the default option was always the worst possible choice. We chose to conduct this treatment without feedback so the results of this treatment are best compared with the results from the no-feedback treatments.

Forty-nine participants took part in this treatment, of which 25 were male and 37 were economics or business students. As with the main treatments, we again look at decisions in the second half where participants in the bad-default and random-default treatments face the same, random defaults. Figure 6 and Tables 6 and 7 show we did not find the predicted debiasing effect, neither on default choices, nor on performance. It is not that participants did not notice that defaults were bad in the first half, they only choose the default option 10.9% of the time much less than in the random-default treatment (p < 0.0001). However, we find no effect on their subsequent decisions. We also do not find evidence for a reduction in the default bias during the first half of the experiment (a linear probability regressions with task fixed effects reveals no significant trend in default choices during the first half



Fig. 6. Simulations and Data on Default Choice No-feedback Treatments and Bad-default Treatment Note. Colour figure can be viewed at wileyonlinelibrary.com.

 $^{37}$  Random defaults might already have had this effect but we did not observe this without feedback and only weakly with feedback.

<sup>© 2016</sup> Royal Economic Society.

| 12 | 03 |
|----|----|
|----|----|

|                                    | First half,<br>no feedback | Second half<br>no feedback |  |
|------------------------------------|----------------------------|----------------------------|--|
| Random-default treatment           | 30.0%                      | 33.9%                      |  |
| Good-default treatment             | 65.3%                      | 45.5%                      |  |
| Bad-default treatment              | 10.9%                      | 31.35%                     |  |
| Mann–Whitney p                     |                            |                            |  |
| Random default versus good default | 0.0000                     | 0.0039                     |  |
| Random default versus bad default  | 0.0000                     | 0.2323                     |  |
| Good default versus bad default    | 0.0000                     | 0.0003                     |  |
|                                    |                            |                            |  |

 Table 6

 Percentage Default Choices No-feedback Treatments

|                                    | First half, no feedback | Second half, no feedback |
|------------------------------------|-------------------------|--------------------------|
| Random-default treatment           | 173.0 (14.1)            | 174.8 (15.1)             |
| Good-default treatment             | 191.5 (11.6)            | 169.1 (17.5)             |
| Bad-default treatment              | 170.6 (17.7)            | 175.7 (12.8)             |
| Mann–Whitney p                     |                         |                          |
| Random default versus good default | 0.0000                  | 0.0811                   |
| Random default versus bad default  | 0.9540                  | 0.9051                   |
| Good default versus bad default    | 0.0000                  | 0.0755                   |

| Table 7 |       |        |        |        |       |        |           |   |
|---------|-------|--------|--------|--------|-------|--------|-----------|---|
| Average | Value | of the | Chosen | Option | No-fe | edback | Treatment | s |

Note. Standard deviations reported in parenthesis.

p = 0.339). Thus, we find evidence that it is possible to increase people's default bias by presenting them with good defaults, but somewhat against our expectations not that it is possible to reduce their default bias with bad-defaults.<sup>38</sup>

## 4. Conclusion

The default bias is one of the most well known and consistently observed behavioural biases. It is also a bias that is often used to influence behaviour, either to help the decision maker (e.g. with libertarian paternalism) or to benefit the person who sets the default (e.g. in marketing). In this study, we investigate whether the default bias itself is affected by the nature of earlier defaults. We find that providing participants with good defaults reinforces their default bias. Compared to participants who always faced random defaults, participants who faced a good default in the first half of the experiment choose the default option significantly more often in the second half were all participants face random defaults. As a result, good-default participants perform

<sup>&</sup>lt;sup>38</sup> Our behavioural model would predict a debiasing effect of first half bad-defaults on second half default choices. However, the model would predict this effect to be smaller than the positive effect of good first half defaults on the choice frequency of second half defaults.

<sup>© 2016</sup> Royal Economic Society.

somewhat worse in the second half.<sup>39</sup> We find this effect both when participants received feedback on the quality of their decisions and when they do not receive any feedback. However, with-feedback participants show a stronger increase in the tendency to choose the default option when facing good defaults and a slower, but eventually larger, decrease in default choices when defaults became random.

To explain our results, we develop a model, which combines choice reinforcement, a tendency to make similar choices as in the past, and reinforcement learning, updating beliefs about the quality of default and non-default options based on past experience, which is more plausible when receiving feedback. The model organises our results quite well, including the observed differences between feedback and no-feedback treatments.

Existing research has studied debiasing decisions (see Larrick, 2004 for a review) but the results presented in this article suggest that we can also, intentionally or not, strengthen existing biases.<sup>40</sup> However, our model does suggest that debiasing people should also be possible: if people face bad defaults their default bias is predicted to decline. Interestingly, results from an additional treatment where participants face bad defaults before facing random defaults, shows no evidence that this works. In the second half of the experiment where all participants face random defaults, these bad-default participants do not choose the default any less often than people who have always faced random defaults. The result suggests that the presented model does not fully capture the cognitive processes which drive our results. A closer look at the cognitive processes and the development of beliefs about the default provide interesting avenues for future research. We believe the results of the bad-default treatment are also interesting in themselves, especially from the perspective of marketing, where consumers are more likely to be confronted with bad defaults.

We conclude from our results that the default bias is not a fixed tendency but something that can be manipulated and, apparently can be increased more easily than that it can be decreased. While this is a single experiment and further studies should assess the robustness of this phenomenon, we do believe this conclusion provides a note of caution to policy makers 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 the way people take decisions, even decisions you do not intend to affect.

Nevertheless, we certainly do not argue that policy makers should never engage in libertarian paternalism. In fact, our experiment shows that the good default we provide helps participants to make better decisions overall.<sup>41</sup> We do believe that our results, especially when taken together with the findings of Brown and Krishna (2004), Campbell

<sup>&</sup>lt;sup>39</sup> This effect on performance is not as strong as the effect on the default bias, which is not surprising as our experiment was not designed to maximise this effects. Other tasks may well reveal a larger, smaller, or even positive effect on performance depending on, for example, the difficulty of the task and the exact nature of the later default.

<sup>&</sup>lt;sup>40</sup> Although he does not consider the effect of a nudge on later decisions Larrick (2004) strikingly calls the libertarian paternalism method of debiasing rebiasing.

<sup>&</sup>lt;sup>41</sup> For a more philosophical discussion on the merits of and objections to libertarian paternalism, which we do not enter in this article, see for example, Sunstein and Thaler (2003), Mitchell (2004) and Sugden (2008).

<sup>© 2016</sup> Royal Economic Society.

(2007), and Caplin and Martin (2013), lend support to the idea of announcing how, to which end and by whom the choice architecture in a particular situation is designed.<sup>42,43</sup>

Another important avenue for future research would be to see if reinforcement of default choices in one choice context could spill over and increase the likelihood of people following defaults or recommended choices in other choice environments. This can be especially relevant in the case of a government actively setting defaults in many choice settings. Our experiment does not address this issue but focuses on establishing the default bias reinforcement result for repeated choice contexts. However, research in the field of behavioural and learning spillovers suggests that such spillovers to other choice environments are not unlikely, because people do, under certain circumstances, take beliefs or even actions learned or reinforced in one strategic context and apply them in another strategic choice environment (Huck *et al.*, 2011; Grimm and Mengel, 2012; Bednar *et al.*, 2012). However, Cooper and Kagel (2008) show an example where people with experience in one environment actually change their strategy quicker towards equilibrium in a subsequent different game.

In light of our results and these related findings, we believe that recognising the effects of a designed choice architecture on subsequent decisions should be considered when designing public policies and might also provide grounds to regulate certain marketing practices. Future (field) experiments on libertarian paternalistic policies can shed further light on the effect of having been nudged with good suggestions by also examining decisions the researchers do not intend to influence, but which could be affected by the choice architecture of earlier decisions.

Norwegian School of Economics, Bergen Maastricht University

Accepted: 18 July 2016

Additional Supporting Information may be found in the online version of this article:

Appendix A. Instructions. Data S1.

#### References

- Abeler, J., Falk, A., Goette, L. and Huffman, D. (2011). 'Reference points and effort provision', American Economic Review, vol. 101(2), pp. 470–92.
- Anderson, C. (2003). 'The psychology of doing nothing: forms of decision avoidance result from reason and emotion', *Psychological Bulletin*, vol. 129(1), pp. 139–67.
- Anufriev, M. and Hommes, C. (2012). 'Evolutionary selection of individual expectations and aggregate outcomes in asset pricing experiments', *American Economic Journal: Microeconomics*, vol. 4(4), pp. 35– 64.

Bednar, J., Chen, Y., Liu, T.X. and Page, S. (2012). 'Behavioral spillovers and cognitive load in multiple games: an experimental study', *Games and Economic Behavior*, vol. 74(1), pp. 12–31.

<sup>42</sup> Although Caplin and Martin show that providing such information can also backfire as they show people can rely on the default too much if it is a good, but not perfect, default.

<sup>43</sup> In light of this the transparency of the British 'Behavioural Insights Team' (http://www.behaviouralin sights.co.uk/) about their interventions and the media attention they receive is a laudable example, although they do not directly inform people subjected to certain interventions.

© 2016 Royal Economic Society.

- Brown, C.L. and Krishna, A. (2004). 'The skeptical shopper: a metacognitive account for the effects of default options on choice', *Journal of Consumer Research*, vol. 31(3), pp. 529–39.
- Cabinet Office Behavioural Insights Team. (2011). 'Behaviour change and energy use', 406537/0711, London.
- 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*, vol. 151(3), pp. 1211–54.
- Campbell, M.C. (2007). "Says Who?!' How the source of price information and affect influence perceived price (un)fairness', *Journal of Marketing Research*, vol. 44(2), pp. 261–71.
- Caplin, A. and Martin, D. (2013). 'Defaults and attention: the drop out effect', Working Paper, New York University.
- Carlin, B.I., Gervais, S. and Manso, G. (2013). 'Libertarian paternalism, information production, and financial decision making', *Review of Financial Studies*, vol. 26(9), pp. 2204–28.
- Choi, J., Laibson, D., Madrian, B. and Metrick, A. (2003). 'Optimal defaults', American Economic Review, vol. 93(2), pp. 180–5.
- Cooper, D.J. and Kagel, J.H. (2008). 'Learning and transfer in signaling games', *Economic Theory*, vol. 34(3), pp. 415–39.
- Erev, I., and Roth, A.E. (1998). 'Predicting how people play games: reinforcement learning in experimental games with unique, mixed strategy equilibria', *American Economic Review*, vol. 88(4), pp. 848–81.
- Frederick, S., Loewenstein, G. and O'Donoghue, T. (2002). 'Time discounting and time preference: a critical review', *Journal of Economic Literature*, vol. 40(2), pp. 351–401.
- Goldstein, D.G., Johnson, E.J., Herrmann, A. and Heitmann, M. (2008). 'Nudge your customers toward better choices', *Harvard Business Review*, vol. 86(12), pp. 99–105.
- Grimm, V. and Mengel, F. (2012). 'An experiment on learning in a multiple games environment', Journal of Economic Theory, vol. 147(6), pp. 2220–59.
- Hoch, S.J. and Loewenstein, G.F. (1991). 'Time-inconsistent preferences and consumer self-control', *Journal of Consumer Research*, vol. 17(4), pp. 492–507.
- Huck, S., Jehiel, Ph. and Rutter, T. (2011). 'Feedback spillover and analogy-based expectations: a multi-game experiment', *Games and Economic Behavior*, vol, 71(2), pp. 351–65.
- Johnson, E. and Goldstein, D. (2003). 'Do defaults save lives?', Science, vol. 302(5649), pp. 1338-9.
- Kahneman, D., Knetsch, J. and Thaler, R. (1991) 'Anomalies: the endowment effect, loss aversion and status quo bias', *Journal of Economic Perspectives*, vol. 5(1), pp. 193–206.
- Kahneman, D. and Miller, D. (1986) 'Norm theory: comparing reality to its alternatives', *Psychological Review*, vol. 93(2), pp. 136–53.
- Kahneman, D. and Tversky, A. (2000). Choices, Values and Frames, New York: Cambridge University Press.
- Kalayci, K. and Potters, J. (2011). 'Buyer confusion and market prices', *International Journal of Industrial Organization*, vol. 29(1), pp. 14–22.
- Larrick, R. (2004). 'Debiasing', in (D.J. Koehler and N. Harvey, eds.), Blackwell Handbook of Judgement and Decision Making, pp. 316–37, Hoboken, NJ: Blackwell Publishing Ltd.
- Levav, J., Heitmann, M., Herrmann, A. and Iyengar, S.S. (2010) 'Order in product customization decisions: evidence from field experiments', *Journal of Political Economy*, vol. 118(2), pp. 274–99.
- Madrian, B. and Shea, D. (2001). 'The power of suggestion: inertia in 401(k) participation and savings behavior', *Quarterly Journal of Economics*, vol. 116(4), pp. 1149–87.
- McFadden, D. (1973). 'Conditional logit analysis of qualitative choice behavior', in (P. Zarembka, ed.), *Frontiers of Econometrics*, pp. 105–42, New York: Academic Press.
- McKelvey, R. and Palfrey, T. (1995). 'Quantal response equilibria for normal form games', Games and Economic Behavior, vol. 10(1), pp. 6–38.
- Mitchell, G. (2004). 'Libertarian paternalism is an oxymoron', Northwestern University Law Review, vol. 99(3), pp. 1245–77.
- Samuelson, W. and Zeckhauser, R. (1988). 'Status quo bias in decision making', Journal of Risk and Uncertainty, vol. 1(1), pp. 7–59.
- Sugden, R. (2008). 'Why incoherent preferences do not justify paternalism', Constitutional Political Economy, vol. 19(3), pp. 226–48.
- Sunstein, C. and Thaler, R. (2003). 'Libertarian paternalism is not an oxymoron', University of Chicago Law Review, vol. 70(4), pp. 1159–202.
- Thaler, R. and Sunstein, C. (2003). 'Libertarian paternalism', American Economic Review, vol. 93(2), pp. 175-9.
- Thaler, R. and Sunstein, C. (2008). Nudge: Improving Decisions about Health, Wealth, and Happiness, New Haven, CT: Yale University Press.
- Tversky, A. and Kahneman, D. (1974). 'Judgement under uncertainty: heuristics and biases', *Science*, vol. 185 (4157), pp. 1124–31.
- Tversky, A. and Kahneman, D. (1981). 'The framing of decisions and the psychology of choice', *Science*, vol. 211(4481), pp. 453–8.
- World Bank. (2015). Mind, Society, and Behavior, Washington DC: World Bank.

© 2016 Royal Economic Society.