

Effects on Knowledge of Nudging Citizens with Information

Morten Jakobsen

&

Søren Serritzlew

Aarhus University

March 2015

Abstract

Nudging¹ is a policy instrument that has proven effective in improving matters such as citizens' health, education, and civic behavior by affecting their decisions. One common strategy is to provide citizens with information that nudges them toward different choices. We argue that this strategy can also affect citizens' knowledge, which can become the basis for future choices. In this manner, providing information can make the nudging initiative stick. Based on a field experiment that solves problems of endogeneity, we test the effect of an information treatment on three types of knowledge. The treatment consists of information on reading strategies designed to nudge parents to read with their children, randomly assigned to a study population of 1438 families with young children. We show how nudging by providing information does in fact increase parents' knowledge. We also show that this effect is equally effective for parents with different socio-economic backgrounds.

¹ The authors would like thank Maiken Kjær Milthers and Line Scott Nesbit for excellent research assistance. Maiken and Line were responsible for data collection, and they contributed during this process with many excellent ideas. They also provided very valuable assistance in developing the booklet. The authors would also like to thank the municipality of Aarhus for an excellent partnership and for great openness and flexibility. We received help and goodwill from many at the Department of Children and Young People. Special thanks to Catharina Damgaard, Jens Møller Hald, Dennis Møller Hansen, and Morten Hjortskov Larsen. Finally, thanks to Solveig Gaarsmand from the National Association of School Parents for valuable help in developing the booklet.

Introduction

Research on how government can affect citizens' choices and behavior has flourished in recent decades. Government often seeks ways to improve citizens' health, education, civic behavior, and so forth. This has been accompanied by research on the effectiveness of different administrative strategies for doing so. In this stream of research, nudging has assumed an important role as an alternative to incentive strategies (see Hood, 2010; John et al., 2011; Thaler & Sunstein, 2008). Hence, nudging is about affecting people's choices and behavior by organizing the context in which they make decisions; for example, by altering default rules or providing information and social cues without forbidding any options or significantly changing economic incentives (Thaler & Sunstein, 2008, pp. 3–6).

One prominent nudging strategy consists of providing people with information that nudges them towards making better choices for themselves and society (John et al., 2011, p. 9). Field experiments have documented that such information nudging can substantially affect how citizens behave (e.g., Calzolari & Nardotto 2011; Hastings, Welden, & Weinstein, 2007; John et al., 2011, pp. 75–84).

However, nudging citizens with information might also have important effects other than immediate changes in behavior. Specifically, we argue that nudging by providing information also affects what people know about the given issue. This is important for two reasons. First, although new knowledge can be forgotten, it can potentially affect behavior for an extended period of time. Citizens must often make repeated choices regarding a matter (e.g., in relation to education). If government follows a nudging strategy, it has to nudge people repeatedly. Yet, if the strategy also enhances knowledge on a matter, such knowledge may work as the basis for future choices. In this sense, providing information can make the nudging initiative stick. Second, increasing citizen knowledge can be a goal in itself; for instance, citizen knowledge may be a valuable input in strategies aimed at improving citizen

choices through deliberation and participation as a supplement or alternative to nudging strategies (see, e.g., the “think” strategies presented by John et al., 2011). In many cases, it is thus valuable for government to increase the knowledge of citizens regarding the matter about which they are making choices. We discuss these arguments in the theoretical section.

We then empirically investigate the effect of information nudging on knowledge. While the empirical studies on how government can nudge people’s behavior is more voluminous, little empirical evidence exists as to whether government can enhance knowledge among citizens. Moreover, although the question of an information-nudging effect on knowledge is simple, it is difficult to settle empirically. Whether it is part of an attempt to nudge citizens or another type of government strategy, governments rarely provide information randomly; they do so with a clear purpose, such as changing the behavior of a specific group of citizens. This creates a methodological problem of endogeneity.

The study at hand uses a randomized field experiment to solve this endogeneity problem. Specifically, we use an experiment on information to families about education services. The experiment was conducted in cooperation with the Danish municipality of Aarhus. 1438 families were randomly assigned to either a treatment group (which received information designed to nudge parents to read with their children) or a control group. This approach has a major advantage compared to the existing studies of the effects of information: It combines an experimental approach to address the endogeneity problem with an examination of the effect in the real world (information is actually provided by government and knowledge measured independently from the treatment). Furthermore, the high number of subjects allows us to examine whether knowledge effects are heterogeneous across different groups of citizens.

Nudging

Nudging as a government strategy for changing citizens' choices and behavior in life has attracted massive attention in recent years, particularly since Thaler and Sunstein's (2008) contribution to this debate. Government often seeks ways to affect people's behavior regarding education, health, transport, household economy, and so forth; and for many years, governments and researchers have devoted much attention to rules and economic incentives as means for producing such effects. Nudging strategies provide a different approach to influencing citizens' behavior. This approach is based on an assumption of bounded rationality (John et al., 2011, pp. 12–14). People have imperfect information, cognitive limitations, and limited time available when making decisions; thus, perfectly rational choices are impossible (Jones, 2001, p. 54ff; Simon, 1957, p. 241 ff.). Nudging takes this into account and designs the choice context in ways that, according to the advocates of nudging, help people overcome these limitations and make better choices for themselves and society.

Following Thaler and Sunstein (2008, p. 3), nudging is about organizing the context in which people make decisions; that is, to be a “choice architect.” Thus, according to Thaler and Sunstein (2008, p. 6):

A nudge, as we will use the term, is any aspect of the choice architecture that alters people's behavior in a predictable way without forbidding any options or significantly changing their economic incentives. To count as a mere nudge, the intervention must be easy and cheap to avoid. Nudges are not mandates. Putting the fruit at eye level counts as a nudge. Banning junk food does not.

An example of a government nudging strategy is to set default rules in ways that increase the likelihood of certain choices (e.g., enrolling people into organ donor registers by default so they must opt out instead of opting in; see Thaler and Sunstein, 2008, pp. 175–182).

Another example is to commit people to a certain behavior by asking them to make a pledge (e.g., a promise to donate books for new school libraries; see John et al., 2011).

A prominent nudging strategy consists in providing citizens with information. Thaler and Sunstein (2008, pp. 202–203) use an example from a study on school choice by Hastings, Welden, & Weinstein (2007) to illustrate this strategy. In 2002, the Charlotte-Mecklenburg Public Schools District (in Charlotte, NC) implemented a school choice plan. Each family received a school choice book, approximately 100 pages, including descriptions of the almost 200 schools (including elementary, middle, and high schools) and their programs. However, previous analyses of the choice program showed that “low-income families place lower weights on academics when choosing schools and that students of such families often had test score losses as a result of attending their first-choice school” (Hastings, Welden, & Weinstein 2007, p. 12). An intervention providing parents with simplified information sheets outlining the schools’ average test scores and odds of admission was therefore undertaken to see if such information influenced school choice. It did. Using a randomized experiment, Hastings, Welden, & Weinstein (2007) show that when provided with the simplified information, the parents put substantially more weight on test scores when choosing schools.

In other cases, information may contain social cues intended to nudge people. For example, John et al. (2011, pp. 75–84) examined whether the choice to sign a petition was higher when people received information indicating that a large number of people had already signed the petition.

As argued in the next section, information nudging may have an important effect not only on citizens’ behavior, which is the main focus of the existing research on government nudging strategies, but also on their knowledge about the matter in question. Moreover, we also argue that, for different reasons, increases in citizen knowledge is important to government.

Nudging and Knowledge

The literature on nudging focuses on citizens' behavior and choices. In many cases, however, it may be valuable for governments to affect not only the immediate behavior of citizens but also their knowledge about a given issue. First, citizens must often make repeated choices.

Education is a good example. An individual's education process can typically span more than a decade, including a vast number of choices, such as the choice of school or daily choice of how to develop skills and how much time to invest. If a government follows a nudging strategy, it has to repeatedly nudge people by building detailed, comprehensive nudging schemes. Yet, if the strategy also enhances knowledge on the given matter (e.g., skill formation), such knowledge may work as the basis for future choices. In that sense, if it affects knowledge, nudging by providing information can make the nudging initiative "stickier".

Second, citizen knowledge may be a valuable input in strategies aimed at improving citizen choices through deliberation and participation, which may be a supplement or an alternative to nudging strategies. As shown by Roberts (2004), there is an ongoing interest among administrators and researchers in using public deliberation and citizen participation in the public administration process. John et al. (2011) use the label "think strategies" to denote strategies in which people's behavior is affected by deliberation processes. Hence, in such strategies, citizens share, discuss, and evaluate their opinions and preferred behavior with others. This deliberation influences their future choices. Hence, in addition to creating institutional spaces for deliberation, government may seek to increase citizens' knowledge on the matters under deliberation as well as knowledge about the deliberative and participatory processes.

Which nudging strategies can potentially increase citizens' knowledge? For some of the strategies, we would expect little if any effect of the strategy on citizens' knowledge. These strategies include nudging by setting default rules, altering the physical context in which people make choices (for an example of the latter, see Thaler and Sunstein 2008, p. 1). Such techniques do not present people with new insights, arguments, or the like. People are not usually even aware that someone may be trying to influence their behavior when such strategies are applied.

For other types of nudging strategies, we would expect a greater potential for an increase in peoples' knowledge, including the provision of information in some form that produces new insight and arguments. The information may be provided in writing, visually, or orally. Information nudging is sometimes used to provide social cues to people in order to push them in a certain direction by means of social pressure. John et al. (2011) present a nudge in which people were given information about how many people had signed a petition (e.g., to introduce a tax on plastic bags) before being asked if they would sign. Such information can produce knowledge about how other people think and behave, but it does not, in itself, produce knowledge on, for instance, damage from plastic bags or the effect of a tax on plastic bags (although knowledge may be increased if people, as a consequence of a nudge, start *seeking* knowledge about the matter).

Nudging by providing information can also in itself transmit insights, arguments, research findings, and experiences that provide knowledge to the people exposed to it. The above-mentioned nudge presented by Hastings, Welden, & Weinstein (2007) concerning school choice is an example of this. This information had an impact on the parents' choices. However, it may also have provided them with knowledge about the local school system, how the different schools performed, and which factors to observe when making educational choices that might also influence their future choices. This end is not reached by providing a

huge mass of information. Instead, the aim is to convert complex information, which is unusable due to its complexity, into information that is simple and can be activated by people. The next sections empirically test whether nudging citizens by providing simple information, which we know from the Hastings, Welden, and Weinstein study (2007), can affect behavior, has a causal effect on knowledge.

Design: How to Address Endogeneity Problems in Empirical Studies of Information Provision

A key methodological problem is endogeneity. First, a public organization's decision to initiate the information provision is likely to be influenced by citizens' current information and knowledge levels (Weiss & Tschirhart, 1994, 85). Hence, the public organization is more likely to provide information if it perceives a need for it. This reverse causal effect, or simultaneity, makes it difficult to estimate the effect of information provision on citizen knowledge by simple comparisons of public organizations that provide information with organizations that do not.

Second, simple correlations between the information provision by public organizations and citizen knowledge are likely to be biased by sorting effects. Citizens with certain characteristics live in certain jurisdictions. For example, a positive effect of information provision may be inflated if high-SES citizens (who are expected to have higher levels of knowledge) sort into jurisdictions in which the public organizations are more likely to send out information, and vice versa.

Laboratory experiments offer an obvious solution to the methodological challenges related to endogeneity but do not say anything about the effect of information provision by actual public organizations. Using survey experiments has a similar drawback, since the information would be provided by the researcher and the effect measured within seconds after

the stimulus is delivered. To meet these methodological challenges, we conducted a randomized field experiment. This approach enables us to handle problems of endogeneity (this is addressed by the random assignment procedure described below) while examining the effects on knowledge in an actual public administration setting.

Primary education was chosen as a test case, as it constitutes an important public service that includes service users (parents and children) with varying backgrounds and characteristics. Hence, we are able to examine whether the effect of nudging by providing information is similar for different types of citizens. This also offers a good example of an area in which service users must make repeated choices, including major solitary choices, such as the choice of school and parents' daily investment of time in children's skill formation (e.g., homework, reading).

To ensure that the information material provided to the service users would be realistic and sent by an appropriate source, the experiment was conducted in collaboration with the municipality of Aarhus.² Ten schools in the municipality were included in the experiment (the inclusion procedure is detailed below), and classes in grades 1–3 were randomly assigned to either treatment or control groups. Members of the treatment group received information directly from the municipality; hence, the information was provided by the municipality and the schools as a normal intervention. Members of the control group did not receive this information.³ The experiment included 1438 families. Next, we describe the information intervention, the random assignment, the estimation approach, and how we measured the dependent variable (knowledge).

² The experiment has also been used (for other purposes) in Nesbit & Milthers (2012), Jakobsen & Thomsen (2015) and Thomsen & Jakobsen (2015).

³ Field experiments can be ethically problematic if they affect citizens adversely. In this field experiment, it was ensured that the information intervention was additional to the information that is normally distributed to parents, meaning that the control group subjects did not receive less information than they normally would have. Moreover, we expected the effect of the extra information to be beneficial to the treatment group citizens (or, in the worst case, have no effect at all). We therefore see no ethical problems in conducting the experiment.

The intervention was simple and consisted of sending a booklet by mail directly to the treatment group parents. The booklet (four pages) contained information on children's reading skills development and advice on how parents can contribute to this development. In accordance with the idea of nudging, the information was chosen to be – and presented in a manner that was – readily accessible and simple. The information and advice came from professionals in reading skills development, based on dialogic reading, which has proven effective (Whitehurst et al., 1988). One type of information concerns reading comprehension: what it is, how it is developed, and its importance in developing skills other than reading. Also included is advice on how to contribute to the development of children's reading skills through various types of dialogic reading. The distributor of the booklet was the Municipality of Aarhus.

The drawback of field experiments is that they usually only allow for a test of one or a few specific types of treatments, as the cluster randomization reduces the power of the test. Our experiment is no exception. This implies that we are only able to estimate the effect of providing the information in one particular way. Additionally, setting up a field experiment requires the permission of the responsible public entity. In practice, this restricts us to examining the question in just one administrative setting. Obviously, this has consequences for the external validity of the study, as discussed in the conclusion.

The municipality has 49 public schools. To allow for cluster randomization (i.e., the allocation of school classes into the treatment and control groups), very small schools (< 400 pupils) were excluded. We used a cluster randomization in which the 81 classes were assigned to either the treatment or control groups in order to reduce the potential risk of experimental contamination or dilution of the treatment group because of control group exposure to the intervention (Donner & Klar, 2000). We also excluded schools that were in the process of closing. This reduces the number of relevant schools to 36. The participation of

schools in the experiment was voluntary. The schools meeting the above-mentioned inclusion criteria were contacted, and 10 agreed to participate. The randomization of classes into treatment and control groups was conducted after the schools had agreed to participate. Two schools that did not meet the inclusion criteria contacted the municipality, asking if they could be included in the project. They were exposed to the intervention but not included in the analyses.

In sum, the study population consisted of 1438 families from 81 school classes (grades 1–3) from 10 schools. Before the random assignment, the classes were stratified with regard to school and grade levels in order to increase the homogeneity of the experimental groups. The classes were then randomly assigned to intervention and control groups, respectively, which ensures that the provision of information is exogenous. Table 1 shows the descriptive statistics for characteristics of subjects in the treatment and control groups. It shows that the stratification and randomization produced a balanced set of experimental groups. Out of 27 comparisons between the treatment and control groups, only two comparisons (two of eight income categories) showed a significant difference at the 0.05 level.

[Table 1 about here]

The effect of the intervention on parents' knowledge was measured using a survey questionnaire. One potential problem is that parents who did not read the booklet in the first place, and for whom the intervention had no effect, could be prompted to find the booklet and read it by receiving the survey itself. We use three strategies to counter this problem. First, the sender of the booklet was the Municipality of Aarhus, while the university was the sender of the survey. Second, we made no reference to the booklet or the intervention in general when introducing the survey and measuring knowledge in the survey. Together, this should reduce

the risk that subjects will see the survey as follow-up to the booklet. Third, the survey was sent to the parents by mail about one week after the information was received. This should ensure that most booklets that were immediately tossed away would be difficult to recover for parents.

The survey could be completed by returning the paper version of the questionnaire (68 percent of the respondents) or filling out a web version (32 percent). The response rate was 64 percent: 63 percent in the control group and 66 percent in the treatment group (this difference is not significant, $p = 0.17$). The parents returned the questionnaire 1–34 days after having received it. The mean return time was about 11 days, or about 18 days after the information was provided.

Three measures of knowledge were used (the item wording is presented in the appendix). Two were objective in the sense that the parents were asked about children's reading skills development and presented with a set of possible answers, some correct and some incorrect. Here, we follow Whitehurst et al. (1988), an influential study documenting the effects of dialogic reading, which is the pedagogical idea behind the booklet. The first objective measure is *Activity Knowledge*. According to the booklet and the ideas behind dialogic reading, it measures knowledge about activities that will help develop reading skills. The second objective measure is *Concept Knowledge*, which captures knowledge about the concept of reading development as explained in the booklet. The activity measure is a 0–3 scale, where 0 indicates that the respondent did not answer any of the activity questions correctly, and 3 indicates that all three questions were answered correctly (1 point is given for every correct answer). The mean is 2.01, and the standard deviation is 0.87. The concept measure is a 0–2 scale (mean: 1.26, standard deviation: 0.66), where 2 indicates that both questions about the concept of reading development were answered correctly.

The third measure was subjective in the sense that it captures the extent to which the parents feel they have knowledge about the development of reading skills. We call this measure *Subjective Knowledge*. Changes in both the objective knowledge and subjective feelings of being knowledgeable can generally be important to behavior; for example, students' own perceptions of the competence to learn affects learning. Similar effects apply to teachers (Bandura, 1993). It is therefore interesting to see whether information nudging affects not only objective knowledge but also subjective feelings of being knowledgeable. The measure was constructed as an additive index of two Likert-scale items. The scale ranges from 0 to 10, 10 indicating the highest degree of perceived knowledge.

Results

We first analyze the effect of the treatment on the two objective measures of knowledge. Table 2 shows the predicted percentages, based on ordered logistic regressions, of the correct answers provided by the parents on each of the two objective knowledge measures in the control and treatment groups, respectively. Each of the regression models include the fixed effects of the 31 school-specific grade levels (the stratification variable) in order to ensure as precise a comparison between the treatment and control groups as possible (i.e., a treatment group at a given school and grade level is compared to a control class at the same school and grade level).

For the first measure of objective knowledge (the left part of Table 2), we find that the treatment group subjects were significantly less likely to provide no or only one correct answer (of three possible). We also find that the treatment group is more likely to provide correct answers only. In the control group, 28 percent were able to do so. In the treatment group, 39 percent answered all three questions correctly, corresponding to a 39 percent increase. Hence, the effect is statistically significant and substantial. The right panel in Table

2 shows similar results for Concept Knowledge. Here, respondents replied to two questions. In the control groups, 31 percent answered both questions correctly; in the treatment groups, 41 percent (a 32 percent increase). Again, we find a substantial effect from the information treatment on citizen knowledge.

[Table 2 about here]

Table 3 shows the results of the subjects' own evaluation of their knowledge. In the control group, subjects on average scored 6.50 on the index measuring Subjective Knowledge. In the treatment group, the mean score is 6.75 (the difference is statistically significant at the 0.05 level). This lends further support to the hypothesis.

[Table 3 about here]

Combined, the results presented in Tables 2 and 3 provide solid evidence that information provided to specific citizens by government can positively affect citizen knowledge.

These results show that the treatment has a considerable effect on knowledge among parents in general. It remains unclear, however, whether this effect is homogeneous or heterogeneous. The following analyses investigate whether the treatment effect depends on the background of the parent. First, we examine the direct effects of SES variables on the three knowledge measures. Table 4 shows the results of OLS models⁴ with fixed effects for the 81 classes (not shown), and cluster-robust standard errors are applied in the significance test. We measure the following dimensions of SES: education, employment status, status as

⁴ Because of the few categories of the dependent variables, ordered logistic regression could be an alternative. For simplicity, however, we present the OLS estimates. Estimating the effects by ordered logistic regressions does not change the conclusions (results available upon request).

immigrant, father or mother of child, and whether or not the responding parent is living with a partner (income is not included in the analysis, because missing values on the variable would reduce our sample considerably). We also examine if knowledge is correlated with response time.

It turns out that immigrants consistently have less knowledge on reading development. This applies to all three measures: Immigrants have significantly less knowledge about effective activities and the concept of reading development, and they also feel less knowledgeable. Education generally has similar effects, although only statistically significant for the first of the objective knowledge measures (Activity Knowledge). Being employed has a positive effect, albeit only significant for Concept Knowledge. Hence, there is considerable variation in knowledge among groups with different SES.

[Table 4 about here]

We investigate heterogeneous treatment effects in Table 5. The table presents the key results of a series of regression analyses that estimate the treatment effect on each of the three knowledge measures while examining whether the treatment effect is moderated by six different respondent characteristics: education, employment status, status as immigrant, father or mother of child, whether the responding parent is living with a partner or not, and the response time of the parent. This requires the estimation of 18 models. The results are presented in Table 5, where we show the estimated coefficients of the interaction term of the 18 models (thus, the full models are not shown in the table).

[Table 5 about here]

In sum, the results of Table 5 are quite clear: The treatment effect does not appear to be heterogeneous across different types of parents. The only case in which we observe a significant interaction effect is when analyzing if the treatment effect on the second objective knowledge measure (Concept Knowledge) is moderated by the father/mother indicator. Thus, the treatment effect on the second objective knowledge measure is found only among the mothers. There are no moderating effects in the other 17 models. The fact that the response time does not moderate the treatment effect suggests that the effect is not just present immediately after receiving the treatment. While, obviously, we cannot know whether the change in knowledge is permanent, there is no evidence in the data that the effect is wearing off during the first month.

It should be noted, however, that for some of the characteristics, especially immigrant status, the group size is rather small. Consequently, if we only analyze the group of immigrant parents (about 11 percent of the sample), the significance test will only capture rather strong interaction effects. For other variables, the respondents are distributed more equally across the different characteristics. For example, roughly one third of the sample is located in the low-education category. This provides a better opportunity to split the sample into low- and high-education parents and examine if we observe an effect in each of these subsamples. If we split the sample into low- and high-education subgroups, we find treatment effects both among parents with low and high educations when analyzing the first measure of objective knowledge. When using the second objective measure as the dependent variable, we find a significant effect for highly educated parents. When analyzing the subgroups using the subjective knowledge measure as the dependent variable, we find no statistically significant effects. This is hardly surprising, since the treatment has the weakest effect on the subjective knowledge, and this effect is not strong enough to produce significant results when analyzing the parents in the subsamples (i.e., there is far less power in each analysis).

Conclusion

Nudging is a policy instrument designed to alter people's behavior without direct regulation or economic incentives. It has attracted considerable scholarly attention, and studies show that it can be quite effective in changing behavior (John et al., 2011; Thaler & Sunstein, 2008).

Nudging by providing information is one prominent type among these instruments. The logic simply is that presenting information can nudge citizens to make other decisions, and field experiments document this effect (Calzolari & Nardotto, 2011; Hastings, Welden, & Weinstein, 2007; John et al., 2011, pp.75–84).

We argue that information nudging can also affect citizens' knowledge. This is important for two reasons. First, new knowledge can become the basis for future choices. In this way, nudging by providing information can make the nudging initiative stick for an extended period of time. Second, citizen knowledge may be a valuable input in strategies aimed at improving citizen choices through deliberation and participation, which may be a supplement or an alternative to nudging strategies.

We test the knowledge effect of information nudging in a randomized field experiment. The randomization breaks the potential reverse effect from knowledge to information provision. Any difference between the treatment and control groups can therefore be taken as evidence of the causal effect of the treatment. We find that providing information by sending a hard-copy booklet on learning strategies in reading has a clear effect on three different measures of parental knowledge about how they can best help their children learn to read. Hence, nudging by providing information does in fact have an impact on knowledge among citizens. The effect on knowledge was not moderated by parents' socioeconomic status.

The cost of sending information directly is difficult to assess. In the present case, the costs of printing the booklet, postage, envelopes, and so forth amounted to about USD 2.2 per

booklet. However, the costs – some of which will be fixed – of deciding, designing, discussing, and handling such a campaign is of course likely to be much higher. So although we cannot provide an estimate of the costs of sending information, we can say that the marginal costs are at least USD 2.2 per participating home or citizen.

These results are based on a field experiment in which the intervention consisted in a hard-copy booklet sent to the parents of children in a number of schools, all located within one administrative setting, the Municipality of Aarhus. This raises the question of external validity. Among experimental designs, the field experiment is often regarded as having the potential of achieving a high degree of external validity because field experiments are carried out with real citizens as subjects, in real settings, and with a treatment that can, as in this case, mirror an intervention that could have been implemented even without the field experiment. In this sense, we can be confident that the results are externally valid in the sense that they are likely to apply to similar settings. This begs the question of generalizability, and two issues are particularly relevant. First, are similar mechanisms likely to occur for other functions of government or in other administrative settings? Second, will other types of interventions produce different results?

Regarding the first question, the intervention was aimed at parents to children in schools, and the information concerned knowledge about reading. In other words, it is a function of government in which citizens will generally be privately interested in securing better outcomes (most parents prefer that their own children learn to read well). We conjecture that weaker effects are likely for government functions that do not affect the private interests of citizens directly. For example, an intervention aimed at promoting recycling may not affect knowledge as strongly, since citizens may find that the private benefits of going through the information material are smaller. Likewise, we expect that an intervention will have stronger effects for the functions of government with very intense and

direct consequences for citizens. Furthermore, some groups are likely more interested in schooling than others. For instance, highly educated parents might have a different view of the importance of reading skills than less-educated parents. In this sense, the present field experiment constitutes a most-likely case for finding heterogeneous effects. Since we find no evidence of heterogeneous effects, we expect that this finding can be generalized to other functions of government.

Regarding the second question, it is relevant to discuss the design of the intervention and the intensity or strength of the treatment. The booklet was designed to be effective, and the wording and layout was designed in collaboration with professionals to ensure this. However, this field experiment does not allow us to conclude whether other designs could be more effective. The answer to such questions would require variation in the design of the intervention. We regard the strength of the intervention as moderate. Sending a booklet by mail meant that participants received a physical object. This indicates that the sender has committed resources to providing it, and citizens may see this as a signal of the importance of the information. It also implies that citizens would actively have to throw out the booklet if uninterested. In other words, it is difficult to ignore it completely. The present design does not allow us to clarify the effect of treatment intensity on knowledge, but it seems very likely that stronger intervention could have even stronger effects on knowledge. Weaker interventions (perhaps from providing information electronically) will possibly not affect knowledge. Again, resolving this question will require further research.

References

- Bandura, A. (1993). Perceived Self-Efficacy in Cognitive Development and Functioning. *Educational Psychologist*, 28(2): 117–148.

- Calzolari, G., & Nardotto, M. (2011). Nudging with Information: A Randomized Field Experiment. Available at SSRN: <http://ssrn.com/abstract=1924901> or <http://dx.doi.org/10.2139/ssrn.1924901>.
- Donner, A., & Klar, N. (2000). *Design and Analysis of Cluster Randomization Trials in Health Research*. London: Arnold.
- Hastings, J. S., Van Weelden, R., & Weinstein, J. M. (2007). Preferences, Information, and Parental Choice Behavior in Public School Choice. NBER Working Paper no. 12995, 2007. <http://www.nber.org/papers/w12995>.
- Hood, C. (2010). Can We? Administrative Limits Revisited. *Public Administration Review*, 70: 527–534.
- Jakobsen, M., & Thomsen, M. K. (2015). Samproduktion, offentlige tiltag og servicebrugernes motivation til at deltage. *Politica*, 47(2), 164–183.
- John, P., Cotterill, S., Lui, H., Richardson, L., Moseley, A., Nomura, H., Smith, G., Stoker, G., & Wales, C. (2011). *Nudge, Nudge, Think, Think: Using Experiments to Change Civic Behavior*. London: Bloomsbury Academic.
- Jones, B. (2001). *Politics and the Architecture of Choice*. Chicago: Chicago University Press.
- Nesbit, L. S., & Milther, M. K. (2012). *Motivation for coproduction*, unpublished MA-thesis, Department of Political Science, Aarhus University.
- Roberts, N. (2004). Public Deliberation in an Age of Direct Citizen Participation. *The American Review of Public Administration*, 34: 315–353.
- Simon, H. A. (1957). *Models of Man*. Oxford: Wiley.
- Thaler, R. H., & Sunstein, C. R. (2008). *Nudge*. New Haven: Yale University Press.
- Thomsen, M. K., & Jakobsen, M. (2015). Influencing Citizen Coproduction by Sending Encouragement and Advice: A Field Experiment. *International Public Management Journal*, 18(2).

Weiss, J. A., & Tschirhart, M. (1994). Public Information Campaigns as Policy Instruments.

Journal of Policy Analysis and Management, 13(1): 82–119.

Whitehurst, G. J., Falco, F. L., Lonigan, C. J., Fischel, J. E., DeBaryshe, B., Valdez-

Menchaca, M. C., & Caulfield, M. (1988). Accelerating Language Development

Through Picture Book Reading. *Developmental Psychology*, 24(4): 552–559.

Tables and Figures

Table 1. Descriptive Statistics of Control and Treatment Groups. Means and Percentages

	Control group	Treatment group
Education		
Primary or secondary	20%	18%
Post-secondary, short	10%	11%
Post-secondary, medium term	36%	35%
Post-secondary, long	34%	37%
Living with a partner	82%	83%
No. of children	2.31	2.35
Immigrant	12%	10%
Employment		
Employed	85%	86%
Unemployed, seeking job	5%	3%
Unemployed, not in job market	5%	6%
Student	5%	5%
Household income (in 1,000 DKK)		
< 200	6%	2% *
200–399	15%	16%
400–599	17%	18%
600–799	24%	25%
800–999	22%	19% *
1,000–1,199	7%	10%
1,200–1,399	5%	5%
> 1,400	4%	4%
Age of child in study population	7.89	7.90
Gender of child in study population, girl	51%	48%
Grade level of child in study population		
Grade level 1	31%	34%
Grade level 2	37%	29%
Grade level 3	31%	37%
Responding parent, mother	74%	75%
Responded by web survey	32%	32%
Response time, no. of days after the intervention was started	17.57	17.68

Note: * $p < .05$, ** $p < .01$, *** $p < .001$, two-tailed hypothesis test of the difference between the control and treatment groups. Except for the examination of grade level, the models used to conduct the significance tests include the fixed effects of 31 school-specific grade levels and cluster-robust standard errors at the class level.

Table 2. Effect of Treatment on Service Users' Knowledge. Predicted Percentages

Knowledge (correct answers)	Objective knowledge I (Activity knowledge)			Objective knowledge II (Concept knowledge)		
	Control group	Treatment group	Difference (T-C)	Control Group	Treatment group	Difference (T-C)
0	5	3	-2**	14	10	-4
1	28	20	-8**	55	49	-6**
2	39	38	-1	31	41	+10***
3	28	39	+11***	-	-	
Total	100	100		100	100	

Note: * $p < .05$, ** $p < .01$, *** $p < .001$, one-tailed hypothesis test. The predicted percentages are based on ordered logistic regressions. The treatment effect is significant in both of the ordered logistic regressions ($p < .001$). The significance tests of the differences are conducted by logistic regressions using the specific number of correct answers (dummy variable) as a dependent variable. All models include the fixed effects of 31 school-specific grade levels and cluster-robust standard errors at the class level.

**Table 3. Effect of Information Intervention on Service Users' Subjective knowledge.
Mean score on 0–10 index**

	Control group	Treatment group	Difference
Mean score on subjective knowledge index	6.50	6.75	0.25*

Note: * $p < .05$, ** $p < .01$, *** $p < .001$, one-tailed hypothesis test. The model includes fixed effects of 31 school-specific grade levels and cluster-robust standard errors at the class level.

Table 4: Explaining Activity Knowledge, Concept Knowledge, and Subjective Knowledge by SES. OLS Regressions.

	Objective knowledge I (Activity knowledge) Model I		Objective knowledge II (Concept knowledge) Model II		Subjective knowledge Model II	
Education						
(ref.: Post-secondary, long)						
Primary or secondary	-0.30**	(0.10)	-0.04	(0.07)	-0.10	(0.25)
Post-secondary, short	-0.37**	(0.11)	-0.12	(0.09)	-0.34	(0.28)
Post-secondary, medium term	0.06	(0.07)	0.05	(0.05)	0.15	(0.18)
Employment (ref.: employed)						
Unemployed, seeking job	-0.27	(0.20)	-0.10	(0.13)	0.15	(0.41)
Unemployed, not in job market	-0.22	(0.13)	-0.28*	(0.11)	0.15	(0.36)
Student	-0.15	(0.15)	0.04	(0.11)	-0.52	(0.37)
Immigrant (ref.: non-immigrant)	-0.35**	(0.13)	-0.18*	(0.07)	-0.89**	(0.27)
Father (ref.: mother)	-0.01	(0.07)	-0.08	(0.05)	-0.13	(0.16)
Living with partner (ref: not living with partner)	-0.04	(0.08)	-0.16*	(0.07)	-0.07	(0.23)
Response time	0.00	(0.00)	-0.00	(0.00)	-0.02	(0.01)
Constant	2.18***	(0.06)	1.35***	(0.06)	7.49***	(0.19)
N	797		807		811	
Classes	81		81		81	

Note: * $p < .05$, ** $p < .01$, *** $p < .001$, two-tailed hypothesis test. All models include the fixed effects of 81 classes and cluster-robust standard errors.

Table 5: Moderator Effects. OLS Regressions.

	Objective knowledge I (Activity Knowledge)	Objective knowledge II (Concept Knowledge)	Subjective Knowledge
Treatment × Low education	0.01	-0.04	-0.18
Treatment × Unemployed	0.00	0.18	0.25
Treatment × Immigrant	0.01	-0.20	-0.11
Treatment × Father responded to survey	-0.03	-0.23**	-0.33
Treatment × Respondent living with partner	0.04	-0.00	0.00
Treatment × Response time	0.01	-0.00	0.00

Note: * $p < .05$, ** $p < .01$, *** $p < .001$ (two-tailed hypothesis test). The coefficients show the interaction effect between the treatment and different characteristics of the respondents. Each coefficient is estimated by a fixed effects regression model (in all, 18 models were estimated). Only the coefficient of the model's interaction term is presented in the table. Each model includes the fixed effects of 31 school-specific grade levels. The reference categories for the first five characteristics of the respondents are high education, employed, non-immigrant, mother, and not living with partner.

Appendix

The survey questions for the three measures of knowledge appear below, translated from Danish.

1. Activity knowledge

Which of the following activities do you believe would be the greatest help in giving your child good reading comprehension and teaching them to read really well? (Tick 3 boxes in total)

- Practicing the spelling of difficult words with the child
- Talking to the child about what you are reading together
- Helping the child stay focused on the reading, e.g., by creating a calm environment for the child and letting it sit alone to read for 15 minutes at a time
- Talking together about how what you read relates to things that the child has experienced
- Telling the child when they have made a mistake and helping them read the text correctly
- Asking questions about what you are reading

2. Concept knowledge

In which school courses would you consider it important to be good at reading? (Tick one or more boxes)

- Danish and language courses

- Mathematics and science courses
- Music and creative courses
- History and religion courses
- All courses

Acquiring reading proficiency requires many different skills. In recent years, good reading comprehension has increasingly been viewed as vital for children’s ability to become good at reading.

Which of the following descriptions do you think fits the word “reading comprehension” best?

- Reading comprehension means understanding the sense of the words you read
- Reading comprehension means understanding the meaning of what you read
- Reading comprehension means understanding how to read
- Reading comprehension means understanding what the word “read” means

3. Subjective knowledge

	Strongly agree	Agree	Neither/ nor	Disagree	Strongly disagree	Don’t know
I don’t know how to help my child become good at reading	<input type="checkbox"/>					
I don’t know how to stimulate my child’s interest in reading	<input type="checkbox"/>					