Ana Caraban, Instituto Superior Técnico
Evangelos Karapanos, Cyprus University of Technology

The '23 Ways to Nudge' Framework

Designing Technologies That Influence Behavior Subtly

Insights

- While cognitive biases tell us *which* problems nudges can solve, they tell us little about *how*.
- The "23 ways to nudge" framework provides interaction designers with a repertoire of nudging mechanisms—the *how* of nudging.
- We have the moral responsibility to be thoughtful about the choice architecture that our products create.

IMAGE BY CIK IN / SHUTTERSTOCK.COM

Twelve years ago, Richard Thaler and and Cass Sunstein [1] introduced the notion of *nudging*. They defined a *nudge* as "any aspect of the choice architecture that alters people's behavior in a predictable way without forbidding any option or significantly changing their economic incentives." They suggested that nudging obeys the principle of *libertarian paternalism*, in that it preserves freedom of choice while at the same time ascribing to the belief that it is legitimate for choice architects to influence people's behaviors for their own good.

Nudging has been applied successfully across a wide range of domains, including public policy, marketing, and technology. For instance, changing from an opt-in to an opt-out organ-donation policy, wherein consent for donation is assumed by default and one retains the right to opt out, can increase the rate of organ donors, thus having a positive impact on societal welfare without forbidding individuals' options or significantly changing their economic incentives. Similarly, replacing cake with fruit in the impulse basket next to the cash register has been found to lead to people buying more fruit and less

cake, while both choices remain available.

Nudging is based on the idea that we humans have two modes of thinking, the *automatic* and the *reflective* [2]. While reflective thinking enables us to make rational decisions by considering all the pros and cons of each choice, automatic thinking enables us to make quick decisions with little effort. We do so through the use of *heuristics*, mental shortcuts that allow us to solve problems by substituting unavailable information with a cue in the environment. For instance, when unsure about how to act in a given situation, we may look at what others do and follow their actions. However, this also makes us susceptible to *cognitive biases*, systematic deviations from rational judgment. Nudging suggest that we can leverage our knowledge on how humans err to design systems and policies that guide people in making better choices.

Figure 1. The Nudge Deck [5] translates the "23 ways to nudge" framework to a set of design cards, thus making it accessible during design meetings. More information about the Nudge Deck can be obtained at http://persuasive.cut.ac.cy/nudgedeck

THE '23 WAYS TO NUDGE' FRAMEWORK

When we started thinking about ways to design nudges in interactive technology, we had a hard time translating cognitive biases into design solutions. We realized that while cognitive biases can motivate the goal of the design effort, they do little in guiding the designer in how to solve the problem. For instance, the *status-quo bias*, also referred to as the power of inertia, denotes our tendency to resist change and to go along with the path of least resistance. This predisposition of "choosing not to choose" leads us to maintain choices already made because the process of searching for a better alternative is often slow, uncertain, or costly. In fighting against this bias, designers can take different approaches: They can make desirable options the default choice. They can rearrange the positioning of the choice (e.g., the rank of a search result). They can also hide undesirable choices, suggest alternatives, pause an action a user just made and provide them with the option to revert it, remind users of the consequences of their action, add inferior alternatives to the set of options with the goal of making some other options seam more appealing (i.e., the decoy effect)—the list goes on.

The "23 ways to nudge" framework [3] provides us with a repertoire of exactly these alternatives for *how to nudge*. The framework emerged out of a systematic review of the use of nudging in HCI literature. We identified 71 papers that presented technology-mediated nudging interventions and attempted to infer the *why* (i.e., which cognitive bias did the nudge attempt to combat?) and the *how* (i.e., what exact mechanism did the nudge employ to incur behavior change?). All in all, we found 23 distinct mechanisms of nudging, grouped in six categories and leveraging 15 different cognitive biases. The six overall categories of nudging mechanisms are:

- *Facilitate:* nudges that reduce the (physical or mental) effort required for an activity, in order to motivate people to pursue it. *Mechanisms:* default options, opt-out policies, positioning, hiding, suggesting alternatives.
- *Confront:* nudges that attempt to pause an unwanted action by instilling doubt. They attempt to break mindless behavior and prompt a reflective choice. *Mechanisms:* throttling mindless activity, reminding of the consequences, creating friction, providing multiple viewpoints.
- *Deceive:* nudges that use deception mechanisms in order to affect how alternatives are perceived, or how activities are experienced, with the goal of promoting particular outcomes. *Mechanisms:* adding inferior alternatives (decoy), biasing the memory of past experiences, placebos, deceptive visualizations.
- *Social influence:* nudges that take advantage of people's desire to conform and comply with what is believed to be expected from them.

While cognitive biases can motivate the goal of the design effort, they do little in guiding the designer in how to solve the problem.

Mechanisms: invoking feelings of reciprocity, leveraging public commitment, raising the visibility of users' actions, enabling social comparisons.

• *Fear:* nudges that evoke feelings of fear, loss, and uncertainty to make the user pursue an activity. *Mechanisms:* making resources scarce, reducing the distance.

• *Reinforce:* nudges that attempt to reinforce behaviors through increasing their presence in individuals' thinking. *Mechanisms:* just-in-time prompts, instigating empathy, subliminal priming, ambient feedback.

Let's consider one example: *reminding the user of consequences*. The *availability heuristic* reflects our tendency to judge the probability of occurrence of an event based on the ease with which it can be recalled. As a result, we might overestimate the probability of events when they are readily available to our cognitive processing (e.g., judging the probability of having cancer as higher than the actual probability of cancer when detecting a lump in our body) while we might be overly optimistic when these events are distant in our minds. Nudges in this category prompt individuals to reflect on the consequences of their actions. For instance, Harbach et al. [4] redesigned the permissions dialogue of the Google Play Store to incorporate personalized scenarios that disclosed potential risks from app permissions. If the app required access to one's storage, the system would randomly select images stored on the phone along with the message "this app can see and delete your photos." Similarly, Wang et al. [4] designed a Web plugin that aims at mitigating impulsive disclosures on social media through reminding users of the audience. The system selects five random contacts from the user's friend list, according to the post's privacy settings, and presents the contacts' profile pictures along with the message "These people and [X] more can see this."

THE NUDGE DECK

To make this framework accessible during design sessions, we translated it to a set of design cards—the *Nudge Deck* [5] (Figure 1). Each of the 23 nudge mechanisms is presented using one card. On its front side, a description of the nudge mechanism is provided along with an example. On the back side of the card, a list of design considerations is presented, with the goal of triggering discussion within the design team and promoting creative thinking.

REFLECTIONS: THE ETHICS OF NUDGING

A common criticism on nudging is that it works through manipulating people's choices. Pelle Hansen and Andreas Jespersen [6] proposed a more complicated picture by classifying nudges against two variables: a) which mode of thinking they primarily engage, automatic or reflective, and b) whether they are transparent, that is, whether people can perceive the intentions of the nudge as well as the means by which the nudge attempts to incur behavior change. By doing this exercise, we found that 78 percent of the nudges presented in HCI literature do not work by manipulation (that is, their intentions and means are transparent to the user), while more than half of the nudges attempt to engage the reflective mind by prompting reflective choice (Figure 2). Some examples are the ones from Harbach et al. and Wang et al. [4] presented earlier, or Laschke et al.'s [4] Keymoment, a key holder that nudges users to choose bike over car by dropping the bike key on the floor when one picks up the car key.

REFLECTIONS: WHEN DO NUDGES FAIL?

Out of the 74 examples of nudging, 49 were empirically studied. We analyzed the reasons why nudges failed, which we present here.

Lack of educational gains. The power of nudges that engage the automatic, rather than the reflective, mind is that they do not rely on individuals' will and capacity to engage in conscious decision making and behavior regulation. This, however, is also their weakness—they provide no educational gains. As such, the effects of the nudge will likely not hold after its removal. For instance, Egebark and Ekstrom [4] changed the default printer option from simplex to double-sided print and found a 15 percent reduction in paper consumption. This default effect was immediate and lasted for more than six months. However, when new printers were introduced, users maintained the new default, a

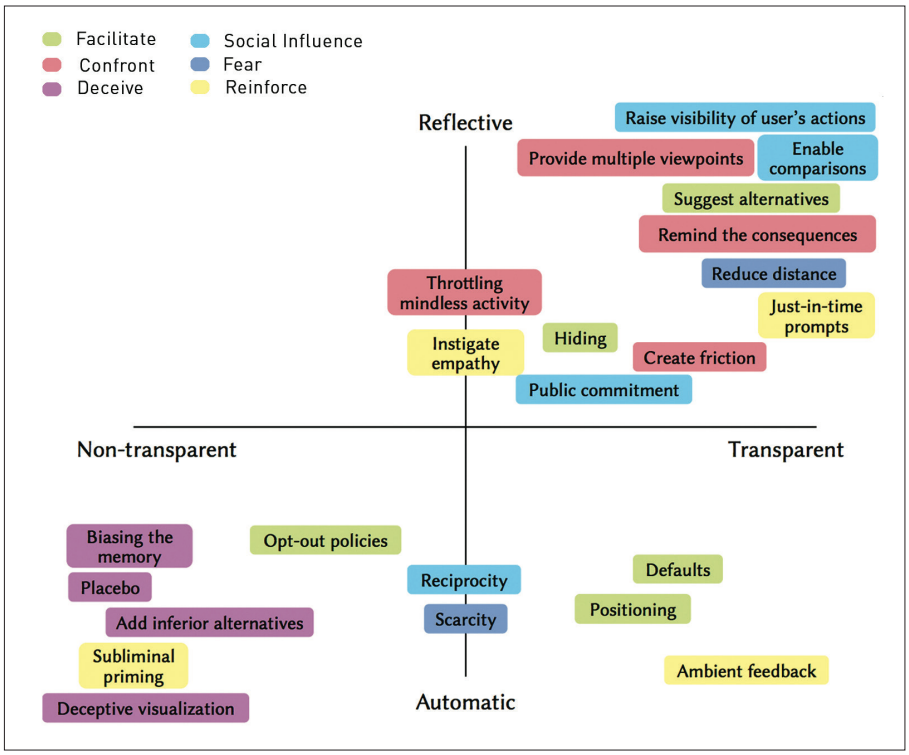

Figure 2. We found that 78 percent of the nudges presented in HCI literature do not work by manipulation—their intentions and means are transparent to the user—while more than half of the nudges attempt to engage the reflective mind by prompting reflective choice.

single-side print option. Interestingly, we observed that only 14 percent of the studies (7 out 49) inquired into whether the effects held after the removal of the nudge.

Nudging effects not sustaining over time. The effects of nudging may degrade over time for a number of reasons, such as habituation or reactance. For instance, any warnings, such as the one about the risks entailed in providing excessive permissions to apps, are bound to lose their effects as we get used to them and start to ignore them. Harbach et al. [4], for instance, suggested altering the feedback techniques to prevent habituation. In another example, Wang et al. [4] observed that their intervention, in some cases, caused reactance. They designed a Web browser plugin that mitigates impulsive disclosures on Facebook by holding a post for 10 seconds, encouraging users to review it. While the countdown could be canceled immediately by clicking a "Post Now" button, some participants were frustrated, as they felt that the plugin reduced their autonomy. Understanding whether initial effects sustain over time is of critical importance in the design of nudges. Yet we were surprised to find that in 65 percent of the studies, effects had a duration of a day or less, while in only 19 percent of the studies the effects lasted over a month.

Unexpected effects and backfiring. Nudges may not always produce the desired effect, due to compensating behaviors, unexpected interpretations, or other reasons. For instance, as Egebark and Ekström discuss, setting double-sided print as the default option may lead individuals to print more as they carry less weight [4]. And providing doctors with a checklist of symptoms to consider in emergency situations might, as Wu et al. write, "become a crutch that de-skills memory" [4]. Inquiring into possible backfires and unexpected effects is of critical importance in the early steps of nudge deployment and evaluation.

Timing and strength of nudges. Fine-tuning the timing and the strength of nudges can be of critical importance. For instance, Forwood et al. [4] explored the power of the peak-end phenomenon in the context of HCI and attributed insignificant results to a weak manipulation of the ending experiences. In another study, Brown et al. [7] altered the default setting of a thermostat room in an attempt to reduce domestic energy consumption. They found that while participants would accept a change in the default temperature from 20°C to 19°C, they would react when this default was further reduced to 17°C.

Strong preferences and established habits. Nudges may not be as effective in the presence of strong preferences or established habits. For instance, Forwood et al. [4], who created a system that nudges individuals toward healthier food choices, suggested that a number of factors, such as the strength of preferences for certain food choices and the extent to which food choices are habitual, can influence the effectiveness of the nudge. Räisänen et al. [4] observed that the less a user smoked, the more affected he was by the smoking-cessation nudges, while in the study of Lehmann et al. [4], a nudge to enroll in a vaccination program was not effective for individuals that held strong negative attitudes toward vaccination. As Sunstein suggests, this ineffectiveness should be seen in a positive light, as it implies that the nudge preserves individuals' freedom of choice: "If choosers ignore or reject it, it is because they know best" [8].

CONCLUSION

Interacting with technology can be seen as a sequence of choices. Some of these choices lead to better outcomes than others. As interaction designers, we have the moral responsibility to be thoughtful of the choice architecture that our products create. The study of cognitive biases enables us to understand the underlying forces that make people susceptible to suboptimal decisions—the *why* of nudging. We hope that the "23 ways to nudge" framework and the Nudge Deck will be a first step in providing interaction designers with a library of mechanisms that can guide them in implementing effective nudges—that is, to enrich our understanding of *how to nudge*. As our physical and our digital worlds increasingly integrate, digital infrastructure is affecting a wider range of choices that impact all facets of human life, from how we spend our time to how we interact with others to how we plan for the future. This gives us, interaction designers, the power and responsibility to thoughtfully design our living environments and to carefully monitor the effect of these environments on our lives.

ENDNOTES

1. Thaler, R.H. and Sunstein, C.R. *Nudge: Improving Decisions About Health, Wealth, and Happiness.* Penguin, 2009.
2. Kahneman, D. Maps of bounded rationality: Psychology for behavioral economics. *American Economic Review 93*, 5 (2003), 1449–1475.
3. Caraban, A., Karapanos, E., Gonçalves, D., and Campos, P. 23 ways to nudge: A review of technology-mediated nudging in human-computer interaction. *Proc. of 2019 CHI Conference on Human Factors in Computing Systems.* ACM, New York, 2019, 1–15.
4. Because of reference limitations, we could not refer to each nudge example separately. All nudges can be found in Appendix B of Caraban et al. [3].
5. Caraban, A., Konstantinou, L., and Karapanos, E. The Nudge Deck: A design support tool for technology-mediated nudging. *Proc. of the 2020 DIS Conference.* ACM, New York, 2020, 395–406.
6. Hansen, P.G. and Jespersen, A.M. Nudge and the manipulation of choice. *European Journal of Risk Regulation 4*, 1 (2013), 3–28.
7. Brown, Z., Johnstone, N., Haščič, I., Vong, L., and Barascud, F. Testing the effect of defaults on the thermostat settings of OECD employees. *Energy Economics 39* (2013), 128–134.
8. Sunstein, C.R. Nudges that fail. *Behavioural Public Policy 1*, 1 (2017), 4–25.

Ana Caraban is a computer science Ph.D. student at Instituto Superior Técnico and a teaching assistant at the University of Madeira. Her research is focused on the design of nudging interventions that promote physical activity.
→ ana.caraban@m-iti.org

Evangelos Karapanos is an associate professor in the Department of Communication and Internet Studies of the Cyprus University of Technology, where he leads the Persuasive Technologies Lab (http://persuasive.cut.ac.cy).
→ evangelos.karapanos@cut.ac.cy

DOI: 10.1145/3414464 © 2020 ACM 1072-5520/20/09 $15.00