

The Double Dividend of Nudges

Steffen Altmann, Andreas Grunewald,
Jonas Radbruch

Discussion Paper #2024/03

DFG Research Unit: “Consumer Preferences, Consumer
Mistakes, and Firms’ Response” (FOR 5392)

THE DOUBLE DIVIDEND OF NUDGES*

Steffen Altmann[†], Andreas Grunewald[‡], Jonas Radbruch[§]

April 12, 2024

Abstract

Nudge-based policies are an important instrument for many policymakers. Based on a laboratory experiment featuring a dual-task paradigm, we examine the effects of two common types of nudge interventions—the simplification of complex decisions and the implementation of high-quality defaults. We find that these interventions do not only improve choices in the targeted domain, but also yield substantial positive indirect effects on non-targeted domains. The latter emerge through a reallocation of cognitive resources. Furthermore, the relative importance of direct and indirect effects varies systematically across the population. Evaluations that focus only on the targeted domain therefore significantly underestimate the interventions’ overall effectiveness and provide a biased assessment of their distributional consequences.

Keywords: nudges, default options, administrative burden, limited attention, limited cognitive resources, behavioral economics, laboratory experiment

*We thank Vincent Selz for excellent research assistance. Financial support by the DFG, grant FOR 5392 project number 462020252 is gratefully acknowledged. We also acknowledge support by the DFG through CRC TRR 190 (project number 280092119).

[†]University of Duisburg-Essen, University of Copenhagen and IZA, Berliner Platz 6-8, 45127 Essen, Germany. Email: steffen.altmann@uni-due.de

[‡]Frankfurt School of Finance and Management, Department of Economics, Adickesallee 32-34, 60322 Frankfurt, Germany. Email: a.grunewald@fs.de

[§]Humboldt University, Spandauer Str. 1, 10178 Berlin, Germany. Email: jonas.radbruch@hu-berlin.de

1 Introduction

Over the past two decades, nudges have become an important new instrument in the toolbox of many policymakers. The establishment of the UK Behavioral Insights Team and similar “nudge units” around the world exemplify this development. The popularity of nudge-based policies—which aim at improving people’s choices by addressing their behavioral biases and cognitive limitations—is also reflected in a large body of academic studies that analyze the impact of nudges in a diverse set of economic applications (see, e.g., Thaler and Sunstein 2008). The considerable academic and practical interest in nudges notwithstanding, there exists a persistent debate on whether nudges are effective in achieving their ultimate goal of improving people’s lives. While some studies point out that nudges compare favorably to other policy tools, especially in terms of cost-effectiveness (Benartzi et al. 2017), recent meta-studies suggest that nudge interventions induce relatively modest changes in the behaviors they target (DellaVigna and Linos 2022, Mertens et al. 2022, Maier et al. 2022), and that their efficacy varies strongly between different groups of individuals (Szasz et al. 2022).

In this paper, we inform the debate about nudges by shedding light on an important, yet commonly overlooked, aspect of nudge-based policies. In particular, we show that nudge interventions do not only have direct effects on targeted behaviors, but they can also cause substantial, positive indirect effects on behavioral domains that are not the primary target of the interventions. These indirect effects are typically neglected in the evaluation of nudge interventions. We demonstrate that this practice can lead analysts to (i) substantially underestimate the overall effectiveness of nudges, (ii) form biased assessments of their impact on different groups of decision makers, and (iii) make erroneous inferences regarding the relative effectiveness of different nudges.

Our empirical analysis is based on a laboratory experiment involving a dual-task paradigm. Participants in the experiment can earn financial rewards by engaging in two distinct cognitively demanding tasks. The first task, a match-pair memory game, remains constant in all treatments and is not targeted by any of our interventions. The second task, a math task, features a pre-selected default option, which governs individuals’ choices if they do not make an active decision. This task is targeted by our experimental interventions. Guided by a simple theoretical framework of optimal allocation of scarce cognitive resources, we study the effects of two nudge interventions that share a common feature: They arguably reduce the cognitive resources required for the targeted domain and thus may induce positive indirect effects on other, non-targeted domains. First, relative to a BASELINE condition, the SIMPLIFICATION treatment reduces the complexity of the targeted task.

Second, in the GOODDEFAULTS condition, the quality of the default option in the targeted task is higher than in the BASELINE condition.

The treatments of our experiment are reminiscent of two commonly implemented nudge interventions. Researchers and policymakers have long argued that overly complex forms, burdensome administrative processes, and the use of convoluted language in public communications (sometimes also referred to as “sludge”) can be a key obstacle for the take-up and effectiveness of social programs, payment compliance, and policy interventions in general (e.g., Abeler and Jäger 2015, Bhargava and Manoli 2015, Sunstein 2021, 2022b). Desmond (2023) estimates that every year more than \$ 140 billion in government benefits go unclaimed.¹ Consequently, simplification interventions, such as the use of simpler forms, the streamlining of administrative processes, and the use of natural language, have gained significant popularity (e.g., Bettinger et al. 2012, Beshears et al. 2013, Finkelstein and Notowidigdo 2019, Ericson et al. 2023). A reduction of administrative burden is also a major goal of current administrations, as documented, for example, in the Burden Reduction Report (OIRA 2023). Similarly, the concept of setting “good defaults” has also been featured very prominently (Thaler and Sunstein 2003, 2008, Carroll et al. 2009, Brot-Goldberg et al. 2023). The strong interest in setting good defaults has been fueled by widespread evidence that defaults are sticky and can substantially alter the fraction of individuals who make a specific decision, for example, when deciding on organ donor status (Johnson and Goldstein 2003), health care (Brot-Goldberg et al. 2023), and retirement savings (Madrian and Shea 2001).

The controlled environment of our laboratory experiment is ideal for studying how SIMPLIFICATION and GOODDEFAULTS alter the quality of individuals’ decisions in the domain that is directly targeted by the interventions as well as in the non-targeted domain. This comprehensive approach allows us to examine not only the direct effects of the interventions, which are typically studied in evaluations of nudge interventions, but also potential indirect effects on the non-targeted domain, which are commonly neglected. We find that GOODDEFAULTS and SIMPLIFICATION have a strong effect on behavior in the targeted domain. In particular, relative to BASELINE the quality of choices in the targeted task increases by 21% and 18% in SIMPLIFICATION and GOODDEFAULTS, respectively. At the same time, both interventions cause a strong and positive indirect effect on individuals’ behavior in the non-targeted task. Compared to BASELINE, the quality of choices in this task increases by 22% in SIMPLIFICATION and also by 22% in GOODDEFAULTS. Both interventions thus yield a double dividend: their *overall* effect on the quality of individual choices is twice as large

¹For instance, only about half of the individuals that are eligible for nutrition assistance from the U.S. Department of Agriculture have actually received that support (OIRA 2023).

as the direct effect observed within the targeted domain alone. Failing to account for the positive indirect effects would, therefore, lead to a substantial underestimation of the overall effectiveness of the interventions.

Furthermore, our experimental setup enables us to investigate how indirect effects vary between different subgroups of participants. For this purpose, we classify subjects according to their available stock of cognitive resources, i.e., their individual “bandwidth”. Our findings reveal substantial indirect effects for both interventions and all subgroups. However, the relative magnitude of the direct and indirect effects depends strongly on individuals’ bandwidth. For participants in the lowest bandwidth quartile, we estimate that approximately 85% (89%) of the overall increase in payoffs due to SIMPLIFICATION (GOODDEFAULTS) can be attributed to the direct effect of the interventions. This proportion decreases to 3% (3%) for the highest bandwidth quartile. Therefore, an evaluation that focuses solely on the targeted domain will reveal almost the entire aggregate effect of the nudge intervention for individuals with strong bandwidth limitations. In contrast, for individuals with ample bandwidth, a narrow focus on the targeted domain would uncover only a negligible portion of the overall effect. Intuitively, the effects of an intervention may remain largely unnoticed within this subgroup of the population, as they make relatively good choices in the targeted domain, even in the absence of the intervention. However, this group may benefit substantially from the positive indirect effects of the intervention, as the intervention allows them to shift attention from the targeted to the non-targeted task. Consequently, assessing the impact of a nudge intervention solely based on its direct effects can lead to a severely biased understanding of the heterogeneity in the overall effects of the intervention and the resulting distributional consequences.

Importantly, our setup also sheds light on the attentional mechanisms underlying the impact of the interventions on choices. In particular, we elicit a detailed measure of how individuals allocate cognitive resources across the targeted and non-targeted domain by recording the time that individuals spend working on each task. This feature allows us to assess the relationship between the overall effects of the interventions and the redistribution of cognitive resources in response to them. Our data reveal that the positive indirect effects observed in the non-targeted domain emerge as a result of an *attention-releasing* effect induced by the interventions within the targeted domain. Both interventions cause participants to devote fewer cognitive resources to the targeted task. Compared to BASELINE, GOODDEFAULTS reduce the attention devoted to the targeted task by 26%, while SIMPLIFICATION reduces the attention allocated to the task by 39%. Consequently, both GOODDEFAULTS and SIMPLIFICATION free up cognitive resources that can be used to make

better choices in other domains.

Although both interventions studied in our paper are, on average, attention-releasing, we also observe noteworthy differences in how they alter individuals' allocation of cognitive resources. In particular, we find that GOODDEFAULTS lead to a rather uniform shift of attention from the targeted to the non-targeted task. Relative to BASELINE, a broad range of subjects devote less attention to the target task and more attention to the non-targeted task. As a result, the cumulative distribution of attention spans observed in GOODDEFAULTS first order stochastically dominates the one in BASELINE. In contrast, SIMPLIFICATION leads to a compression of the attention spans devoted to the targeted task, relative to BASELINE. In particular, we observe that in 16% of all cases, participants in BASELINE do not devote any cognitive resources to the targeted task. As a result of SIMPLIFICATION, this fraction decreases to only 8%. Thus, the intervention encourages subjects to engage with the (simplified) task. At the same time SIMPLIFICATION allows those individuals, who would devote substantial attention to the targeted task in BASELINE, to withdraw some attention from the task. In general, SIMPLIFICATION therefore induces different participants to devote more homogeneous amounts of cognitive resources on the targeted task, leading to a compression of the attention distribution.

These differences in attentional responses to the interventions are crucial for understanding how nudge interventions affect choice quality across different subgroups of the population. Consistent with the uniform shift in attention, we observe that GOODDEFAULTS induce relatively homogeneous positive indirect effects on choice quality in the non-targeted domain. In contrast, the compression in the attention distribution due to SIMPLIFICATION also implies that the indirect effects differ substantially across the population. Neglecting indirect effects when deciding which nudge to implement may therefore lead to an erroneous inference of which nudge is most effective in reducing inequalities between individuals.

Our paper offers important new insights for the evaluation of nudges and other behavioral policy interventions. The effectiveness of such policies has been extensively debated. Scholars have argued that the current emphasis on nudges may lead to a neglect of potentially more effective policy instruments (Loewenstein and Chater 2017, Chater and Loewenstein 2022) and that the impact of nudges on behavior in the targeted domains is relatively small (DellaVigna and Linos 2022, Mertens et al. 2022, Maier et al. 2022). Our findings demonstrate a simple effect, which is of first-order importance for this debate, yet has been widely neglected so far: the existence of positive indirect effects of nudge interventions on other domains that are not directly targeted by

the interventions. Our results show that these indirect effects arise systematically and can be as substantial as the direct effects of the interventions. In fact, certain subgroups of the population may exclusively benefit in non-targeted domains, implying that the absence of direct effects can not necessarily be interpreted as an ineffectiveness of the policy instrument. For instance, simplifications of the enrollment process for social aid or transfer programs may not affect outcomes in the targeted domain for individuals who would have been enrolled in the program regardless. Nevertheless, these individuals may clearly benefit from the intervention, as it eases their cognitive resource constraints and enables them to utilize these released resources in other domains.

While the majority of studies has focused on evaluating the effects on targeted domains, there is a growing interest in understanding the indirect effects of nudges and other behavioral policy interventions. To the best of our knowledge, however, the existing literature has exclusively discussed negative side effects. For instance, reminders may crowd out non-reminded actions (Medina 2020, Koch et al. 2023), messages to increase compliance with rental income taxes can reduce tax payments in other domains (Castro et al. 2022), and encouraging healthy behavior in one domain can reduce health promoting activities in other domains (Trachtman 2023). Nafziger (2020) and Altmann et al. (2022) highlight that these negative indirect effects may occur because policy interventions interfere with individuals' deliberate allocation of cognitive resources, directing some resources towards the targeted domain. In contrast, our study focuses on the existence and characteristics of positive indirect effects of simplifications and high-quality defaults. By identifying shifts in attention as a main driver of indirect effects, we contribute to a better understanding of how nudges can be classified into those with negative and those with positive indirect effects: Attention-grabbing nudges (e.g., reminders, active choice) induce a reallocation of scarce cognitive resources from non-targeted domains to the targeted domain and are, therefore, likely to cause negative indirect effects. On the other hand, attention-releasing nudges (good defaults, simplification, reduction of sludge) rather soften individuals' cognitive resource constraint and, therefore, induce positive indirect effects.

One important aspect in the evaluation of policy interventions is the question which group of the population benefits the most. Research suggests that high-income and highly educated individuals tend to disproportionately profit from social benefit programs due to their greater access and utilization of available resources (Heckman and Landersø 2021). Whether a similar pattern emerges in the realm of behavioral policy interventions is the subject of intense debate. On the one hand, scholars argue that behavioral interventions which simplify decision-making processes or improve default options are inherently likely to help those who are most in need (Mrkva et al.

2021, Sunstein 2021, 2022a). Specifically, by alleviating cognitive resource constraints, they can benefit individuals who face significant bandwidth limitations, such as the poor or other vulnerable subgroups of the population (Mani et al. 2013, Bertrand et al. 2004, Sharafi 2023, Herd et al. 2023). On the other hand, Finkelstein and Notowidigdo (2019) demonstrate that individuals who sign up for the Supplemental Assistance Nutrition Program in response to a simplified procedure are generally less needy than those who would have signed up without the simplification. This finding supports the argument of Roberts (2017) that the progressive effects of behavioral policy measures may be commonly overstated, as the most vulnerable individuals may be resistant to nudges due to structural factors that restrict their autonomy in decision-making. Our findings highlight the importance of considering indirect effects when assessing the distributional impacts of nudges. For both interventions, we find that subgroups with minimal direct effects actually benefits the most from indirect effects. Therefore, ignoring indirect effects when evaluating the distributional consequences of behavioral interventions may lead to a systematically biased assessment.

We also contribute to a body of work in industrial organization that examines the optimal design of regulatory interventions when consumers face limited cognitive resources (see, e.g., Barr et al. 2008, Heidhues et al. 2021, Johnen and Leung 2022). An important theme in this literature is that harmonizing or simplifying the features of secondary products can promote competition (Ericson and Starc 2016, Heidhues et al. 2021, Fehr and Wu 2023). Specifically, simplifying secondary contract terms may allow consumers to allocate their cognitive resources towards more extensive cross-product search, rather than spending time scrutinizing the secondary terms for one product. Our findings support this notion. Simplifications in our setting prompt individuals to redirect their cognitive resources from the simplified task and effectively employ them in other domains.

By providing evidence on the underlying attentional mechanism, our paper is also related to studies that use choice data from laboratory experiments (Caplin and Dean 2013, Dertwinkel-Kalt et al. 2022, Dean and Neligh 2023, Martin 2017, Altmann et al. 2022) and field experiments (Bartoš et al. 2016, Bronchetti et al. 2023) to test models featuring limited cognitive resources. We complement this body of research by studying how the fact that cognitive resources are limited influences the design and evaluation of nudges. This research question requires the elicitation of detailed choice process data in order to track subjects' allocation of cognitive resources. Methodologically, our paper is therefore also related to studies that use various methods to track the choice process instead of choice outcomes such as Mouselab (e.g., Johnson et al. 1989, Gabaix et al. 2006), eye-tracking (e.g., Wang et al. 2010), or click data (Caplin et al. 2011).

2 Design of the Experiment

The goal of our experiment is to study the direct and indirect effects of nudge interventions. Towards this end, we set up a stylized decision environment. This environment allows us to precisely measure how the interventions alter individuals' choices in the domain they directly target as well as in other, non-targeted domains. Participants in our experiment face two cognitively demanding tasks. Our treatments vary whether one of the tasks is targeted by a nudge intervention. We focus on two interventions that share a common feature. Specifically, they aim to help participants make better decisions in the targeted choice domain by reducing the cognitive resources required to solve the targeted task. We investigate how the interventions affect participants' allocation of cognitive resources across tasks and the quality of their choices in both the targeted and non-targeted tasks.

Figure 1: The Targeted Task

Please choose the option with the highest sum.

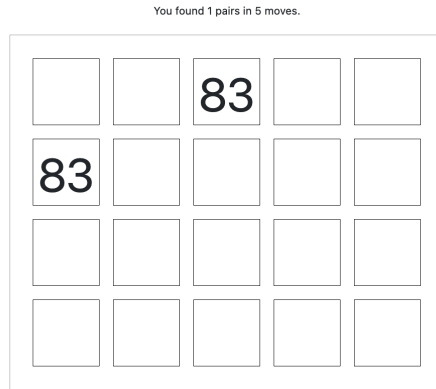
- $8 + 5 + 7 + 5$
- $8 + 1 + 9 + 5$
- $1 + 3 + 7 + 7$
- $1 + 4 + 7 + 6$
- $9 + 6 + 2 + 6$

Participants' task in the *targeted choice domain* consisted of solving simple math problems. In each round of the experiment, subjects faced 5 summations (see Figure 1 for an example). Their goal was to find out which of the 5 options yields the highest sum.² The targeted task featured a default option that was implemented if subjects did not make an active decision. The quality of the default, i.e., the probability that the default option was correct, differed across treatments (see Section 2.1). Subjects were informed about the existence and the quality of the default in the instructions of the experiment.

Participants' task in the second, *non-targeted choice domain* was a matching-pairs memory game. Specifically, subjects had to find matching pairs of two-digit numbers on a 4×5 memory grid (see Figure 2). At the beginning of each round, the cards laid face down. Subjects could then flip two cards per move by clicking on the corresponding cards. If both cards showed the same number,

²To hold the difficulty of the task roughly constant across different rounds of the experiment, each option resulted in a sum between 15 and 25.

Figure 2: The Non-targeted Task



subjects found a pair and the cards stayed flipped over.

The financial rewards for solving the two tasks were as follows. Subjects received €0.30 if they correctly solved the math task in a given round. In the memory task, subjects received €0.03 for every pair of cards they found. Hence, they received €0.30 if they found all 10 pairs in the 4x5 memory grid. In each round of the experiment, subjects had a total of 60 seconds to work on the targeted and non-targeted task. Within this time span, subjects could freely navigate between the tasks by pressing a “Switch Task” button at the bottom of the screen, or by using tabs at the top of the screen (Figure D.5 in the appendix displays an example of subjects’ decision screens).

2.1 Treatments

We implemented three treatments to study how nudge interventions affect cognitive resource allocation and the quality of individuals’ choices. Between treatments, we exogenously varied the characteristics of the targeted task, while keeping all features of the non-targeted task constant. In our first treatment, denoted as BASELINE, subjects worked on the targeted and non-targeted task, as depicted in Figures 1 and 2. The BASELINE environment reflects situations in which nudge interventions can potentially help people to make better choices. Individuals face two demanding tasks that compete for their scarce cognitive resources, the representation of the math task is rather complex, and paying little attention to it may lead to wrong decisions, as the task features a low-quality default option. In particular, one of the five options of the math task was randomly selected and displayed as default. Therefore, the default option in BASELINE is the correct choice only in 20% of the cases.

Our second treatment—the GOODDEFAULTS environment—was identical to BASELINE except for the quality of the default option. In GOODDEFAULTS, the default option in the math task

Figure 3: The Targeted Task in SIMPLIFICATION

Please choose the option with the highest sum.

- 20 + 5
- 18 + 5
- 11 + 7
- 12 + 6
- 17 + 6

corresponded to the correct choice in 80% of cases. This intervention aims to improve the quality of choices by individuals who pay little or no attention to the targeted task. It reflects the idea that, in light of the widely observed “stickiness” of defaults, policy makers should set defaults that are (likely to be) a good choice for individuals (Johnson and Goldstein 2003, Thaler and Sunstein 2003, 2008, Brot-Goldberg et al. 2023). Furthermore, in line with our interest in a possible attention-releasing effect of our interventions, some authors have suggested that high-quality defaults have the potential to alleviate cognitive burden, thus helping individuals who face binding bandwidth limitations (e.g., Duflo 2012).

In our final treatment—the SIMPLIFICATION environment—subjects faced the memory task and the math task with a low-quality default, which was correct with 20% probability, as in the BASELINE environment. However, we simplified the representation of the math task relative to BASELINE. Specifically, while each option of the math task in BASELINE contained 4 summands, we reduced the number of summands to 2 in SIMPLIFICATION. We did so by summing up the first three summands of each option. Figures 1 and 3 provide an illustration of the math task in BASELINE and SIMPLIFICATION, respectively. Although SIMPLIFICATION does not change the nature of the targeted task, it directly targets the amount of cognitive resources required to solve the task. Therefore, the intervention SIMPLIFICATION is reminiscent of policy initiatives to reduce “sludge” or administrative burden, by simplifying overly complex forms or administrative processes (see e.g. Bhargava and Manoli 2015, Sunstein 2021, OIRA 2023, Ericson et al. 2023).

2.2 Procedures

The experiment consisted of three parts. In the first two parts, we familiarized participants with the memory task, the math task, and the corresponding payoffs. Subjects first received detailed

instructions on the memory task.³ Subsequently, subjects engaged in two incentivized practice rounds of 60 seconds each, focusing exclusively on the memory task without the possibility of working on the math task. In the second part, participants received instructions on the math task and subsequently completed two incentivized practice periods in which they worked on the math task alone. The third and main part of the experiment consisted of ten rounds. In each round, participants could simultaneously work on both tasks for 60 seconds. Subjects were not able to proceed to the next round of the experiment before this time had elapsed. The order and content of the memory task in different rounds was identical across all subjects and treatments. Similarly, the math task was kept constant between subjects and treatments in every given round, except for the quality of the default and the simpler vs. more complex representation of the task (see Section 2.1). The experiment ended with a short post-experimental questionnaire and a short test on fluid intelligence. Table B.1 in the appendix provides descriptive statistics and balancing checks. As the randomization checks yield some minor imbalances across treatments, we always report estimation results with and without controls for additional covariates when discussing our empirical results in Section 4.

The experiment was conducted online with subjects from the participant pool of the BonnEcon-Lab at the University of Bonn. It was implemented with oTree (Chen et al. 2016) and the online recruitment system by Bock et al. (2014). When registering for the experiment, participants were told that they could only participate on a laptop or personal computer, but not on mobile devices. We ensured compliance by screening device types on the first screen of the experiment. A total of 840 subjects participated in our experiment, distributed across five sessions on five consecutive days. Within each session, subjects were randomly assigned to the BASELINE, SIMPLIFICATION and GOODDEFAULTS environment. On the day of their experimental session, subjects received an email with the link to the experimental software around noon. The link remained active until 07:00 p.m. and subjects had to complete the experiment until that point to qualify for payment. On average, participants spent approximately 30 minutes to complete the experiment, including the time needed for instructions and the post-experimental questionnaire. The mean earnings of subjects were €7.73, including a show-up fee of €1.

³A translated version of the instructions for all parts and all treatments of the experiment can be found in Appendix C.

3 Behavioral Predictions

In what follows, we develop our behavioral predictions. To fix ideas, we make use of a model in which decision makers have a limited stock of cognitive resources but optimally allocate their limited resources across tasks. This assumption is clearly a simplification, as individuals might make mistakes when determining the allocation of cognitive resources, both in terms of whether and how much attention they devote to different tasks. Nevertheless, the approach is useful to illustrate how GOODDEFAULTS and SIMPLIFICATION can have an attention-releasing effect on individuals and how this, in turn, may influence the quality of decisions in different choice domains. We build upon the frameworks and ideas introduced in Alonso et al. (2014) and Altmann et al. (2022) and relegate the analysis of the model and all formal results to Appendix A. In what follows, we illustrate the main theoretical intuitions that give rise to our hypotheses.

We assume that each decision maker i is equipped with an individual-specific stock of cognitive resources or bandwidth X .⁴ The decision maker faces two tasks, T and NT, which simultaneously require cognitive resources to be solved. Task T features a default option and can be targeted by a nudge intervention; task NT denotes the non-targeted task. Allocating resources x_T to task T and x_{NT} to task NT results in the likelihoods $\pi_{NT}(x_{NT})$ and $\pi_T(x_T, \lambda, \phi)$ to solve the tasks correctly and obtain payoffs u_T and u_{NT} , respectively. To make explicit that the likelihood of solving task T depends on the treatment, π_T has two additional arguments, λ , denoting the quality of the default option, and ϕ , denoting the difficulty of task T. The decision problem of the individual is given by:

$$\begin{aligned} \max_{x_{NT}, x_T} u(x_{NT}, x_T) &= \pi_{NT}(x_{NT})u_{NT} + \pi_T(x_T, \lambda, \phi)u_T & (1) \\ \text{s.t. } x_T + x_{NT} &\leq X \end{aligned}$$

Decision makers in BASELINE optimally allocate their cognitive resources between the two tasks. The optimal allocation of resources depends on how scarce these resources are for a given individual. Specifically, individuals whose stock of cognitive resources X lies below a threshold X^* will optimally allocate all their resources to Task NT. This result arises because allocating minimal resources to task T will not allow decision makers to exceed the payoff they can expect if they simply stick to the default in this task. Therefore, it is optimal for decision makers with a low bandwidth to rely on

⁴Alternatively, one can also interpret X as efficiency units of time that participants in the experiment have available. Suppose, for example, that all participants have the same amount of time T to work on the tasks but differ in their ability a , which describes how effective they are in solving the tasks with a given amount of time. Then $X = aT$ describes their available (effective) resources that they can allocate across tasks.

the default while devoting all their cognitive resources to Task NT. Conversely, individuals whose bandwidth X is above the threshold X^* will attend to both tasks. The larger their stock of X , the more resources they will devote to each of the tasks, resulting in better choices and higher payoffs in both tasks. Building on this intuition, the appendix derives a formal condition on the shapes of π_{NT} and π_T such that the threshold X^* exists and is unique. Based on this insight about the optimal allocation of scarce cognitive resources in BASELINE, we now analyze how GOODDEFAULTS and SIMPLIFICATION change individuals’ cognitive resource allocation and the quality of their choices.

3.1 Good Defaults

The implications of GOODDEFAULTS for the allocation of cognitive resources and the quality of individuals’ choices depend on their bandwidth X . Following the intuition for the BASELINE environment, it will be optimal for decision makers in GOODDEFAULTS to devote no resources to Task T if their bandwidth is below the threshold X^* . Individuals with higher X devote at least some attention to the targeted task in BASELINE. For some of these individuals—i.e., individuals whose bandwidth lies just above X^* —the higher default quality in GOODDEFAULTS and the resulting higher expected payoff from sticking to the default will make it optimal to pay no attention to task T. In a similar vein, individuals with higher bandwidth X will benefit from GOODDEFAULTS as they can withdraw some resources from the targeted task and utilize these resources to solve the non-targeted task. These insights yield the following three predictions for treatment differences between BASELINE and GOODDEFAULTS.

First, GOODDEFAULTS have an attention-releasing effect. Specifically, all individuals should weakly decrease the cognitive resources allocated to the targeted task and increase the resources allocated to the non-targeted task. To formalize this argument, denote by $G_C(\cdot)$ the cumulative distribution of cognitive resources allocated to task T in condition $C \in \{Base, Good, Simple\}$. Then the above arguments yield that $G_{Base}(\cdot)$ should first-order stochastically dominate $G_{Good}(\cdot)$.

Second, the reallocation of attention has immediate consequences for the quality of individuals’ decisions. The higher amount of cognitive resources devoted to the non-targeted choice domain should improve the quality of individuals’ decisions in this domain. At the same time, the impact of GOODDEFAULTS on the quality of individuals’ choices in the targeted domain is ambiguous. On the one hand, individuals who pay relatively little attention to the targeted domain in BASELINE, such that their likelihood of solving task T lies below 80% (i.e., the likelihood that the default is correct in GOODDEFAULTS), will benefit from GOODDEFAULTS. On the other hand, individuals who devote high amounts of cognitive resources to the targeted task in BASELINE, such that they

solve the task in 80% or more of the cases, may experience a drop in choice quality in the targeted task, due to the shift in attention from task T to task NT. Which of the two effects dominates depends on the distribution of cognitive resources X in the population. Regardless of the sign of the direct effect of GOODDEFAULTS on choice quality in the targeted domain, GOODDEFAULTS will increase the *overall* average payoffs.

Third, our framework also speaks towards the relative size of the effects of GOODDEFAULTS in the targeted and non-targeted domain. For participants with relatively small X , GOODDEFAULTS will have larger direct effects on the quality of choices in the targeted domain than indirect effects on the non-targeted task. This is the case as low-bandwidth subjects concentrate (almost) all their resources on the non-targeted task also in BASELINE. GOODDEFAULTS generically increase the quality of choices in the targeted domain for these subjects, but the attention-releasing effect and, thus, the improvement of choices in the non-targeted domain is relatively small. The relative importance of the direct and indirect effects of GOODDEFAULTS, however, reverses for participants with high bandwidth X . These subjects devote substantial amounts of cognitive resources to the targeted task in BASELINE. With GOODDEFAULTS, high-bandwidth subjects can devote more attention to the non-targeted task, resulting in a relatively large, positive indirect effect on the quality of choices in the non-targeted domain. In contrast, the direct effect of GOODDEFAULTS on choice quality in the targeted domain will be relatively small for these individuals, as they already exhibit a high likelihood of solving the targeted task in BASELINE. The following hypothesis summarizes the effects of GOODDEFAULTS. Their formal derivation is relegated to Proposition 1 in the Appendix.

Hypothesis 1. *GOODDEFAULTS lead to the following effects:*

(i) Choice quality: *In the non-targeted domain, the average quality of decisions are higher in GOODDEFAULTS than in BASELINE. The effect on choice quality and payoffs in the targeted domain is ambiguous. Average overall payoffs of participants are higher in GOODDEFAULTS than in BASELINE.*

(ii) Relative Payoff Effect: *For participants with small X , GOODDEFAULTS have a positive direct effect on participants' payoffs in the targeted domain, which is larger than the indirect effect on payoffs in the non-targeted domain. For participants with large X , the positive indirect effect on payoffs in the non-targeted domain is larger than the direct effect.*

(iii) Cognitive Resources: *GOODDEFAULTS have an attention-releasing effect. In particular, the*

cumulative distribution of cognitive resources devoted to task T in BASELINE first-order stochastically dominates the distribution in GOODDEFAULTS.

3.2 Simplification

We next turn to analyzing the expected treatment differences between SIMPLIFICATION and BASELINE. To conceptualize SIMPLIFICATION, we assume that the marginal likelihood of solving the task increases if the task becomes simpler, i.e., $\frac{\partial^2 \pi_T}{\partial x_T \partial \phi} < 0$ for all x_T . How this change in the slope of π_T affects the optimal allocation of cognitive resources again crucially depends on individuals' stock of cognitive resources.

Individuals with a relatively large stock of cognitive resources will devote large amounts of resources to the targeted task in BASELINE and solve this task (almost) with certainty. The SIMPLIFICATION of the targeted task has an attention-releasing effect on these individuals. As solving the targeted task requires fewer resources, they reallocate attention to the non-targeted task. Empirically, we should therefore observe fewer instances of subjects devoting comparatively large amounts of resources to task T. In contrast, individuals whose bandwidth is below the threshold X^* , do not attend to task T in BASELINE (see above). However, in SIMPLIFICATION, the allocation of resources to the targeted task becomes more attractive for these individuals, as the likelihood of solving the task increases. Therefore, we should observe fewer instances of subjects devoting no or only negligible amounts of attention to task T in SIMPLIFICATION. Formally, these arguments imply that there should exist a threshold \bar{x} such that $G_{Base}(x_T) > G_{Simple}(x_T)$ if and only if $x_T < \bar{x}$ (see Appendix A). In other words, SIMPLIFICATION should induce a squeeze in the distribution of cognitive resources allocated to task T such that the cumulative distributions $G_{Base}(\cdot)$ and $G_{Simple}(\cdot)$ intersect exactly once.

Although the effects of SIMPLIFICATION on the allocation of cognitive resources, therefore, depend on X , the effect on choice quality and payoffs in the targeted domain is weakly positive for all subjects. In particular, individuals with large bandwidth do not sacrifice payoffs in the targeted domain, they simply need fewer resources to reach the same (or even higher) choice quality in the targeted task, relative to BASELINE. Conversely, individuals with low bandwidth devote more cognitive resources to the simplified targeted task in SIMPLIFICATION. As a result, their choice quality in the targeted choice domain increases. Since some subjects increase and some decrease the amount of cognitive resources devoted to the targeted task, the impact of SIMPLIFICATION on choice quality in the non-targeted domain is ambiguous and depends on the distribution of X . However, in general, SIMPLIFICATION should again result in an increase in average overall payoffs relative to BASELINE.

Finally, we consider the relative importance of the direct and indirect effect of SIMPLIFICATION. As for GOODDEFAULTS, the direct effect will be larger than the indirect effect for individuals with small X . These individuals will increase the amount of cognitive resources devoted to the targeted task and additionally benefit from the fact that it is more likely to solve the task for any given amount of cognitive resources, due to the simpler representation of the task. However, since individuals with sufficiently low bandwidth shift cognitive resources from the non-targeted task to the targeted task in response to simplification, the indirect effect on the non-targeted domain should be absent or even negative for this group of individuals. In contrast, individuals with large X have a high likelihood to solve the targeted task in both BASELINE and SIMPLIFICATION, such that there is little or no direct effect on the quality of the choice in the targeted task. However, as a result of the attention-releasing effect of SIMPLIFICATION, these individuals devote more attention to the non-targeted task, resulting in relatively large positive indirect effects on the quality of choices in the non-targeted domain. The following hypothesis summarizes our predictions for the effects of SIMPLIFICATION. Their formal derivation is relegated to Proposition 2 in the appendix.

Hypothesis 2. SIMPLIFICATION leads to the following effects:

(i) Choice quality: In the targeted domain, the average quality of decisions is higher in SIMPLIFICATION than in BASELINE. The effect on choice quality in the non-targeted domain is ambiguous. Average overall payoffs of participants are higher in SIMPLIFICATION than in BASELINE.

(ii) Relative Payoff Effect: For participants with small X , the direct positive effect on participants' payoffs in the targeted domain is larger than the indirect effect on payoffs in the non-targeted domain. For participants with large X , the positive indirect effect on payoffs in the non-targeted domain is larger than the direct effect.

(iii) Cognitive resources: SIMPLIFICATION has an ambiguous effect on the allocation of cognitive resources. More specifically, the cumulative distributions of cognitive resources allocated to task T $G_{Base}(\cdot)$ and $G_{Simple}(\cdot)$ intersect exactly once.

4 Results

This section describes the results of our experiment. We first examine the impact of the interventions on the quality of individuals' choices in the targeted and non-targeted domain. In a second step, we shed further light on the mechanisms behind the observed treatment differences in behavior, by analyzing how the interventions alter individuals' allocation of scarce cognitive resources.

4.1 Choice Quality in the Targeted and Non-Targeted Domain

In a first step, we study how SIMPLIFICATION and GOODDEFAULTS affect the quality of individuals' choices. As our measure of choice quality in the targeted and non-targeted domain, we examine individuals' payoffs in the corresponding task at the participant-round level. Columns (1) and (2) of Table 1 display treatment differences in individuals' payoffs in the task that is targeted by the interventions. We find that SIMPLIFICATION and GOODDEFAULTS substantially improve the quality of participants' choices in the targeted domain. While participants in BASELINE earn, on average, 22.1 cents per round in the targeted task, this number increases significantly by 4.0 cents in GOODDEFAULTS and by 4.7 cents in SIMPLIFICATION. As the rewards for correctly solving the task in a given round were €0.30, these numbers imply that the subjects, on average, solve the targeted task in 73.6% of the cases in BASELINE, 87.0% of cases in GOODDEFAULTS, and 89.3% of cases in SIMPLIFICATION.⁵ Hence, both nudge interventions improve the quality of individuals' choices in the domain that they target.

Table 1: Direct and Indirect Effects of the Interventions

	Targeted Task		Non-Targeted Task		Overall	
	(1)	(2)	(3)	(4)	(5)	(6)
GOOD DEFAULTS	4.034*** (0.529)	3.807*** (0.497)	4.157*** (0.557)	3.725*** (0.515)	8.190*** (0.707)	7.532*** (0.608)
SIMPLIFICATION	4.716*** (0.563)	4.544*** (0.529)	4.099*** (0.521)	3.695*** (0.476)	8.815*** (0.730)	8.239*** (0.627)
Controls	No	Yes	No	Yes	No	Yes
Mean BASELINE	22.077	22.077	18.911	18.911	40.988	40.988
N	8400	8400	8400	8400	8400	8400
No. Subjects	840	840	840	840	840	840
R^2	0.035	0.091	0.042	0.191	0.075	0.226
SIMPLIFICATION= GOOD DEFAULTS	0.071	0.058	0.910	0.947	0.327	0.211

Note: The table reports results of OLS regressions of differences in participants' payoffs in different treatments of the experiment. The dependent variable in Columns (1)–(2) is the payoff in the targeted task in a given round of the experiment. The dependent variable in Columns (3)–(4) is the payoff in the non-targeted task in a given round of the experiment. The dependent variable in Columns (5)–(6) is the overall payoff in both tasks in a given round of the experiment. Control variables in Columns (2), (4), and (6) include subjects' age, gender, performance in the trial rounds in the targeted and non-targeted task, task-round fixed effects, and Raven score. The lower part of the table reports p-values from post-estimation tests of the equality of selected coefficients (Wald tests). Robust standard errors, clustered at the subject level, are reported in parentheses. ***, **, * indicates significance at the 1%, 5%, and 10% level, respectively.

⁵When comparing payoffs in the targeted task between GOODDEFAULTS and SIMPLIFICATION, we find that the payoffs are slightly higher in SIMPLIFICATION; cp. post-estimation tests in columns (1) and (2) of Table 1.

However, the direct positive effects in the targeted domain do not fully capture the behavioral consequences of SIMPLIFICATION and GOODDEFAULTS. In particular, columns (3) and (4) of Table 1 demonstrate that both interventions also affect the quality of individuals’ choices in the second, non-targeted task. We find that participants’ payoffs in the non-targeted domain increase from 18.9 cents per round in BASELINE to 23.1 cents in GOODDEFAULTS and 23.0 cents in SIMPLIFICATION. Hence, both interventions increase the quality of individuals’ choices in the non-targeted domain by about 22%, relative to BASELINE. While the differences between both treatments and the BASELINE condition are highly statistically significant, the payoffs in the non-targeted domain do not differ significantly between SIMPLIFICATION and GOODDEFAULTS.

The positive indirect effects of the interventions in the non-targeted choice domain are similar in magnitude to the interventions’ direct effects in the targeted domain. Columns (5) and (6) of Table 1 document the impact of the interventions on subjects’ overall payoffs. The estimates illustrate that both interventions lead to substantive and statistically significant increases in participants’ overall payoffs. Most importantly, the size of the coefficients demonstrates that the overall benefits of SIMPLIFICATION and GOODDEFAULTS are almost twice as large as the direct benefits that arise from subjects making better decisions in the targeted choice domain.

Result 1. *Two important types of nudges—SIMPLIFICATION and GOODDEFAULTS—improve individuals’ decisions in the choice domain that the interventions target. In addition, both interventions cause substantial, positive indirect effects on the quality of individuals’ choices in other domains. These findings are in line with part (i) of Hypotheses 1 and 2.*

Direct and indirect effects for different subgroups

Our first result demonstrates that SIMPLIFICATION and GOODDEFAULTS yield a ‘double dividend’. They do not only improve individuals’ decisions in the domain that they target, but also have a positive indirect effect on choices in other, non-targeted tasks. Our theoretical analysis in Section 3 indicated that the relative importance of the direct and indirect effect should differ systematically across the population. Specifically, we hypothesized that the positive direct effects of the interventions should be particularly pronounced for individuals with a relatively small stock of cognitive resources who, in the absence of the interventions, make relatively poor choices in the targeted domain. In contrast, subjects with a relatively large stock of cognitive resources should already exhibit a high likelihood of solving the targeted task in BASELINE. For these high-bandwidth subjects, the direct effect on the targeted domain is relatively less important than the positive indirect effect of the interventions (cp. part (ii) of Hypotheses 1 and 2).

Since we do not readily observe participants’ stock of cognitive resources, we cannot provide a direct test of this hypothesis. However, the theoretical analysis in Section 3 implies that, within each treatment, there should be a monotonic relationship between an individual’s stock of cognitive resources and the quality of her choices in the targeted task. Hence, we can use the performance percentile of each subject in the targeted domain as a measure of her stock of cognitive resources relative to the other participants. Alternatively, we can also measure subjects’ stock of cognitive resources or bandwidth by the test of fluid intelligence that we administered after the experiment, their performance in the first two parts of the experiment, or a combination of both (see Appendix B).⁶ The hypothesis implies that the relative importance of the direct effect of the interventions should be higher (smaller) for subjects that rank lower (higher) in the bandwidth distribution. To test this conjecture, we divide the sample within each treatment into four quartiles, based on the average quality of each participant’s choices in the targeted task. We then analyze the direct and indirect effect of the interventions for each quartile, following the logic of Table 1.

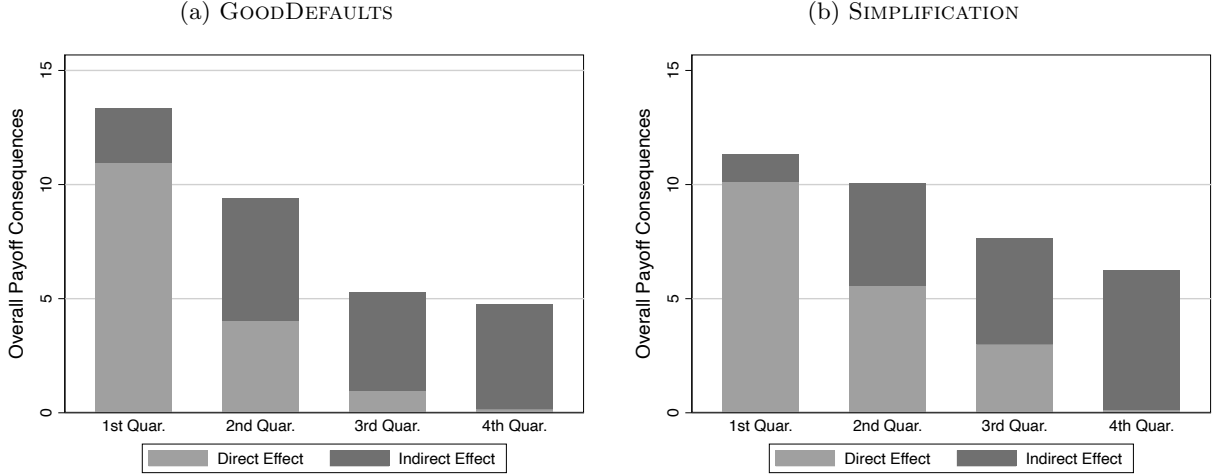
Figure 4 depicts the direct and indirect effects of GOODDEFAULTS (Panel a) and SIMPLIFICATION (Panel b) across the population. For each subgroup, the figure shows the increase in payoffs in the targeted task (light gray bars) and non-targeted task (dark gray bars). The total height of the bars in Figure 4 shows the overall increase in payoff due to GOODDEFAULTS and SIMPLIFICATION, relative to the corresponding subgroup in BASELINE. To estimate these payoff effects, Table B.2 in the appendix provides subgroup specific analyses paralleling the estimations in Table 1.

Two key findings emerge from Figure 4. First, GOODDEFAULTS and SIMPLIFICATION yield positive, statistically significant and sizable indirect effects for a wide range of participants (see Appendix Table B.2 for details). The payoff increase in the non-targeted domain ranges from approximately 1 Cent (SIMPLIFICATION, bottom quartile) to approximately 6 Cents (SIMPLIFICATION, top quartile).

Second, the relative importance of the direct and indirect effects differs substantially between different subgroups of participants. In both treatments, individuals who make relatively poor choices predominantly benefit from the direct effect, as illustrated by their strong increase in payoffs in the targeted task. The indirect effect of the interventions on choice quality in the non-targeted task is also positive for this group, but relatively modest. Indeed, when comparing the direct and indirect

⁶As we show in Figure B.1 and Table B.3 of the appendix, the results we present in what follows are qualitatively similar for these alternative measures. The measures provide a less direct test of the hypothesis, since they do not directly emerge from our model. However, they do not depend on the assumption that performance percentiles are not affected by the treatment. To combine the data from the intelligence test and the first two parts of the experiment (i.e., the trial rounds in which subjects worked on only one of the tasks), we extract a common factor from a factor analysis.

Figure 4: Relative Importance of the Direct and Indirect Effect



Note: The figure illustrates the payoff consequences of GOODDEFAULTS (panel (a)) and SIMPLIFICATION (panel (b)) in the targeted and non-targeted domain. The bars depict the average payoff differences between the treatments and BASELINE. The groups are categorized based on their performance quartiles in the targeted task within each treatment. 1st Quartile refers to the lowest performance level. To ensure that groups are of equal size, ties in performance are broken randomly. The tie-breaking procedure is repeated 500 times and the average payoff differences across repetitions is plotted.

effect for this subgroup in Table B.2, we find that the indirect effect is significantly smaller than the direct effect for both interventions (p-value < 0.0001 for both treatments).⁷ As a consequence, the benefits of GOODDEFAULTS and SIMPLIFICATION for this group of participants almost exclusively accrue from an improvement of choices in the targeted choice domain.

This is in stark contrast to the relative importance of the direct and indirect effect for the top quartiles of participants. For these subjects, the interventions only have a minimal direct effect on the quality of choices in the targeted domain. Since this subgroup already exhibits a high likelihood of solving the targeted task in BASELINE, GOODDEFAULTS and SIMPLIFICATION have only very little influence on the quality of their choices in this task. Yet, Figure 4 demonstrates that the interventions, nevertheless, have substantial value for these participants. In particular, as a result of SIMPLIFICATION and GOODDEFAULTS in the targeted choice domain, these subjects make better choices in the *non-targeted* domain. Indeed, when comparing the direct and indirect effect for this subgroup in Table B.2, we find that the indirect effect is significantly larger than the direct effect for both interventions (p-value < 0.0001 for both treatments).⁸ Hence, the top quartiles of participants

⁷This result is robust to using alternative measures of bandwidth. Considering the group of subjects for which we estimate the lowest bandwidth, Appendix Table B.3 shows that the direct effect is significantly larger than the indirect effect for both interventions and for all three alternative measures (cp. Footnote 6).

⁸This result is qualitatively robust to using alternative measures of bandwidth. Considering the group of subjects

benefit from SIMPLIFICATION and GOODDEFAULTS primarily as a result of the indirect effect of the interventions on the quality of choices in the non-targeted domain.

Result 2. SIMPLIFICATION and GOODDEFAULTS increase the quality of choices in the non-targeted domain for a wide range of individuals. Subgroups with a relatively low bandwidth primarily benefit from an improvement of choices in the targeted task but exhibit a rather small indirect effect. In contrast, subgroups with a higher bandwidth benefit primarily from an improvement of choices in the non-targeted choice domain. These findings are in line with part (ii) of Hypotheses 1 and 2.

4.2 Reallocation of Cognitive Resources

Our results so far demonstrate that SIMPLIFICATION and GOODDEFAULTS improve the quality of individuals' choices in both the targeted and non-targeted choice domain. In our theoretical analysis in Section 3, we have argued that the interventions can yield these double dividends, since they allow subjects to withdraw some cognitive resources from the targeted task and use these resources to make better decisions in the non-targeted task. Our experiment allows us to directly test whether the interventions indeed have such an attention-releasing effect. As our measure for how individuals allocate cognitive resources between the targeted and non-targeted task, we use detailed data on subjects' attention spans. Specifically, for each subject and each round of the experiment, we measure the precise length of the time spans during which a subject attends to the screens for the targeted task and non-targeted task, respectively (cp. Figure D.5 in the appendix). Based on these attention spans, we can examine how SIMPLIFICATION and GOODDEFAULTS affect individuals' allocation of attention across the two choice domains. We focus on two main questions: First, do the interventions have an attention-releasing effect, i.e., do they lead to a reallocation of attention from the targeted to the non-targeted task? Second, does this reallocation of attention differ between the interventions?

Columns (1) and (2) of Table 2 depict treatment differences in the average length of attention spans for the targeted task. We find that both interventions cause participants to devote less attention to the targeted task. On average, subjects in BASELINE enter the targeted task for 17.7 seconds in each round of the experiment. This value decreases to 13.2 and 10.9 in GOODDEFAULTS and SIMPLIFICATION, respectively. Differences in average attention spans are statistically significant

for which we estimate the highest bandwidth, Appendix Table B.3 shows that the indirect effect is larger than the direct effect for both interventions and for all three alternative measures (cp. Footnote 6). While being qualitatively identical for all three alternative measures, the differences between the direct and indirect effects is only statistically significant for two of the measures (see post-estimation tests in Table B.3).

Table 2: Allocation of Attention

	Average Attention		No Attention	
	(1)	(2)	(3)	(4)
GOOD DEFAULTS	-4.553*** (0.727)	-4.484*** (0.719)	0.066*** (0.023)	0.070*** (0.020)
SIMPLIFICATION	-6.874*** (0.643)	-6.894*** (0.637)	-0.074*** (0.018)	-0.068*** (0.017)
Controls	No	Yes	No	Yes
Mean BASELINE	17.730	17.730	0.158	0.158
N	8400	8400	8400	8400
No. Subjects	840	840	840	840
R^2	0.065	0.103	0.025	0.110
SIMPLIFICATION= GOOD DEFAULTS	.0001	<0.0001	<0.0001	<0.0001

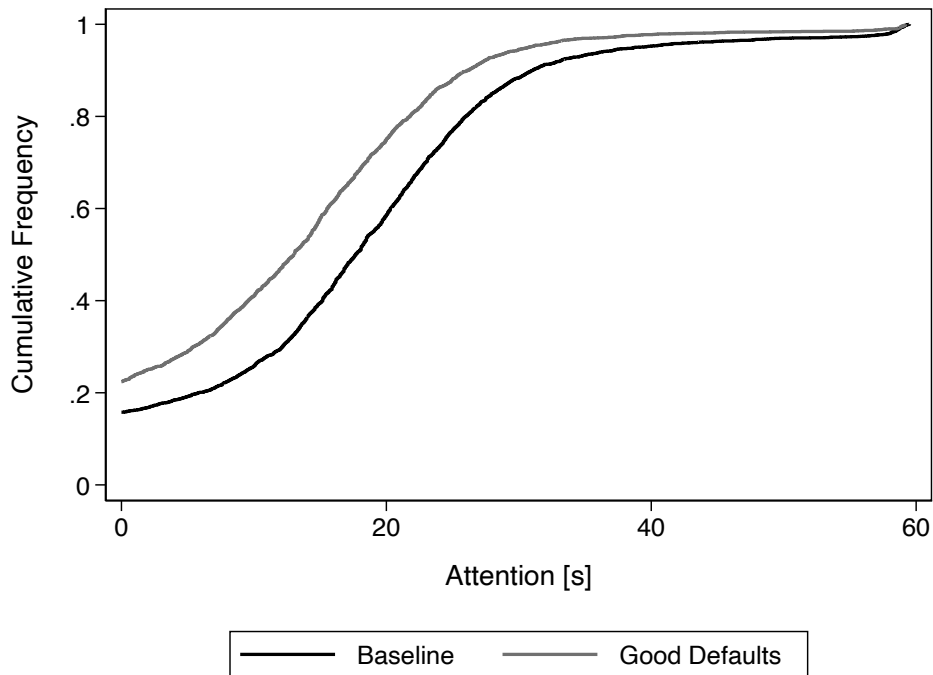
Note: The table reports results of OLS regressions of treatment differences in attention allocation. The dependent variable in Columns (1)–(2) is the average visual attention devoted to the targeted task in a given round of the experiment. The dependent variable in Columns (3)–(4) is an indicator equal to 1 if a subject did not pay any attention to the targeted task in a given round of the experiment and 0 otherwise. Control variables in Columns (2) and (4) include subjects’ age, gender, performance in the trial rounds in the targeted and non-targeted task, round-task fixed effects, and Raven score. The lower part of the table reports the p-values of post-estimation tests of the equality of selected coefficients (Wald tests). Robust standard errors, clustered at the subject level, are reported in parentheses. ***, **, * indicates significance at the 1%, 5%, and 10% level, respectively.

for all pairwise treatment comparisons. The results in Table 2, thus, demonstrate that SIMPLIFICATION or GOODDEFAULTS indeed have an attention-releasing effect. Since the overall time span is limited to 60 seconds in each round of the experiment, the withdrawal of attention from the targeted task directly implies an increase in attention devoted to the non-targeted task. Hence, subjects devote 4.6 seconds and 6.9 seconds more to the non-targeted task in GOODDEFAULTS and SIMPLIFICATION, respectively. This increase of resources devoted to the non-targeted task enables participants to make better choices in this task—resulting in the positive indirect effect documented in columns (3) and (4) of Table 1.

Result 3. *GOODDEFAULTS and SIMPLIFICATION have an attention-releasing effect: On average, individuals reduce the amount of attention devoted to the domain that is targeted and increase attention to the non-targeted domain.*

Although both interventions, on average, lead to a shift of attention from the targeted to the non-targeted task, we also observe systematic differences in how GOODDEFAULTS and SIMPLIFICATION affect individuals’ allocation of attention across tasks. A first important difference between the interventions is depicted in columns (3) and (4) of Table 2. These columns show treatment

Figure 5: Attention Allocation in GOODDEFAULTS and BASELINE



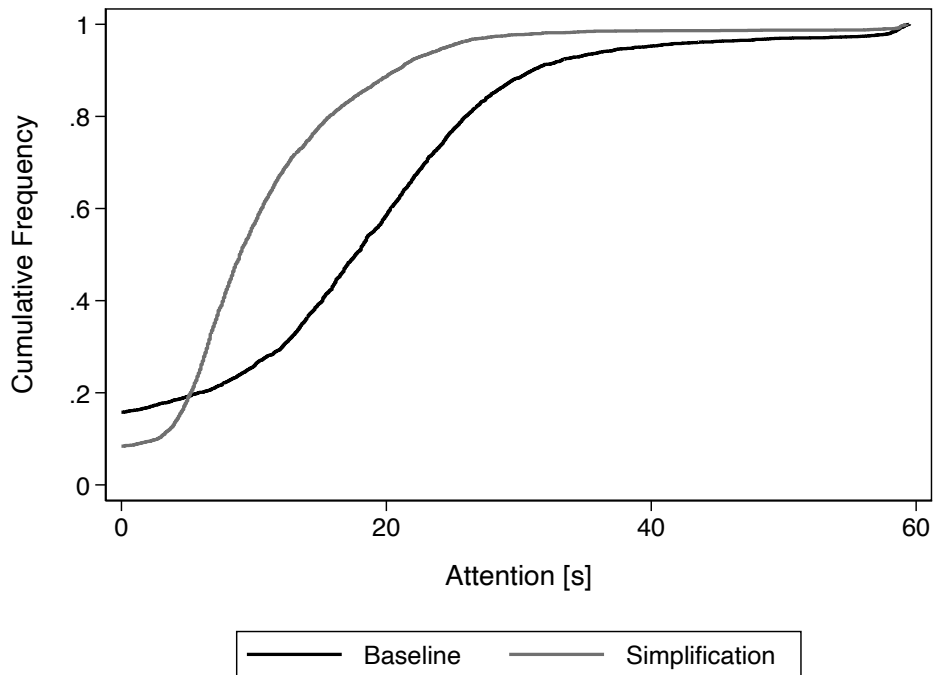
Note: The figure depicts cumulative distribution functions of the amount of attention devoted to the targeted task in GOODDEFAULTS and BASELINE (based on measures of attention spans at the subject-round level). Appendix Figure B.3 plots the corresponding figure using subject averages.

differences in individuals' propensity not to pay *any* attention to the targeted task. For GOODDEFAULTS, we find that the fraction of cases in which participants allocate zero attention to the targeted task increases from 15.8% in BASELINE to 22.4%. Hence, in 6.6% of the cases, subjects respond to GOODDEFAULTS by completely withdrawing their attention from the targeted task.

This contrasts sharply with the evidence for the SIMPLIFICATION intervention, for which we observe that the frequency of cases in which individuals pay zero attention to the targeted task is 7 percentage points *lower* than in BASELINE. Put differently, while GOODDEFAULTS cause an extensive-margin shift away from the targeted task, SIMPLIFICATION increases the likelihood that individuals engage with the task.

To shed further light on the patterns of attention reallocation, we examine treatment-specific distributions of attention spans. Figure 5 compares the cumulative distributions of attention devoted to the targeted task in BASELINE and GOODDEFAULTS. The figure documents that GOODDEFAULTS induce a shift in the entire distribution of attention, away from the targeted task. As a result, the distribution of attention spans observed in GOODDEFAULTS first-order stochastically dominates the one in BASELINE. Overall, GOODDEFAULTS lead to a rather uniform shift of cognitive resources

Figure 6: Attention Allocation in SIMPLIFICATION and BASELINE



Note: The figure depicts cumulative distribution functions of the amount of attention devoted to the targeted task in SIMPLIFICATION and BASELINE (based on measures of attention spans at the subject-round level). Appendix Figure B.4 plots the corresponding figure using subject averages.

from the targeted to the non-targeted domain.

Result 4. GOODDEFAULTS lead to a rather uniform shift of attention from the targeted task to the non-targeted task. In line with part (iii) of Hypothesis 1, this shift implies that the cumulative distribution of attention spans in BASELINE first-order stochastically dominates the one in GOODDEFAULTS.

Figure 6 compares the cumulative distributions of attention spans in SIMPLIFICATION and BASELINE. Unlike the uniform shift observed in Figure 5, we find that the two distributions intersect. In particular, as already documented in Columns (3) and (4) of Table 2, SIMPLIFICATION increases the likelihood that individuals engage with the targeted task at all. Hence, we observe fewer cases in which individuals devote zero or very little attention to the targeted task, and the c.d.f. of the attention spans in SIMPLIFICATION lies below the one from BASELINE for small attention spans. At the same time, SIMPLIFICATION allows other participants who thoroughly engage with the targeted task to shift attention from the targeted to the non-targeted task. As a consequence, the cumulative distribution of attention for SIMPLIFICATION is above the one for BASELINE for larger

attention spans. Therefore, the average reduction in attention in response to SIMPLIFICATION, which we documented in Table 2, is composed of two countervailing effects. In about 20% of the cases, participants increase their attention relative to BASELINE. In the remaining 80% of the cases, participants withdraw cognitive resources from the targeted task.

In sum, the analysis of attention distributions reveals that SIMPLIFICATION does not lead to a uniform shift of cognitive resources from the targeted to the non-targeted domain, as observed for GOODDEFAULTS. Instead, SIMPLIFICATION induces a compression of the attention distribution.⁹ Hence, while SIMPLIFICATION has, on average, an attention-releasing effect, it is an *attention-grabbing* intervention for individuals who would otherwise not pay any attention or devote only minimal amounts of attention to the targeted task.

Result 5. SIMPLIFICATION *leads to a compression of the attention spans devoted to the targeted task, relative to BASELINE. In line with part (iii) of Hypothesis 2, the cumulative distribution of the attention spans in BASELINE intersects with the one from SIMPLIFICATION exactly once.*

Although SIMPLIFICATION and GOODDEFAULTS have almost identical effects on average choice quality in the targeted and the non-targeted domain, Results 4 and 5 document rather different effects on how participants (re)allocate scarce cognitive resources. GOODDEFAULTS induce a rather uniform shift of cognitive resources from the targeted to the non-targeted domain. In contrast, SIMPLIFICATION leads to a compression in the observed attention spans, with an attention-grabbing effect for some individuals and an attention-releasing effect for others. As a result of this compression in the attention distribution, the cognitive resources devoted to the targeted task are less dispersed in SIMPLIFICATION. For instance, when comparing the distribution of attention spans in SIMPLIFICATION, BASELINE, and GOODDEFAULTS, we observe that the standard deviation of attention spans in SIMPLIFICATION is about 33% lower than in BASELINE and about 24% lower than in GOODDEFAULTS.¹⁰ This observation also helps to understand why the positive indirect effects of SIMPLIFICATION on the quality of choices differ across the population more strongly than those of GOODDEFAULTS (see Figure 4). The compression in the attention distribution implies that subjects with poor choice quality benefit much less from the positive indirect effects than their counterparts that exhibit higher choice quality. On the contrary, the relatively uniform shift in attention due to

⁹The histograms of attention underlying the cumulative density functions, depicted in Figure B.2 in the appendix, also illustrates this effect. While SIMPLIFICATION strongly increases the number of cases in which participants devote about 5-10 seconds to the targeted task compared to BASELINE, the intervention decreases both the participants' propensity to pay zero attention to the task and their propensity to engage with the task for more than ~15 seconds.

¹⁰The reduction in standard deviations is highly statistically significant (sd-test based on subject averages, $p < 0.01$ for SIMPLIFICATION vs. BASELINE, $p < 0.01$ for SIMPLIFICATION vs. GOODDEFAULTS).

GOODDEFAULTS also yields rather uniform indirect effects across the population.

5 Implications for the Evaluation of Nudge Interventions

Several of our results have immediate consequences for the evaluation of nudge interventions. This section illustrates these consequences and discusses how our findings help to enhance our understanding of the average treatment effects and the distributional consequences of nudges.

Results 1 and 2 show the critical role of both direct and indirect effects in understanding the overall impact and distributional consequences of interventions such as SIMPLIFICATION and GOODDEFAULTS. In particular, the magnitude of the indirect effect is non-negligible. In fact, in our setting the indirect effect on choice quality in the non-targeted domains is as large as the direct effect in the targeted domain (cp. Table 1). An evaluation that focuses solely on the domain that is targeted by the intervention is thus prone to underestimate the overall value of attention-releasing interventions. Notably, this observation also implies that if a policy analyst does not detect a positive direct effect of an (attention-releasing) intervention, this does not necessarily mean that the intervention is ineffective. For instance, while reducing administrative burden in social programs might not further increase take-up rates due to existing high baseline levels, the indirect benefits in other domains of people’s lives, accruing from the reallocation of cognitive resources, can be substantial.

Second, policy evaluations that focus exclusively on the targeted domain could lead to a biased assessment of the distributional effects of an intervention (cp. Result 2). In particular, the direct effects depicted in Figure 4 suggest that the policies considered in our experiment primarily help those who are prone to make rather poor decisions. However, the figure also shows that the indirect effects of the interventions—which are primarily reaped by individuals with higher bandwidth—have a dampening effect on these distributional effects. Put differently, attention-releasing interventions may be less progressive than it appears when only assessing the targeted domain. Furthermore, this dampening effect should be especially pronounced for policies that reduce cognitive burden and simplify administrative processes. Such policies lead to a compression in the attention distribution on non-targeted domains which may bias an evaluation of distributional effects. For instance, simplifying application processes for social programs might appear to disproportionately benefit lower-ability individuals. However, when also considering indirect effects accruing to higher-ability individuals, the overall policy’s impact on reducing income disparities may be less pronounced than initially thought.

Third, Results 3–5 demonstrate how nudge interventions affect the allocation of cognitive resources across domains. This analysis is useful to distinguish two broad categories of interventions

according to the attentional reactions they trigger. *Attention-grabbing interventions* like reminders or active choice policies cause individuals to devote more attention to the task targeted by the intervention. In doing so, they lead to a shift of cognitive resources away from other choice domains, with potentially negative indirect effects on choice quality in non-targeted domains. In contrast, *attention-releasing* interventions including good defaults, simplification, and measures to reduce bureaucratic barriers, free up scarce cognitive resources, and can produce positive indirect effects on non-targeted choice domains. In situations where both intervention types offer similar direct effects, attention-releasing nudges might therefore be the preferable option.

6 Conclusion

Nudge interventions have become an important policy instrument to improve choices in various domains of people’s lives. At the same time, the effectiveness of nudges is the subject of an intense, ongoing debate. This paper uses a simple experimental paradigm to analyze the existence and nature of indirect positive effects of two popular classes of nudges—simplifications of choice processes and the implementation of high-quality defaults—on choice domains that are not directly targeted by the interventions. Applications of these interventions include reductions in administrative burdens (OIRA 2023), the use of natural language, simplifying retirement savings decisions (Beshears et al. 2013), assistance for applying for social benefits (Finkelstein and Notowidigdo 2019), and improving default health care plans (Brot-Goldberg et al. 2023). We document that such interventions can induce not only direct effects, but also substantial positive indirect effects. Neglecting these indirect effects can lead policy evaluations to substantially underestimate the overall effectiveness of such attention-releasing nudges. Moreover, we demonstrate that the relative importance of indirect and direct effects differ systematically across subgroups. As a consequence, neglecting indirect effects in the evaluation of interventions can also lead to a biased assessment of their distributional consequences.

The implications for the evaluation of nudge interventions extend to a broad set of policies and therefore may serve as a catalyst for further research. In particular, it seems crucial to develop a more comprehensive understanding of the estimation of indirect, attention-driven effects of policy interventions in field settings. One approach could be to conceptualize how we can use additional data to estimate lower and upper bounds for indirect effects of interventions. While it seems impossible to elicit detailed data on the entire set of relevant non-targeted domains, it might be feasible to obtain measures for those non-targeted domains, which are most likely to be affected (e.g., focus on other health-related behaviors when evaluating a health intervention as in Trachtman 2023).

Alternatively, more detailed survey measures on the attentional consequences of an intervention, i.e., the extent to which a specific policy is attention-releasing versus attention-grabbing, could help to bound or approximate aggregate indirect effects as well as subgroup-specific indirect effects. A first step in this direction would be the elicitation of procedural or survey measures that examine how strongly individuals engage with a specific choice process, and whether this engagement increases or decreases with a particular policy intervention.

References

- ABELER, J. AND S. JÄGER (2015): “Complex tax incentives,” *American Economic Journal: Economic Policy*, 7, 1–28.
- ALONSO, R., I. BROCAS, AND J. CARILLO (2014): “Resource Allocation in the Brain,” *Review of Economic Studies*, 81, 501–534.
- ALTMANN, S., A. GRUNEWALD, AND J. RADBRUCH (2022): “Interventions and Cognitive Spillovers,” *The Review of Economic Studies*.
- BARR, M. S., S. MULLAINATHAN, AND E. SHAFIR (2008): “Behaviorally informed financial services regulation,” *Asset Building Program Policy Paper*.
- BARTOŠ, V., M. BAUER, J. CHYTILOVÁ, AND F. MATĚJKA (2016): “Attention Discrimination: Theory and Field Experiments with Monitoring Information Acquisition,” *American Economic Review*, 106, 1437–75.
- BENARTZI, S., J. BESHEARS, K. L. MILKMAN, C. R. SUNSTEIN, R. H. THALER, M. SHANKAR, W. TUCKER-RAY, W. J. CONGDON, AND S. GALING (2017): “Should Governments Invest More in Nudging?” *Psychological Science*, 28, 1041–1055.
- BERTRAND, M., S. MULLAINATHAN, AND E. SHAFIR (2004): “A behavioral-economics view of poverty,” *American Economic Review*, 94, 419–423.
- BESHEARS, J., J. J. CHOI, D. LAIBSON, AND B. C. MADRIAN (2013): “Simplification and saving,” *Journal of Economic Behavior and Organization*, 95, 130–145.
- BETTINGER, E. P., B. T. LONG, P. OREOPOULOS, AND L. SANBONMATSU (2012): “The Role of Application Assistance and Information in College Decisions: Results from the H&R Block Fafsa Experiment,” *The Quarterly Journal of Economics*, 127, 1205–1242.
- BHARGAVA, S. AND D. MANOLI (2015): “Psychological frictions and the incomplete take-up of social benefits: Evidence from an IRS field experiment,” *American Economic Review*, 105, 3489–3529.
- BOCK, O., I. BAETGE, AND A. NICKLISCH (2014): “hroot: Hamburg Registration and Organization Online Tool,” *European Economic Review*, 71, 117–120.

- BRONCHETTI, E. T., J. B. KESSLER, E. B. MAGENHEIM, D. TAUBINSKY, AND E. ZWICK (2023): “Is Attention Produced Optimally? Theory and Evidence From Experiments With Bandwidth Enhancements,” *Econometrica*, 91, 669–707.
- BROT-GOLDBERG, Z., T. LAYTON, B. VABSON, AND A. Y. WANG (2023): “The Behavioral Foundations of Default Effects: Theory and Evidence from Medicare Part D,” *American Economic Review*, 113, 2718–2758.
- CAPLIN, A. AND M. DEAN (2013): “Behavioral Implications of Rational Inattention with Shannon entropy,” *NBER Working Paper No. 19318*.
- CAPLIN, A., M. DEAN, AND D. MARTIN (2011): “Search and Satisficing,” *American Economic Review*, 101, 2899–2922.
- CARROLL, G. D., J. J. CHOI, D. LAIBSON, B. C. MADRIAN, AND A. METRICK (2009): “Optimal Defaults and Active Decisions,” *The Quarterly Journal of Economics*, 124, 1639–1674.
- CASTRO, J. F., D. VELÁSQUEZ, A. BELTRÁN, AND G. YAMADA (2022): “The Direct and Indirect Effects of Messages on Tax Compliance: Experimental Evidence from Peru,” *Journal of Economic Behavior and Organization*, 203, 483–518.
- CHATER, N. AND G. LOEWENSTEIN (2022): “The i-frame and the s-frame: How focusing on individual-level solutions has led behavioral public policy astray,” *Behavioral and Brain Sciences*, 1–60.
- CHEN, D., M. SCHONGER, AND C. WICKENS (2016): “oTree – An Open-source Platform for Laboratory, Online, and Field Experiments,” *Journal of Behavioral and Experimental Finance*, 9, 88 – 97.
- DEAN, M. AND N. NELIGH (2023): “Experimental Tests of Rational Inattention,” *Journal of Political Economy*, 131, 3415–3461.
- DELLAVIGNA, S. AND E. LINOS (2022): “RCTs to Scale: Comprehensive Evidence From Two Nudge Units,” *Econometrica*, 90, 81–116.
- DERTWINKEL-KALT, M., H. GERHARDT, G. RIENER, F. SCHWERTER, AND L. STRANG (2022): “Concentration Bias in Intertemporal Choice,” *The Review of Economic Studies*, 89, 1314–1334.
- DESMOND, M. (2023): *Poverty, by America*, Crown.
- DUFLO, E. (2012): “Human values and the design of the fight against poverty,” *Tanner Lectures*, 2012, 1–55.
- ERICSON, K. M., T. LAYTON, A. MCINTYRE, AND A. SACARNY (2023): “Reducing Administrative Barriers Increases Take-up of Subsidized Health Insurance Coverage: Evidence from a Field Experiment,” *NBER Working Paper No. 30885*.

- ERICSON, K. M. AND A. STARC (2016): “How product standardization affects choice: Evidence from the Massachusetts Health Insurance Exchange,” *Journal of Health Economics*, 50, 71–85.
- FEHR, E. AND K. WU (2023): “Obfuscation in competitive markets,” Tech. rep., Working Paper.
- FINKELSTEIN, A. AND M. J. NOTOWIDIGDO (2019): “Take-up and targeting: Experimental evidence from SNAP,” *The Quarterly Journal of Economics*, 134, 1505–1556.
- GABAIX, X., D. LAIBSON, G. MOLOCHE, AND S. WEINBERG (2006): “Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model,” *American Economic Review*, 96, 1043–1068.
- HECKMAN, J. J. AND R. LANDERSØ (2021): “Lessons from Denmark about inequality and social mobility,” Tech. rep., National Bureau of Economic Research.
- HEIDHUES, P., J. JOHNNEN, AND B. KŐSZEGI (2021): “Browsing versus studying: A pro-market case for regulation,” *The Review of Economic Studies*, 88, 708–729.
- HERD, P., H. HOYNES, J. MICHENER, AND D. MOYNIHAN (2023): “Introduction: Administrative Burden as a Mechanism of Inequality in Policy Implementation,” *The Russell Sage Foundation Journal of the Social Sciences*, 9, 1–30.
- JOHNNEN, J. AND B. T. K. LEUNG (2022): “Distracted from Comparison: Product Design and Advertisement with Limited Attention,” *CEPR Discussion Paper No. DP17234*.
- JOHNSON, E. J. AND D. GOLDSTEIN (2003): “Do Defaults Save Lives?” *Science*, 302, 1338–1339.
- JOHNSON, E. J., J. W. PAYNE, D. A. SCHKADE, AND J. R. BETTMAN (1989): “Monitoring Information Processing and Decisions: The Mouselab System,” *Working Paper 89-4, Center for Decision Studies, Fuqua School of Business*.
- KOCH, A., D. MØNSTER, AND J. NAFZIGER (2023): “Nudging in complex environments,” *IZA Discussion Paper No. 16137*.
- LOEWENSTEIN, G. AND N. CHATER (2017): “Putting Nudges in Perspective,” *Behavioral Public Policy*, 1, 26–53.
- MADRIAN, B. C. AND D. F. SHEA (2001): “The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior,” *The Quarterly Journal of Economics*, 116, 1149–1187.
- MAIER, M., F. BARTOS, T. STANDLEY, D. R. SHANKS, A. J. L. HARRIS, AND E.-J. WAGENMAKERS (2022): “No Evidence for Nudging after Adjusting for Publication Bias,” *Proceedings of the National Academy of Sciences of the United States of America*, 119, e2200300119.
- MANI, A., S. MULLAINATHAN, E. SHAFIR, AND J. ZHAO (2013): “Poverty impedes cognitive function,” *Science*, 341, 976–980.
- MARTIN, D. (2017): “Rational Inattention in Games: Experimental Evidence,” *Working Paper*.

- MEDINA, P. C. (2020): “Side Effects of Nudging: Evidence from a Randomized Intervention in the Credit Card Market,” *The Review of Financial Studies*, forthcoming.
- MERTENS, S., M. HERBERZ, U. J. J. HAHNEL, AND T. BROSCHE (2022): “The effectiveness of nudging: A meta-analysis of choice architecture interventions across behavioral domains,” *Proceedings of the National Academy of Sciences*, 119, e2107346118.
- MRKVA, K., N. A. POSNER, C. REECK, AND E. J. JOHNSON (2021): “Do nudges reduce disparities? Choice architecture compensates for low consumer knowledge,” *Journal of Marketing*, 85, 67–84.
- NAFZIGER, J. (2020): “Spillover effects of nudges,” *Economics Letters*, 190, 109086.
- OIRA (2023): “Burden Reduction Report,” *The Office of Information and Regulatory Affairs*, <https://www.whitehouse.gov/wp-content/uploads/2023/07/OIRA-2023-Burden-Reduction-Report.pdf>.
- ROBERTS, J. L. (2017): “Nudge-proof: Distributive justice and the ethics of nudging,” *Michigan Law Review*, 116, 1045.
- SHARAFI, Z. (2023): “Poverty and perseverance: The detrimental effect of poverty on effort provision,” *Journal of Development Economics*, 162, 103040.
- SUNSTEIN, C. (2021): *Sludge: What Stops Us from Getting Things Done and What to Do about It*, MIT Press.
- SUNSTEIN, C. R. (2022a): “The distributional effects of nudges,” *Nature human behaviour*, 6, 9–10.
- (2022b): “Sludge audits,” *Behavioural Public Policy*, 6, 654–673.
- SZASZI, B., A. HIGNEY, A. CHARLTON, A. GELMAN, I. ZIANO, B. ACZEL, D. G. GOLDSTEIN, D. S. YEAGER, AND E. TIPTOP (2022): “No Reasons to Expect Large and Consistent Effects of Nudge Interventions,” *Proceedings of the National Academy of Sciences of the United States of America*, 119, e2200732119.
- THALER, R. H. AND C. R. SUNSTEIN (2003): “Libertarian Paternalism,” *American Economic Review*, 93, 175–179.
- (2008): *Nudge*, New Haven & London: Yale University Press.
- TRACHTMAN, H. (2023): “Does Promoting one Healthy Behavior Detract from Others? Evidence from a Field Experiment,” *American Economic Journal: Applied Economics*, forthcoming.
- WANG, J. T., M. SPEZIO, AND C. F. CAMERER (2010): “Pinocchio’s Pupil: Using Eyetracking and Pupil Dilation to Understand Truth Telling and Deception in Sender-Receiver Games,” *American Economic Review*, 100, 984–1007.

Appendix A Theoretical Setup and Derivation of Hypotheses

Our behavioral predictions in Section 3 are informed by an illustrative model that builds on the framework presented in Altmann et al. (2022). This section presents the formal arguments proofs of the propositions underlying the behavioral predictions.

Appendix A.1 The Model

Suppose there is a mass one of individuals. Each individual is endowed with a fixed stock of cognitive resources $X \in \mathbb{R}_{>0}$ such that the distribution of X in the population is described by the cumulative distribution function $F(X)$, which we assume to have full support and no mass points. Individuals face two tasks, a non-targeted task NT , and a targeted task T to which they can allocate resources $x_{NT}, x_T \geq 0$, such that $x_{NT} + x_T \leq X$. Allocating an amount of resources x_{NT} to the task NT results in the likelihood of $\pi_{NT}(x_{NT})$ to correctly solve the task and obtain a payoff of u_{NT} . We assume that π_{NT} is strictly increasing, strictly concave, and continuously differentiable. Allocating resources x_T to Task T results in a likelihood $\pi_T(x_T, \phi, \lambda) = \max\{\lambda, \tilde{\pi}_T(x_T, \phi)\}$ to solve the task and obtain a payoff u_T , where $\lambda \in (0, 1]$ corresponds to the quality of the default option, $\phi \in \mathbb{R}^+$ to the difficulty of the task, and $\tilde{\pi}_T(0, \phi) = 0 \forall \phi$. Moreover, we assume that there exists a threshold $\bar{x}(\phi)$ such that $\tilde{\pi}_T(x_T, \phi) = 1 \forall x_T \geq \bar{x}(\phi)$. Therefore, the targeted task can be perfectly solved if an individual dedicates enough resources to it. For all ϕ , function $\tilde{\pi}_T$ is continuous for all x_T and strictly increasing, differentiable and strictly concave in x_T for all $x_T < \bar{x}(\phi)$ with $\lim_{x_T \rightarrow \bar{x}(\phi)^-} \frac{\partial \tilde{\pi}_T(x_T, \phi)}{\partial x_T} > 0$. Finally, we assume that $\frac{\partial^2 \tilde{\pi}_T(x_T, \phi)}{\partial \phi \partial x_T} < 0 \forall x_T, \phi$. Therefore, the difficulty of the task affects the slope of π_T with respect to x_T . If the difficulty of the task ϕ increases, this slope decreases, which means that a marginal increase in cognitive resources devoted to the targeted task is less effective in solving this task. For an individual with cognitive resources X the decision problem is given by:

$$\begin{aligned} \max_{x_{NT}, x_T} u(x_{NT}, x_T) &= \pi_{NT}(x_{NT})u_{NT} + \pi_T(x_T, \phi, \lambda)u_T \\ \text{s.t. } x_{NT} + x_T &\leq X \end{aligned} \quad (2)$$

The optimal allocation of cognitive resources, which we denote by $(x_{NT}^*(\lambda, \phi, X), x_T^*(\lambda, \phi, X))$, depends on the shape of π_T and π_{NT} , the relative payoffs for the tasks (u_T/u_{NT}), the default quality λ , the difficulty of the task ϕ , and the stock of cognitive resources X . To derive clear-cut predictions, we assume the following:

Condition 1. For all ϕ , the functions π_T and π_{NT} satisfy :

$$(i) u_{NT} \geq \frac{\frac{\partial \tilde{\pi}_T(0, \phi)}{\partial x_T}}{\pi_{NT}(0)} u_T$$

$$(ii) \lim_{x \rightarrow \infty} \pi'_{NT}(x) = 0$$

The condition imposes two restrictions on the profit functions. The first part is a sufficient condition to ensure that the incentives in Task NT are strong enough so that it is never optimal for subjects to allocate all their resources to Task T . The second part implies that at some point allocating further resources to Task NT only yields a negligible increase in the likelihood of correctly solving this task. In analogy to the main text, we will start by analyzing the behavioral changes induced by GOODDEFAULTS. The distribution of resources allocated to the targeted task implied by the distribution F is denoted by $G_{\phi,\lambda}(\cdot)$. The following proposition derives the results underlying Hypothesis 1.

Proposition 1. *Suppose that Condition 1 holds and let $\lambda_G > \lambda_B$. Then,*

(i) $x^*_{NT}(\lambda, \phi, X)$ and thus $\pi_{NT}(x^*_{NT}(\lambda, \phi, X))$ are weakly increasing with λ for all X . Moreover, $\frac{du(x^*_{NT}(\lambda, \phi, X), x^*_T(\lambda, \phi, X))}{d\lambda} \geq 0$ for all X .

(ii) there exist two thresholds $X^1(\lambda_B, \lambda_G, \phi) < X^2(\lambda_B, \lambda_G, \phi)$ such that:

$$\pi_{NT}(x^*_{NT}(\lambda_G, \phi, X)) - \pi_{NT}(x^*_{NT}(\lambda_B, \phi, X)) \geq \pi_T(x^*_T(\lambda_G, \phi, X)) - \pi_T(x^*_T(\lambda_B, \phi, X)) \text{ if } X > X^2(\lambda_B, \lambda_G, \phi)$$

and

$$\pi_{NT}(x^*_{NT}(\lambda_G, \phi, X)) - \pi_{NT}(x^*_{NT}(\lambda_B, \phi, X)) \leq \pi_T(x^*_T(\lambda_G, \phi, X)) - \pi_T(x^*_T(\lambda_B, \phi, X)) \text{ if } X < X^1(\lambda_B, \lambda_G, \phi).$$

(iii) $G_{\phi,\lambda_B}(x) \leq G_{\phi,\lambda_G}(x)$ for all x .

The different parts of Proposition 1 inform the respective parts of Hypothesis 1. Note that the ambiguity of the treatment effect with respect to the average payoffs in the targeted domain (part (i) of Hypothesis 1) can also be seen from the arguments presented in Appendix A.2 below. Equation 3 implies that a group of individuals should reduce the resources allocated to Task T . The overall extent of this effect depends on unobserved shape of π_{NT} and the unobserved distribution $F(X)$. Whether this effect dominates the positive effect of higher quality default options on the choice quality in the targeted domain is thus not determined in our model.

Next, we turn to the implications of SIMPLIFICATION. The following proposition derives the results underlying Hypothesis 2.

Proposition 2. *Suppose that Condition 1 holds and let $\phi_S < \phi_B$. Then,*

(i) $\pi_T(x^*_T(\lambda, \phi, X))$ and $u(x^*_{NT}(\lambda, \phi, X), x^*_T(\lambda, \phi, X))$ are weakly decreasing with ϕ for all X .

(ii) there exist two thresholds $X^3(\lambda, \phi_B, \phi_S) < X^4(\lambda, \phi_B, \phi_S)$ such that:

$$\pi_{NT}(x^*_{NT}(\lambda, \phi_S, X)) - \pi_{NT}(x^*_{NT}(\lambda, \phi_B, X)) \geq \pi_T(x^*_T(\lambda, \phi_S, X)) - \pi_T(x^*_T(\lambda, \phi_B, X)) \text{ if } X > X^4(\lambda, \phi_B, \phi_S)$$

and

$$\pi_{NT}(x^*_{NT}(\lambda, \phi_S, X)) - \pi_{NT}(x^*_{NT}(\lambda, \phi_B, X)) \leq \pi_T(x^*_T(\lambda, \phi_S, X)) - \pi_T(x^*_T(\lambda, \phi_B, X)) \text{ if } X < X^3(\lambda, \phi_B, \phi_S).$$

(iii) there exist two thresholds $x^1(\lambda, \phi_B, \phi_S)$ and $x^2(\lambda, \phi_B, \phi_S)$ with $x^1(\lambda, \phi_B, \phi_S) \leq x^2(\lambda, \phi_B, \phi_S)$ such that $G_{\phi_B,\lambda}(x) > G_{\phi_S,\lambda}(x)$ if $x < x^1(\lambda, \phi_B, \phi_S)$ and $G_{\phi_B,\lambda}(x) < G_{\phi_S,\lambda}(x)$ if $x > x^2(\lambda, \phi_B, \phi_S)$.

Each part of Proposition 2 informs the respective part of Hypothesis 2 in the main body of the paper. There is one point worth discussing. The formal result in part (iii) of Proposition 2 states that there are two thresholds. These thresholds will be identical for a wide range of distributions. We therefore decided to include the somewhat stricter statement in the Hypothesis that they have to be identical, which simplifies the exposition in the main text. Clearly, if we find evidence in favor of part (iii) of Hypothesis 2 it also corroborates the weaker statement in Proposition 2. Finally, note that the ambiguous payoff effect on the targeted domain (part (i)) again follows from the arguments presented in Appendix A.2 below, because some individuals will reduce their attention to that task, while others will increase their attention to the targeted domain.

Appendix A.2 Proofs

Proof of Proposition 1

We start by deriving the optimal allocation of cognitive resources. Since $\pi_{NT}(x_{NT})$ is strictly increasing, all cognitive resources will be exhausted. Consider first the auxiliary maximization problem:

$$\begin{aligned} \max_{x_T} \pi_{NT}(X - x_T)u_{NT} + \tilde{\pi}_T(x_T, \phi)u_T \\ \text{s.t. } x_T \in [0, X] \end{aligned}$$

This objective function is strictly concave in x_T . Denote its maximizer by $\tilde{x}_T, \tilde{x}_{NT}$, where we omit the arguments to simplify notation. Due to part (i) of Condition 1, \tilde{x}_{NT} will be strictly positive and $\tilde{x}_T < X$ for all levels of X . The maximizer of the individual's maximization problem:

$$\begin{aligned} \max_{x_T} \pi_{NT}(X - x_T)u_{NT} + \pi_T(x_T, \phi, \lambda)u_T \\ \text{s.t. } x_T \in [0, X] \end{aligned}$$

is then given by $\tilde{x}_T, \tilde{x}_{NT}$ if and only if

$$\pi_{NT}(X - \tilde{x}_T)u_{NT} + \tilde{\pi}_T(\tilde{x}_T, \phi)u_T \geq \pi_{NT}(X)u_{NT} + \lambda u_T$$

and by $x_T = 0, x_{NT} = X$ otherwise. Note first that the right-hand side is strictly larger than the left-hand side for $X = 0$, because $\lambda u_T > 0$. Moreover, the optimality of \tilde{x}_T implies that the slope of the left-hand side with respect to X is weakly larger than the one of the right-hand side. Formally,

the derivative of the left hand side is given by:

$$\begin{aligned} & \pi'_{NT}(X - \tilde{x}_T)u_{NT} + \frac{d\tilde{x}_T}{dX} \left[\frac{d\tilde{\pi}_T(\tilde{x}_T, \phi)}{d\tilde{x}_T} u_T - \pi'_{NT}(X - \tilde{x}_T)u_{NT} \right] \\ & = \pi'_{NT}(X - \tilde{x}_T)u_{NT}. \end{aligned}$$

This in turn is larger than $\pi'_{NT}(X)u_{NT}$, which is the derivative of the right-hand side, due to the concavity of π_{NT} . Therefore, there exists a strictly positive threshold $X^*(\phi, \lambda) \in \mathbb{R}^+$ such that an individual will abstain from devoting any cognitive resources to the targeted task if and only if $X \leq X^*(\phi, \lambda)$. Note that this cut-off is also smaller than infinity due to Condition 1 part (ii). $X^*(\phi, \lambda)$ is implicitly defined by:

$$\pi_{NT}(X^* - \tilde{x}_T)u_{NT} + \tilde{\pi}_T(\tilde{x}_T, \phi)u_T = \pi_{NT}(X^*)u_{NT} + \lambda u_T$$

Moreover, there is a second threshold $X^{**}(\phi, \lambda)$ such that all subjects with $X \geq X^{**}(\phi, \lambda)$ dedicate $x_T = \bar{x}(\phi)$ to the targeted task and solve it perfectly. This threshold exists because $\lim_{x \rightarrow \infty} \pi'_{NT}(x) = 0$, while we have that π_T is strictly increasing for all $x_T < \bar{x}(\phi)$.

Next, consider how $X^*(\phi, \lambda)$ varies in λ . Implicitly differentiating the cutoff X^* yields:

$$\begin{aligned} & \frac{\partial X^*(\phi, \lambda)}{\partial \lambda} > 0 \\ \Leftrightarrow & \frac{u_T - \frac{d\tilde{x}_T}{d\lambda} \left[\frac{d\tilde{\pi}_T(\tilde{x}_T, \phi)}{d\tilde{x}_T} u_T - \pi'_{NT}(X^* - \tilde{x}_T)u_{NT} \right]}{\pi'_{NT}(X^* - \tilde{x}_T)u_{NT} + \frac{d\tilde{x}_T}{dX} \left[\frac{d\tilde{\pi}_T(\tilde{x}_T, \phi)}{d\tilde{x}_T} u_T - \pi'_{NT}(X^* - \tilde{x}_T)u_{NT} \right] - \pi'_{NT}(X^*)u_{NT}} > 0 \\ \Leftrightarrow & \frac{u_T}{[\pi'_{NT}(X^* - \tilde{x}_T) - \pi'_{NT}(X^*)]u_{NT}} > 0, \end{aligned} \tag{3}$$

where the terms in brackets in the second line are zero due to the optimality of \tilde{x}_T . The last line is given due to the concavity of π_{NT} . For subjects with $X \neq X^*(\phi, \lambda)$ a marginal shift in λ does not affect $x_T^*(\lambda, \phi, X)$. Overall, $x_T^*(\lambda, \phi, X)$ is thus weakly decreasing in λ for all X . As a consequence, $G_{\phi, \lambda_B}(x) \leq G_{\phi, \lambda_G}(x)$ holds for all x , which proves part (iii) of Proposition 1.

The first half of part (i) of Proposition 1 follows immediately from the above arguments, since all resources that are withdrawn from Task T will be allocated to Task NT . The second half of part (i) is implied since individuals allocate their resources optimally.

For part (ii) of the Proposition 1 consider first the case of small X . Let $X^1(\lambda_B, \lambda_G, \phi)$ be equal to $X^*(\phi, \lambda_B)$ such that consumers with $X \leq X^1(\lambda_B, \lambda_G, \phi)$ dedicate zero resources to the targeted task under λ_B and λ_G . For these subjects $\pi_T(x_T^*(\lambda_G, \phi, X)) > \pi_T(x_T^*(\lambda_B, \phi, X))$ holds because $\lambda_G > \lambda_B$. At the same time, there is no indirect effect because the allocation of cognitive resources remains unaffected, i.e., $\pi_{NT}(x_{NT}^*(\lambda_G, \phi, X)) = \pi_{NT}(x_{NT}^*(\lambda_B, \phi, X)) \quad \forall X \leq X^1(\lambda_B, \lambda_G, \phi)$.

Second, consider the case of sufficiently large X . Define $X^2(\lambda_B, \lambda_G, \phi)$ implicitly such that $\pi_T(x_T^*(\lambda_B, \phi, X^2)) = \lambda_G$. $X^2(\lambda_B, \lambda_G, \phi)$ exists because of the mean value theorem and it holds that $X^2(\lambda_B, \lambda_G, \phi) < X^*(\phi, \lambda_G)$. For all subjects with $X \in (X^2(\lambda_B, \lambda_G, \phi), X^*(\phi, \lambda_G)]$ we have $x_T^*(\lambda_G, \phi, X) < x_T^*(\lambda_B, \phi, X)$. At the same time $x_T^*(\lambda_G, \phi, X) = x_T^*(\lambda_B, \phi, X)$ holds for $X > X^*(\phi, \lambda_G)$. Hence, the effect on choice quality is non-positive in the targeted domain for $X \geq X^2(\lambda_B, \lambda_G, \phi)$. However, there will be a positive indirect effect for $X \in (X^2(\lambda_B, \lambda_G, \phi), X^*(\phi, \lambda_G)]$ because they will reallocate resources to the non-targeted task. Hence, the indirect effect is weakly larger than the direct effect for all individuals with $X \geq X^2(\lambda_B, \lambda_G, \phi)$.

Proof of Proposition 2

We start from the optimal allocation $x_T^*(\lambda, \phi, X), x_{NT}^*(\lambda, \phi, X)$, which is derived in the proof of Proposition 1. First, consider subjects with $X < X^*(\phi, \lambda)$. These individuals ignore Task T and a marginal shift in the difficulty will not affect $x_T^*(\lambda, \phi, X)$. Implicit differentiation yields:

$$\begin{aligned} \frac{\partial X^*(\phi, \lambda)}{\partial \phi} &> 0 \\ \Leftrightarrow \frac{-\frac{\partial \bar{\pi}_T(\tilde{x}_T, \phi)}{\partial \phi} u_T - \frac{d\tilde{x}_T}{d\lambda} \left[\frac{d\bar{\pi}_T(\tilde{x}_T, \phi)}{d\tilde{x}_T} u_T - \pi'_{NT}(X^* - \tilde{x}_T) u_{NT} \right]}{\pi'_{NT}(X^* - \tilde{x}_T) u_{NT} + \frac{d\tilde{x}_T}{dX} \left[\frac{d\bar{\pi}_T(\tilde{x}_T, \phi)}{d\tilde{x}_T} u_T - \pi'_{NT}(X^* - \tilde{x}_T) u_{NT} \right] - \pi'_{NT}(X^*) u_{NT}} &> 0 \\ \Leftrightarrow \frac{-\frac{\partial \bar{\pi}_T(\tilde{x}_T, \phi)}{\partial \phi} u_T}{[\pi'_{NT}(X^* - \tilde{x}_T) - \pi'_{NT}(X^*)] u_{NT}} &> 0 \end{aligned}$$

where the terms in brackets in the second line are zero due to the optimality of \tilde{x}_T . The last line is given due to the concavity of π_{NT} . Hence, the fraction of individuals that disregard task T is higher in BASELINE than in SIMPLIFICATION, i.e., $G_{\phi_B, \lambda}(0) > G_{\phi_S, \lambda}(0)$. All subjects with $X > X^{**}(\phi_S, \lambda)$ allocate $\bar{x}(\phi_S)$ to Task T in SIMPLIFICATION. Since $F(X)$ has full support and $\bar{x}(\phi_S) < \bar{x}(\phi_B)$, $G_{\phi_B, \lambda}(\bar{x}(\phi_S)) < 1 = G_{\phi_S, \lambda}(\bar{x}(\phi_S))$. As a consequence of the mean value theorem, there exist two thresholds $x^1(\lambda, \phi_B, \phi_S)$ and $x^2(\lambda, \phi_B, \phi_S)$ with $x^1(\lambda, \phi_B, \phi_S) \leq x^2(\lambda, \phi_B, \phi_S)$ such that $G_{\phi_B, \lambda}(x) > G_{\phi_S, \lambda}(x)$ if $x < x^1(\lambda, \phi_B, \phi_S)$ and $G_{\phi_B, \lambda}(x) < G_{\phi_S, \lambda}(x)$ if $x > x^2(\lambda, \phi_B, \phi_S)$. This concludes the proof of part (ii) of Proposition 2.

Part (i) of Proposition 2 also follows from the arguments above. In particular, the payoffs in the targeted domain for individuals with $X \notin [X^*(\phi, \lambda), X^{**}(\phi, \lambda)]$ will not be affected by a marginal decrease in ϕ . However, $X^*(\phi, \lambda)$ and $X^{**}(\phi, \lambda)$ decrease if the task becomes simpler. Moreover, all individuals with $X \in (X^*(\phi, \lambda), X^{**}(\phi, \lambda))$ also increase x_T^* due to $\frac{\partial^2 \bar{\pi}_T(x_T, \phi)}{\partial \phi \partial x_T} < 0$. Hence, $\pi_T(x_T^*(\lambda, \phi, X))$ decreases with ϕ for all X . The second statement in part(i) follows immediately from the optimality of resource allocation.

For part (ii) of Proposition 2 consider first the case of small X . Let $X^3(\lambda_B, \lambda_G, \phi)$ be equal

to $X^*(\phi_B, \lambda)$ such that consumers with $X \leq X^3(\lambda_B, \lambda_G, \phi)$ dedicate zero resources to the targeted task in BASELINE. Since $X^*(\phi_B, \lambda) > X^*(\phi_S, \lambda)$, we have $\pi_T(x_T^*(\lambda, \phi_S, X)) \geq \pi_T(x_T^*(\lambda, \phi_B, X))$ for $X \in [X^*(\phi_S, \lambda), X^3(\lambda_B, \lambda_G, \phi)]$. Hence, $\pi_{NT}(x_{NT}^*(\lambda, \phi_S, X)) \leq \pi_{NT}(x_{NT}^*(\lambda, \phi_B, X))$ for all $X < X^3(\lambda_B, \lambda_G, \phi)$. Second consider the case of sufficiently large X . Define $X^4 = X^{**}(\phi_B, \lambda)$. For all individuals with $X \geq X^4(\lambda_B, \lambda_G, \phi)$ it holds that $x_T^*(\lambda, \phi_S, X) = \bar{x}(\phi_S) < x_T^*(\lambda, \phi_B, X) = \bar{x}(\phi_B)$. However, $\pi_T(x_T^*(\lambda, \phi_S, X)) = \pi_T(x_T^*(\lambda, \phi_B, X))$ for $X \geq X^4(\lambda_B, \lambda_G, \phi)$, because they solve the targeted task perfectly in both conditions. Hence, there is no positive effect on choice quality in the targeted domain but a positive indirect effect.

Appendix B Additional Tables and Figures

Table B.1: Descriptives

	Treatments			Kruskal-Wallis
	BASELINE	GOOD DEFAULTS	SIMPLIFICATION	p-value
Age of Subject	25.18 (4.98)	27.07 (7.40)	25.75 (6.08)	0.00
Female	0.69 (0.46)	0.60 (0.50)	0.65 (0.51)	0.06
Overall Grade of Abitur	2.04 (0.60)	2.12 (0.65)	2.03 (0.65)	0.28
Last Math Grade in School	2.11 (0.90)	2.29 (0.98)	2.21 (1.02)	0.14
Ability Math Task	0.89 (0.22)	0.90 (0.22)	0.91 (0.21)	0.73
Ability Memory Task	0.84 (0.20)	0.88 (0.18)	0.87 (0.19)	0.02
Raven Test Score	6.81 (2.53)	6.74 (2.35)	6.64 (2.45)	0.77
N	287	263	290	

Note: The table depicts mean values of basic subject-level characteristics across treatments. Standard deviations are reported in parentheses.

Table B.2: Direct and Indirect Effects of the Interventions: Heterogeneity

	Rank within Treat x Targeted Task	
	Targeted	Non-Targeted
	(1)	(2)
Good Defaults × 1st Quar.	10.899*** (0.628)	2.630*** (0.896)
Good Defaults × 2nd Quar.	3.829*** (0.421)	5.240*** (1.034)
Good Defaults × 3rd Quar.	1.359*** (0.228)	4.302*** (1.003)
Good Defaults × 4th Quar.	0.042 (0.042)	4.456*** (1.082)
Simplification × 1st Quar.	9.958*** (0.897)	1.152 (0.958)
Simplification × 2nd Quar.	5.764*** (0.317)	4.584*** (1.075)
Simplification × 3rd Quar.	3.092*** (0.083)	4.734*** (0.916)
Simplification × 4th Quar.	0.042 (0.042)	5.927*** (0.894)
N	8400	8400
No. Subjects	840	840
R^2	0.268	0.054
<i>Statistical Tests of Differences between Direct and Indirect Effect within Group</i>		
	Rank within Treat x Targeted Task	
	Diff.	p-value
GOOD DEFAULTS, 1st Quart.	-8.2692	<0.0001
GOOD DEFAULTS, 4th Quart.	4.4137	<0.0001
SIMPLIFICATION, 1st Quart.	-8.8062	<0.0001
SIMPLIFICATION, 4th Quart.	5.8855	<0.0001

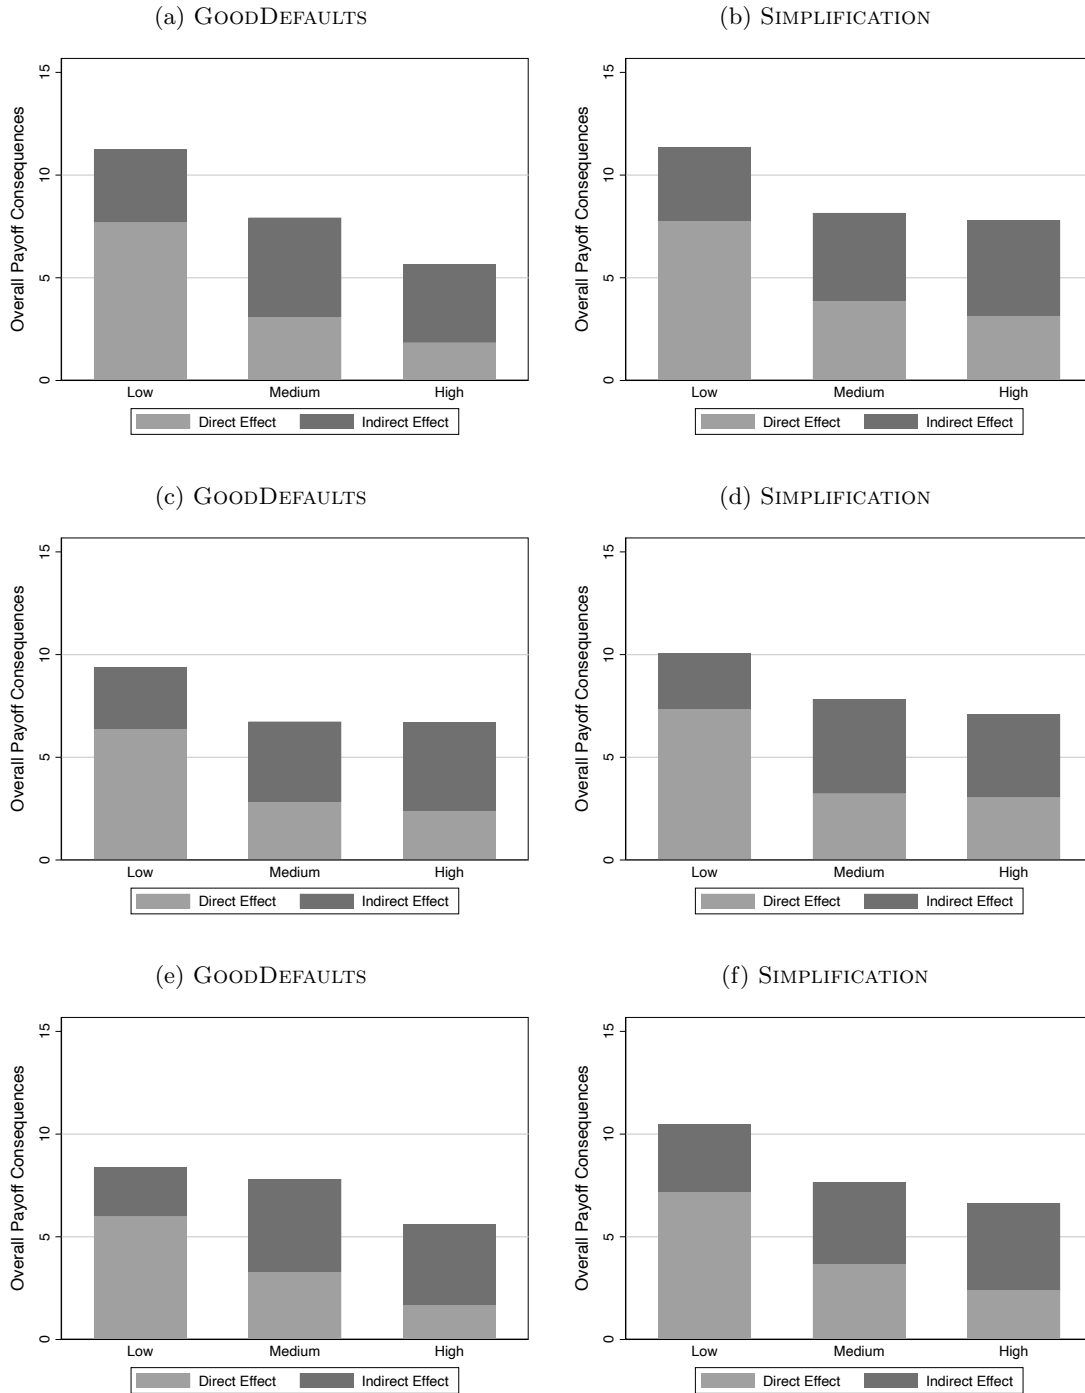
Note: The table reports results of OLS regressions of treatment differences in payoffs in the targeted and non-targeted task for different groups. The dependent variable in Column (1) is the payoff in the targeted task in a given round of the experiment. The dependent variable in Column (2) is the payoff in the non-targeted task in a given round of the experiment. The groups are categorized based on their performance quartiles in the targeted task within each treatment. 1st Quartile refers to the lowest performance level. To ensure that groups are of equal size, ties in performance are broken randomly. Robust standard errors, clustered at the subject level, are reported in parentheses. ***, **, * indicates significance at the 1%, 5%, and 10% level, respectively. The lower part of the table reports the differences between coefficients in column (1) and (2) for selected groups. We also report p-values from post-estimation tests for the equality of the coefficients that represent the direct effect and the indirect effect within these groups (Wald tests).

Table B.3: Direct and Indirect Effects of the Interventions: Alternative Measures of Bandwidth

	Raven Test		Trial Rounds		Ressource Factor	
	Targeted	Non-Targeted	Targeted	Non-Targeted	Targeted	Non-Targeted
	(1)	(2)	(3)	(4)	(5)	(6)
Good Defaults \times Low	7.718*** (1.158)	3.510*** (1.107)	6.392*** (1.004)	2.980*** (1.070)	6.014*** (0.988)	2.378** (1.003)
Good Defaults \times Medium	3.114*** (0.727)	4.810*** (0.827)	2.857*** (1.024)	3.876*** (1.084)	3.323*** (0.866)	4.487*** (0.909)
Good Defaults \times High	1.858** (0.754)	3.784*** (0.989)	2.416*** (0.687)	4.289*** (0.711)	1.694** (0.687)	3.919*** (0.759)
Simplification \times Low	7.770*** (1.239)	3.555*** (0.962)	7.363*** (1.100)	2.693*** (0.980)	7.188*** (1.087)	3.289*** (0.839)
Simplification \times Medium	3.874*** (0.790)	4.278*** (0.817)	3.254*** (1.115)	4.572*** (0.970)	3.683*** (0.903)	3.968*** (0.921)
Simplification \times High	3.144*** (0.735)	4.656*** (0.891)	3.069*** (0.720)	4.008*** (0.693)	2.414*** (0.729)	4.196*** (0.671)
N	8400	8400	8400	8400	8400	8400
No. Subjects	840	840	840	840	840	840
R^2	0.059	0.055	0.050	0.099	0.056	0.135
<i>Statistical Tests of Differences between Direct and Indirect Effect within Group</i>						
	Raven Test		Trial Rounds		Ressource Factor	
	Diff.	p-value	Diff.	p-value	Diff.	p-value
GOOD DEFAULTS, Low	-4.2074	.0133	-3.4119	.0368	-3.6355	.0248
GOOD DEFAULTS, High	1.9259	.1557	1.8727	.0808	2.2251	.0477
SIMPLIFICATION, Low	-4.2147	.0081	-4.6696	.0027	-3.8990	.0093
SIMPLIFICATION, High	1.5120	.2241	0.9388	.3718	1.7823	.0915

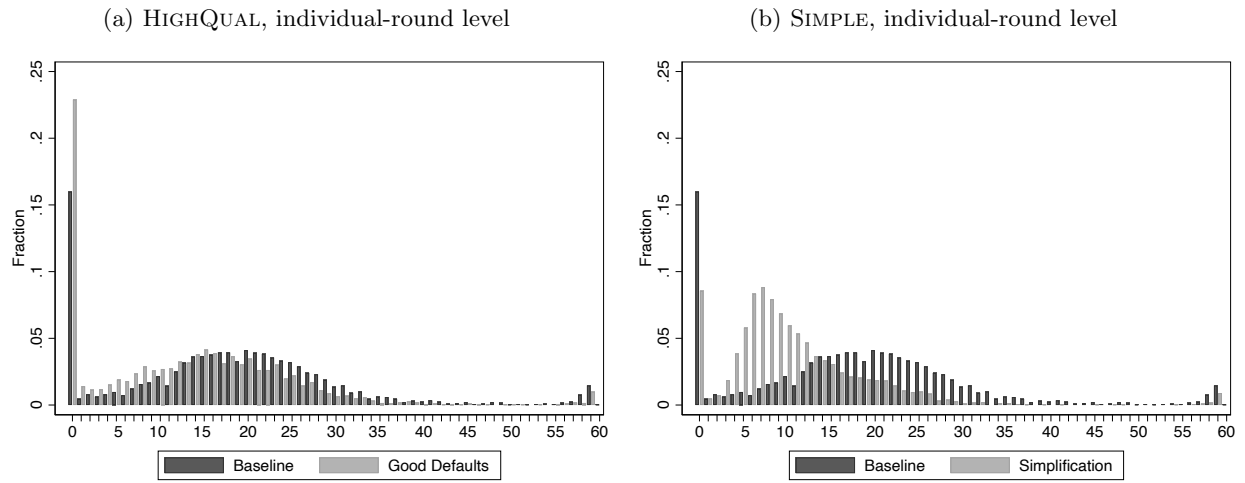
Note: The table reports results of OLS regressions of treatment differences in payoffs in the targeted and non-targeted task for different groups. The dependent variable in Columns (1), (3), and (5) is the payoff in the targeted task in a given round of the experiment. The dependent variable in Columns (2), (4), and (6) is the payoff in the non-targeted task in a given round of the experiment. The groups are categorized based on their performance in a Raven test in columns (1) and (2), their performance in the trial rounds in columns (3) and (4), and using a factor analysis extracting a common factor from all available measures from the trial rounds and Raven score (including Raven score, performance in trial rounds, time needed in trial rounds, number of moves needed in trial rounds) in columns (5) and (6). For each variable we form groups which are approximately of similar size. Robust standard errors, clustered at the subject level, are reported in parentheses. ***, **, * indicates significance at the 1%, 5%, and 10% level, respectively. The lower part of the table reports the differences between coefficients in the first and second column for each measure for selected groups. We also report p-values from post-estimation tests for the equality of the coefficients that represent the direct effect and the indirect effect within these groups (Wald tests).

Figure B.1: Relative Importance of the Direct and Indirect Effect: Alternative Measures of Bandwidth



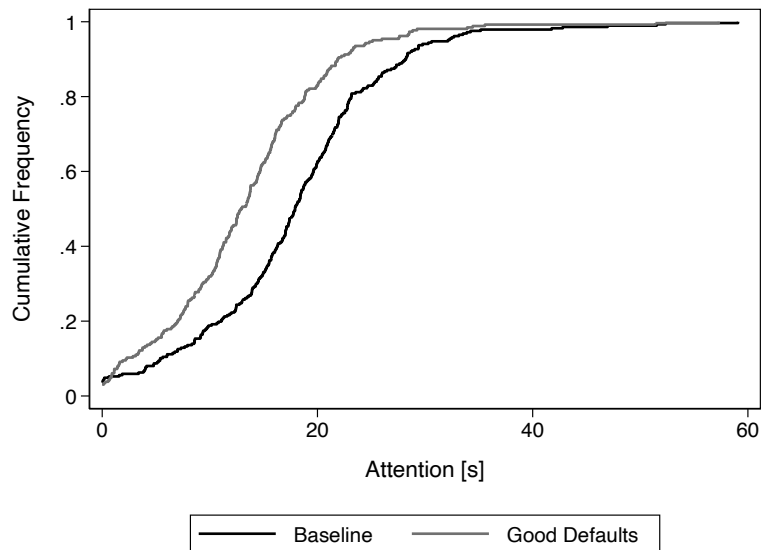
Note: The figure illustrates the payoff consequences of GOODDEFAULTS (panel (a),(c), and (e)) and SIMPLIFICATION (panel (b), (d), and (f)) in the targeted and non-targeted domain. The bars depict the average payoff differences between the treatments and the BASELINE for the respective groups. The groups are categorized based on three alternative measures of individuals' stock of cognitive resources. In particular, we use subjects' relative performance in a Raven test in panel (a) and (b) and their relative performance in the trial rounds in panel (c) and (d). In panel (e) and (f), we base the categorization on a factor analysis extracting a common factor from all available measures from the trial rounds and the Raven score (including Raven score, performance in trial rounds, time needed in trial rounds, number of moves needed in trial rounds).

Figure B.2: Histograms of Attention in BASELINE and Treatments



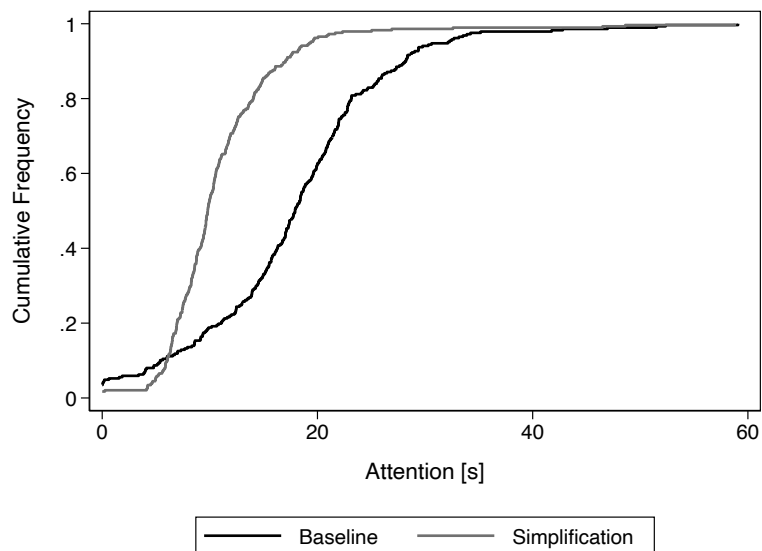
Note: The figure shows the probability distribution of attention spans in GOODDEFAULTS and BASELINE in panel (a) and SIMPLICATION and BASELINE in panel (b). The unit of observation is the subject-round level.

Figure B.3: Attention Allocation in GOODDEFAULTS and BASELINE



Note: The figure shows the cumulative distribution of attention spans in GOODDEFAULTS and BASELINE. The unit of observation is the subject level average.

Figure B.4: Attention Allocation in SIMPLIFICATION and BASELINE



Note: The figure shows the cumulative distribution of attention spans in SIMPLIFICATION and BASELINE. The unit of observation is the subject level average.

Appendix C Instructions of the Experiment

First Part

Please read the following information carefully. On the next page you will find a comprehension question that you have to answer correctly in order to start part 1 of the experiment.

The first part of the experiment consists of two rounds. In each round, you work on a memory game. You have 60 seconds per round to work on the memory game.

The memory consists of ten pairs of cards, each with two-digit numbers, as shown in the illustration below. At the beginning of a round, the cards are face down, and you can flip them over with a mouse click. You can always turn over two cards per turn. If you find a pair - i.e., both cards show the same number - the cards remain turned over.

The aim of the task is to find as many pairs as possible. **For each pair found you will receive 3 Cents.**

[Screenshot of Task]

Comprehension Question [on second screen]:

What are your earnings in Euro, if you find all 10 pairs in a round?

Second Part

Please read the following information carefully. On the next page you will find a comprehension question that you have to answer correctly in order to start part 2 of the experiment.

The second part of the experiment consists of two rounds. In each round, you will be shown 5 options as in the figure below. Each option corresponds to a math problem. Your task is to choose the option with the **highest sum**. To select an option, please click on the field in front of the respective option. To do so, you will have 60 seconds. If you have selected the correct option, your earnings in this round will be 30 Cents. If you have not selected an option or have selected a wrong option, your earnings in this round will be 0 cents.

[Screenshot of Task]

Comprehension Question [on second screen]:

What are your earnings in Euro, if you select the correct option in the math problem?

Third Part

Please read the following information carefully. On the next page you will find comprehension questions that you have to answer correctly in order to start part 3 of the experiment.

The third part of the experiment consists of **ten rounds**. In each round you can work on both the task from part 1 and the task from part 2. You have 60 seconds to work on both tasks.

As in Part 1 of the experiment, the **memory** in each round consists of 10 pairs of cards. Your task is to find as many pairs as possible.

In the **math problem** from part 2 of the experiment, your goal is again to select the option with the highest sum. [next sentence only in BASELINE and GOODDEFAULTS:] Each option of the math problem again consists of **4 numbers** that you have to add up. [next sentence only in SIMPLIFICATION:] In contrast to part 2, each option option of the math problem **now only consists of 2 numbers** that you have to add up.

For the math problem from part 2, one of the 5 options is already preselected in each round. If you do not work on this task, the preselected option will count as your answer. The preselection is determined by the computer with the help of a random generator. The probability that the preselected option also corresponds to the correct option, is [next part only in BASELINE and SIMPLIFICATION:] 20%. This means that **on average in 2 out of 10 rounds** the preselected option is correct, and on average in 8 out of 10 rounds it is wrong. [next part only in GOODDEFAULTS:] 80%. This means that **on average in 8 out of 10 rounds** the preselected option is correct, and on average in 2 out of 10 rounds it is wrong.

In the course of a round, you can decide for yourself when to work on which task. There are two ways to switch from one task to another. You can either click on the tabs in the upper part of the screen or on the “**Switch task**” button at the bottom of the screen. At the top of the screen you

can see how much working time is left for the current round.

Your earnings in a round is the sum of the earnings achieved for both tasks. For each pair found in the memory task, you get 3 Cents. If the selected option is correct in the math problem, you earn 30 Cents for that task. If a wrong option is selected, you will receive 0 cents for this task.

Example: Suppose you choose the option with the highest sum in the math problem and find 8 pairs of cards in the memory. In this case, you earn 30 Cent for the math problem and 8×3 Cent = 24 Cent for the memory. So your earnings for this round would be 54 Cent.

Comprehension Questions [on second screen]:

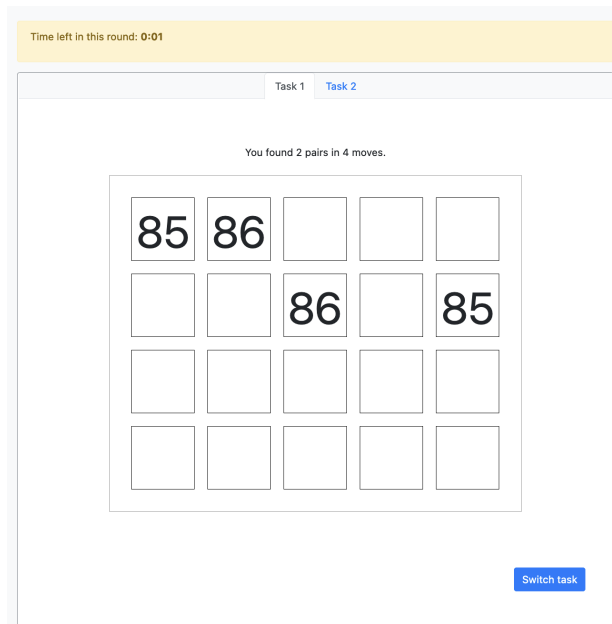
What is the likelihood that the preselected option in the math problem is the correct option?

How many numbers do you have to add up for each option in the math problem?

Appendix D Screenshots of the Experiment

Figure D.5: Screenshots of the Experiment (BASELINE condition)

(a) Screen 1: Non-targeted Task



(b) Screen 2: Targeted Task

