THE PINCHAS SAPIR CENTER

FOR DEVELOPMENT

Tel Aviv University

המרכז לפיתוח ע"ש פנחס ספיר
ליד אוניברסיטת תל אביב
"עמותה רשומה"
580010221

THE PINHAS SAPIR CENTER FOR DEVELOPMENT
TEL AVIV UNIVERSITY

Behavioral Spillovers from Promoting Healthier Consumer Choices

Mathias Wagner Barløse, Kfir Eliaz, Neil Thakral, and Sarit Weisburd

Discussion Paper No. 12-2022

אוניברסיטת תל-אביב, תל אביב יפו 6997801 | טל': 03-6409202 | פקס: 03-6405815
Tel Aviv University, Tel Aviv Yafo 6997801 Israel | Tel. 972-3-6409202 | Fax 972-3-6405815 | email: sapircen@tauex.tau.ac.il

Behavioral Spillovers from Promoting Healthier Consumer Choices

Mathias Wagner Barløse[1], Kfir Eliaz[2], Neil Thakral[3], and Sarit Weisburd[4]

[1]Aarhus University

[2]Tel Aviv University and the University of Utah

[3]Brown University

[4]The Hebrew University

September 2022*

Abstract

This project analyzes data from a large-scale randomized controlled experiment run on about 10,000 shoppers on an online grocery platform. During the four-month intervention period of the Swap and Be Healthy experiment, 50 per cent of the users were randomly allocated to the Swap and be Healthy treatment nudge. When these treatment shoppers added a less healthy product to their basket, a popup appeared with the option to swap to a healthier substitute (lower fat, less salt, etc.) with the click of an icon. We find that the nudge had a significant effect on the purchase of healthier products and that this effect persists over time even after the intervention period ends. Interestingly, the nudge also tended to decrease the price sensitivity of shoppers, reduce their

*Barløse: Department of Economics and Business Economics, Aarhus University (email: mathias@barlose.dk). Eliaz: Eitan Berglas School of Economics, Tel Aviv University and David Eccles School of Business, the University of Utah (email: kfire@tauex.tau.ac.il). Thakral: Department of Economics, Brown University (email: neil_thakral@brown.edu). Weisburd: The Hebrew University Business School (email: sarit.weisburd@mail.huji.ac.il). We thank Elior Cohen, Maya Fuks, and Isabella Kuhr for excellent research assistance. Special thanks to Bernard Feldman for his assistance with the data. Research funding from the Pinhas Sapir Center, and the Israel Science Foundation is gratefully acknowledged.

demand for produce, and cause them to shorten both the length of their trip and the size of their purchase. We apply machine learning analysis to examine what types of consumers are most impacted by the intervention, as well as the relationship between a shopper's propensity to respond to the nudge (by adding a healthier product) and the overall effect of the nudge on shopping behaviour. We find that this overall effect is present across all treatment shoppers, even those that were unlikely to respond to the nudge directly by adding a healthier product to their basket. We examine multiple mechanisms that could explain this result and conclude that guilt-costs may have played an important role in the impact of the healthy nudge on overall shopping behavior.

1 Introduction

Obesity has been one of the major public health challenges of the past several decades, affecting both adults and children, and households in both developed and developing countries. The World Health Organization (WHO) formally recognized obesity as a global epidemic in 1997 as it began "replacing the more traditional public health concerns, including undernutrition and infectious disease as one of the most significant contributors to ill health" (WHO, 2000), with over 1 billion adults now being overweight and at least 300 million clinically obese. The "fundamental causes" are sedentary lifestyles and unhealthy diets, with diet playing a particularly important role as specific nutrients—high intakes of saturated fat, sodium, added sugar, and low intakes of fibre, fruits, and vegetables—contribute to disease and mortality beyond their impact on body weight (WHO, 2000). Nearly 8 million deaths worldwide are attributable to dietary risk factors (Murray et al., 2020) and direct medical costs in the U.S. alone exceed $250 billion (Cawley et al., 2021).

The predominant public policy response to promote healthier consumer choices is information provision. Most notably, nutrition education and food labels reach large audiences. The vast majority of Americans indicate that they have heard of the U.S. Food Guide Pyramid, even though few follow its recommendations (Guthrie, Mancino and Lin, 2015). About 80 percent of U.S. adults report using nutrition facts labels and 72 percent report using nutritional content claims (e.g., "low fat") when making purchasing decisions (Choinière and Lando, 2008; Guthrie, Mancino and Lin, 2015). Puzzlingly, however, excessive intakes of saturated fat, sodium, and added sugar persist despite the significant consequences and the availability of information.

More recent food guidelines, motivated by behavioural insights, come with "clear and actionable" suggestions, e.g., "make half your plate fruits and vegetables," "drink water instead of sugary drinks," and "switch to fat-free or low-fat (1%) milk" (Sunstein, 2013*a*). The latest version of *Dietary Guidelines for Americans*, published by the U.S. Department of Agriculture (USDA) and Department of Health and Human Services (HHS), recommends to "start with small changes to make healthier choices you can enjoy." One of the developers of U.S. dietary guidelines emphasizes that "You can still choose foods that you enjoy, but you need to align them with healthy eating patterns: less sugar, sodium, and saturated fat... Making small changes in your diet over time... can pay off in the long run" (NIH, 2016). Similar recommendations appear in dietary guidelines outside the U.S. as well.[1] However, making dietary choices that moderate the consumption of saturated fat, sodium, and added sugar remains difficult as significant quantities are "hidden" in everyday foods, with numerous examples cited in scientific research (Ponzo et al., 2021), the media (ABC, 2012), and messaging from governmental agencies (CHP, 2017). Insights from marketing suggest the need to go a step further than the existing guidelines-based approach by developing new technologies that provide consumers with customized, personally relevant nutritional information during their shopping trips (Lowe, de Souza-Monteiro and Fraser, 2013).

This paper uses comprehensive data on over 14,000 online shoppers to study a novel intervention that provides consumers with information when they add less healthy versions of certain food items to their shopping baskets. The intervention, which we refer to as "Swap and be Healthy" (SABH), provides specific, actionable, brief health information at the time of decision. The intervention is *informative* (shoppers learn about an item's nutritional content), *shopper-specific* (information pertains to a product the shopper is actively considering), *actionable* (shoppers receive specific suggestions for healthier alternatives), *brief* (information is conveyed clearly and concisely), and *timely* (shoppers receive information while they are making decisions).

In partnership with an established online platform for supermarket shopping across all major retailers in Israel, we conducted a randomized field experiment to evaluate the SABH intervention. Shoppers randomly assigned to the intervention group receive information about particular nutrients and alternative products when they add to their shopping basket one of 57 common food items that a registered dietitian identified to have healthier alternatives. Over 9,588 shoppers added at least one of these products to their baskets during the three months of the experiment. The information consists of a simple statement about the improved nutrient profile of the alternative(s): less sugar, less saturated fat, less sodium, more fiber, lower glycemic index, and more iodine.[2] About 16 percent of shoppers receiving SABH information switched to a

[1]For example, the National Health and Medical Research Council in Australia notes (emphasis added), "Many of the health problems in Australia today are linked to poor eating habits. Too many people eat too much saturated fat, added salt, added sugars and alcohol. *Even reducing these by small amounts can make us healthier.* It can help us manage our weight better and reduce our risk of chronic diseases like heart disease, stroke, Type 2 diabetes, some cancers and chronic kidney disease" (NHMRC, 2015).

[2]Iodine deficiency, while uncommon in the U.S. is a "major public health problem worldwide" that affects approximately 2 billion people, with about 50 million exhibiting clinical symptoms, and is particularly concerning for pregnant women and children (Biban and Lichiardopol, 2017). In Israel, a national survey of iodine intake in 2017 revealed that 85 percent of pregnant women and 62 percent

healthier alternative during the intervention period.

However, a change in the percentage of accepted swap-offers does not capture the full extent of the impact of the intervention on demand. This is because a shopper who accepted the swap during their first trip but in later baskets chose the healthier product will not appear in this statistic. We, therefore, apply a difference-in-differences analysis comparing demand across treatment and control shoppers during the intervention in contrast to pre-intervention. We find that the intervention increased demand for the healthier products by 26 percent, with a corresponding decrease in demand for the less healthy products. The difference-in-differences analysis also provides evidence that the SABH had a persistent effect on demand even after the intervention had ended. During the three months after the experiment, purchases of the healthier alternatives (those items recommended by the SABH prompts) continued to be 7 percent higher for treated shoppers.

A unique aspect of our data is that it allows us to examine the impact of the SABH pop-up on the demand for other products. This enables us to estimate the overall consequences of our intervention on shoppers. Our results suggest that the pop-up impacts the type of products that are added to baskets, as well as the cost and size of baskets, and the probability of switching to an alternative supermarket. In particular, we find that after a shopper receives the first SABH prompt, he decreases the purchases of items in the fruit and vegetable category by 2.5 percentage points (s.e. 0.7) at an average purchase rate of 26 percent. He also chooses products that have 5 percent more saturated fat and 10 percent more cholesterol per serving (a 0.17 (s.e. 0.06) increase in saturated-fat at an average rate of 4, and a 0.002 (s.e. 0.0005) increase in cholesterol at an average cholesterol rate of 0.02).

We also observe a 2 percent increase (s.e. 0.8) in the price of products that are added to the basket, where the average observed price of a product is 10 NIS. Relatedly, there is a 5 percentage point decrease (s.e. 0.9) in the probability of switching a supermarket during a shopping trip at an average switch rate of 10 per cent. In addition, shoppers spend on average 37 NIS less on a basket of goods costing on average 572 NIS and decrease their baskets (with an average of 77 items) by 6 items (s.e. 1.2). These findings suggest that while the pop-up promoted healthier products, it also had

of school-age children have insufficient iodine levels, putting adults at risk for thyroid disease, and putting fetuses, infants, and children at risk for impaired neurocognitive development (Ovadia et al., 2017; Lazarus, 2020).

negative externalities regarding the consumption of produce, overall nutritional value, price-per-unit costs, and the likelihood of switching to a cheaper supermarket.

To examine whether our intervention had a differential effect on different shopper types, we use the recently developed causal forest technique (see Athey and Imbens, (2016), Wager and Athey (2018) and Athey et al. (2019)), which searches across covariates for good predictors of heterogeneous treatment effects. These are described by Conditional Average Treatment Effect (CATE) estimates, which are the expected effects of a treatment for individuals in subpopulations defined by covariates. Following Chernozhukov et al. (2018), we characterize those shoppers who are most and least responsive to our intervention.

Not surprisingly, consumers who are more likely to purchase relevant products are more likely to be impacted by the intervention. We find that the consumers who were most impacted by the intervention were those consumers who seemed to be more prone to less healthy purchasing habits in the pre-intervention period. As such, the most responsive shoppers are more likely to have added something unhealthy to their basket, and more likely to have purchased alcohol during the pre-intervention period. Further, children seem to make consumers more prone to embrace the healthy nudges provided by the intervention. Lastly, the most affected consumers were less likely to buy on-sale products or check for cheaper alternatives and generally spent more on their shopping trips in the pre-intervention period.

Indeed, when we run our difference-in-differences analysis on the shoppers most-likely to respond to the intervention based on their CATE estimates, we find a larger response to the intervention, namely a 6 percentage point increase in purchases of healthier alternatives during the intervention (we estimate an effect of less than 1 percentage point on the full sample), and an increase of 1.5 percentage points in the months following the intervention (we estimate an effect of 0.25 percentage points for the full sample). Shoppers predicted to be least-likely to respond to the intervention, have the opposite response to the SABH nudge - and are 3.6 percentage points less likely to purchase a healthier product during the intervention than untreated shoppers in the control group. This effect persists even after the intervention, when these shoppers continue to be half a percentage point less likely to purchase healthy than those in the control group.

Interestingly, when we consider the overall consequences of the SABH pop-up, we find that both shoppers most likely and least likely to respond to the intervention by buying the healthier alternative respond similarly in their demand for products outside of the intervention. In other words, even shoppers who were very unlikely to respond to the healthy pop-up by purchasing the healthier product, bought less produce, and purchased products with higher levels of saturated-fat and cholesterol. Similarly, both types of shoppers were less likely to change supermarkets, purchased less products, and spent less on their shopping basket.

These similar outcomes for those shoppers most and least likely to respond to the intervention help to shed light on the mechanism driving these behavioural changes. One explanation for the decrease in consumption of produce could be that consumers have a target for healthy consumption; however, the SABH intervention crowds out fruits and vegetables.[3] This would not explain why shoppers who don't respond to the intervention still decrease their fruit and vegetable consumption. An alternative explanation is that processing the additional information provided by the pop-ups, and solving the decision problem induced by these pop-ups (to switch or not to switch) may deplete shoppers' cognitive resources. This interpretation draws on previous studies that showed people's ability to exert self-control, or to make careful deliberate choices is hampered when following a prior decision-making task that requires some cognitive effort (see Muraven and Baumeister (2000) with regards to self-control, and Pocheptsova et al. (2009) with regards to decision-making). This mechanism would suggest that the effect should be strongest for shopper's who we observe making more decisions prior to receiving the health popup. However, we don't find a stronger effect for shoppers who have already considered alternative products to those added to their shopping trip prior to receiving the health nudge, or on shoppers who are further along in their shopping trip.

Our results are most in line with a scenario where the SABH popup triggers negative feelings of guilt that cause shoppers to avoid the shopping experience. We present a simple model whereby consumers with high guilt costs related to nutrition, or who practice self-control try to avoid situations that will trigger guilt or the need for self-control. If the health nudge caused shoppers with high guilt-costs to feel bad when they didn't choose the healthier alternative, this could cause shoppers to end their shopping trip more quickly to avoid another negative experience. In their haste to exit the site, they may be less likely to look for cheaper alternatives, or make healthy

[3] See work on taxi drivers, who quit working when reaching a certain threshold of earnings (Camerer et al. (1997) and the more recent study by To and Thakral (in press)).

product choices. We identify consumers with higher guilt costs in the pre-intervention period as those who added junk-food to their basket, and later removed it prior to purchase.[4] Indeed, we find the strongest effects among shoppers who exhibited guilt or self-control costs in the pre-intervention period. This raises the concern that even a positive intervention like the SABH can make shoppers worse off.

Our paper contributes to three broad areas of existing work, including policy-oriented, experimental, and theoretical research. The first area of related work focuses on food policy. Existing approaches for promoting healthier food consumption such as information provision (Balasubramanian and Cole, 2002), incentives (Loewenstein, Price and Volpp, 2016; Olsho et al., 2016; Griffith, von Hinke and Smith, 2018), and nudges (Wilson et al., 2016; Cadario and Chandon, 2020) face several challenges. Complexity poses a significant barrier for information-based approaches such as descriptive nutritional labelling (Guthrie, Mancino and Lin, 2015).[5] Simpler forms of information such as evaluative nutritional labelling may also be ineffectual due to tastes or behavioural factors such as habits, present focus, or inattention (Horgen and Brownell, 2002). More importantly, any attempt to promote healthier food consumption in one area may result in *behavioural spillovers* or compensatory changes in other areas. Such spillovers can undermine attempts to promote healthier food consumption via information, incentives, or nudges, and possibly even lead to the opposite conclusion about their efficacy. An important advantage of our data is that we observe the decision maker's entire consumption bundle, beyond the set of products considered in the experiment.

Our analysis contributes to a broader literature on behavioural spillovers (Dolan and Galizzi, 2015; Galizzi and Whitmarsh, 2019). Several studies document spillover effects within the domain of food choice on a small scale.[6] Wilcox et al. (2009) documents that the presence of a healthy option can lead to greater indulgence. Wisdom, Downs and Loewenstein (2010) find that providing calorie information at a

[4] We define junk food as chocolate, candies, chewing gum, dough, pastries, pizza, bread, cereal with added sugar, cookies, jams, sweet spreads, chips, pretzels, ice cream, drinks with added sugar, cakes, waffles, and popcorn.

[5] Perhaps not surprisingly then, the literature shows mixed results on information about calories (Downs et al., 2013; Cawley, Susskind and Willage, 2020).

[6] Another set of papers examines spillover effects across domains, such as food choice and exercise or effort (Polivy, Herman and McFarlane, 1994; Urbszat, Herman and Polivy, 2002; Chiou, Yang and Wan, 2011; De Witt Huberts, Evers and De Ridder, 2012; Dolan and Galizzi, 2014; Buyalskaya and Shum, 2020).

fast-food sandwich chain causes a reduction in sandwich calories but does not reduce total calories. Griffith, von Hinke and Smith (2018) show that government-provided vouchers for fruit, vegetables, and milk do not lead to offsetting changes in spending on unhealthy foods (i.e., added sugar and saturated fats) but slightly increase total calories. Relative to this set of papers, our work not only measures behavioural spillovers using a large-scale experiment with detailed data but also characterizes the consumers who exhibit such compensatory effects. Moreover, previous theoretical (Nafziger, 2020), empirical (Trachtman, 2022), and experimental (Altmann, Grunewald and Radbruch, forthcoming) work proposes interpretations of behavioral spillovers based on attention. We provide tests of attention-based explanations against a novel guilt-based mechanism for the spillover effects we observe and find support for the latter.

Our study also contributes to research on the design of information interventions. We develop a novel information-provision intervention by combining insights about effective design from previous experiments. In the marketing literature, Lowe, de Souza-Monteiro and Fraser (2013) estimate using a discrete choice experiment that 88 percent of respondents are willing to pay for hypothetical new technologies that would provide customized information such as nutrition alerts during their shopping trips.[7] Our intervention provides information on products the shopper is actively considering, following research on the importance of providing personalized rather than generalized information (Kling et al., 2012; Hoxby and Turner, 2013; Herber, 2018; Arteaga et al., 2022) as well as research emphasizing the timeliness of information provision (Hulshof and De Jong, 2006; Nahum-Shani et al., 2018).[8] The information, unlike that of nutrition labels, is simple (Sunstein, 2013*b*; Bhargava and Manoli, 2015) and consists of "clear and actionable" (Sunstein, 2013*b*) suggestions for alternative products. To maximize the usefulness of the information, we do not provide information about products in categories where such information would be unlikely to change consumers' prior expectations, such as chocolates or cookies (Araya et al., 2022). The subsequent

[7]While this stated preference approach is widely used in marketing, see Drake, Thakral and Tô (2022); Drake et al. (2022) and the references therein for examples across a broad set of fields within economics as well as recent methodological developments.

[8]Other recent work emphasizes that providing information in advance of making decisions can lead to healthier food choices (Brownback, Imas and Kuhn, 2021) or generally more forward-looking behaviour (Thakral and Tô, 2020; Imas, Kuhn and Mironova, forthcoming; Thakral and Tô, 2022); however, in these settings, information consists of upcoming choice sets, whereas in our setting, information consists of specific attributes of the choice alternatives.

adoption of our intervention as a permanent feature of the online platform corroborates the practical value of this combination of design features.

Our work also contributes to recent literature that uses field experiments in the context of online shopping to evaluate strategies for promoting healthier food choices. These papers tend to analyze the effects of different labelling schemes on product demand, finding mixed results (Sacks et al., 2011; Finkelstein et al., 2019; Finkelstein, Ang and Doble, 2020; Shin, van Dam and Finkelstein, 2020). The few papers that consider the impacts of suggested product alternatives or default options (Huang et al., 2006; Coffino, Udo and Hormes, 2020) do not capture possible spillover effects.[9] In comparison, our paper characterizes the mechanisms through which suggested product alternatives affect food-related decision-making more broadly.

Finally, we contribute to the extensive literature on intertemporal decision-making, particularly in the context of health-related choices. This literature primarily analyzes present-focused preferences (Ericson and Laibson, 2019), emphasizing that choices about immediate consumption tend to be inconsistent with longer-term goals such as health (Read and Van Leeuwen, 1998). For example, the survey article by Wilson et al. (2016) notes that field interventions for increasing healthier food choices take place almost exclusively in cafeterias, laboratories, and restaurants. The majority of food spending, however, occurs online or in supermarkets and therefore reflects decisions regarding future consumption.[10] Understanding the behavioural phenomena underlying unhealthy consumption requires richer theoretical frameworks that capture the basic features of making choices over bundles of goods in shopping environments, occurring in advance of actual consumption. A conceptual contribution of our work is to view supermarket shopping as a choice over menus, where actual consumption takes place in the future. By contrast, frameworks on immediate gratification would predict that making choices about future consumption largely eliminates self-control issues.

The remainder of the paper is organized as follows. Section 2 describes the experimental design and the data, and presents summary statistics on shopping

[9]In simulated online supermarket environments, treatments in which subjects are prompted with healthier product alternatives also find mixed effects (Forwood et al., 2015; Koutoukidis et al., 2019; Riches et al., 2019; Bunten et al., 2021; Jansen, van Kleef and Van Loo, 2021).

[10]In fact, for populations with the highest obesity risks, food away from home shows no association with Body Mass Index (Drichoutis, Nayga and Lazaridis, 2012; Crespo-Bellido et al., 2021), further highlighting the importance of studying decisions on the food that is not immediately consumed.

behaviour. Section 3 introduces our empirical strategies. Section 4 presents our main analysis, where we estimate the impact of the intervention. Section 5 considers heterogeneous treatment effects and identifies shopper types who are most and least likely to be impacted by the intervention. Section 6 discusses possible mechanisms, and Section 7 concludes.

2 Swap and Be Healthy Intervention

In this section, we describe the platform where we conducted our field experiment, the nature of the experiment, and the data collected.

2.1 The environment

Our field experiment was conducted on a popular platform for online grocery shopping in Israel. The platform provides the consumer with a "smart" shopping experience that allows easy comparison of purchase costs for the given basket across alternative online supermarket chains. Additionally, when adding a product to the basket, the platform automatically prompts the consumer with cheaper alternatives if available. Thus, our intervention focused on incorporating health into the "smart" shopping experience.

The platform offers two unique features aimed at helping shoppers make more informed decisions. One feature, which does not exist on any online supermarket website, allows shoppers to compare prices across supermarkets. When a shopper creates his basket he can observe how much his current basket would cost in competing supermarkets. For this calculation, if some items - like generic brands - do not exist in the competing supermarket, the platform replaces them with close substitutes. A shopper can switch supermarkets (in which case his entire current basket is transferred) as many times as he likes before checking out. Any product that is not available in the new supermarket is removed from the basket.

A second feature called "Swap and Save", is meant to help shoppers save money within each supermarket. Whenever a shopper adds to his basket an item for which either there is a cheaper close substitute (as determined by a proprietary algorithm developed by the platform), or there is a quantity discount, a "Swap and Save" button appears on the screen. If the button is pressed, the shopper is shown options that can lower the unit price for the product that he added to his basket (i.e., either cheaper brands or higher quantity). The shopper can then decide whether to swap the item he chose with one of the recommended options, or to leave the item and continue shopping.

Shoppers need to register to use the platform, but registration is free. Registered shoppers who log in to this platform (either from a mobile device or a laptop/desktop) can choose to start their shopping trip in one of several supermarket chains (our data covers the six major supermarket chains in Israel).[11] Once a shopper selects a supermarket, he can start adding items to his basket. Items can be added from previous shopping trips, from promotional banners, from a menu of categories (e.g., dairy, meat, fish, produce, etc.) or from or from a direct search for specific items.

All the activities that shoppers perform on the platform are recorded. This includes the source of each item added to the basket (e.g., from the shopper's previous basket or a promotional banner), items that were added but then removed, the sequence in which items were added, swap and save prompts that were observed and those that were accepted, the time between each item that was added and the total duration of the shopping trip, supermarket switches, and the final basket that was sent to the retailer site.

2.2 The experiment

Together with a registered dietitian, we identified 65 food items in 15 categories of staple foods (e.g., milk, pretzels, pudding, and soup) that had healthier alternatives in terms of having less sugar, less saturated fat, less sodium, lower glycemic index, more fiber, and added iodine. Table 1 provides the full list of selected items, their healthier alternatives and the reason each alternative was healthier.

In our selection of items, we had two objectives in mind. First, we focused on nutritional components for which overconsumption (or underconsumption in the case of iodine) is associated with an increased likelihood of diseases. In particular, according to the Center for Disease Control and Prevention (CDC) in the U.S., excess consumption of sodium can increase blood pressure and the risk for heart disease, while added sugars in food contribute to health problems such as weight gain and obesity, type 2

[11] While the platform could be to examine prices in any local supermarket, our analysis focuses on baskets created for on-line orders as these were sent directly to the retailer site upon completion.

diabetes, and heart disease.[12] Saturated fats (fats found in dairy products) have twice as many calories per gram as carbohydrates and protein, and studies have shown that they can raise the risk for heart disease and stroke. Because of this, the American Heart Association recommends limiting intake of saturated fats.[13] Diets with low glycemic index values improve the prevention of coronary heart disease in diabetic and healthy subjects, while in obese or overweight individuals, low-glycemic index meals increase satiety and facilitate the control of food intake (Rizkalla, Bellisle and Slama, 2002). Iodine deficiency has multiple adverse effects on growth and development and is the most common cause of preventable intellectual disability in the world. The primary source of iodine in food is iodized salt.[14]

Second, we chose healthier alternatives that are sufficiently similar in characteristics and price to the unhealthy target item such that shoppers would be more likely to switch to the healthier alternatives. Thus, for instance, we did not pair jasmine rice (which has a high glycemic index) with brown rice, but with another type of white rice that has a lower glycemic index (basmati rice). The average price difference between the unhealthy product and its alternative was 1.7 (the healthy being the most expensive). For roughly half of the selected categories, the healthier alternative was even cheaper than the unhealthy target item.

We then randomly allocated 14,282 shoppers into treatment and control groups. For the treatment group, the platform displayed the SABH pop-up whenever they added to their basket one of the 57 unhealthy target items (unlike the "swap and save" feature that appeared only if a shopper pressed a button). As Figure 1 illustrates, each pop-up showed a menu of healthier options including the price of each option and the reason for being healthier (e.g., lower fat, no added sugar). Shoppers were then able to substitute for the healthier option with the click of a button.

The experiment lasted for three months, and for each shopper in our sample, we have data on any basket created on the platform in the four months preceding the intervention and also during three months after the intervention ended.

[12]See https://www.cdc.gov/heartdisease/sodium.htm and https://www.cdc.gov/nutrition/data-statistics/added-sugars.html

[13]See https://www.heart.org/en/healthy-living/healthy-eating/eat-smart/fats/saturated-fats

[14]NIH, Office of Dietary Supplements.

2.3 Data

We focus our analysis on five databases from the online platform that provide information on the basket choices of individual shoppers, as well as the prices of alternative products both within the supermarket they shopped, and among alternative online suppliers. The "Products-Added" database records the order in which each product was added to the basket over each shopping trip including where the product was added from (free search, previous purchase, favourites, etc.), the price of that product, whether or not it is discounted, quantity purchased, whether the shopper considered a cheaper alternative when prompted, whether the shopper accepted a cheaper alternative, and whether a shopper accepted a healthy alternative for the product. The "Retailer" database provides the price of each purchased basket at each of the online retailers. The "Basket" database provides information on shopping time as well as the device used for shopping (app, computer, mobile-web). The "Swap" database provides information on the original item added to the basket and its price as well as the offered swap item and its price, including the reason for swap (health, price, or because the item is unavailable) as well as whether the swap was accepted. The "Prices" database provides data on all prices of all items at all stores during the experiment period to calculate prices of unchosen alternative items.

The SABH nudge was applied across the following 15 food categories: rice, canned corn, chocolate milk, chocolate pudding, yogurt, date-honey, jam, mayonnaise, salt, soy sauce, milk, nuts, pretzels, soup powder, and soy milk. Figure 3 highlights the differences in popularity across the different categories where the nudge was applied. The categories where shoppers are most likely to select a product are milk and rice, where products such as date-honey and nuts tend to be less frequent purchases.

Using these data, we construct a set of consumer characteristics based on their shopping behavior in the pre-intervention period. A shopper who did not make a purchase in this period is dropped from the data. These characteristics fall under three main categories: price sensitivity (thriftiness), health preferences, and lifestyle (Sacco et al., 2017).

It seems reasonable that general health awareness could play a role in a shopper's response to the SABH intervention. Thus, one might expect healthy types to be most responsive to the intervention as health is a priority for them. Alternatively, healthier shoppers may already be aware of these healthier alternatives and therefore face little benefit of the added nudge. To measure health awareness we create the following measures from shopping behavior in the pre-intervention period: the fraction of healthier alternatives purchased out of all relevant products, the fraction of baskets that include at least one junk-food or alcohol or cigarette product, the fraction of trips where the shopper purchased produce before junk food, the fraction of trips where junk-food was the last product purchased, and the fraction of junk food products that were added to basket and removed before purchase.

Lifestyle may also play an important role in determining response to the intervention. Thus, perhaps families with young children may benefit more from SABH if they have limited time and are concerned about health. Degree of religiosity, employment, or amount of free time may impact whether users benefit from the health nudge. Lastly, the degree to which shoppers purchase products relevant to the intervention will directly impact the intensity of treatment. We therefore include the following lifestyle variables in the analysis: the day of week when purchases were generally made, whether the shopper shopped during business hours, the fraction of products in basket that were relevant for the intervention, whether the shopper purchased products that have ultra orthodox (kosher) supervision, and whether the basket contained baby products.

Lastly, it is important to consider the correlation between price sensitivity and health sensitivity. Interestingly, in roughly half of the cases the healthier product was also cheaper, in which cause more price sensitive consumers may be more likely to respond to the SABH nudge. Additionally, more price sensitive shoppers may be more open to trying new products. However, when the alternative product is more expensive, these same consumers may be more hesitant to respond to the nudge.[15]. Price Sensitivity is determined based on: whether the average shopping cart is above or below the median value, the average price paid per product, the fraction of Swap & Save opportunities embraced by the shopper, the fraction of products that were purchased at full price, and the fraction of trips where the shopper chose to make her purchase at the supermarket with the lowest basket price. The full list of features and their definition can be found in Table 4.

In Table 2 we provide some summary statistics for the 8,625 shoppers in the treatment and control groups observed in the pre-intervention period. Generally, shoppers spend roughly 600 NIS per basket and purchase about 85 products in 30

[15]In order to account for price differences (and other factors that may impact SABH response) across product categories we also include category fixed effects in the analysis

minutes. The average shopper shops on the site 1.5 times per month. Shoppers take advantage of sales for about 10 percent of their purchases, 15 percent of shoppers choose to purchase their supermarket basket at the cheapest supermarket. Roughly 20 percent of the average supermarket basket is made up of fruits and vegetables, and 12 percent of products fall under the junk-food definition. It is common for supermarket baskets to include at least one junk-food item, 20 percent of baskets include baby products. Since consumers were randomly allocated to treatment, it is not surprising that their characteristics are almost identical to those of control shoppers.

We construct a balanced panel to analyze the direct effect of the intervention on the probability of purchasing healthier products. This dataset is organized by shopper, period, and, category. Thus, each of the 8,625 shoppers included in the database appears in 15 categories in 3 periods (pre-intervention, during-intervention, and post-intervention).[16] In cases where the shopper did not buy any products in that category during that period, we set the outcome to 0.

In addition to the SABH data set, we also put together a comprehensive data-set with product-level nutritional information. This data contains information on total calories, total fat, saturated fat, sodium, dietary fibres, sugar, protein, iron, carbohydrates, and cholesterol for each product in our data set. The primary source for this information is the same website that we used to collect the rest of the data. However, several gaps exist, either because the product is no longer sold or because nutritional information is simply not available for the product on that site. In these cases, we first attempt to manually collect the information by searching for the products online. When this is not feasible, we use broader categories (such as canned peaches instead of the particular brand) and scrape the information from wolfram alpha. Table 3 gives a summary of nutritional components of an average supermarket basket in the pre-intervention period for both treatment and control shoppers.

[16]We focus on this subset of shoppers because part of our analysis is conducted with the causal forest technique which uses pre-intervention characteristics (see Section 3.)

3 Empirical strategy

3.1 Across trips

We begin our analysis by examining the impact of the intervention on the purchase of the nudged healthier products that appeared in the popup for treatment shoppers. Thus, we run a difference-in-differences analysis where the outcome variable $buyhealthy_{ict}$ is equal to 1 if shopper i purchased a healthier alternative from category c during period t:

$$buyhealthy_{ict} = \alpha_0 + \alpha_1 \text{Treat} \times \text{During}_{it} + \alpha_2 \text{Treat} \times \text{Post}_{it} + \theta_i + \rho_t + \varepsilon_{ict} \qquad (1)$$

We separately explore both intention-to-treat outcomes using a balanced panel of all shoppers in the treatment and control groups during the pre,during, and post periods of the intervention and a sub-sample of consumers who made purchases in the relevant categories during the intervention.

3.2 Within trips

A unique aspect of our data is that it allows us to examine the impact of the SABH pop-up on the demand for other products. For this analysis we focus on a subset of consumers who added unhealthy relevant products to their basket and then consider how the behavior of treated shoppers (who view the SABH popup) changes relative to that of control shoppers.

We estimate this effect by examining differences between treatment and control shoppers who add a relevant product to their basket during the intervention period and measuring within-trip changes in behaviour across these groups after receiving the SABH prompt. Thus, we model outcomes $Y_{ibt} = (\text{price}_{ibt}, \text{produce}_{ibt}, \text{nutritionalcontent}_{ibt}, \text{totpurchases}_{ibt}, \text{totspending}_{ibt}, \text{timeshopping}_{ibt}, \text{storechange}_{ibt})$ as a function of whether shopper i (who was assigned to the treatment or control group) adds a product at the time, t in basket b is before or after selecting a relevant product.

$$Y_{ibt} = \beta_0 + \beta_1 \text{Treat} \times \text{Post}_{ibt} + \beta_2 \text{Post}_{ibt} + \theta_i + \rho_b + \varepsilon_{ibt} \qquad (2)$$

To analyze the impact of the prompt on price sensitivity (price$_{ibt}$), the nutritional content of products added to basket (nutritionalcontent$_{ibt}$), and the demand for fruits and vegetables (produce$_{ibt}$) we consider how shoppers change behavior regarding each product added to the basket after receiving the SABH nudge. Thus, we examine the probability of an added product being in the "fruit or vegetable category," or the relative change in the price of any product being added after the nudge. We measure nutritional content as a separate count of calories, saturated-fat, cholesterol, sugar, sodium, dietary-fibers, protein, and iron. To examine outcomes at the aggregate level (e.g., total spending, total products purchased), we collapse each basket into purchases pre and post-adding the first relevant product.

This analysis excludes control shoppers who never added unhealthy relevant products to their basket, and treated shoppers who either never added an unhealthy product or did not receive an SABH popup after adding an unhealthy product. This could raise concerns that shoppers in the treated group are no longer identical to those in the control. To further ensure that selection is not driving these results, we include data from the pre-intervention period to calculate these estimates using a triple-difference framework.

3.3 Heterogeneity

To identify which type of shoppers are most and least likely to respond to the intervention, we use a causal forest[17]. Causal forest works by partitioning a data set into sub-samples and then estimating a separate treatment effect per sub-sample to minimize a loss function. Following the potential outcome of Rubin's causal model (see Wager and Athey, 2018, p. 3), we let y_{its} be the outcome for individual i at time t when in group s. When $|s| = 2$, as is the case for this intervention, we can write this as $(y_{it}(0), y_{it}(1)$ for each i. Using this definition, we define the unit level causal effect as the difference in potential outcome, i.e. $\tau_{it} = y_{it}(1) - y_{it}(1)$.

As the causal forest model cannot yet handle multiple periods, we will divide the sample into three periods, pre-, during and post-intervention. The pre-intervention period will then be used to define the set of features as these should be unaffected by the treatment for them to be valid as inputs to the model. The effects can then be estimated on the two remaining periods separately. For this reason, we remove

[17] A variation of a random forest

the time subscript and write: $\tau_i = y_i(1) - y_i(1)$. As we have a randomized controlled trial, we get unconfoundedness (see Athey and Imbens, 2016, p. 4) for free and so can define the conditional average treatment effect as

$$\tau(z) = \mathbb{E}[y_i(1) - y_i(0) \mid Z_i = z] \tag{3}$$

which is estimated as $\hat{\tau}(z)$. Note that estimating Equation (3) without conditioning on $Z_i = z$ is the population average treatment effect which is the effect we find in a standard DID model. Now let Π correspond to a partition of the feature space $\mathbb{Z}$ with $\#(\Pi)$ being the number of partitions as

$$\Pi = \{l_1, ..., l_{\#(\Pi)}\}, \quad \text{with} \bigcup_{j=1}^{\#(\Pi)} l_j = \mathbb{Z} \tag{4}$$

The model then works by selecting a partition Π to minimize the mean squared error, defined as

$$\begin{aligned} MSE(\mathcal{S}^{te}, \mathcal{S}^{est}, \Pi) &= \frac{1}{\mid \mathcal{S}^{te} \mid} \sum_{i \in \mathcal{S}^{te}} ((y_i - \hat{\mu}(Z_i; \mathcal{S}^{est}, \Pi))^2 - y_i^2) \\ \hat{\mu}(Z_i; \mathcal{S}, \Pi) &= \frac{1}{\#\{i \in \mathcal{S} : Z_i \in l(z; \Pi)\}} \sum_{\{i \in \mathcal{S}: Z_i \in l(z;\Pi)\}} y_i \end{aligned} \tag{5}$$

where $\mathcal{S}$ is a sample, $l(z; \Pi)$ is a leaf $l \in \Pi$ and "te" and "est" indicates two parts of the total sample $\mathcal{S}$, $\mathcal{S}^{te} \cup \mathcal{S}^{est} = \mathcal{S}$.

When building the model, the algorithm recursively partitions the observations of the training sample into subgroups. For each of these groups, the algorithm calculates all candidate splits and chooses the one that minimizes the MSE as defined in Equation (5). Thus, for given samples $\mathcal{S}^{tr}, \mathcal{S}^{est}$ we define the conditional average treatment effect as:

$$\begin{aligned} \hat{\tau}(z) = &\frac{1}{\#\{i \in \mathcal{S}^{est} : Z_i \in l(z; \Pi), s(i) = 1\}} \sum_{\{i \in \mathcal{S}^{est}: Z_i \in l(z;\Pi), s(i)=1)\}} y_i \\ &- \frac{1}{\#\{i \in \mathcal{S}^{est} : Z_i \in l(z; \Pi), s(i) = 0\}} \sum_{\{i \in \mathcal{S}^{est}: Z_i \in l(z;\Pi), s(i)=0\}} y_i \end{aligned} \tag{6}$$

The splitting of the tree is called "honest" by Athey and Imbens (2016), which is an important property to get unbiased estimates. This means that the tree has to use a different set of points for constructing the splits and for estimating the treatment effect. In other words, the likelihood must be the same as if we had two data points with covariates x in our data (and responses drawn separately, i.i.d.), and we used one of them to determine the splits, and then the other one to output the estimation. We ensure this by using the same algorithm as Athey and Imbens (2016), which is available through the r-package `grf`.

The extension from a "causal tree" to the causal forest is straightforward. Let Q be the desired size of the forest. Then for $q = 1, 2, ..., Q$ randomly split $\mathcal{S}$ into an estimation and a training part. Then use the training part to grow a tree and use the estimation sample to get estimate $\hat{\tau}(z)$ using Equation (6). Lastly, average over the estimates to get the final single CATE estimate per shopper. Note that this method is using the entire sample as the training and estimation samples are re-sampled for each q.

To perform the partitioning, the model is given a set of shopper-level features, along with an indicator for the user being treated and the outcome variable. This set of features is presented in Table 4 and can generally be divided into categorising shoppers based on their healthiness, their thriftiness, and their general demographics. Note that each feature is defined in the pre-period to satisfy the stable unit treatment value assumption (Rubin, 1978). The outcome is defined as the probability of selecting the healthy products in each of the relevant categories. The data contains 15 distinct categories, with multiple observations within each category for each shopper, corresponding to each time the shopper added a product from the category in the intervention period. Using the fraction of healthy choices within categories collapses the data into one observation per category for each shopper, mitigating the over-weighting of shoppers who purchases many products. However, this still means some shoppers will be naturally over-weighted in the sample if they purchase in all the relevant categories. Conversely, some shoppers will be under-weighted if they purchase in fewer categories. This makes CATE estimates harder to interpret. To further correct this, we change the format of the data into a balanced version by making the following transformations. We balance the panel on the category level such that each shopper has one observation per category. Every feature is constant within the shopper and so no transformation of these is needed. The category fixed effect is also easily transformed. Recall that the target variable is the average of an indicator of choosing the healthy variety within each category such that the interpretation is the probability of selecting the healthy alternative at any time during the intervention period for each category for each shopper. In cases where the shopper has not bought any products in a category, the outcome is, therefore, naturally missing. In these cases, we set the outcome to 0. This gives 15 observations per user id in each period which we use for estimating the causal forest and the DID.

We use these estimates to get several interesting results. First, we attempt to label shoppers as having either a high or low effect of the intervention. Secondly, we use this labelling to investigate which characteristics are best at describing the difference between the top (high effect) and bottom (low effect) shoppers. Thirdly, we investigate the partial effect of the most important shopper characteristics in determining the effect of the intervention.

The first of these analyses are done using the expected conditional average treatment effect (CATE) on the shopper level. We then use the distribution of CATEs to identify the predicted top and bottom 25 percent. We label each shopper as belonging either to the top, the bottom, or the middle, and use this information in subsequent analysis. To get at the second analysis, We then perform the same exercise in a simulation study inspired by Chernozhukov et al. (2018). Here we again predict the top and bottom 25 percent, but this time repeating the study 250 times, each time running a regression of several shopper characteristics on top and then reporting the median coefficient with the median standard error. This way, we can identify significant differences between the top (shoppers expected to have the most positive response to the SABH nudge) and bottom shoppers (those expected to have the smallest, or even negative, response to the SABH nudge).

For the third analysis, we fix a seed and estimate a single causal forest of 6,000 trees. For each tree in the forest, we then return the first splitting value of each variable selected in that tree and then take the median value across trees. For each sub-sample defined by splitting the sample at the median value for each variable in the forest, we then predict the CATE for each shopper. With this, we comment on the effect of each variable on the predicted CATE to obtain something akin to a partial effect. All results are presented in Section 5.1.

4 How Do Shoppers Respond to the SABH Nudge?

In this section, we consider both the immediate impact of the intervention on shoppers choices regarding healthier products and how the nudge affects decisions made throughout the remainder of the shopping trip.

4.1 Across trips

Table 5 Panel A summarizes our results when running our difference-in-differences analysis on a balanced panel of all treatment and control shoppers. This balanced panel summarizes each shopper's purchases of the healthier product in each of the 17 categories during the pre, during, and post periods of the intervention. Thus, our intention-to-treat outcomes (which include treatment shoppers who never purchased in relevent categories and thus, never received the healthy nudge) reveal that the SABH nudge increased purchasing rates of the healthier products by 2.2 percentage points during the intervention and a noisily measured .6 percentage points after the intervention. This estimate is robust to including category and shopper fixed effects.

Table 5 Panel B focuses the analysis specifically on categories where shoppers made a purchase during the intervention. We would expect a larger impact on these treatment shoppers who actually received the shopping nudge.[18] Indeed, we find that treatment shoppers who received the nudge increase their probability of purchasing the healthy product by 4.5 percentage points during the intervention, and even after the intervention are 1.35 percentage points more likely to buy the healthier product.

4.2 Within trips

In this section we examine how receiving the SABH nudge promoting healthier products changes the way people shop. We find that after receiving an SABH nudge, shoppers are more likely to cut their trip short, tend to purchase less produce and vegetables, and become less price sensitive. Our analysis focuses on a within-basket comparison so that we examine changes in behavior within the same shopping trip before vs. after a treatment shopper adds a product that prompts a SABH nudge in contrast to a

[18]Not all shoppers in the treatment group receiving the SABH nudge on all relevant purchases during the intervention. This could be a result of product shortages at the supermarket (if the healthier product was not available) or ad blockers applied on individual computers. We continue to treat these shoppers as treated in order to avoid selection concerns.

shopper in the control group who adds a relevant product but who's trip continues uninterrupted. See empirical specification discussed in part 2 of Section 3.

The identifying assumption of this difference-in-differences analysis is that shoppers' post-nudge behaviour would have mimicked that of the control group if they hadn't received the nudge. This is because the they were randomly assigned to treatment or control group. We check the validity of this assumption by running this same DD analysis during the pre-intervention period. Lastly, we invoke a standard triple difference design (DDD) where we use the pre-intervention period to correct for any pre-existing differences between the treatment and control groups in terms of their supermarket purchasing trends.

The estimates of β_1 from Equation (2) are reported in column (2) of Table 6. We find that products added by shoppers in the treatment group after receiving the SABH prompt are 3 percentage points (s.e. 1) less likely to be from the produce category at an average purchase rate of 26 percent. Moreover, these products tend to be more nutritionally dense. Thus products added post SABH nudge have more saturated fat (an increase of 5 percent), cholesterol (an increase of 10 percent), protein (an increase of 6 percent), and iron (an increase of 2 percent). These products are also 2.5 percent (s.e. 1) more expensive than products added by shoppers who did not receive the SABH prompt (the average product price is 10 NIS).

A decrease in price sensitivity is also observed when examining shopping behavior at the aggregate level as shoppers are 5 percentage points (s.e. 0.9) less likely to move their current purchases in basket b to an alternative, cheaper supermarket after receiving the SABH prompt (on average 10 percent of shoppers in our sample change supermarkets mid-basket). In total, shoppers in the treatment group spend 25 percent (s.e. 7.5) less time shopping after receiving the SABH prompt and purchase 8 fewer products per basket (s.e. 1) than shoppers in the control group after receiving the SABH nudge. This decrease in basket size results in a decrease in spending of 56 NIS (s.e. 8) per basket.

Column (3) of Table 6 repeats this analysis when focusing on these same consumers in the pre-intervention period. As expected, the estimates are much smaller, and mostly not statistically significant, when comparing treatment shoppers to control shoppers when neither of them received a SABH prompt. However, we do find that shoppers in the treatment group purchased 2 products (s.e. 1) less than shoppers in the control group during the pre-intervention period, leading to a decrease in spending of 19 NIS (s.e. 7).

Column (4) of Table 6 reports the DDD results, in essence providing a further control for any pre-existing differences in shopping trends between shoppers assigned to the treatment group who received the SABH nudge and shoppers in the control group. We continue to find that shoppers who received the SABH prompt purchase products that are more expensive and are less likely to be from the produce category. The DDD analysis strengthens our results regarding the increase in demand for products with higher levels of cholesterol, saturate fat, protein, and iron. Treated shoppers spend 21 per cent (s.e. 7.2) less time shopping and are less likely to move their basket to an alternative supermarket. They purchase 6 (s.e. 1) fewer items than shoppers in the control group, decreasing their spending per basket by 37 NIS (s.e. 8).

The analysis so far suggests that while the SABH nudge was successful in causing shoppers to buy healthier alternatives, it also had a broader impact on shopping choices. Specifically, the reduction in both the price sensitivity of shoppers and the probability of purchasing fruits and vegetables could raise concerns about a general welfare reduction for shoppers. Moreover, an increase in the saturated-fat and cholesterol content of products added post nudge may undue the nutrition benefits created when substituting intervention products with healthier alternatives. Additionally, supermarkets may hesitate to implement a healthy shopping nudge that decreases total basket size and spending of shoppers. In the next section, we apply machine learning methods to categorize the types of shoppers who respond to the SABH nudge by switching to healthier products. This will allow us to both provide a summary of shopper-types who benefit from this type of healthy nudge, and examine whether the change in shopping behaviour outside of the direct response to the nudge is driven by the selection of a new healthier product (e.g. reaching a healthy threshold), or alternatively, an adverse response to the nudge either due to a depletion of cognitive resources or a guilt-response that causes shoppers to decrease their shopping time regardless of selecting a healthier product.

5 What types of shoppers respond to the SABH nudge?

5.1 Defining top and bottom types

Figure 6 displays the differences between the top and bottom CATE consumers. The bottom half of the plot shows the average value for the various features for the bottom type consumers. The top half of the plot shows the difference between the bottom and the top type consumers, with confidence intervals. The colour displays whether the difference is estimated to be positive or negative. The features are ordered by their absolute value so that the first one shows the largest absolute difference between the top and bottom.

We see that the average top shopper is more likely to have purchased alcohol at some point in the pre-intervention period. A similar effect in size can be seen for the indicator baskets_per_period_am which indicates the shopper has more baskets per month (on average) than the median shopper. Thus, this indicates that consumers who are more often exposed to the site are more likely to have a larger treatment effect. Conversely, we see a large negative difference in the average time between items in the basket. This indicates that consumers who spend a lot of time per item in their basket are likely to be less affected by the treatment, possibly indicating that they have a lower cognitive capacity to contemplate another decision. Summarizing the remaining variables, we see that top shoppers are more likely to have purchased some unhealthy items, more likely to have children, have a larger average fraction of their purchases in the junk food category, and more likely to have bought products associated with being ultra-orthodox, more likely to buy more expensive products and dedicate a larger share of their average basket to cigarettes or alcohol (the last two effects are small but statistically significant). Conversely, we see that top shoppers have a smaller average fraction of their basket dedicated to the fruit and vegetable category, are less likely to shop on Saturdays and less likely to accept a swap and save offer.

Table 7 shows predicted CATE for various sub-samples. The first column gives the variable name which is split on. The second gives the variable importance in determining the CATE and, the third is the median split value across trees in the forest. The fourth and the fifth column give the CATE for the sub-sample above and below the median split value, respectively. The last column gives out the categorisation of the variable.

As can be seen, the most important variable for splitting the sample is the average price paid. Note that the variable is defined based on the pre-intervention data and therefore indicates the general price level paid by the customer. We see that consumers who paid on average more than 9.897 shekels have a CATE of .011, while those who paid less have a CATE of only .007 (an increase of more than 50%). We also find that consumers with higher guilt costs, as measured by their pre-intervention tendency to remove already added junk food from their basket prior to purchase (see Section 2), are less likely to respond to the intervention by buying the healthier alternative. Generally, we find that variables related to the healthiness of the consumer tend to be good predictors of the intervention effect (four out of the 6 strongest prediction variables come from the healthiness category). Our results suggest that the intervention has a larger effect on healthy consumers than on less healthy ones. While category of purchase and shopping day are generally weaker predictors of intervention outcomes, we see the strongest effect in the milk category and among shoppers who tend to shop on Saturdays (after peak demand for pre-weekend shopping).

5.2 Across trips

In **??** we re-estimate equation Equation (1) for top and bottom shoppers separately as defined by their CATE estimates. When running the analysis on the top shoppers we observe a large 6 percentage point increase in the probability of purchasing a healthier alternative during the intervention period. These treated shoppers continue to be 1.5 percentage points more likely to purchase a healthier product after the intervention. With an average purchase rate of roughly 4 percent for this sample, this demonstrates the extent to which these consumers changed their shopping behavior, presumably for the better. However, not all shoppers, benefited from the nudge. Indeed, the treated shoppers least likely to respond to the SABH nudge actually decreased their consumption of the healthier alternatives by 3.7 percentage points relative to the control group (perhaps by simply avoiding the product category). This response shrinks significantly to half a percentage point after the intervention period.

5.3 Within trips

Figure 4 and Figure 5 plot the DDD results when running the analysis separately on top and bottom shoppers to examine how receiving the SABH nudge impacts the price of products added to the basket, their willingness to switch supermarkets, their demand for produce, the nutritional content of products added to basket, and the length, size, and expense of their shopping trip. This allows us to measure the difference in shopping behavior between those in the treatment and control group after adding a product that prompts a SABH nudge. The pre-intervention data allows us to control for any pre-existing differences in shopping trends between these groups. Despite our finding that shoppers in the bottom group were much less likely to respond directly to the SABH nudge, Figure 4 and Figure 5 suggest that their overall shopping behaviour was still affected. While noisily measured, our results suggest that even bottom shoppers tended to buy fewer products, were less likely to swap to a different supermarket, decreased their consumption of produce, and increased their consumption of saturated fat, and iron. The similar response across treatment shoppers,regardless of whether they actually added the healthier product to their basket, suggests that the healthy nudge has a direct effect on shopping behavior. The nudge may deplete cognitive resources, or induce guilt costs for all shoppers, even for those who choose not to accept the nudge.

6 Mechanisms

One of our main findings is somewhat puzzling: the intervention caused *both* responders and non-responders to make fewer decisions (they shorten their trip and reduce the number of products bought), forego opportunities to save (they reduce the likelihood of switching to a cheaper supermarket and buy more expensive products), purchase food with more cholesterol and saturated-fat, and reduce their purchases of produce. In this section, we propose two possible explanations for this finding and report the results of empirical tests of these explanations.

6.1 Possible explanations

One plausible explanation is that our intervention, which imposed additional decision-making on shoppers, depleted these shoppers' cognitive resources to make further decisions. Consequently, they either exited their shopping trip earlier (to avoid facing more decisions), or made poorer decisions in terms of savings or nutrition. This explanation draws on previous studies that showed people's ability to exert self-control, or to make careful deliberate choices is hampered when following a prior decision-making task that requires some cognitive effort (see Muraven and Baumeister (2000) with regards to self-control, and Pocheptsova et al. (2009) with regards to decision-making). Note that this explanation applies to both responders and non-responders since the SABH pop-up forced both groups to make a decision.

An alternative explanation is that shoppers would like to avoid anticipated "guilt" from choosing an unhealthy product when they know a healthier substitute is available. Consequently, these shoppers end their trip earlier, and as a result, make fewer purchasing decisions, buy fewer items (including produce), and have fewer opportunities to switch supermarkets. This explanation is motivated by studies that document feelings of anticipated guilt from indulging in unhealthy food (see, e.g., Steenhuis, 2009; Gonzalez and Vitousek, 2004; Hur and Jang, 2015).

To see how avoidance of anticipated guilt costs can lead to a decision to exit the shopping trip by both responders and non-responders, consider the following example, which adopts a simple version of a model of anticipated guilt costs by Kopylov (2012).

Suppose there are only two food categories, X and Y, and consider a shopper, who before the intervention, is aware of only a single item from each category: x in X and y in Y. Suppose this shopper purchases the one item he is aware of in each category. For simplicity, assume these items are the most unhealthy in their respective category. Because he is unaware of the healthier option, this shopper does not incur any self-control costs associated with resisting temptations and also does not suffer from guilt associated with succumbing to temptation.

Now suppose the intervention makes the shopper aware of healthier options. To describe how this affects the shopper's payoff we now introduce self-control and guilt costs. These are modelled in the spirit of the utility representation axiomatized in Kopylov (2012).

First, a shopper has two consumption utilities, a *commitment* utility $u(\cdot)$ and a *temptation* utility $v(\cdot)$. The function $u(\cdot)$ represents his long-term considerations, which take into account the healthiness of the food item. This is the utility he would get if he could avoid the temptation of unhealthy food by having a singleton menu with only a single option (e.g., if he could commit to going to a restaurant that serves only vegetable salads). In contrast, the function $v(\cdot)$ captures the shopper's visceral urges by representing the strength of the temptation of unhealthy items. Thus, an individual who cares about taste and health needs to balance these two types of utilities.

Thus, a shopper who buys item a when he is aware that there is a set of alternatives A gets a payoff equal to (his total utility from a basket is equal to the sum of his payoffs from all the categories):

$$u(a) - c(a, A) - g(a, A) - p(a) \tag{7}$$

where $p(a)$ is the price of a and $c(a, A)$ and $g(a, A)$ denote the self-control cost and guilt cost, respectively, associated with choosing item a when there is a known set of alternatives A. These costs take the following form:

$$c(a, A) = \kappa[\max_{b \in A} v(b) - v(a)] \tag{8}$$

and

$$g(a, A) = \gamma[\max_{b \in A} u(b) - u(a)] \tag{9}$$

These cost specifications capture the idea that the cost associated with exerting self-control when choosing from a menu of options is proportional to the difference in temptation utility between the most tempting item on the menu and the item that is chosen, while the guilt cost is proportional to the difference between the "healthiness" of the healthiest item on the menu and the healthiness of the item that is actually chosen. The scalars κ and γ capture the relative intensity of temptation and guilt. In particular, if $\kappa = 0$ the individual has no self-control problems, and if $\gamma = 0$, the the individual does not experience guilty feelings.

For the sake of illustration, suppose that x and y are the most unhealthy options in their respective categories, and the intervention only makes the shopper aware that y has a healthier alternative y^*. The shopper has the following options: (i) continue to buy y, (ii) switch to the healthier option y^*, or (iii) end the trip and only buy x.

If the shopper does not respond to the intervention (i.e., he continues to purchase the unhealthy y), his payoff from category Y is equal to

$$u(y) - g(y, Y) - p(y) \tag{10}$$

while if he responds and switches to y^*, his total payoff from category Y is given by

$$u(y^*) - c(y^*, Y) - p(y^*) \tag{11}$$

If both of these expressions are negative, but $u(x) - p(x)$ (the payoff from X) is positive, a shopper who anticipates a SABH pop-up in category Y would decide to finish his shopping after buying x.

Suppose next that the intervention also makes a shopper aware that x has a healthier alternative x^*, and that whether a shopper switches (responds to the intervention) or not he still gets a positive surplus. In this case, both shoppers who respond to the intervention (those who switch to x^*) and those who don't (those who still buy x), end their shopping early. Thus, the anticipated guilt cost can explain why non-responders end their trip early, while the anticipated self-control cost can explain why responders end their trip early.

The intervention can lead the shopper to end their trip early by impacting the purchase of products both within and outside the intervention. This may occur because any future product added to the basket has some probability of appearing with a popup. Alternatively, the intervention may raise awareness of healthy alternatives more generally, which can increase shoppers' feelings of guilt or self-control costs regardless of whether a popup appears. In either case, the shopper would respond to the intervention by ending their shopping trip prematurely.

6.2 Tests

In an attempt to examine which of the above explanations has more support in our data, we conducted the following tests. While neither of the tests are perfect, they do seem to suggest that an explanation based on guilt costs may be the driver of the findings described at the beginning of this section.

To investigate whether cognitive overload is the likely explanation, we conducted the following two tests. One test checked whether shoppers who clicked on a Swap & Save option before receiving the first SABH pop-up were more likely to exhibit the behavioural patterns described above. The idea behind this test is the following: if shoppers' cognitive resources are depleted by the intervention because it forced them to make more decisions (choose healthily or not), then those shoppers who also had to decide whether to save money or not via the Swap & Save options depleted even more cognitive resources.

The second test was based on a similar idea. As the shopping trip progresses, a shopper depletes more cognitive resources, so a shopper who faced the SABH pop-up late in his shopping trip should be more likely to be cognitively exhausted and hence, more likely to exhibit the behavioural patterns described above.

In columns (2) & (3) of Table 9 we separately examine the impact of the SABH nudge on shoppers who already considered a Swap & Save alternative on the site (i.e., may have depleted some of their cognitive resources) versus those who have not. If anything, the SABH popup seemed to have more of an effect on shoppers who had not yet viewed a Swap & Save alternative. Similarly, in columns (5) & (6), we generally don't find significant differences in behavior between shoppers who receive the popup at the beginning of their trip versus shoppers who receive the popup at the end of their trip. However, while we observe a similar decrease in total quantity purchased, we do find that the decrease in shopping time (see coefficient on Log(Time Shopping)) is being driven by shoppers who receive the popup later on in their trip. Thus, Table 9 makes it difficult to conclude that the depletion of cognitive resources as measured by engaging in Swap & Save decisions, or facing the SABH pop-up later in the shopping trip, are the mechanisms driving our results.

If guilt and self-control costs are the main driver of our findings, then the behavioural patterns described above would be mostly driven by shoppers with higher costs. To test this hypothesis, we define an indicator for whether a shopper faces guilt or self-control costs based on whether the shopper ever adds and subsequently removes junk items from their basket during the pre-intervention period. In columns (2) & (3) of Table 10 we separately examine the impact of the SABH nudge on shopper types who face high guilt or self-control costs versus those that do not. Indeed, we find that shoppers with higher guilt costs drive the behavioral change in almost all of the outcome variables (price, produce consumption, saturated fat, iron, basket price, and number of products purchased). We therefore conclude that this is a more plausible explanation for the mechanism driving our results.

7 Conclusion

In the majority of OECD countries, over 50 percent of the population is considered overweight, with obesity rates more than doubling in the US in the last 30 years to close to 40 percent of the population (Indicators, 2015). While one would expect that information on nutritional values of products could result in better eating habits, most of the literature finds little to no effect of nutrition labeling on consumer choice (see Balasubramanian and Cole, 2002; Downs, Loewenstein and Wisdom, 2009; Van Herpen and Van Trijp, 2011)). A more promising direction seems to be simplifying nutrition information with a nutrition grade or star system (see Cawley et al., 2015; Nikolova and Inman, 2015; Hobin et al., 2017)).

The SABH intervention allows us to take nutritional labeling one step further into the future. It provides an opportunity to make healthier choices as simple as clicking a button. We find that the SABH nudge resulted in a significant 2-5 percentage point increase in purchasing rates of healthier products for treated shoppers. This effect continued even after treated shoppers no longer received the SABH prompts so that treated shoppers remained 1 percentage point more likely to purchase healthier products even 1 month after the intervention. It also presented us with data from online shopping that can provide a much more detailed understanding of the shopping experience, and how shoppers respond to outside interference with their shopping choices. Despite the success of the SABH nudge in encouraging shoppers to make healthier choices within the product categories of the intervention, we observe changes in shopping behavior that raise concerns regarding the implications of this type of "forced involvement" on consumer choice. Generally, treated shoppers purchased 3 percentage points less fruits and vegetables after receiving the SABH prompt. These shoppers also purchased products with 5 percent more saturated fat and 10 percent more cholesterol. They all exhibited a decrease in price sensitivity, purchasing products that were on average 2.5 percent more expensive.

We apply machine learning analysis to better understand which types of shoppers respond most to the SABH nudge and find that the shoppers most likely to respond are those who frequently shopped on the site, who spend less time selecting products, who buy more junk food and/or baby products, and generally buy more expensive products. Interestingly, we find that the changes in shopping behavior discussed above are observed very similarly for both shoppers that are most likely to respond to the intervention and shoppers that are least likely to respond to the intervention. While an increase in cognitive overload could explain why the SABH nudge caused shoppers who did not respond to the nudge to still exhibit a change in behavior, our results suggest that it is “guilt” or “self-control” costs that drive the externalities associated with the SABH nudge. Specifically, we find that shoppers who exhibit higher guilt-costs or self-control in the pre-intervention period are the shoppers driving the behavioral changes observed for treated shoppers after receiving the SABH nudge (decrease in number of products purchased, and specifically fruits and vegetables, alongside an increase in average product price, and saturate fat).

These results point to the importance of focusing not only on the health intervention but on its timing within the shopping trip. Thus, if the SABH popup only appeared at the end of the shopping trip this could help to mitigate some of the consequences we document occurring within the trip. Also, targeting the order in which shoppers add products to basket (as opposed to which products are added to basket) could have important policy implications - as targeted nudges may result in an early end to trip, and then the order is crucial for determining which products are purchased.

References

ABC. 2012. “Foods that have hidden fat, sugar, salt.” *https://abc7news.com/archive/8589311/*. 2

Altmann, Steffen, Andreas Grunewald, and Jonas Radbruch. forthcoming. “Passive choices and cognitive spillovers.” 7

Araya, Sebastian, Andrés Elberg, Carlos Noton, and Daniel Schwartz. 2022. “Identifying food labeling effects on consumer behavior.” *Available at SSRN 3195500.* 7

Arteaga, Felipe, Adam J Kapor, Christopher A Neilson, and Seth D Zimmerman. 2022. “Smart Matching Platforms and Heterogeneous Beliefs in Centralized School Choice.” *The Quarterly Journal of Economics.* 7

Athey, Susan, and Guido Imbens. 2016. “Recursive partitioning for heterogeneous causal effects.” *Proceedings of the National Academy of Sciences*, 113(27): 7353–7360. 17, 18

Balasubramanian, Siva K, and Catherine Cole. 2002. "Consumers' search and use of nutrition information: The challenge and promise of the nutrition labeling and education act." *Journal of marketing*, 66(3): 112–127. 6, 30

Bhargava, Saurabh, and Dayanand Manoli. 2015. "Psychological frictions and the incomplete take-up of social benefits: Evidence from an IRS field experiment." *American Economic Review*, 105(11): 3489–3529. 7

Biban, Bianca Georgiana, and Corina Lichiardopol. 2017. "Iodine deficiency, still a global problem?" *Current health sciences journal*, 43(2): 103. 2

Brownback, Andy, Alex Imas, and Michael Kuhn. 2021. "Behavioral Food Subsidies." *Available at SSRN 3422272*. 7

Bunten, Amanda, Lucy Porter, Jet G Sanders, Anna Sallis, Sarah Payne Riches, Paul Van Schaik, Marta González-Iraizoz, Tim Chadborn, and Suzanna Forwood. 2021. "A randomised experiment of health, cost and social norm message frames to encourage acceptance of swaps in a simulation online supermarket." *PloS one*, 16(2): e0246455. 8

Buyalskaya, Anastasia, and Matthew Shum. 2020. "Self-Regulation Failures in the Field: Investigating Spillovers from Morning Exercise to Food Choice." *Mimeo*. 6

Cadario, Romain, and Pierre Chandon. 2020. "Which healthy eating nudges work best? A meta-analysis of field experiments." *Marketing Science*, 39(3): 465–486. 6

Cawley, John, Adam Biener, Chad Meyerhoefer, Yuchen Ding, Tracy Zvenyach, B Gabriel Smolarz, and Abhilasha Ramasamy. 2021. "Direct medical costs of obesity in the United States and the most populous states." *Journal of Managed Care & Specialty Pharmacy*, 27(3): 354–366. 1

Cawley, John, Alex Susskind, and Barton Willage. 2020. "The impact of information disclosure on consumer behavior: evidence from a randomized field experiment of calorie labels on restaurant menus." *Journal of Policy Analysis and Management*, 39(4): 1020–1042. 6

Cawley, John, Matthew J Sweeney, Jeffery Sobal, David R Just, Harry M Kaiser, William D Schulze, Elaine Wethington, and Brian Wansink. 2015. "The impact of a supermarket nutrition rating system on purchases of nutritious and less nutritious foods." *Public health nutrition*, 18(1): 8–14. 30

Chernozhukov, Victor, Mert Demirer, Esther Duflo, and Ivan Fernandez-Val. 2018. "Generic machine learning inference on heterogeneous treatment effects in randomized experiments, with an application to immunization in India." National Bureau of Economic Research. 4, 19

Chiou, Wen-Bin, Chao-Chin Yang, and Chin-Sheng Wan. 2011. "Ironic effects of dietary supplementation: illusory invulnerability created by taking dietary supplements licenses health-risk behaviors." *Psychological science*, 22(8): 1081–1086. 6

Choinière, CJ, and A Lando. 2008. "Health and diet survey." *Food and Drug Administration: Silver Spring, MD, USA*. 1

CHP. 2017. "Centre for Health Protection, Department of Health, The Government of the Hong Kong Special Administrative Region." *https://www.chp.gov.hk/en/static/90037.html*. 2

Coffino, Jaime A, Tomoko Udo, and Julia M Hormes. 2020. "Nudging while online grocery shopping: A randomized feasibility trial to enhance nutrition in individuals with food insecurity." *Appetite*, 152: 104714. 8

Crespo-Bellido, Mayra S, Stephanie K Grutzmacher, Yumie Takata, and Ellen Smit. 2021. "The Association Between Food-Away-From-Home Frequency and a Higher BMI Varies by Food Security Status in US Adults." *The Journal of Nutrition*, 151(2): 387–394. 8

De Witt Huberts, Jessie C, Catharine Evers, and Denise TD De Ridder. 2012. "License to sin: Self-licensing as a mechanism underlying hedonic consumption." *European Journal of Social Psychology*, 42(4): 490–496. 6

Dolan, Paul, and Matteo M Galizzi. 2014. "Because I'm worth it: a lab-field experiment on the spillover effects of incentives in health." 6

Dolan, Paul, and Matteo M Galizzi. 2015. "Like ripples on a pond: behavioral spillovers and their implications for research and policy." *Journal of Economic Psychology*, 47: 1–16. 6

Downs, Julie S, George Loewenstein, and Jessica Wisdom. 2009. "Strategies for promoting healthier food choices." *American Economic Review*, 99(2): 159–64. 30

Downs, Julie S, Jessica Wisdom, Brian Wansink, and George Loewenstein. 2013. "Supplementing menu labeling with calorie recommendations to test for facilitation effects." *American Journal of Public Health*, 103(9): 1604–1609. 6

Drake, Marshall, Fernando Payró, Neil Thakral, and Linh T Tô. 2022. "Bayesian Adaptive Choice Experiments." *Mimeo*. 7

Drake, Marshall, Neil Thakral, and Linh T Tô. 2022. "Wage Differentials and the Price of Workplace Flexibility." *Mimeo.* 7

Drichoutis, Andreas C, Rodolfo M Nayga, and Panagiotis Lazaridis. 2012. "Food away from home expenditures and obesity among older Europeans: are there gender differences?" *Empirical Economics*, 42(3): 1051–1078. 8

Ericson, Keith Marzilli, and David Laibson. 2019. "Intertemporal choice." In *Handbook of Behavioral Economics: Applications and Foundations 1.* Vol. 2, 1–67. Elsevier. 8

Finkelstein, Eric A, Felicia Jia Ler Ang, Brett Doble, Wei Han Melvin Wong, and Rob M van Dam. 2019. "A randomized controlled trial evaluating the relative effectiveness of the Multiple Traffic Light and Nutri-Score front of package nutrition labels." *Nutrients*, 11(9): 2236. 8

Finkelstein, Eric Andrew, Felicia Jia Ler Ang, and Brett Doble. 2020. "Randomized trial evaluating the effectiveness of within versus across-category front-of-package lower-calorie labelling on food demand." *BMC public health*, 20(1): 1–10. 8

Forwood, Suzanna E, Amy L Ahern, Theresa M Marteau, and Susan A Jebb. 2015. "Offering within-category food swaps to reduce energy density of food purchases: a study using an experimental online supermarket." *International Journal of Behavioral Nutrition and Physical Activity*, 12(1): 1–10. 8

Galizzi, Matteo M, and Lorraine Whitmarsh. 2019. "How to measure behavioral spillovers: a methodological review and checklist." *Frontiers in psychology*, 10: 342. 6

Gonzalez, Vivian MM, and Kelly M Vitousek. 2004. "Feared food in dieting and non-dieting young women: a preliminary validation of the Food Phobia Survey." *Appetite*, 43(2): 155–173. 26

Griffith, Rachel, Stephanie von Hinke, and Sarah Smith. 2018. "Getting a healthy start: The effectiveness of targeted benefits for improving dietary choices." *Journal of health economics*, 58: 176–187. 6, 7

Guthrie, Joanne, Lisa Mancino, and Chung-Tung Jordan Lin. 2015. "Nudging consumers toward better food choices: Policy approaches to changing food consumption behaviors." *Psychology & Marketing*, 32(5): 501–511. 1, 6

Herber, Stefanie P. 2018. "The role of information in the application for highly selective scholarships: Evidence from a randomized field experiment." *Economics of Education Review*, 62: 287–301. 7

Hobin, Erin, Bryan Bollinger, Jocelyn Sacco, ELI Liebman, Lana Vanderlee, FEI Zuo, Laura Rosella, Mary L'abbe, Heather Manson, and David Hammond. 2017. "Consumers' response to an on-shelf nutrition labelling system in supermarkets: evidence to inform policy and practice." *The Milbank Quarterly*, 95(3): 494–534. 30

Horgen, Katherine Battle, and Kelly D Brownell. 2002. "Comparison of price change and health message interventions in promoting healthy food choices." *Health Psychology*, 21(5): 505. 6

Hoxby, Caroline M, and Sarah Turner. 2013. "Informing students about their college options: A proposal for broadening the expanding college opportunities project." 7

Huang, Amy, Federica Barzi, Rachel Huxley, Gareth Denyer, Beth Rohrlach, Kathy Jayne, and Bruce Neal. 2006. "The effects on saturated fat purchases of providing internet shoppers with purchase-specific dietary advice: a randomised trial." *PLoS clinical trials*, 1(5): e22. 8

Hulshof, Casper D, and Ton De Jong. 2006. "Using just-in-time information to support scientific discovery learning in a computer-based simulation." *Interactive Learning Environments*, 14(1): 79–94. 7

Hur, JungYun, and SooCheong Shawn Jang. 2015. "Anticipated guilt and pleasure in a healthy food consumption context." *International Journal of Hospitality Management*, 48: 113–123. 26

Imas, Alex, Michael Kuhn, and Vera Mironova. forthcoming. "Waiting to Choose: The Role of Deliberation in Intertemporal Choice." 7

Indicators, OECD. 2015. "Health at a Glance 2011." *OECD Indicators, OECD Publishing, Paris DOI: https://doi org/101787/health_glance-2015-en Accessed February*, 15: 2016. 30

Jansen, Laura, Ellen van Kleef, and Ellen J Van Loo. 2021. "The use of food swaps to encourage healthier online food choices: a randomized controlled trial." *International journal of behavioral nutrition and physical activity*, 18(1): 1–16. 8

Kling, Jeffrey R, Sendhil Mullainathan, Eldar Shafir, Lee C Vermeulen, and Marian V Wrobel. 2012. "Comparison friction: Experimental evidence from Medicare drug plans." *The quarterly journal of economics*, 127(1): 199–235. 7

Kopylov, Igor. 2012. "Perfectionism and choice." *Econometrica*, 80(5): 1819–1843. 26

Koutoukidis, Dimitrios A, Susan A Jebb, José M Ordóñez-Mena, Michaela Noreik, Melina Tsiountsioura, Sarah Kennedy, Sarah Payne-Riches, Paul Aveyard, and Carmen Piernas. 2019. "Prominent positioning and food swaps are effective interventions to reduce the saturated fat content of the shopping basket in an experimental online supermarket: a randomized controlled trial." *International Journal of Behavioral Nutrition and Physical Activity*, 16(1): 1–14. 8

Lazarus, John H. 2020. "Iodine deficiency in Israeli pregnant women–a time for action." *Israel Journal of Health Policy Research*, 9(1): 1–3. 3

Loewenstein, George, Joseph Price, and Kevin Volpp. 2016. "Habit formation in children: Evidence from incentives for healthy eating." *Journal of health economics*, 45: 47–54. 6

Lowe, Ben, Diogo M de Souza-Monteiro, and Iain Fraser. 2013. "Nutritional labelling information: Utilisation of new technologies." *Journal of Marketing Management*, 29(11-12): 1337–1366. 2, 7

Murray, Christopher JL, Aleksandr Y Aravkin, Peng Zheng, Cristiana Abbafati, Kaja M Abbas, Mohsen Abbasi-Kangevari, Foad Abd-Allah, Ahmed Abdelalim, Mohammad Abdollahi, Ibrahim Abdollahpour, et al. 2020. "Global burden of 87 risk factors in 204 countries and territories, 1990–2019: a systematic analysis for the Global Burden of Disease Study 2019." *The Lancet*, 396(10258): 1223–1249. 1

Nafziger, Julia. 2020. "Spillover effects of nudges." *Economics Letters*, 190: 109086. 7

Nahum-Shani, Inbal, Shawna N Smith, Bonnie J Spring, Linda M Collins, Katie Witkiewitz, Ambuj Tewari, and Susan A Murphy. 2018. "Just-in-time adaptive interventions (JITAIs) in mobile health: key components and design principles for ongoing health behavior support." *Annals of Behavioral Medicine*, 52(6): 446–462. 7

NHMRC. 2015. "Fat, Salt, Sugars and Alcohol." *https://www.eatforhealth.gov.au/food-essentials/fat-salt-sugars-and-alcohol*. 2

NIH. 2016. "New dietary guidelines urge Americans to eat less added sugars, saturated fat, and sodium." *https://www.nhlbi.nih.gov/news/2016/new-dietary-guidelines-urge-americans-eat-less-added-sugars-saturated-fat-and-so* 1

Nikolova, Hristina Dzhogleva, and J Jeffrey Inman. 2015. "Healthy choice: the effect of simplified point-of-sale nutritional information on consumer food choice behavior." *Journal of Marketing Research*, 52(6): 817–835. 30

Olsho, Lauren EW, Jacob A Klerman, Parke E Wilde, and Susan Bartlett. 2016. "Financial incentives increase fruit and vegetable intake among Supplemental Nutrition Assistance Program participants: a randomized controlled trial of the USDA Healthy Incentives Pilot." *The American journal of clinical nutrition*, 104(2): 423–435. 6

Ovadia, Yaniv S, Jonathan E Arbelle, Dov Gefel, Hadassah Brik, Tamar Wolf, Varda Nadler, Sandra Hunziker, Michael B Zimmermann, and Aron M Troen. 2017. "First Israeli national iodine survey demonstrates iodine deficiency among school-aged children and pregnant women." *Thyroid*, 27(8): 1083–1091. 3

Polivy, Janet, C Peter Herman, and Traci McFarlane. 1994. "Effects of anxiety on eating: Does palatability moderate distress-induced overeating in dieters?" *Journal of Abnormal Psychology*, 103(3): 505. 6

Ponzo, Valentina, Marianna Pellegrini, Paola Costelli, Laura Vázquez-Araújo, Lucía Gayoso, Chiara D'Eusebio, Ezio Ghigo, and Simona Bo. 2021. "Strategies for reducing salt and sugar intakes in individuals at increased cardiometabolic risk." *Nutrients*, 13(1): 279. 2

Read, Daniel, and Barbara Van Leeuwen. 1998. "Predicting hunger: The effects of appetite and delay on choice." *Organizational behavior and human decision processes*, 76(2): 189–205. 8

Riches, Sarah Payne, Paul Aveyard, Carmen Piernas, Mike Rayner, and Susan A Jebb. 2019. "Optimising swaps to reduce the salt content of food purchases in a virtual online supermarket: A randomised controlled trial." *Appetite*, 133: 378–386. 8

Rizkalla, Salwa W, F Bellisle, and G Slama. 2002. "Health benefits of low glycaemic index foods, such as pulses, in diabetic patients and healthy individuals." *British Journal of Nutrition*, 88(S3): 255–262. 11

Rubin, Donald B. 1978. "Bayesian inference for causal effects: The role of randomization." *The Annals of statistics*, 34–58. 18

Sacco, Jocelyn, Heather G Lillico, Emily Chen, and Erin Hobin. 2017. "The influence of menu labelling on food choices among children and adolescents: a systematic review of the literature." *Perspectives in public health*, 137(3): 173–181. 12

Sacks, Gary, Kim Tikellis, Lynne Millar, and Boyd Swinburn. 2011. "Impact of 'traffic-light' nutrition information on online food purchases in Australia." *Australian and New Zealand journal of public health*, 35(2): 122–126. 8

Shin, Soye, Rob M van Dam, and Eric A Finkelstein. 2020. “The Effect of Dynamic Food Labels with Real-Time Feedback on Diet Quality: Results from a Randomized Controlled Trial.” *Nutrients*, 12(7): 2158. 8

Steenhuis, Ingrid. 2009. “Guilty or not? Feelings of guilt about food among college women.” *Appetite*, 52(2): 531–534. 26

Sunstein, Cass R. 2013*a*. “Nudges.gov: Behavioral economics and regulation.” *Oxford Handbook of Behavioral Economics and the Law (Eyal Zamir and Doron Teichman eds.).* 1

Sunstein, Cass R. 2013*b*. *Simpler: The future of government.* Simon and Schuster. 7

Thakral, Neil, and Linh T Tô. 2020. “Anticipation and Temptation.” *Mimeo.* 7

Thakral, Neil, and Linh T Tô. 2022. “Anticipation and Consumption.” *Mimeo.* 7

Trachtman, Hannah. 2022. “Does promoting one behavior distract from others? Evidence from a field experiment.” *Mimeo.* 7

Urbszat, Dax, C Peter Herman, and Janet Polivy. 2002. “Eat, drink, and be merry, for tomorrow we diet: effects of anticipated deprivation on food intake in restrained and unrestrained eaters.” *Journal of abnormal psychology*, 111(2): 396. 6

Van Herpen, Erica, and Hans CM Van Trijp. 2011. “Front-of-pack nutrition labels. Their effect on attention and choices when consumers have varying goals and time constraints.” *Appetite*, 57(1): 148–160. 30

Wager, Stefan, and Susan Athey. 2018. “Estimation and inference of heterogeneous treatment effects using random forests.” *Journal of the American Statistical Association*, 113(523): 1228–1242. 16

WHO. 2000. “Obesity: preventing and managing the global epidemic.” 1

Wilcox, Keith, Beth Vallen, Lauren Block, and Gavan J Fitzsimons. 2009. “Vicarious goal fulfillment: When the mere presence of a healthy option leads to an ironically indulgent decision.” *Journal of Consumer Research*, 36(3): 380–393. 6

Wilson, Amy L, Elizabeth Buckley, Jonathan D Buckley, and Svetlana Bogomolova. 2016. “Nudging healthier food and beverage choices through salience and priming. Evidence from a systematic review.” *Food Quality and Preference*, 51: 47–64. 6, 8

Wisdom, Jessica, Julie S Downs, and George Loewenstein. 2010. “Promoting healthy choices: Information versus convenience.” *American Economic Journal: Applied Economics*, 2(2): 164–78. 6

Figure 1: Screenshot of SABH

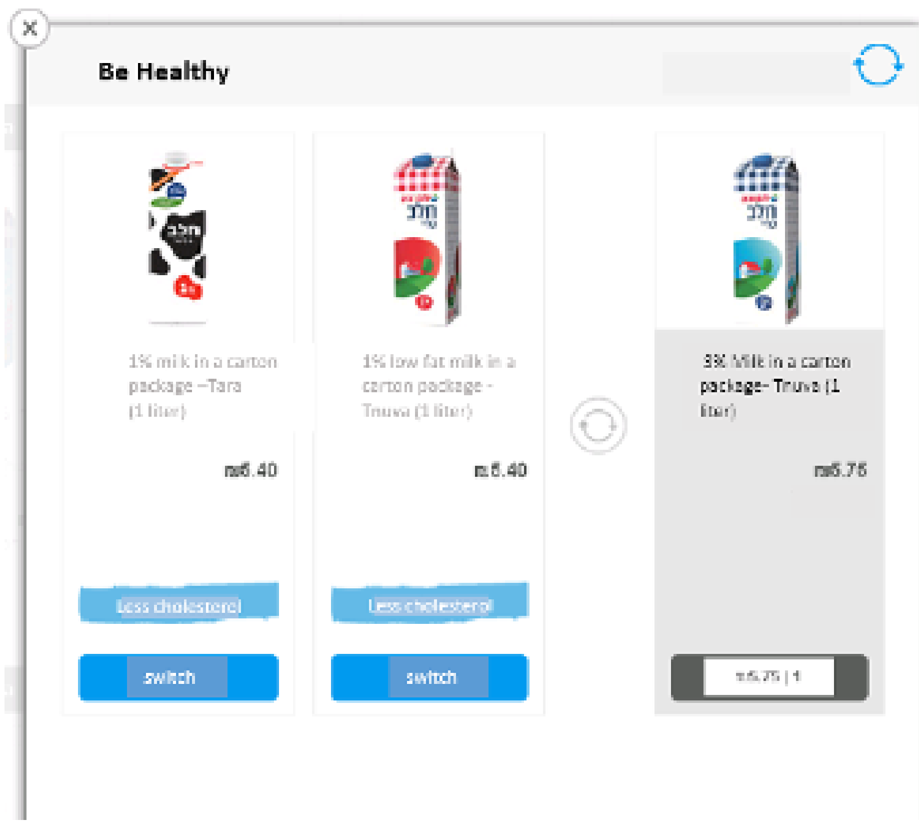

Note: This is a screenshot of the milk SABH nudge. This popup appeared on the screen for any treatment shopper who added a 3 percent milk carton to their basket during the intervention period.

Figure 2: Effect of treatment by price difference

(a) Treatment group

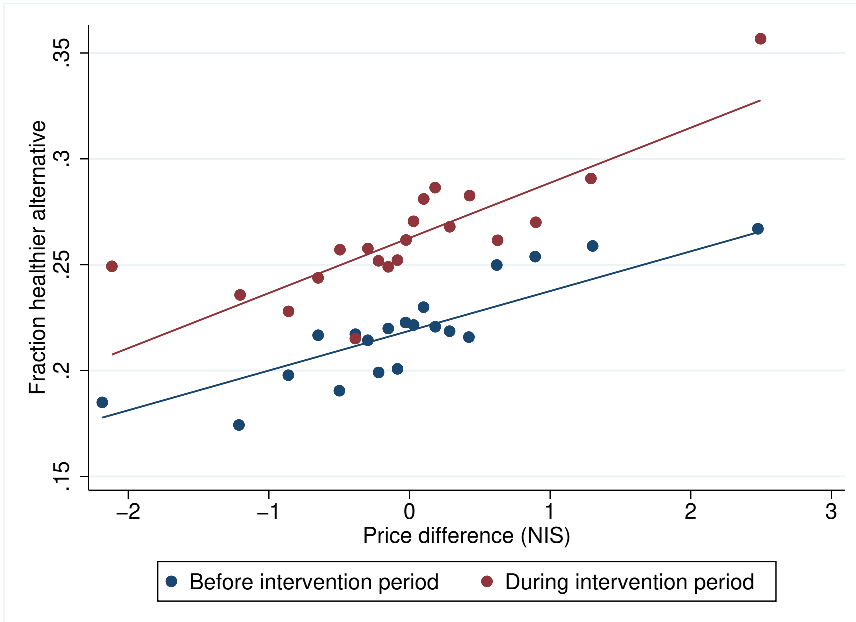

(b) Control group

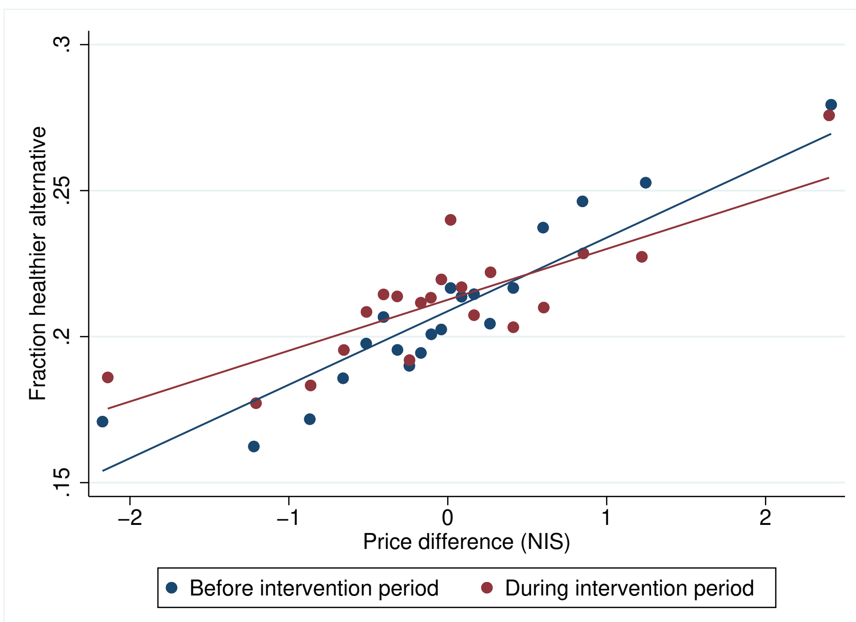

Note: Each panel shows a binned scatterplot of the relationship between the probability of purchasing the healthier alternative and the price difference, controlling for shopper and product category fixed effects. Price difference is defined as the difference between the price of the less healthy product and the healthier alternative. The sample in the top panel consists of shoppers randomly assigned to the treatment group, and the sample in the bottom panel consists of shoppers randomly assigned to the control group.

Figure 3

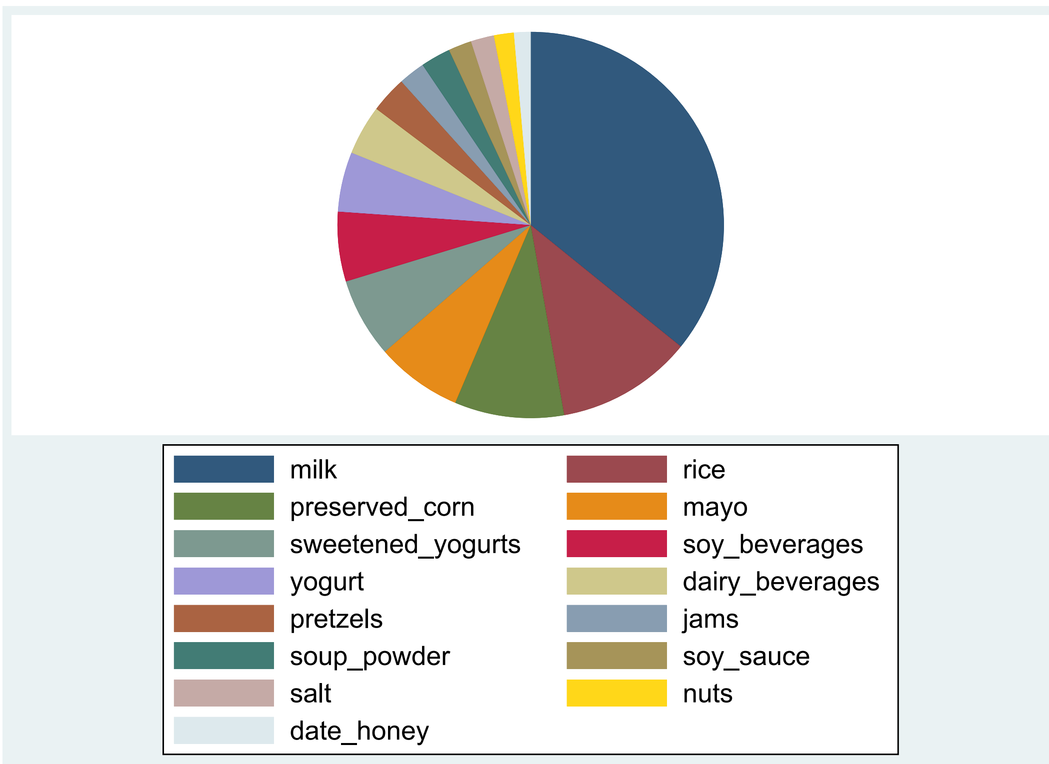

Note: This figure maps the composition of purchases made in the pre-period across all categories impacted by the intervention.

Figure 4: These figures map the point estimate and 95 percent confidence interval surrounding the DDD estimates when the analysis is run separately on consumers who are defined as top and bottom types.

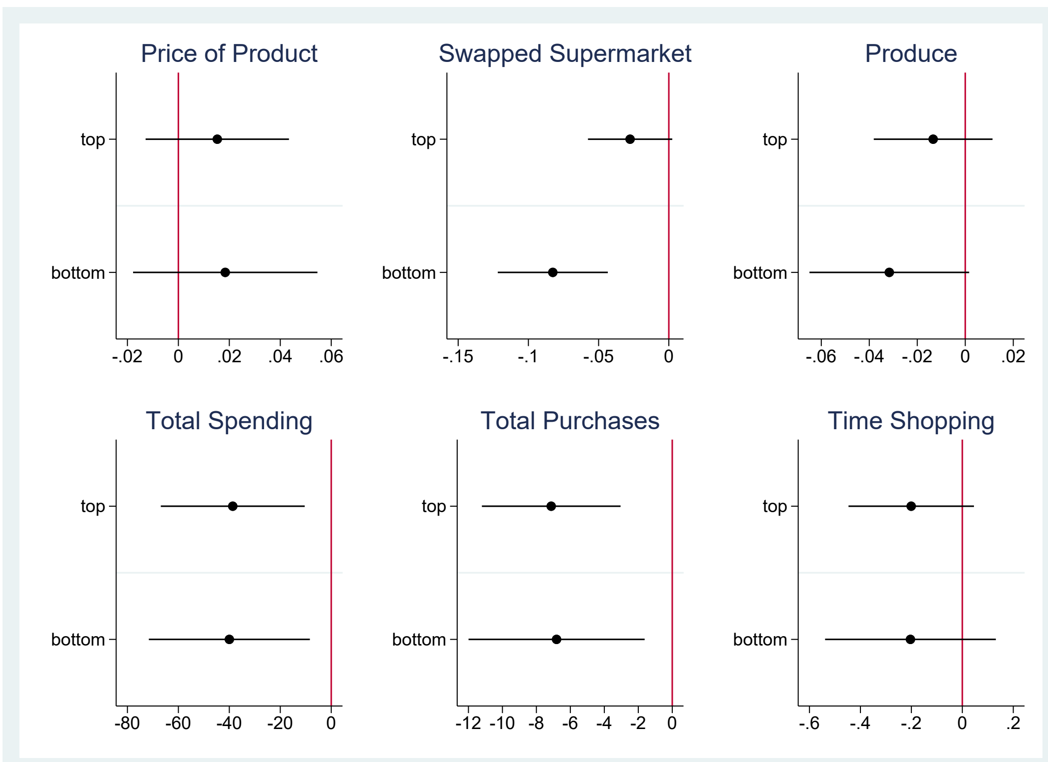

Figure 5: These figures map the point estimate and 95 percent confidence interval surrounding the DDD estimates when the analysis is run separately on consumers who are defined as top and bottom types.

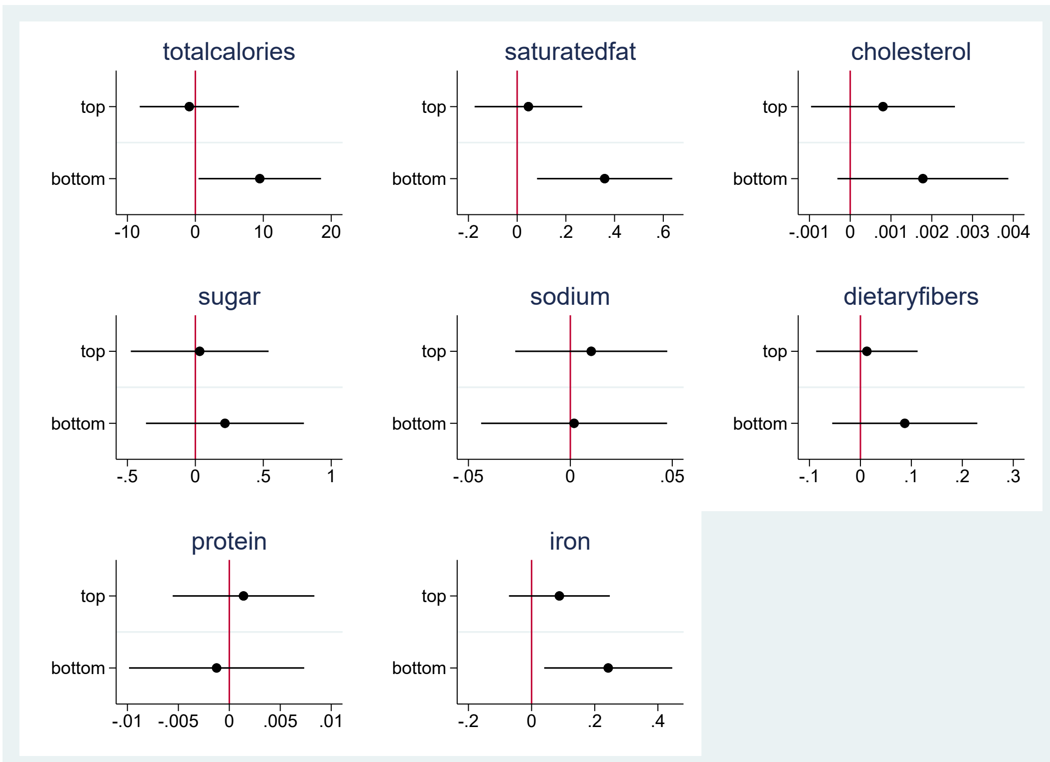

Figure 6: The bottom histogram shows the average variable value for the bottom type while the top dotplot shows the difference between the bottom type and top type with confidence bands. The red coloured dots display positive difference while the blue display negative.

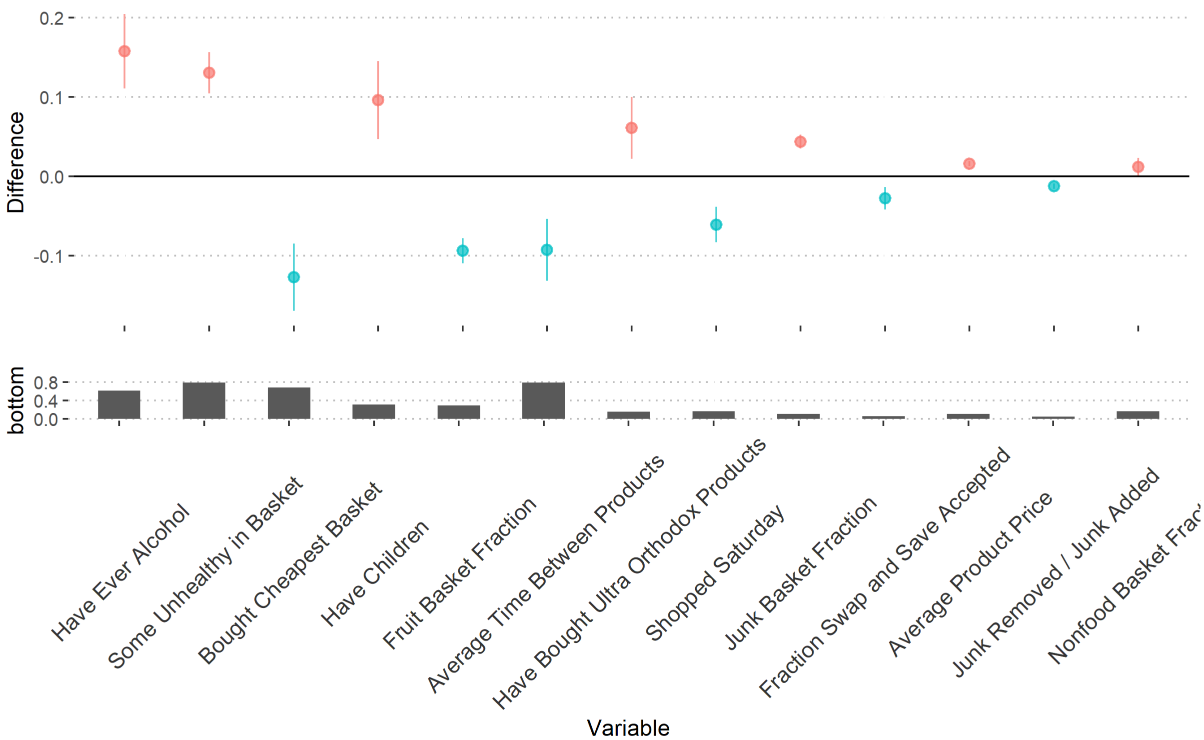

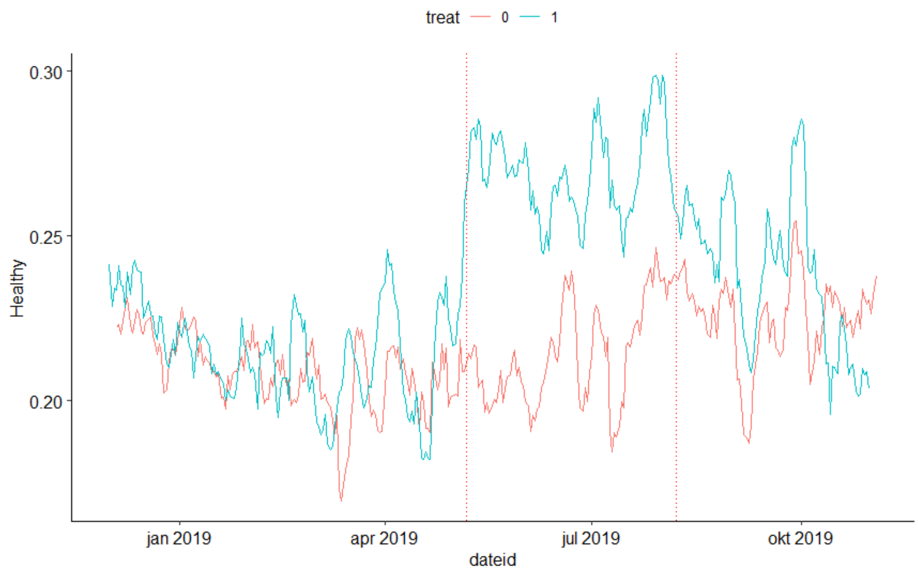

Note: This event study figure maps the 5 day moving average of the fraction of healthy-alternatives out of all relevant purchases for shoppers in the treatment and control groups.

Table 1: List of experimental products and healthier alternatives

Product	Healthier alternative	Price unhealthy	Price healthy	Unit
Canned corn	No added sugar	10.30	7.40	600 g
Chocolate milk	No added sugar	10.20	10.40	1 L
Yogurts	No added sugar	9.40	11.10	100 g
Date honey	No added sugar	8.20	13.90	400 g
Nuts	No added sugar	17.20	15.70	200 g
Soy milk	No added sugar	10.90	13.50	1 L
Jam	No added sugar	10.10	16.50	300 g
Salt	Low sodium	3.30	10.00	200 g
	Iodine fortified		5.70	200 g
Soy Sauce	Low sodium	14.90	19.00	400 g
Soup powder	Low sodium	14.40	16.60	400 g
Chocolate pudding	Low sodium, more calcium	7.10	8.40	100 g
Pretzels	Low sodium		8.30	300 g
	Low fat	11.80	11.80	400 g
Mayo	Low saturated fat	11.60	11.30	500 g
3% milk	1% milk (Low cholesterol)	6.20	5.80	1 L
Jasmine rice	Lower glycemic index	11.40	13.90	1 Kg
Persian rice	High fiber	9.20	8.40	1 Kg

Table 2: Summary

	Control	Treat	Difference
Basket Price	645.547	638.921	-6.62
	(518)	(465)	(3.824)
Number of products	38.105	37.736	-0.368
	(17.8)	(17.6)	(0.137)
Time Spent Shopping	23.683	23.874	0.191
	(16.7)	(17.2)	(0.132)
Business Hours	0.566	0.562	-0.004
	(0.496)	(0.496)	(0.004)
Fraction on sale	0.059	0.058	-0.001
	(0.143)	(0.14)	(0.001)
Swap & Save Accepts / Swap Save Offers	0.026	0.028	0.002
	(0.0771)	(0.079)	(0.001)
Produce / All Food	0.258	0.254	-0.004
	(0.17)	(0.168)	(0.001)
Junk / All Food	0.123	0.126	0.003
	(0.103)	(0.105)	(0.001)
Alcohol & Cigarettes / All Food	0.019	0.019	-0.001
	(0.0558)	(0.0539)	(0)
Contains Junk	0.866	0.877	0.011
	(0.22)	(0.207)	(0.002)
Contains Baby Products	0.39	0.383	-0.007
	(0.488)	(0.486)	(0.004)
Bought Cheapest Supermarket	0.141	0.14	0
	(0.183)	(0.18)	(0.001)
N	32872	33342	470

Basket Level

Table 3: Summary - Bought healthy during intervention

	Control	Treat	Difference
Total Calories	91.9690	92.3170	0.347
	(36.7)	(36.9)	(0.286)
Sugar	3.2200	3.1930	-0.027
	(2.31)	(2.31)	(0.018)
Sodium	0.2350	0.2380	0.003
	(0.327)	(0.316)	(0.002)
Saturated Fat	2.4210	2.4660	0.045
	(1.47)	(1.47)	(0.011)
Dietary Fibers	1.1500	1.1520	0.002
	(0.68)	(0.675)	(0.005)
Cholesterol	0.0130	0.0130	0
	(0.0101)	(0.0103)	(0)
Protein	4.9100	4.9460	0.035
	(2.94)	(2.94)	(0.023)
Iron	0.7950	0.7790	-0.015
	(0.701)	(0.623)	(0.005)
N	32872	33342	470

Basket Level

Table 4: Features with definition

Variable	Definition
Contains Junk	Likelihood of adding some unhealthy item to basket in pre-period
Business hours (Sunday to Thursday 8-17) (0/1)	Likelihood of shopping between 9-17 on working days in pre-period
Low Average Basket value	Likelihood of have a basket value below the median basket value in pre-period
Swap & Save Not accepted/Swap & Save Offers	Fraction of ´Swap N' save´ opened but not accepted in pre-period
Swap & Save Accepts/Swap & Save Offers	Fraction of ´Swap N' save´ accepted conditional on opening
Produce before junk	Likelihood of adding a long sequence of produce before a long sequence of junk
Products bought at full price / All products bought	Average fraction of products in basket not on sale
Contains Ultra Orthodox Product	Have ever bought item identified as kocher
Contains Baby Product	Have ever bought items associated with having children
Ended basket on junk	Likelihood of last item in basket being a junk item in pre-period
Have ever added alcohol	Have ever added alcohol or cigarettes to basket
Healthy/Relevent Products pre-intervention	Likelihood of adding the healthy alternative in pre-period
Junk removed / Junk added	Average fraction of added junk removed again
Time between items	Average time passed between adding items
Distribution of shopping days	Fraction of baskets ended on day i
Junk fraction	Number of junk items bought / all products
Produce fraction	Number of produce items bought / all products
Nonfood fraction	Number of nonfood items bought / all products
Alcohol/cigarettes fraction	Number of alcohol and cigarettes items bought / all products
mean product price	Average Price paid for products
Cheapest basket selected	The fraction of times where the cheapest version of the basket was selected

Table 5: Effect of SABH on healthy purchases—Difference in differences

	Balanced panel			Relevant purchases		
	(1)	(2)	(3)	(4)	(5)	(6)
Treat	-0.0003	-0.0003	0.0000	0.0092	0.0075	0.0000
	(0.0012)	(0.0012)	(0.0000)	(0.0077)	(0.0074)	(0.0000)
During experiment	-0.0096	-0.0096	-0.0096	0.0065	0.0101	0.0054
	(0.0008)	(0.0008)	(0.0008)	(0.0042)	(0.0040)	(0.0036)
After experiment	-0.0149	-0.0149	-0.0149	0.0138	0.0140	0.0076
	(0.0008)	(0.0008)	(0.0008)	(0.0052)	(0.0051)	(0.0041)
Treat × During experiment	0.0087	0.0087	0.0087	0.0441	0.0427	0.0457
	(0.0012)	(0.0012)	(0.0012)	(0.0061)	(0.0059)	(0.0053)
Treat × After experiment	0.0025	0.0025	0.0025	0.0074	0.0074	0.0133
	(0.0012)	(0.0012)	(0.0012)	(0.0074)	(0.0072)	(0.0062)
Outcome mean	0.0403	0.0403	0.0403	0.2116	0.2116	0.2116
Sample size	373050	373050	373050	136415	136415	136415
Category FE		X	X		X	X
Shopper FE			X			X

Note: The dependent variable is an indicator for purchasing the healthy product. Column (1) contains a treatment group indicator, time period fixed effects (before, during, and after the intervention period), and their interactions. Column (2) adds product category fixed effects. Column (3) adds shopper fixed effects. Columns (4) to (6) are analogous but restrict the analysis to purchases of relevant products in the experiment. Standard errors reported in parentheses are adjusted for clustering at the shopper level.

Table 6: The Effect of the SABH Nudge on Shopping Behavior

Panel A: Product-level effects

	(1) Mean	(2) During Intervention	(3) Pre-Intervention	(4) DDD
Log (price)	10.3376	0.0248	-0.0038	0.0220
	(8.7799)	(0.0101)	(0.0092)	(0.0082)
Produce	0.2588	-0.0304	0.0050	-0.0250
	(0.4380)	(0.0104)	(0.0093)	(0.0074)
Calories	199.7691	4.5685	-1.2989	2.9970
	(186.7672)	(2.5976)	(2.2302)	(2.0896)
Saturated Fat	3.7201	0.1693	-0.0612	0.1665
	(6.4739)	(0.0752)	(0.0628)	(0.0626)
Cholesterol	0.0190	0.0019	-0.0005	0.0021
	(0.0575)	(0.0006)	(0.0005)	(0.0005)
Sugar	8.8449	0.0938	0.1346	-0.1158
	(15.8148)	(0.1606)	(0.1295)	(0.1381)
Sodium	0.3163	0.0080	-0.0014	0.0031
	(1.4591)	(0.0100)	(0.0073)	(0.0103)
Dietary Fibers	3.2849	0.0554	-0.0024	0.0399
	(3.7355)	(0.0305)	(0.0235)	(0.0301)
Protein	0.0624	0.0039	-0.0011	0.0045
	(0.2419)	(0.0024)	(0.0018)	(0.0020)
Iron	4.7473	0.1087	-0.0440	0.1100
	(4.8110)	(0.0463)	(0.0402)	(0.0446)
N	943790	694585	943790	2144192

Panel B: Basket-level effects

	(1) Mean	(2) During Intervention	(3) Pre-Intervention	(4) DDD
Basket size	38.2138	-7.9948	-2.0993	-5.9887
	(38.2177)	(1.1012)	(0.8813)	(1.1517)
Basket price	284.6531	-55.5129	-19.4482	-36.7780
	(232.8622)	(8.1569)	(6.3129)	(7.4650)
Log (time shopping)	3.3764	-0.2509	-0.0520	-0.2105
	(2.6517)	(0.0754)	(0.0572)	(0.0717)
Supermarket swap	0.1038	-0.0531	-0.0044	-0.0490
	(0.3050)	(0.0094)	(0.0060)	(0.0089)
N	70168	34030	70168	104198

Note: Standard errors reported in parentheses are adjusted for clustering at the shopper level. All specifications control for user fixed effects. In Panel A the analysis is run on a full database where observations include all products added to each basket by shoppers in the treatment and control groups. In Panel B the database is collapsed to 1 observation prior to adding the first relevant product to basket, and 1 observation post adding this product (if the relevant product was the first or last product added to basket then this basket will only have one observation. Column (1) provides summary statistics for each variable in the Pre-intervention period. Columns (2) and (3) provide the coefficient on the interaction term postxtreat during each of the relevant periods. Column (4) provides the coefficient on the DDD interaction term postxtreatxDuring.

Table 7: Exploring heterogeneity in response to the Swap and Be Healthy intervention. The table shows the result for variables with importance higher than 0.01 percent. The median split value is calculated as the median over the average for each tree and S refers to this value.

	Variable	Importance	Median Split Value	X >S	X <= S	Category
1	Average Product Price	0.1190	9.8970	0.011	0.007	thriftiness
2	Junk Removed / Junk Added	0.1070	0.0350	0.008	0.011	Healthy
3	Healthy / Relevant products	0.1070	0.1250	0.035	0.008	Healthy
4	Category-Rice	0.0650	0	0.02	0.009	Category
5	Fruit / All products	0.0540	0.2400	0.008	0.011	Healthy
6	Ended Basket on Junk	0.0510	0.1500	0.012	0.008	Healthy
7	Shopped Saturday	0.0460	0.2000	0.006	0.01	Day
8	Category-Milk	0.0430	0	0.021	0.008	Category
9	Junk / All products	0.0430	0.1070	0.011	0.008	Healthy
10	Nonfood / All products	0.0400	0.1490	0.01	0.009	Basket
11	Contains junk	0.0330	0.8180	0.011	0.004	Healthy
12	Shopped Tuesday	0.0320	0.2110	0.011	0.009	Day
13	Started with Produce	0.0280	0.2170	0.01	0.009	Healthy
14	Average Time Between Products	0.0280	37.7190	0.009	0.01	Basket
15	Alcohol and Cigarettes / All products	0.0250	0.0160	0.011	0.009	Healthy
16	Swap and Save Accepted / Swap and Save Opened	0.0230	0.0320	0.011	0.009	thriftiness
17	Shopped Friday	0.0230	0.1000	0.007	0.01	Day
18	Shopped Sunday	0.0220	0.2220	0.011	0.008	Day
19	Shopped Monday	0.0170	0.2000	0.009	0.009	Day
20	Shopped Wednesday	0.0170	0.2000	0.01	0.009	Day
21	Bought Cheapest Supermarket	0.0150	0.4470	0.01	0.009	thriftiness
22	Category-Mayo	0.0140	0	0.012	0.009	Category
23	Category-Jams	0.0120	0	0.003	0.01	Category

Table 8: Effect of SABH on healthy purchases—Difference in differences (Top and Bottom groups)

	Balanced panel			Relevant purchases		
	(1)	(2)	(3)	(4)	(5)	(6)
Treat	-0.0014	-0.0014	0.0000	0.0428	0.0424	0.0000
	(0.0023)	(0.0023)	(0.0000)	(0.0135)	(0.0129)	(0.0000)
During experiment	-0.0390	-0.0390	-0.0390	-0.0767	-0.0616	-0.0690
	(0.0014)	(0.0014)	(0.0014)	(0.0072)	(0.0069)	(0.0062)
After experiment	-0.0283	-0.0283	-0.0283	-0.0275	-0.0268	-0.0319
	(0.0015)	(0.0015)	(0.0015)	(0.0084)	(0.0083)	(0.0072)
Treat × During experiment	0.0596	0.0596	0.0596	0.2276	0.2068	0.2081
	(0.0022)	(0.0022)	(0.0022)	(0.0113)	(0.0109)	(0.0101)
Treat × After experiment	0.0153	0.0153	0.0153	0.0674	0.0639	0.0720
	(0.0022)	(0.0022)	(0.0022)	(0.0128)	(0.0125)	(0.0115)
Outcome mean	0.0509	0.0509	0.0509	0.2498	0.2498	0.2498
Sample size	128160	128160	128160	48007	48007	48007
Category FE		X	X		X	X
Shopper FE			X			X
Treat	-0.0073	-0.0073	0.0000	-0.0638	-0.0528	0.0000
	(0.0017)	(0.0017)	(0.0000)	(0.0153)	(0.0145)	(0.0000)
During experiment	0.0178	0.0178	0.0178	0.1508	0.1416	0.1260
	(0.0016)	(0.0016)	(0.0016)	(0.0100)	(0.0098)	(0.0093)
After experiment	-0.0054	-0.0054	-0.0054	0.0642	0.0592	0.0511
	(0.0015)	(0.0015)	(0.0015)	(0.0120)	(0.0118)	(0.0101)
Treat × During experiment	-0.0366	-0.0366	-0.0366	-0.2210	-0.1995	-0.1692
	(0.0019)	(0.0019)	(0.0019)	(0.0122)	(0.0120)	(0.0112)
Treat × After experiment	-0.0046	-0.0046	-0.0046	-0.0621	-0.0607	-0.0419
	(0.0019)	(0.0019)	(0.0019)	(0.0150)	(0.0147)	(0.0128)
Outcome mean	0.0283	0.0283	0.0283	0.1878	0.1878	0.1878
Sample size	128115	128115	128115	31526	31526	31526
Category FE		X	X		X	X
Shopper FE			X			X

Note: The dependent variable is an indicator for purchasing the healthy product. The sample in Panels A and B consist of households in the top and bottom 25 percent, respectively, of the distribution of predicted responses to the treatment. Column (1) contains a treatment group indicator, time period fixed effects (before, during, and after the intervention period), and their interactions. Column (2) adds product category fixed effects. Column (3) adds shopper fixed effects. Columns (4) to (6) are analogous but restrict the analysis to purchases of relevant products in the experiment. Standard errors reported in parentheses are adjusted for clustering at the shopper level.

Table 9: The Effect of the SABH Nudge on Shopping Behavior—Does Cognitive Overload Play a Role?

Panel A: Product-level effects

	(1) Mean	(2) Had S&S	(3) No S&S	(4) Diff	(5) Prod. 1–5	(6) Prod. 6	(7) Diff
Log (price)	2.0578	-0.0474	0.0279	-0.0752	0.0473	0.0171	0.0302
	(0.7618)	(0.0307)	(0.0095)	(0.0322)	(0.0320)	(0.0094)	(0.0331)
Produce	0.2559	0.0207	-0.0310	0.0517	-0.0490	-0.0253	-0.0238
	(0.4364)	(0.0249)	(0.0081)	(0.0261)	(0.0219)	(0.0080)	(0.0234)
Calories	198.3879	-1.6465	3.1510	-4.7975	10.2911	2.0361	8.2549
	(186.3036)	(8.1613)	(2.2987)	(8.4490)	(7.5898)	(2.3495)	(7.9240)
Saturated Fat	3.7038	-0.0920	0.1979	-0.2899	0.4200	0.1786	0.2413
	(6.4600)	(0.2230)	(0.0685)	(0.2310)	(0.2511)	(0.0703)	(0.2593)
Cholesterol	0.0190	0.0024	0.0021	0.0003	0.0037	0.0022	0.0015
	(0.0572)	(0.0018)	(0.0005)	(0.0019)	(0.0025)	(0.0005)	(0.0026)
Sugar	8.7793	-0.8309	-0.0920	-0.7389	-0.1407	-0.1504	0.0096
	(15.6095)	(0.5506)	(0.1501)	(0.5670)	(0.5595)	(0.1578)	(0.5817)
Sodium	0.3123	0.0318	0.0080	0.0237	0.0487	0.0023	0.0464
	(1.4272)	(0.0398)	(0.0106)	(0.0412)	(0.0341)	(0.0112)	(0.0359)
Dietary Fibers	3.2071	-0.0426	0.0534	-0.0960	-0.0498	0.0372	-0.0871
	(3.6576)	(0.1178)	(0.0310)	(0.1207)	(0.1294)	(0.0330)	(0.1327)
Protein	0.0628	-0.0024	0.0042	-0.0066	0.0202	0.0026	0.0176
	(0.2426)	(0.0081)	(0.0022)	(0.0084)	(0.0086)	(0.0023)	(0.0089)
Iron	4.6044	0.0908	0.1203	-0.0296	-0.2388	0.1362	-0.3750
	(4.5905)	(0.1541)	(0.0478)	(0.1607)	(0.2526)	(0.0486)	(0.2553)
N	1385260	143490	2000702	1385260	327666	1816526	1385260

Panel B: Basket-level effects

	(1) Mean	(2) Had S&S	(3) No S&S	(4) Diff	(5) Prod. 1–5	(6) Prod. 6–	(7) Diff
Basket size	37.8873	-1.4489	-5.8285	4.3796	-3.6248	-4.3700	0.7452
	(39.0682)	(3.9137)	(1.1365)	(4.0240)	(2.0026)	(1.1043)	(2.1549)
Basket price	285.2282	-32.0953	-33.9148	1.8196	-7.1316	-29.5888	22.4572
	(236.9177)	(26.6876)	(7.3409)	(27.4346)	(11.8568)	(7.4859)	(13.5634)
Log(duration)	3.3491	-0.3779	-0.2268	-0.1511	0.1147	-0.2893	0.4039
	(2.6317)	(0.2427)	(0.0741)	(0.2535)	(0.1832)	(0.0722)	(0.1956)
Sup. swap	0.1208	0.0304	-0.0534	0.0837	-0.0511	-0.0469	-0.0042
	(0.3259)	(0.0354)	(0.0089)	(0.0366)	(0.0241)	(0.0092)	(0.0260)
N	104198	6337	97861	104198	18352	85846	104198

Note: Standard errors reported in parentheses are adjusted for clustering at the shopper level. In Panel A the analysis is run on a full database where observations include all products added to each basket by shoppers in the treatment and control groups. In Panel B the database is collapsed to 1 observation prior to adding the first relevant product to basket, and 1 observation post adding this product (if the relevant product was the first or last product added to basket then this basket will only have one observation. Column (1) provides summary statistics for each variable in the full sample. Columns (2) and (3) provide the coefficient on the DDD interaction term postxtreatxDuring for shoppers based on whether or not they examined a Swap & Save opportunity prior to receiving the SABH popup. Columns (5) and (6) provide this same coefficient based on whether or not they received the SABH popup at the beginning of their shopping trip (within the first 5 products added to basket).

Table 10: The Effect of the SABH Nudge on Shopping Behavior—Do Guilt/Self Control Costs Play a Role?

Panel A: Product-level effects	(1) Mean	(2) Guilty Types	(3) Non-Guilty Types	(4) Diff
Log (price)	2.0578	0.0388	0.0076	0.0312
	(0.7618)	(0.0160)	(0.0109)	(0.0194)
Produce	0.2559	-0.0434	-0.0130	-0.0303
	(0.4364)	(0.0133)	(0.0093)	(0.0162)
Calories	198.3879	12.5115	-2.6703	15.1818
	(186.3036)	(4.1235)	(2.6103)	(4.8798)
Saturated Fat	3.7038	0.2746	0.0856	0.1890
	(6.4600)	(0.1184)	(0.0791)	(0.1424)
Cholesterol	0.0190	0.0003	0.0025	-0.0022
	(0.0572)	(0.0008)	(0.0006)	(0.0010)
Sugar	8.7793	0.0838	-0.2650	0.3489
	(15.6095)	(0.2880)	(0.1672)	(0.3329)
Sodium	0.3123	0.0043	0.0078	-0.0035
	(1.4272)	(0.0185)	(0.0123)	(0.0222)
Dietary Fibers	3.2071	0.0472	0.0398	0.0074
	(3.6576)	(0.0546)	(0.0364)	(0.0656)
Protein	0.0628	-0.0001	0.0044	-0.0045
	(0.2426)	(0.0039)	(0.0026)	(0.0047)
Iron	4.6044	0.2257	0.0702	0.1555
	(4.5905)	(0.0819)	(0.0552)	(0.0987)
N	1385260	630791	1513401	1385260

Panel B: Basket-level effects	(1) Mean	(2) Guilty Types	(3) Non-Guilty Types	(4) Diff
Basket size	37.8873	-11.9441	-3.5487	-8.3954
	(39.0682)	(2.4745)	(1.1961)	(2.7478)
Basket price	285.2282	-73.8924	-23.3991	-50.4933
	(236.9177)	(15.4335)	(7.9098)	(17.3390)
Log (time shopping)	3.3491	-0.2200	-0.2071	-0.0129
	(2.6317)	(0.1572)	(0.0781)	(0.1755)
Supermarket swap	0.1208	-0.0687	-0.0423	-0.0265
	(0.3259)	(0.0178)	(0.0097)	(0.0203)
N	104198	30435	73763	104198

Note: Standard errors reported in parentheses are adjusted for clustering at the shopper level. In Panel A the analysis is run on a full database where observations include all products added to each basket by shoppers in the treatment and control groups. In Panel B the database is collapsed to 1 observation prior to adding the first relevant product to basket, and 1 observation post adding this product (if the relevant product was the first or last product added to basket then this basket will only have one observation. Column (1) provides summary statistics for each variable in the full sample. Columns (2) and (3) provide the coefficient on the DDD interaction term postxtreatxDuring for shoppers based on whether or not they struggled with self-control/guilt costs in the pre-period. A shopper is defined as a guilty type if she added a junk food item to her basket, but removed it prior to purchase.

Table 11: Effect of SABH on nutrition—Difference in differences

	All		Top 25%		Bottom 25%	
	(1)	(2)	(3)	(4)	(5)	(6)
Total calories	-35.5326	-35.5920	119.5089	81.2771	-107.9633	-52.3974
	(26.4421)	(23.6024)	(45.1268)	(41.1050)	(51.2909)	(47.0726)
Total fat	-3.6149	-2.8061	4.9111	3.4645	-5.3835	0.8296
	(2.5703)	(2.3054)	(4.5098)	(4.0365)	(4.8502)	(4.6979)
Saturated fat	-1.2693	-0.6997	1.4106	1.3397	-0.8446	1.8343
	(0.9185)	(0.8382)	(1.6396)	(1.4823)	(1.6998)	(1.7095)
Cholesterol	-0.0058	-0.0037	0.0045	0.0045	-0.0150	-0.0011
	(0.0046)	(0.0042)	(0.0077)	(0.0071)	(0.0092)	(0.0082)
Sugar	-1.3387	-1.7735	7.2756	5.3806	-4.7735	-3.8420
	(1.2186)	(1.1540)	(2.0795)	(1.9611)	(2.2829)	(2.2194)
Sodium	-0.0810	-0.0420	0.2591	0.2228	-0.6528	-0.4625
	(0.1432)	(0.1463)	(0.2537)	(0.2642)	(0.2771)	(0.2877)
Dietary fiber	-0.4608	-0.3968	0.4350	0.1120	-1.5205	-0.4438
	(0.3840)	(0.3431)	(0.6441)	(0.5896)	(0.7681)	(0.6654)
Protein	-1.7608	-0.9785	5.1009	3.7358	-5.1202	-2.7131
	(1.6577)	(1.5517)	(2.6657)	(2.5969)	(3.1952)	(3.2117)
Carbohydrates	-12.2444	-9.6376	15.4698	10.9140	-30.8538	-17.3916
	(5.9544)	(5.2261)	(10.2850)	(9.1957)	(11.3244)	(10.4955)
Shopper FE		X		X		X

Note: Each row reports the effect of being in the treated group compared to the control group on the total quantity of the specified nutrient in the shopper's basket. Columns (1) and (2) display results for the full sample of shoppers, Columns (3) and (4) display results for shoppers in the top 25 percent of the distribution of predicted responses to the treatment, and Columns (5) and (6) display results for shoppers in the bottom 25 percent of the distribution of predicted responses to the treatment. Standard errors reported in parentheses are adjusted for clustering at the shopper level.

Online Appendix