Precision Nudging: The Future of Behaviour Change?

Themistocleous Christos
Lecturer of Consumer Behaviour,
Department of Communication and Marketing,
Cyprus University of Technology,
Email: c.themistocleous@cut.ac.cy
ORCID: 0000-0001-9074-7711

Evangelos Karapanos
Associate Professor of Human-Computer Interaction
Department of Communication and Internet Studies
Cyprus University of Technology,
Email: evangelos.karapanos@cut.ac.cy
ORCID: 0000-0001-5910-4996

Abstract

Behavioural science literature has seen the recent emergence of Nudge theory, a proposition that behaviour change can be achieved by altering the "choice architecture" of our physical and digital environments while considering the boundaries of human rationality. The chapter advocates the term *Precision Nudging*, the use of Artificial Intelligence to advance the efficacy of nudges while mitigating their adverse effects. It discusses two methods of Precision Nudging, *tailoring* and *timing*, and their potential in addressing adverse effects, in particular, behavioural spillovers and misfires. In doing so the authors categorise nudges in *reactive* ones that aim at improving adverse effects and *proactive* ones that aim to shield individuals from future spillovers and misfires.

1. Introduction

Artificial Intelligence (AI) has recently gained much attention for its prospect to advance the efficacy of behaviour change technology. A prominent tool in behavioural sciences with the aim of steering individuals' decisions towards more beneficial, healthy alternatives is *Nudge*, proposed in 2008 by Thaler and Sunstein. Since its conception nudges have attracted much attention in terms of their types, contextual applications, and effectiveness. Remaining under academic scrutiny, nudges have the potential of reshaping behavioural and social sciences. We attempt to make a case of how AI mediated nudges can help address some of their downsides while enhancing their effectiveness.

The chapter is structured as follows. First, we provide an overview of what Nudge theory entails, what qualifies as a nudge and review its latest developments in social, behavioural and computer science literature. Secondly, we identify new ways in which AI can enhance the use of nudges by improving their effectiveness through two elements, namely, *tailoring* and *timing,* thus leading to *Precision Nudging*. We present occasions where AI can identify nudge failures in the forms of misfires and spillovers. Finally, we make recommendations on how AI can provide corrective courses for nudge failures while covering implications for behavioural, social and HCI research. Throughout the chapter we employ *Behavioural Pathways*, visual representations of trajectories that demonstrate intended and unintended behavioural changes of nudges. In these pathways we highlight AI-mediated corrective courses for when nudges create unintended adverse effects.

In terms of contextual application for our ideas, we use the example of smartwatches as fitness and wellbeing devices that provide visual, sound and haptic information that nudge users towards more beneficial behaviours either through reminders, social comparison statistics and more. We note that beneficial behaviours are defined as behaviours that enhance individual and societal welfare. Concrete examples of the latter include exercising, more movement after large intervals of being idle for improved blood circulation as well as sugar intake reduction in cases where average intake significantly exceeds recommended levels by physicians.

2. Nudging for Behavioural Change

Nudge Theory serves as the behavioural birthchild of the Heuristics and Biases program named after Kahneman and Tversky's respective *Science* paper (Tversky and Kahneman, 1974). They posit that Heuristics (i.e. availability, representation, and anchoring) frugal as they might be, can give rise to biases that hinder individuals from optimal decisions. This perspective on human rationality generated research streams that descended from their work and isolated occasions in which heuristics can lead to systematic errors. Albeit difficult to unsee these biases, Kahneman proposes that slow and analytical thinking processes can lead to better decisions compared to ones made by heuristics (Kahneman, 2011). Nudges are based on the assumption that human decisions aim to maximise expected utility in the most efficient, heuristic-based, way. As such, the effort required to analytically evaluate the presented alternatives in a decision-making context is significant and can have ego depleting effects (Baumeister et al., 1998). In this context, nudging can help individuals make better, more efficient decisions in a chosen environment.

Thaler and Sunstein (2008, p. 6) define nudges as:

> "Any aspect of the choice architecture that alters people's behaviour in a predictable way without forbidding any options or significant changing their economic incentives. To count as a mere nudge, the intervention must be easy and cheap to avoid. Nudges are not mandates."

Choice architecture refers to the design of the different ways in which choices are presented to decision-makers. A choice architect aiming at bettering the individual's wellbeing, can present choices in ways that steer individuals towards beneficial outcomes, without restricting choices. Examples of nudges take the form of smart disclosures, default options, alternative suggestions, reminders, sequencing effects and opt-out policies, among others. Since its introduction, the concept of nudging has been eagerly explored in behavioural sciences and beyond, in a variety of contexts including welfare improvement, law compliance and disclosure decisions (i.e. Caraban et al., 2019; Steffel et al., 2016; Loewenstein et al., 2014; Themistocleous et al., 2014).

Nudges are set to lead individuals into more beneficial behaviours compared to the trajectory of behaviour where such nudges are absent. For example, an individual that has been idle for a prolonged period of time can be nudged towards a more beneficial behaviour – to stand up and move - compared to the continuation of the existing behaviour which is remaining sedentary. *Diagram 1* visualises the hypothesised use of nudges. Embedded in their design and philosophical positioning, nudges are not mandates, thus one can choose to ignore them. The epitome of a nudge is that the expected utility of the nudged behaviour will exceed that of the projected behaviour.

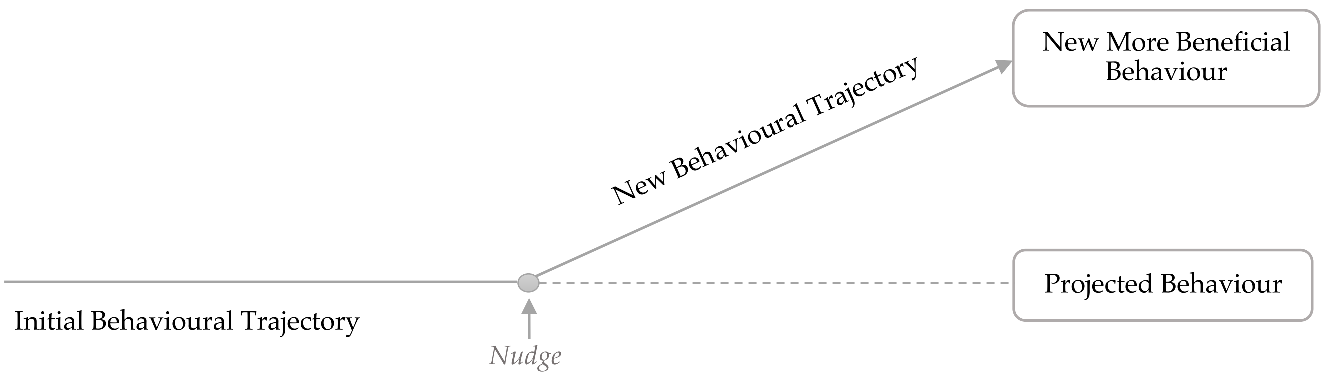

Diagram 1. *A nudge behavioural pathway. An individual can ignore the nudge and continue with the projected behaviour. (Source: Themistocleous and Karapanos, 2024)*

Several empirical studies have shown nudging to be an effective behaviour change strategy (i.e Hummel and Maedche, 2019). For instance, research found that a simple change to a default of double-sided printing led to a 15% paper reduction compared to the default counterpart of single-sided printing (Egebark and Ekstrom, 2016). In another example, Opower found their social comparison nudges, such as comparing one's energy consumption to that of her neighbours, to lead to a 2% reduction in energy consumption (Loewenstein et al., 2014)

Nudges, nevertheless, may fail and even produce adverse effects for a number of reasons – for example, due to the provision of complex or confusing information, due to repeated exposure, or due to compensating behaviours (Sunstein, 2017). In HCI, Caraban et al. (2019) reviewed the use of nudging and identified seven reasons of nudge failures, such as their lack of educational effects (i.e., nudges working without depleting individuals' cognitive resources but their effects becoming obsolete once nudges are removed), or failing due to individuals' strong preferences and established habits. Choosing a nudge type in a given situation, or for a given individual, is also important. For instance, some individuals may be more prone to follow a social comparison nudge, while others might be more prone to follow a nudge that explains the reasons, or the benefits acquired from following the target behaviour, instead of what others are doing on the subject matter.

AI can help here. The aim of the following section is to present how AI can improve the effectiveness of nudges, focusing on two main dimensions: *tailoring* (*i.e.,* selecting the right nudge type for a given individual or occasion) and *timing* (*i.e.,* selecting the right point in time to deliver a nudge).

3. AI-Mediated Nudges

ML has enabled social, behavioural and data scientists to migrate from a one-size-fits-all approach in technology design towards a more tailored one, founded on psychographic insights and repeated behavioural patterns of users in a given environment. In the case of Nudging, we suggest that AI can increase effectiveness by addressing two main questions: *how to nudge (tailoring)* and *when to nudge (timing)?* Here both timing and tailoring can enhance nudging in terms of its precision. Thus we use the term *Precision Nudging* to describe AI-mediated nudges that are enhanced in terms of the point in time that are presented and the form they take dependent on the individual and the environment the behaviour occurs.

3.1 Tailoring

Tailoring refers to adapting the type of nudge, mode of delivery, or any other aspect of the nudge's function, with the goal of increasing its effectiveness in a given situation, or for a given individual. As accurately pointed out by Kaptein et al., (2015) the following two statements "*70% of users run at least twice a week"* and *"Fitness Experts recommend that you run at least twice a week"* are fundamentally different in the way they nudge. Some individuals might be more prone to follow social comparisons (first statement), while others might prefer prompts of exercise that stem from experts' opinions (second statement).

To delineate tailoring, consider an example where we are designing alternative nudges that aim at breaking individuals' sedentary behaviour through the individuals' smartwatches (see Gouveia et al., 2016 as an example). Which heuristic do we attempt to tap into, and which type of nudge do we leverage? Individuals may vary in their susceptibility to different persuasion techniques, and as such, some nudges may be suitable for some individuals but not others. Individual differences in terms of need for cognition, scientific reasoning, numeracy, personality traits, or general decision-making style, will influence the effectiveness of different types of nudge for different individuals (see Warberg et al., 2019).

While measuring individuals' susceptibility to different cognitive biases psychometrically may be unrealistic, AI offers a pragmatic approach, that of Reinforcement Learning (RL). In its simplest form RL is a technique that identifies the optimal behaviour in order to obtain the maximum reward. In essence an AI agent is trained through interactions with the environment and the accumulation of observations and responses. Agents, through RL, learn as users interact with the model. A particular implementation of RL that is applicable to nudging, is Multi-Armed Bandits (MABs). In the example given above, assume that we have designed five different types of nudging to break sedentary behaviour. Each time we decide to nudge the individual to stand up and move, we select one of the five nudges. The individual is completely unknown to us - no prior data describe her behavioural patterns, no information highlight persuasion tactics that are more likely to be effective and no nudge type is more favourable over another. Because of this each nudge has equal change of being delivered (i.e., 20%). At each trial, we monitor a proximal behavioural variable (e.g., did the individual stand up and move for 50 steps or more during the five minutes following the delivery of the nudge?). This informs the model and thus, in each subsequent delivery, we have more knowledge about which of the five types of nudges are more likely to steer this individual to break sedentary behaviour. The greater the sample of observations, the greater the certainty of nudge type effectiveness for each individual.

In the early steps, it may be more fruitful to *explore* by trying out different types of nudges to acquire knowledge about the individual. Once sufficient knowledge is gathered (*e.g.,* we know that a given type of nudge is significantly more likely that then others to steer the individual to the beneficial behaviour) we can exploit (use a particular nudge type over the others). MABs define the strategy and the mechanisms for this exploration versus exploitation dilemma and have been used effectively both in research and in industry. For instance, Kaptein et al (2015) leveraged Cialdini's principles of persuasion (Cialdini, 2009) to suggest the notion of *persuasion profiles*: while some of us may be more susceptible to the principle of *reciprocity* (i.e., to return a favour made to us), others may be more susceptible to the principle of *scarcity* (i.e., to be attracted by an offering whose availability is limited). Kaptein et al (2015), aiming at tailoring, suggested the use of Multi-Armed Bandits to personalise messages Cialdini's six principles of influence.

3.2 Timing

We propose that *timing*, or the selection of an appropriate point in time to deliver a nudge, is the second method of Precision Nudging. A time-to-stand-up nudge might be irrelevant when one engages in long drives and a nudge for taking a walk outdoors might be more effective on a sunny instead of a rainy day. Timing should aim at inferring variables from the individuals' internal state and environment conditions so that the nudge is delivered at opportune moments. Interventions provided just-in-time will in essence maximise both their effectiveness and make a change in behaviour more attractive and beneficial. Consider having worked for 3hs straight up with no break from sedentary activity. This not only implies that a break is much needed but also that a nudge to have a break has high likelihood of being successful.

Timing is a central tenet in *Just-In-Time-Adaptive-Interventions* (JITAIs), defined as *"intervention design(s) aiming to provide just-in-time support, by adapting the dynamics of an individual's internal state and context"* (Nahum-Shani et al., 2018, p. 448). In their work, they describe how inferring the individuals' internal state and behaviour, such as their mood, medication adherence, sleep, hallucination coping and social functioning, can help in providing schizophrenia patients a timely and tailored intervention on how to cope when at a psychological peak, a time when help is most needed. Beyond JITAIs, inferring opportune moments for interventions has been a topic of considerable interest across various domains. For instance, Pielot et al (2015) aimed at optimising users' engagement with mobile phone notifications. They did so by inferring moments of boredom from a wide range of mobile phone sensor data and behavioural logs, and envisioned "boredom-triggered proactive recommender systems that attune their users' level of attention and need for stimulation." (p. 1). They in essence aim to influence and persuade an individual at a time that he is more receptive of a recommendation.

In a different line of work, Psychologist Mihaly Csikszentmihalyi employed Experience Sampling, aka Ecological Momentary Assessment, where individuals were beeped around eight times a day in order to complete a short questionnaire about their feelings, thereby mitigating recall and other biases present in global measures of wellbeing (Schwatrz et al., 2009; Csikszentmihalyi, 2002). The tool aimed to address the challenge of the remembering self and the experiencing self. Prompting you to think how happy you felt the last time you played guitar can create a misalignment compared to true emotions at the point of playing and the emotions you were just prompted to remember. By identifying opportune moments for triggering a survey during the day, Experience Sampling aimed at reducing the chronological difference between the actual experience and the remembering of that experience increasing recall accuracy.

While tackling these measurement challenges, Experience Sampling introduces new problems, particularly relating to the intrusiveness of the prompt, as participants need to stop their ongoing activity and complete the questionnaire. Recent efforts have attempted to utilise sensor data from smartphones to identify opportune moments for trigger Experience Sampling surveys, thereby reducing intrusiveness and maximising the likelihood of a response yet abiding to the chronological proximity requirements of Experience Sampling (Mehrotra et al., 2015).

All in all, Precision Nudging can draw on these past efforts to *tailor* and *time* the delivery of nudges. However, these two methods of Precision Nudging, *tailoring* and *timing*, can be further used in mitigating the unintended consequences of nudging. What if someone follows the nudge yet a follow-up, unintended adverse behaviour is triggered? What if there is a misfire and the nudge does not trigger the right behaviour? What if the intervention nudges an individual to a worse off situation? In the next section we elaborate on two frequent unintended consequences, *spillovers* and *misfires*, and provide recommendations for the use of AI mediated tailoring and timing to address them.

4. Mitigating Nudge failures through AI

4.1 Nudge Spillovers and AI

Spillovers in psychology relate to a person's behaviours, emotions and attitudes in one domain that can unintendedly flow into another (Zedeck 1992) and can have either adverse or positive effects (Bell et al., 2012). Detecting spillovers, in the context of Nudging, is particularly important as a seemingly beneficial behaviour triggered from a nudge, can lead to an unintended adverse behaviour. Dolan and Galizzi (2015, p.2) echo this argument and make an analogy between behavioural spillovers and ripples in a pond, prompting researchers to "capture all ripples of behaviour when a pebble of intervention is thrown in the pond". In essence, once following a nudge, the new behavioural trajectory might lead the individual to a subsequent unintended adverse behaviour. We call this a *Nudge Spillover* with diagram 2 depicting the respective behavioural pathway. The two behaviours, both intended and unintended, can cancel out the overall expected utility ultimately leading to a worse-off situation for the individual. In this case the nudge might ultimately do more harm than good.

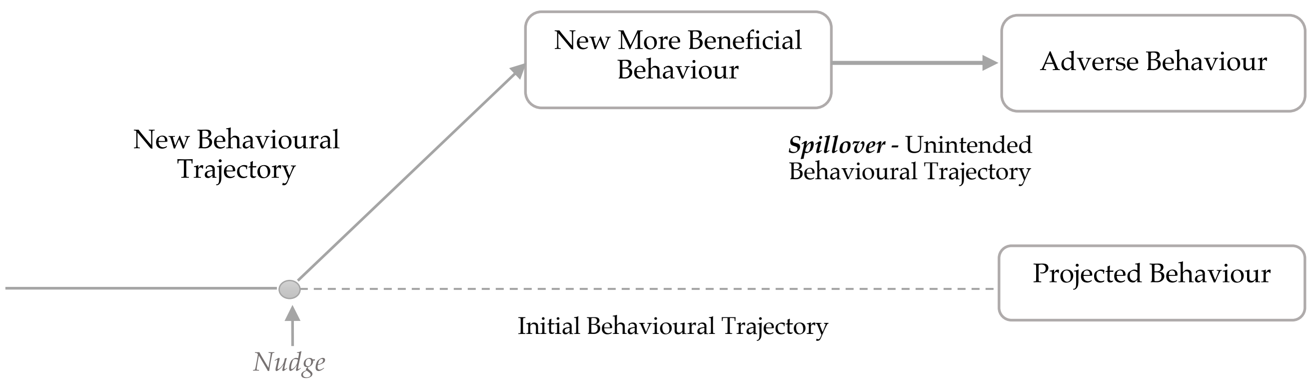

Diagram 2. *A Nudge Spillover behavioural pathway. (Source: Themistocleous and Karapanos, 2024)*

As an example, imagine that, following a prolonged period of sitting, an individual is nudged, by her smartwatch, to stand up. Adopting this behaviour is beneficial for the individual as prolonged periods of sedentarism have detrimental effects on blood circulation (Pedromo et al., 2019), and correlate with all-cause mortality independent of leisure time (Katzmarzyk et al., 2009). Imagine that the user responds to the nudge, stands up for her 1-minute walk, however, she walks into the office's kitchen and consumes a few cookies. A repetition of this behaviour creates a pattern with adverse effects: The individual uses the stand-up nudge as a cookie reminder, making it in the process, a habit. Assuming that the calorie intake from the cookie is not desirable for the user, the nudge is seemingly effective in its purpose (to prompt movement) yet it results to an unexpected negative effect, the unwanted increase of calorie intake.

AI can undertake a mitigating role here. First, assuming the presence of a continuous glucose monitoring device, machine learning models can detect emerging causal relationship between the nudge behaviour (*i.e.,* standing up) and the spillover (*i.e.,* calorie intake as inferred from the glucose monitor), and the contextual factors that play a role in such causal relationships. For example, this spillover might take place only in specific locations, such as the office due to the availability of cookies in the office's kitchen, or at specific points in time, such as after lunch due to reactive hypoglycemia. Second, once a spillover is identified, we propose two main courses of action: one that aims at preventing similar adverse effects in the future, a *Proactive Nudge* correction, and one that corrects current adverse behaviour, a *Reactive Nudge* correction. The two main tools for these adjustments are, as previously mentioned, timing and tailoring.

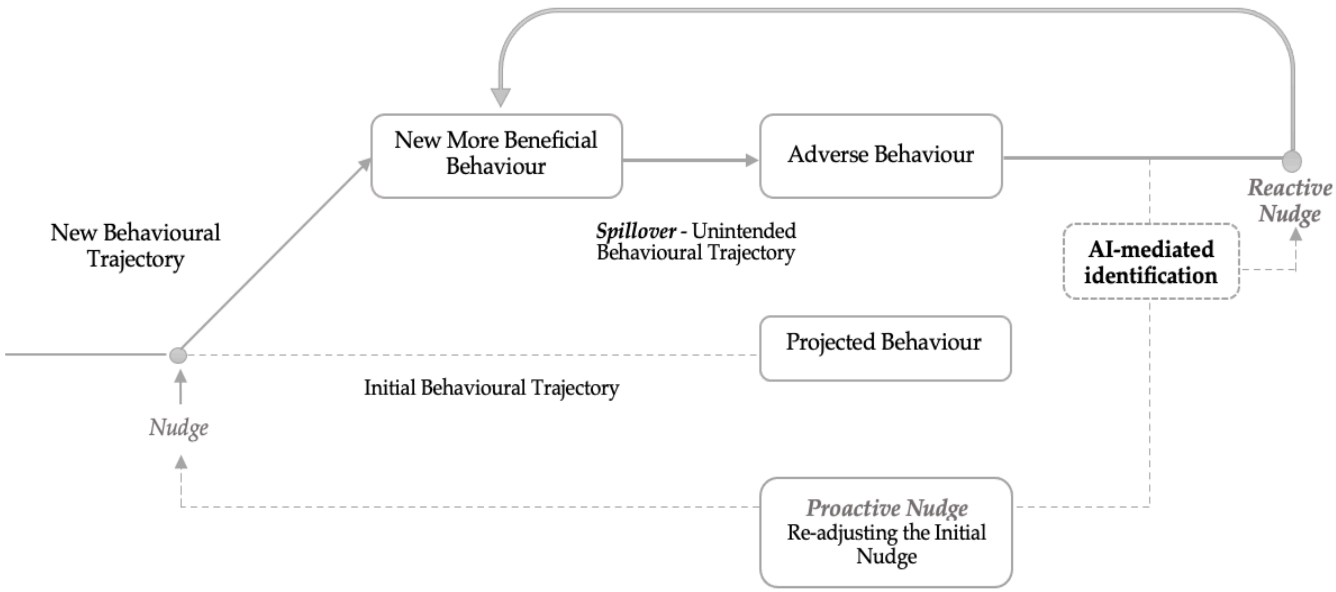

Diagram 3. *AI-mediated identification and corrective courses for Nudge Spillovers. (Source: Themistocleous and Karapanos, 2025).*

Diagram 3 reflects how following the AI-mediated identification of the spillover the two types of nudge corrections are triggered, *proaction* and *reaction*. A Proactive Nudge aims at future spillover prevention. Since the present type of nudge was ineffective in preventing the spillover, what is needed is to identify a different type that is equally effective in achieving the intended behaviour yet avoids the negative effect. In this case, a smart disclosure nudge that explains the benefits of standing up while simultaneously highlighting the negatives of sugar intake can reinforce the intended behaviour of standing and avoiding the cookie consumption. Timing, on the other hand, aims at triggering this type of nudge at times where the probability of a spillover is low, especially if the utility of standing up is overshadowed by the sugar intake. This timing readjustment for precision nudging can be used tactically to break bad habits while steering the reinforcement of beneficial ones.

A *Reactive Nudge* aims at correcting the current adverse behaviour caused from the spillover. This requires a new nudge and relates to the stage after the cookie consumption and upon the user's return to her desk. The user can be met with a new nudge prompting her to measure her sugar levels, contrasting those figures with the desirable daily intake tailored to her case. The sugar level nudge timed after cookie consumption can prompt the user to both re-adjust the sugar intake for the remainder of the day, reduce further consumption and simultaneously passively call for a behaviour re-evaluation that can affect the next behavioural cycle. Alternatives can also be offered in relation to other nudge types that mitigate this spillover effect or by bypassing the stand-up nudge during certain hours that cookie consumption seems imminent.

4.2 Nudge Misfires and AI

Nudges need to be tested to ensure they do what they are supposed to do (Thaler, 2015) yet due to reasons pertaining to emotional states of individuals at the time of the nudge or the frequency of nudges that can lead to reactance. Certain adverse behaviours might be triggered that deviate from the intended nudged trajectory. We call this a *Nudge Misfire* as the initial nudge can lead directly to an unintended adverse behaviour putting the individual at a deficit when compared to their initial behavioural trajectory. *Diagram 4* depicts the nudge misfire pathway.

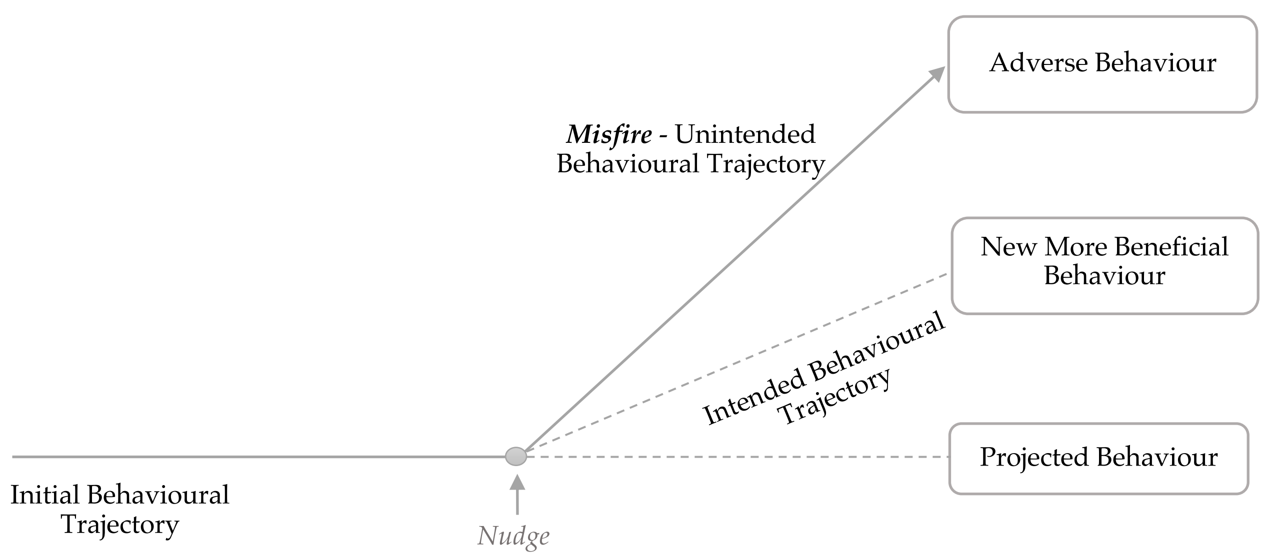

Diagram 4. *A Nudge Misfire behavioural pathway. (Source: Themistocleous and Karapanos, 2025)*

Using an example, this time a smartwatch-wearing user receives a nudge with the aim of prompting him to exercise. Here the type of nudge was a Social Comparison Nudge (similar to the one used by Opower mentioned in section 2) where the targeted individual is nudged with information that facilitate comparisons with others in a similar situations. As such the user is met with information indicating that his calorie-burning performance is in the lowest quartile (bottom 25%) from a similar pool of other smartwatch users - *"Out of 1000 similar users you are ranked*

920th for your calorie-burning performance". Albeit some findings report favourable comparative effects (i.e. Loewenstein et al., 2014) others report adverse ones (Samra et al., 2022) as comparisons can lead to low self-esteem and depressing states (White et al 2006).

The social comparison nudge here could have been effective for individuals in the top quartile putting them in a favourable, motivated position, yet individuals being informed of their position in the lowest tiers might not only lead them to disregard the exercise nudge, but also enter into a negative mood swing. In this moment the misfire took effect, making the user worse-off. Specifically, the individual might balance out the negative emotional effects with consumption of unhealthy sugar-rich foods, an association that is systematically documented in psychological and nutritional research (Paans et al., 2018; Heatherton and Baumeister, 1991). Here the nudge failed to direct the user into a better behaviour making things worse off for him in the process.

Similar with spillovers, triangulating the effects from multiple sources including the smartwatch and sugar meter is essential. In detecting phases of low self-esteem multiple markers can be employed. Two examples include sugar intake and heart rate. Low self-esteem and depressive states are associated with sugary comfort food consumption (Paans et al., 2018). Cardiovascular activity and specifically low heart-rate, under conditions, is linked with low self-esteem levels (O'Donnell et al., 2008). The inclusion of these and more parameters in the Machine Learning process can allow an AI to detect occasions where adverse effects have been recorded as a result of the initial nudge, primarily in the immediate phase after the nudge. After the identification Proactive and a Reactive Nudge Corrections are required. *Diagram 5* summarises the two AI-mediated nudge corrections.

The proactive correction requires a re-evaluation of the nudge type and its timing. In the previous case the social comparison nudge was responsible for the misfire thus in need to be substituted with a different nudge type. A nudge in the form of a reminder *"Let's Exercise"* can be more effective in facilitating a beneficial behaviour and mitigating the chance for an adverse unintended one. Timing wise, a social comparison nudge can be resumed when individuals have made their way in the first quartile (Top 25%) of similar performers where comparisons with others can likely work favourably in reinforcing intended behaviours.

The Reactive Nudge aims at correcting the adverse effects caused by the first nudge and complement the Proactive one. Once an individual enters in the low self-esteem stage as a result of the nudge misfire, albeit challenging, a new nudge needs to reinstate self-esteem and confidence. In this case a nudge prompting the user for taking a moment to reflect on past successes and practising *mindfulness* can improve the adverse effects caused. A follow-up prompt to exercise describing the benefits of exercising can lead back to the initial beneficial behaviour and correct the adverse course caused (Randal, Pratt and Bucci, 2015).

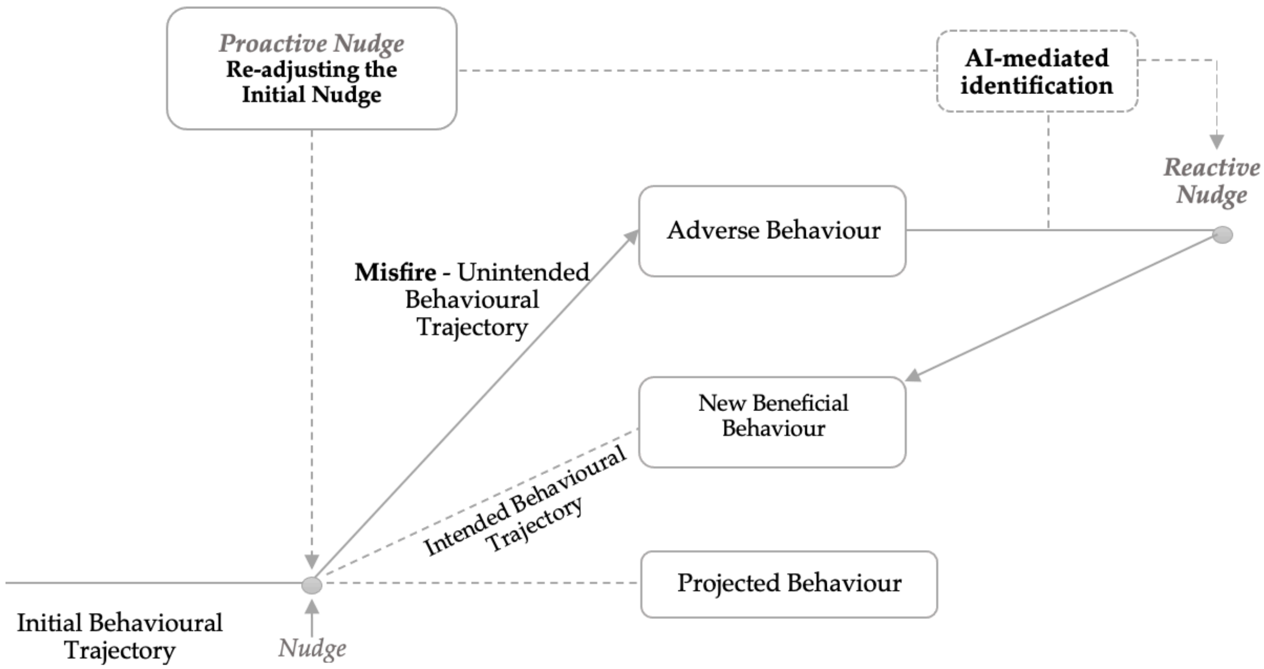

Diagram 5. *AI-mediated identification and corrective courses for Nudge Misfires. (Source: Themistocleous and Karapanos, 2025)*

5. Conclusions

Offering the ability to analyse behavioural data, identify patterns, and provide tailored interventions at the right time, AI offers the promise of more effective and precise behavioural interventions. In this chapter we discussed how AI, and in particular, reinforcement learning, can improve the effectiveness of nudging while mitigating some of its downfalls, namely *spillovers* and *misfires*. We analyse the use of two methods, *tailoring* and *timing,* and exemplified their application, using the conceptual tool of *behavioural pathways*.

As with all behavioural tools, nudges can fail. An individual might heed the call of a nudge yet end up in a worse off situation as a consequence of the heed. Likewise, a nudge might be ignored but its type or timing might lead to a miscalculated adverse effect generating a misfire. Better nudging serves as a seemingly straightforward solution to both problems but the implementation of the latter is not. Identifying the exact moment a nudge might be needed and heeded is challenging with conventional means. Here we discussed a pragmatic approach offered by Reinforcement Learning, and particularly, Multi-Armed Bandits, for the delivery of *tailored* and *just-in-time* nudging interventions. The implications of the proposed ideas are linked to persuasive technology designs and can further advance policy making that heighten an individual's benefit. Furthermore, the proposed precision nudges can revolutionise sustainable marketing practises for the promotion of environmental, economic and societal wellbeing by tailoring nudge types based on consumer data and psychological profiling (Themistocleous, 2023).

Our analysis highlights a number of implications and directions for future research. We conceptualise the terms of *Reactive* and *Proactive Nudges* when handling misfires and spillovers. Future empirical research should explore how reactive, corrective nudges can help mitigate adverse effects of a previously delivered nudge. When individuals enter induced nudge spillover and misfire situations, tests can be conducted to examine the effectiveness of Reactive Nudges for resolving the caused issues. Mapping certain nudge types that equalise the adverse effects of others can provide social and computer researchers with a nudge recovery toolbox for when nudges fail while MABs can generate precision nudges that provide solutions regarding when and in which form the Proactive Nudges can assist with future spillover and misfire prevention.

6. References

Baumeister, R.E. Bratslavsky, E. Muraven, M. and Tice, D.M. (1998). Ego Depletion: Is the Active Self a Limited Resource? *Journal of personality and social psychology*, 74 (5), 1252-1265.

Baumeister, R. F., & Leary, M. R. (2017). The need to belong: Desire for interpersonal attachments as a fundamental human motivation. *Interpersonal Development*, 57–89.

Bell, A.S., Rajendran, D., Theiler S. (2012). Job stress, wellbeing, work-life balance and work-life conflict among Australian academics, *Journal of Applied Psychology*, 8 (1), 25-37.

Caraban A., Karapanos E., Gonçalves D., and Campos R. (2019). 23 Ways to Nudge: A Review of Technology- Mediated Nudging in Human-Computer Interaction. CHI Conference on Human Factors in Computing Systems Proceedings (CHI 2019), May 4–9, 2019, Glasgow, Scotland UK.

Cialdini, R. B. (2009). *Influence: Science and practice.* Boston: Pearson education.

Csikszentmihalyi M. (2002). *Flow: The classic work on how to achieve happiness*. Rider.

Dolan O. and Galizzi M. (2015). Like ripples on a pond: Behavioural spillovers and their implications for research and policy, *Journal of Economic Psychology*, 47, 1-16.

Egebark, J. and Ekstrom M. (2016). Can indifference make the world greener? *Journal of Environmental Economics and Management,* 76, 1–13.

Gouveia R. Pereira F. Karapanos E. Munson S.A. Hassenzahl M. (2016). Exploring the Design Space of Glanceable Feedback for Physical Activity Trackers. In UbiComp 2016, September 12-16, Heidelberg, GER, 144-155.

Heatherton, T. F. and Baumeister, R. F. (1991). Binge eating as escape from self-awareness. *Psychological Bulletin,* 110 (1), 86–108.

Hummel, D., & Maedche, A. (2019). How effective is nudging? A quantitative review on the effect sizes and limits of empirical nudging studies. *Journal of Behavioural and Experimental Economics*, 80, 47-58.

Kahneman, D. (2011). *Thinking, Fast and Slow.* Penguin Psychology.

Kaptein, M., Markopoulos, P., De Ruyter, B., & Aarts, E. (2015). Personalizing persuasive technologies: Explicit and implicit personalization using persuasion profiles. *International Journal of Human-Computer Studies*, 77, 38-51.

Katzmarzyk, P. T., Church, T. S., Craig, C. L., & Bouchard, C. (2009). Sitting time and mortality from all causes, cardiovascular disease, and cancer. *Medicine & Science in Sports & Exercise*, 41 (5), 998-1005.

Loewenstein, G., Sunstein, C. R. and Golman, R. (2014). 'Disclosure: Psychology changes everything', *Annual Review of Economics*, 6, 391–419.

Mehrotra A., Vermeulen J. Pejovic V. Musolesi M. (2015). Ask, but don't interrupt: the case for interruptibility-aware mobile experience sampling. Adjunct Proceedings of the 2015 ACM International Symposium on Wearable Computers, 723-732.

Nahum-Shani, I., Smith, S. N., Spring, B. J., Collins, L. M., Witkiewitz, K., Tewari, A., & Murphy, S. A. (2018). Just-in-time adaptive interventions (JITAIs) in mobile health: key components and design principles for ongoing health behavior support. *Annals of Behavioural Medicine*, 1-17.

O'Donnell K. Brydon L. Wright C. aSteptoe A. (2008). Self-esteem levels and cardiovascular and inflammatory responses to acute stress, *Brain gf and Immunity*, 22 (8), 1241-1247.

Paans, S.P Bot, M. Brouwer, I.A. Visser, M. Roca, M. Kohls E. Watckins, E. and Penninx B. (2018). The association between depression and eating styles in four European countries: The MooDFOOD prevention study. *Journal of Psychomatic Research*, 108, 85-92.

Perdomo, S.J. Gibbs, B.B. Kowalsky, R.L., Taormina, J.M. Balzer J.R. (2019). Effects of Alternating Standing and Sitting Compared to Prolonged Sitting on Cerebrovascular Hemodynamics, *Sports Sci Health*, 15(2), 375-383.

Pielot, M., Dingler, T., Pedro, J. S., and Oliver, N. (2015). When attention is not scarce-detecting boredom from mobile phone usage. In *Proceedings of the 2015 ACM international joint conference on pervasive and ubiquitous computing*, 825-836.

Randal, C. Pratt, D. and Bucci, S. (2015). Mindfulness and Self-esteem: A Systematic Review, *Mindfulness*, 6, 1366–1378.

Raafat, R. M., Chater, N., & Frith, C. (2009). Corrigendum: Herding in humans. *Trends in Cognitive Sciences*, *13*(12), 420–428. https://doi.org/10.1016/j.tics.2009.09.002

Samra, A. Warburtin, W.A. and Collins, A.M. (2022). “Social Comparisons: A potential Mechanism linking problematic social media use with depression”. *Journal of Behavioural Addiction*, 11(2), 607-614.

Schwarz, N., Kahneman, D., Xu, J., Belli, R., Stafford, F., & Alwin, D. (2009). Global and episodic reports of hedonic experience. *Using calendar and diary methods in life events research*, 157-174.

Steffel, M., Williams, E. F. and Pogacar, R. (2016). Ethically deployed defaults: Transparency and consumer protection through disclosure and preference articulation, *Journal of Marketing Research*, 53(5), 865–880.

Sunstein, C.R. (2017). Nudges that fail. *Behavioural Public Policy*, 1(1), 4-25.

Thaler R. and Sunstein C. (2008). *Nudge Improving Decisions About Wealth, and Happiness*. Penguin Publications.

Thaler, R. (2015) *Misbehaving*, New York: W.W. Norton & Company, 309–345.

Thaler, R. H. (2018). Nudge, not sludge. *Science*, *361*(6401), 431. https://doi.org/10.1126/science.aau9241

Themistocleous, C. Wagner, C. and Smith, A. (2014). The ethical dilemma of implicit vs explicit data collection: Examining the factors that influence the voluntary disclosure of information by consumers to commercial organizations. Proceedings of the *2014 IEEE on Ethics in Social Sciences and Technology*. ISBN: 978-1-4799-4992-2.

Themistocleous C., (2023), *Sustainable Marketing*, in: Mathews L., Bianchi L. and Ingram C. (Ed.), *Concise Encyclopaedia of Corporate Social Responsibility*. Edward Elgar Publishing.

Tversky A. and Kahneman D. (1974). Judgment under uncertainty: Heuristics and biases. *Science,* 185(4157), 1124–1131.

Warberg L., Acquisti. A., and Sicker D. (2019). Can Privacy Nudges be Tailored to Individuals' Decision Making and Personality Traits?. In 18th Workshop on Privacy in the Electronic Society (WPES'19), November 11, 2019, London, UK https://doi.org/10.1145/3338498.3358656

White, J.B. Langer, E.J. Yariv, L. Welch J.C. (2006). Frequent Social Comparisons and Destructive Emotions and Behaviors: The Dark Side of Social Comparisons. *Journal of Adult Development*, 13(1), 36-44.

Zedeck, S. (Ed.). (1992) *Work, families, and organizations.* Jossey-Bass/Wiley.

BIOS

Dr Christos Themistocleous is a Lecturer of Consumer Behaviour at the Department of Communication and Marketing, Cyprus University of Technology (CUT). He holds a PhD in Consumer Behaviour (University of Nottingham), MSc in Corporate Strategy and Governance (University of Nottingham) and BA(Hons) in Finance, Accounting and Management (University of Nottingham). Having worked previously in faculty positions at the University of Reading, his research focuses on Disclosure Decision-Making, Consumer Privacy and Data Collection Systems.

Dr. Evangelos Karapanos is an Associate Professor in the Department of Communication and Internet Studies and the Director of the Persuasive Technologies Lab at the Department of Communication and Internet Studies of the Cyprus University of Technology (CUT). Dr. Karapanos holds a PhD (cum laude) in Human-Computer Interaction from Eindhoven University of Technology, a MSc in Human-Computer Interaction from University College London, and a BSc in Physics with a specialisation in microelectronics from the University of Patras. His expertise is in experience-centred design of interaction with technology. His work is regularly published at top-tier conferences and journals in Human-Computer Interaction and has received more than 5000 citations in the literature and a number of awards, including a best paper and a best poster award at Persuasive Technology 2017, an honourable mention for best paper at Ubicomp 2015, the tech challenge for atrocity prevention 2013 Award, and the cum laude distinction for his PhD thesis (given to the top 5% of PhD theses defended annually at the Eindhoven University of Technology).