

A new year, a new you?  
A two-selves model of within-individual  
variation in food purchases

Laurens Cherchye, Bram De Rock, Rachel Griffith,  
Martin O'Connell, Kate Smith and Frederic Vermeulen\*

September, 2019

**Abstract**

We document that *within*-individual variation in food choices is substantial and has potentially important consequences for nutrition, and hence well-being. We develop an approach that allows us to study the determinants of this within-individual variation within an economic framework and allow for across-individual preference heterogeneity. We show that around one-quarter of within-individual fluctuations in diet quality are explained by standard economic variables (prices and budgets), along with advertising and weather. The residual fluctuations are important and are larger for lower income and younger people, and individuals who state they are impulsive.

**Keywords:** Two-selves model, collective model, revealed preferences, diet quality  
**JEL classification:** C14, D12, D90, I12.

---

\*Cherchye is at the University of Leuven, De Rock is at Université Libre de Bruxelles and the University of Leuven, Griffith is at the Institute for Fiscal Studies and University of Manchester, O'Connell is at the Institute for Fiscal Studies, Smith is at the Institute for Fiscal Studies and University College London and Vermeulen is at the University of Leuven and the Universidad del Pacífico. **Correspondence:** laurens.cherchye@kuleuven.be, bderock@ulb.ac.be, rgriffith@ifs.org.uk, martin\_o@ifs.org.uk, kate\_s@ifs.org.uk and frederic.vermeulen@kuleuven.be. **Acknowledgements:** The authors gratefully acknowledge financial support under ERC-2015-AdG-694822 and ERC-2013-CoG-614221, the ESRC grant number ES/M010147/1, the British Academy under pf160093, the FWO and FNRS. Data supplied by Kantar. The use of Kantar data in this work does not imply the endorsement of Kantar in relation to the interpretation or analysis of the data. All errors and omissions remained the responsibility of the authors.

“New Year’s Day ... now is the accepted time to make your regular annual good resolutions. Next week you can begin paving hell with them as usual.” – *Mark Twain*

## 1 Introduction

The consequences of poor dietary choices for well-being and welfare are potentially profound, and governments around the world are struggling with how to address rising rates of obesity and diet related non-communicable diseases.<sup>1</sup> The relation between inequalities in diet quality and associated health outcomes *across* individuals is well documented.<sup>2</sup> We focus, however, on the empirical determinants of *within*-individual variation in dietary choices. It is informative for policy design to understand whether successfully encouraging people to more consistently behave as they themselves do in their relatively healthy periods could be a significantly more fruitful strategy than policies that focus more abstractly on persuading people to adopt more healthy choice behavior.

In this paper we make three contributions to the existing literature. First, we exploit new longitudinal data to document substantial within-individual variation in diet quality using information on individuals’ entire shopping baskets. We show that this variation is large and important from a nutritional perspective, and that it is of a similar magnitude to the well documented variation we see across people. Second, we develop a modelling approach that allows us to study the determinants of within-individual variation in food choices. This includes the impact of standard economic variables such as prices and budgets, as well as other factors, while also allowing for rich cross-section heterogeneity in choice behavior. We model individuals’ food choices as a bargaining process between two selves – a “healthy” and “unhealthy” self. This approach relies on insights from the collective household literature, and allows us to adapt revealed preference methods to show that our data are consistent with a two-selves model of behavior. Third, we empirically quantify the drivers of within-individual fluctuations in diet quality. We show that standard economic variables (prices, budgets), along with advertising and weather, are important and explain about one-quarter of within-individual variation in choices between healthy and unhealthy foods. The remaining three-quarters of variation is not explained by these variables. We show that the extent of this residual vari-

---

<sup>1</sup>See Finkelstein et al. (2005), Currie (2009), Currie et al. (2010) and Finkelstein et al. (2013).

<sup>2</sup>See Baum and Ruhm (2009), Cutler and Lleras-Muney (2010), Bhattacharya and Sood (2011) and Allcott et al. (2017).

ation is correlated with observable characteristics of individuals, such as age and income, and with stated attitudes that reflect impulsiveness.

We use longitudinal data on the entire shopping basket of a sample of British individuals (the Kantar Worldpanel) to document substantial within-individual variation in diet quality over time. Looking over several years, we show that within-individual variation in diet quality is substantial from a nutritional perspective. For example, the within-individual standard deviation of the share of calories from added sugar (across months) is 4 percentage points; in comparison the UK government’s sugar reduction strategy targets a 5% reduction in sugar consumption in its first year (Public Health England (2017)). We show that the variation in diet quality that we see in our data can be parsimoniously described by shifts in the share of foods that are contained in a healthy (e.g. fruits, vegetables, whole grains, etc.) versus an unhealthy (e.g. soda, crisps, confectionery, etc.) basket. There is a sharp increase in the healthiness of foods purchased at the beginning of January, followed by a steady decline in healthiness over the calendar year. These declines are substantial: the median decline in the share of calories from healthy foods from the first to the final quarter of the calendar year is 5 percentage points, which is the same as the difference between the average purchases of a normal weight individual and an obese individual.

We are interested in quantifying what drives this within-individual variation in food choices. We want to retain the ability of standard economic variables, such as prices and budgets, to affect choices, and to allow for cross-sectional heterogeneity in preferences, but also to relax the usual assumptions about the stability of preferences within-individual over time in an economically intuitive way. The ability to describe variation as shifts between a healthy and unhealthy basket is appealing because it accords well with two-selves models of choice behavior.<sup>3</sup> We model an individual’s food choices as the outcome of a bargaining process between a healthy and an unhealthy self, each with stable preferences over separate food baskets. In order to analyze this model empirically we draw on insights from the collective household model.<sup>4</sup> Most importantly, this allows us to map the empirically observed share of spending on the healthy and unhealthy baskets to the theoretical bargaining weight between the selves; this bargaining weight can be affected by both standard and non-standard economic variables.

---

<sup>3</sup>See Strotz (1955) and Peleg and Yaari (1973), Thaler and Shefrin (1981), Laibson (1997), O’Donoghue and Rabin (1999a, 1999b), Gul and Pesendorfer (2001, 2004), Bénabou and Pycia (2002) and Fudenberg and Levine (2006).

<sup>4</sup>See Chiappori (1988), Chiappori (1992), Browning and Chiappori (1998), Chiappori and Ekeland (2009), Dunbar et al. (2013) and Browning et al. (2013).

Another advantage of this approach is that it allows us to make use of revealed preference methods to check whether the theoretical implications of our two-selves model are satisfied for our data, in the tradition of Samuelson (1938), Afriat (1967) and Varian (1982). We build on the work of Cherchye et al. (2007, 2011), who develop revealed preference methods to analyze collective choice behavior, to evaluate the consistency of the two-selves model with the data for each individual separately, and thus avoid assuming that individuals are characterized by homogeneous preferences. The results suggest that the two-selves model does a good job at explaining variation in the data, and a better job than the more standard single-self model, in which each individual is characterized by a single, stable utility function.

We combine the data and the model to provide an empirical quantification of the impact of standard economic variables and other factors in driving the healthiness of individuals' food choices. We find that, in the cross-section, the mean bargaining weight of the healthy self is higher for higher income and older individuals, people with a lower body mass index, and those that state they try to eat a healthy diet. Our long panel allows us to estimate the relationship between the bargaining weight and explanatory factors separately for each individual, avoiding the need to impose homogeneous preferences. Looking within-individual over time, we see that, on average, around 25% of the within-individual variation in the bargaining weight is explained by responses to changes in standard economic variables such as prices, budgets, along with advertising and weather variation.

Nonetheless, there remains substantial within-individual variation that is not explained by these variables, which is broadly characterized by a decline in diet quality over the calendar year. We show how these fluctuations vary with demographics, with particular focus on the correlation with age and income, as policy interventions are often targeted specifically at young people and lower income households. We find that individuals with lower income experience greater variation in the bargaining weight of the healthy self. This is true even after we control for potentially greater variation in the prices and food budgets that lower income individuals face,<sup>5</sup> variation in their responsiveness to these changes, as well as the influence of advertising and weather. We find a similar relationship with age; the average within-individual variation in the bargaining weight is higher for younger people than older people.

An extensive psychological literature shows that individual choice behavior varies with context and time, and that individuals sometimes use self-regulation

---

<sup>5</sup>See Kaplan and Menzio (2015) and Kaplan and Schulhofer-Wohl (2017).

and behavior modification in an attempt to mitigate these influences (see references and discussion in Rabin (1998) and DellaVigna (2009)). For example, experimental evidence suggests that individuals may be willing to impose (sometimes costly) commitments on themselves.<sup>6</sup> New Years' resolutions to eat a more healthy diet are a common form of self-regulation and behavior modification (Dai et al. (2014, 2015)). We use information on individuals' stated preferences and attitudes to investigate whether greater fluctuations in the bargaining weight of the healthy self reflect impulsive behavior. We find that fluctuations are larger for individuals who state that they are more impulsive (e.g. spend money without thinking). Our results therefore build on a literature that finds empirical evidence of considerable within-individual variation in choice behavior in other settings,<sup>7</sup> as well as in grocery purchases using alternative identification strategies.<sup>8</sup>

The rest of the paper is structured as follows. In the next section we provide evidence that there is substantial variation in food purchasing behavior within individuals across time. In Section 3 we describe the model and show that it can rationalize observed purchasing patterns. In Section 4 we analyze how the bargaining weight of the healthy self varies with standard economic factors, age and income. In Section 5 we discuss possible interpretations of our results and, in a final section, conclude and discuss some avenues for future work.

## 2 Food purchasing behavior

Poor dietary choices are thought to be a leading cause of rising obesity and diet related disease, with profound consequences for well-being and welfare,<sup>9</sup> and of particular concern for low income households (Currie (2009)). The socio-economic gradient in diet quality and associated health outcomes *across* individuals is well documented. We bring new data to bear on the extent and importance of *within*-individual variation in diet quality, which is less well understood.

---

<sup>6</sup>See Read and Van Leeuwen (1998), Read et al. (1999), Trope and Fishbach (2000), Ariely and Wertenbroch (2002) and Gilbert et al. (2002).

<sup>7</sup>See Ashraf et al. (2006), DellaVigna and Malmendier (2006), Oster and Morton (2005), Bucciol (2012) and Hinnsaar (2016).

<sup>8</sup>See Shapiro (2005), Milkman et al. (2010) and Sadoff et al. (2015).

<sup>9</sup>See Cutler et al. (2003), Bleich et al. (2008), Finkelstein and Zuckerman (2008) and World Health Organization (2015).

## 2.1 Data and measurement

We use data from the Kantar Worldpanel over the period 2005-2011, which we describe in detail in Appendix A.1. These data contain information on a representative sample of over 25,000 households. The data are longitudinal, with households remaining in the sample for on average two years. In this paper we focus on a sample of 3,645 individuals in Britain who live on their own. We do this to avoid the confounding issue of intra-household allocation mechanisms; we show in Appendix B that the basic patterns we describe also hold in the full sample for all household types.

Individuals record (using handheld scanners in the home) all grocery purchases made and brought into the home. The data are at the transaction (i.e. the barcode or UPC) level and include all foods and drinks, as well as household goods such as cleaning supplies and toiletries. We know the exact products purchased, the price paid (including discounts and special offers), and we have information on the nutritional characteristics of each product. This type of data is increasingly widely used in research.<sup>10</sup> From now on we use “foods” as shorthand for foods and non-alcoholic drinks.<sup>11</sup>

Diet quality is a complex multi-dimensional object; whether a diet is “healthy” depends on the consumption of a whole range of nutrients. Policy in the UK, US and many other developed countries has focused on foods that are high in fat, salt and sugar.<sup>12</sup> Eating too much of these types of foods is linked to increased risk of non-communicable diseases, such as heart disease, diabetes, and cancer, as well as to obesity.<sup>13</sup> In the UK, 88% of adults consume more than recommended levels of sugar, 45% more than recommended levels of saturated fat, and 35% more than recommended levels of salt.<sup>14</sup> There are other nutrients, such as protein and fibre, which are good for health and are under-consumed. We describe variation in purchases of these key nutrients below.

We show that this variation is well captured by dividing foods into “healthy” and “unhealthy” on the basis of an index used by the UK government. The nutrient profiling score (NPS) converts the multi-dimensional nutrient profile of

---

<sup>10</sup>See, for example, Hausman and Leibtag (2007), Aguiar and Hurst (2007), Broda and Weinstein (2010), Kaplan et al. (2016) and Dubois et al. (2014).

<sup>11</sup>We exclude alcoholic drinks from our main analysis; we have information on alcoholic drinks purchased and brought into the home. Our main conclusions are robust to their inclusion.

<sup>12</sup>See Public Health England (2017) and US Department of Agriculture (2015).

<sup>13</sup>See Ascherio et al. (1999) and World Health Organization (2015).

<sup>14</sup>Calculated using the National Diet and Nutrition Survey. UK government recommendations specify: no more than 5% of calories from added sugar, no more than 30g of saturated fat per day for men and 20g for women (aged 19-64), and no more than 6g salt per day.

each individual food product into a single-dimensional score.<sup>15</sup> The NPS is used by the UK government, for example, to regulate advertising of food and drinks on children’s television programming. It is similar in spirit to the Healthy Eating Index used by the US government to measure how well a basket of foods align with key government recommendations. The popular “MyPlate” tool is designed to encourage people to consume diets in line with these recommendations. The NPS ranges from -15 to 40, with a lower score indicating that the product is more healthy.<sup>16</sup> The products that have the lowest score are pulses and vegetables, with scores of around -10, and those with highest scores are solid fats, chocolates and biscuits, with scores over 20. We describe how nutrients are recorded in the data and provide details on the construction of the NPS in Appendix A.2.

Food products with an NPS below 4 and drinks products with an NPS below 1 are deemed “healthy” – this threshold is used by the UK Government to restrict advertising of unhealthy products to children. In this section, in order to describe variation in food choices, we use the government’s cutoff to classify foods as healthy or unhealthy. In Sections 3 and 4 we allow for the fact that individuals might have heterogeneous views about what counts as healthy by endogenizing the classification of foods as healthy or unhealthy.

## 2.2 Extent of within-individual variation in diet quality

We use the panel element of the data to look at the degree of within-individual (over time) variation in diet quality, in comparison with the cross-sectional variation. Table 2.1 shows that when purchases are aggregated to the weekly level, the within-individual intertemporal variation in diet quality is larger than the cross-sectional variation. This falls when we aggregate purchases to the monthly level, but within-individual intertemporal variation is still roughly the same as the well documented cross-sectional variation. For example, the within-individual standard deviation of the share of calories from added sugar (computed across months) is 4%, which is similar to the between individual standard deviation, and compares with a mean share of 12%. The bottom panel of the table shows that this pattern holds for the share of calories from healthy foods. In the rest of the paper we use monthly aggregation. This means we analyze broad changes in the healthiness of people’s diets rather than day-to-day fluctuations that may

---

<sup>15</sup>For more details see Rayner et al. (2005), Arambepola et al. (2008) and Rayner et al. (2009).

<sup>16</sup>The NPS gives positive points for the amount of saturated fat, sodium, sugar and calories per 100g and negative points for the amount of fibers, proteins and fruit, vegetables and nuts per 100g.

be driven by people eating cake on one day and compensating on the following day by abstaining, or by diets that entail high frequency changes (e.g. the 5:2 diet, which prescribes eating as normal for five days and eating a very restricted quantity of calories for two days).

### 2.3 Within-individual variation through time

One particularly salient dimension in which diet quality varies is that most people tend to make, and fail to keep, New Years' resolutions to lead more healthy lifestyles, and specifically to eat a healthier diet; for example we can see it in Google trends data in searches for "diet" (see Dai et al. (2014, 2015)). We show that this pattern is evident in our observational data, and has important implications for diet quality. However, it is important to note that this is only one of the dimensions in which diet quality varies within-individual.

In Figure 2.1, we show that, on average, the share of calories from healthy foods is highest in January and declines steadily over the year reaching a low in December; in Appendix B we show this pattern is evident for key nutrients, and for the full sample of households (which includes multiple as well as single occupancy households). In Figure 2.1(b) we plot the distribution of changes in the share of calories from healthy food from the first to the final quarter of the year, across individuals.

These figure shows two things. First, there is considerable heterogeneity across individuals in the change in share of calories from healthy food. The majority of individuals see their diet quality decline from Q1 to Q4, but these declines vary from a 20 percentage point decline in the share of calories from healthy food to no change, and some individuals actually seeing an increase in the healthiness of their diets over the calendar year. This illustrates the importance of allowing for heterogeneity in purchasing patterns across individuals in order to understand the determinants of within-individual variation. Second, these observed changes over the year are substantial. For example, the median decline in the share of calories from healthy foods from Q1 to Q4 is approximately 5 percentage points. This is the same as the difference between the average share of calories from healthy foods bought by a normal weight individual and an obese individual.<sup>17</sup>

---

<sup>17</sup>We have information on the body mass index (BMI) for individuals in our sample. Individuals with a BMI of between 18 and 20 (normal weight is classed as a BMI of between 18.5 to 25) get, on average, 51% of their calories from healthy foods, in comparison with obese individuals with a BMI between 35 and 40, who get, on average, 46% of their calories from healthy foods.




Table 2.1: *Variation in diet quality*

	(1)	(2)
	Purchases aggregated to:	
	Week	Month
<i>% calories from added sugar</i>		
Mean	12.19	12.02
Standard deviation	11.93	5.51
within individuals	11.25	3.90
between individuals	4.04	3.94
<i>% calories from saturated fat</i>		
Mean	14.34	14.40
Standard deviation	6.60	3.95
within individuals	6.01	2.85
between individuals	2.75	2.75
<i>% calories from protein</i>		
Mean	14.96	14.26
Standard deviation	9.68	3.85
within individuals	9.24	2.76
between individuals	2.97	2.74
<i>Salt intensity (g per 100g)</i>		
Mean	0.48	0.47
Standard deviation	0.51	0.32
within individuals	0.49	0.28
between individuals	0.15	0.17
<i>Fibre intensity (g per 100g)</i>		
Mean	1.36	1.33
Standard deviation	0.77	0.44
within individuals	0.70	0.31
between individuals	0.33	0.32
<i>% calories from healthy foods</i>		
Mean	51.02	48.34
Standard deviation	22.64	13.77
within individuals	20.50	10.00
between individuals	9.60	9.46

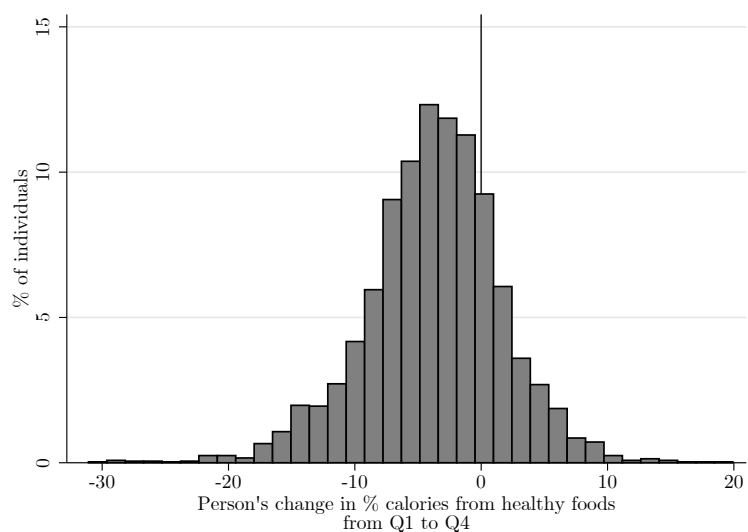
Notes: Column (1) shows the mean and variance decomposition for shopping baskets aggregated to the weekly level, column (2) shows the same for shopping baskets aggregated to the monthly level.

Figure 2.1: *Variation over time in % calories from healthy foods*

(a) Mean across individuals



(b) Within-individual change from Q1 to Q4



*Notes: Panel (a) shows the percentage of calories from healthy foods purchased on each shopping trip for a sample of 3,645 individuals over 2005–2011. Healthy foods are defined as those with an NPS of less than 4 (or drinks with an NPS below 1), see the text and Appendix A.2. The grey dotted line shows the mean across individuals on each day, the solid black line shows the mean across individuals in each week. Panel (b) shows the distribution of each individual's mean change from Q1 to Q4.*

A striking feature of the data is that there is substantial variation in individuals' diet quality over time, which is driven by changes in the choices they make over which foods to purchase. A second key feature of the data is heterogeneity in this behavior across individuals, in terms of both their average behavior and the

variation in their choices over time. Our objective is to understand what drives the within-individual variation in behavior. In doing this it is important to allow for the rich heterogeneity in behavior so clearly apparent in the data.

Our approach is to model agents' food purchasing behavior in a way that allows for some flexibility in agents' food preferences over time. This relaxes, in an intuitive way, the usual stability of preferences imposed by the standard single-self model of rational choice behavior. This standard model provides a useful starting point for addressing the question of whether the variation in individuals' purchasing patterns across time can be rationalized by responses to changes in the economic environment i.e. prices and food budgets. As we show at the end of Section 3, this standard rational choice model cannot fully explain individuals' food purchase behavior over time. Our model provides an economic interpretation to the additional variation in the share of spending on healthy foods; and we quantify the determinants of this variation in Section 4. A useful feature of our approach is that we can use a revealed preference methodology (Samuelson (1938), Afriat (1967) and Varian (1982)) to investigate whether the observed purchasing behavior can be rationalized by our model. This allows us to analyze each individual separately, so avoiding the assumption that different individuals are characterized by homogeneous preferences.

### **3 A two-selves model of food choice**

We present a two-selves model in which individuals' food choices are driven by the influence of a healthy self and an unhealthy self; both selves are characterized by their own stable and well-behaved preferences. The influence of the unhealthy relative to the healthy self may vary over time and determines individuals' food choices. The model draws on the literature on multiple-selves models.<sup>18</sup> The main feature that we take away from these multi-selves models is that consumer decision making can be thought of as a process in which two selves each have an influence on decision making.

---

<sup>18</sup>This dates back to Strotz (1955) and Peleg and Yaari (1973), with the main intuition underpinning these models remaining popular in the theoretical literature, e.g., Thaler and Shefrin (1981), Laibson (1997), O'Donoghue and Rabin (1999a, 1999b), Gul and Pesendorfer (2001, 2004), Bénabou and Pycia (2002) and Fudenberg and Levine (2006).

### 3.1 Data structure and notation

Consider  $I$  individuals, indexed  $i \in \{1, \dots, I\}$ , with food consumption bundles that contain both healthy and unhealthy foods. A first set of goods,  $H_i$ , is associated with a healthy lifestyle and contains items such as fruits, vegetables and whole grains. A second set,  $L_i$ , contains unhealthy foods, such as soda, crisps and confectionery. The  $i$  subscripts reflect the fact that individuals may have different views on the nutritional quality of specific food items. In Section 3.4, we explain how we empirically define the individual specific categorization of healthy and unhealthy sets of goods. In our empirical application we consider 85 goods that together make up the entire food and non-alcoholic drink grocery basket; these are each an aggregate of nutritionally similar food or drink products (UPCs or barcodes).

We observe  $T_i$  grocery baskets purchased by each individual  $i$ . The number of observations is individual specific, which does not pose a problem since the application of our model is individual-by-individual. For each observation  $t \in \{1, \dots, T_i\}$ , denote the quantities of the healthy food items by  $\mathbf{q}_{it}^h \in \mathbb{R}_+^{H_i}$  and the quantities of the less healthy food items by  $\mathbf{q}_{it}^l \in \mathbb{R}_+^{L_i}$ . The market prices associated with these are denoted by  $\mathbf{p}_t^h \in \mathbb{R}_{++}^{H_i}$  and  $\mathbf{p}_t^l \in \mathbb{R}_{++}^{L_i}$ , respectively. The individual's food budget spent on healthy food items is denoted by  $x_{it}^h$  and is equal to  $\mathbf{p}_t^h \mathbf{q}_{it}^h$ ; the food budget spent on less healthy food is denoted by  $x_{it}^l$  and it equals  $\mathbf{p}_t^l \mathbf{q}_{it}^l$ . The food budget of consumer  $i$  at time  $t$  is denoted by  $x_{it}$ , where  $x_{it} = x_{it}^h + x_{it}^l$ . To summarize, for each individual  $i \in \{1, \dots, I\}$  we observe the data  $S^i = \{(\mathbf{p}_t^h, \mathbf{p}_t^l; \mathbf{q}_{it}^h, \mathbf{q}_{it}^l), t = 1, \dots, T_i\}$ .

### 3.2 Two-selves model

We propose a two-selves model in which we assume that individual  $i$  is characterized by two selves, each with stable preferences. The first self is associated with a healthy lifestyle and derives utility from only the healthy food items  $\mathbf{q}_{it}^h$ . The second self derives utility from only the unhealthy food items  $\mathbf{q}_{it}^l$ . The preferences of each self are represented by the well-behaved utility functions  $u^{ih}(\mathbf{q}_i^h)$  and  $u^{il}(\mathbf{q}_i^l)$ . The two selves enter into a bargaining process that is different for every individual and that may not be stable over time.

One useful way to quantify the influence of both selves is to make use of the sharing rule concept, which we borrow from the literature on collective models.<sup>19</sup>

---

<sup>19</sup>See, for example, Chiappori (1988), Browning and Chiappori (1998) and Chiappori and Ekeland (2009).

The sharing rule distributes the food budget of individual  $i$ ,  $x_{it}$ , to the budget spent on healthy food items by the healthy self,  $x_{it}^h$ , and the budget spent on less healthy food items by the unhealthy self,  $x_{it}^l$ . The results from Chiappori (1988) imply that, under the assumption that each self has stable preferences and both selves choose Pareto efficient allocations, the sharing rule is a direct indication of the bargaining power of each self.<sup>20</sup>

More formally, Pareto efficiency implies that individual  $i$ 's observed food purchase behavior  $(\mathbf{q}_{it}^h, \mathbf{q}_{it}^l)$  can be represented as the solution of the following maximization problem:

$$\max_{\mathbf{q}_i^h, \mathbf{q}_i^l} \mu_{it} u^{ih}(\mathbf{q}_i^h) + (1 - \mu_{it}) u^{il}(\mathbf{q}_i^l) \quad (3.1)$$

subject to

$$\mathbf{p}_t^h \mathbf{q}_i^h + \mathbf{p}_t^l \mathbf{q}_i^l \leq x_{it}.$$

In this representation the parameter  $\mu_{it} \in [0, 1]$  is a Pareto weight that represents the bargaining weight of the healthy self in consumer  $i$ 's optimization problem in period  $t$ . If  $\mu_{it}$  equals one, then the individual behaves according to the healthy self's preferences, while if  $\mu_{it}$  equals zero the allocation of the food budget is determined by the unhealthy self's preferences.

The two-selves model is a direct generalization of a rational choice model in which healthy and unhealthy foods are strongly separable: this is characterized by the case where  $\mu_{it} = \mu_i \in ]0, 1[$ . The rational choice model with strongly separable preferences is a special case of the standard single-self model without any such constraints. In Section 3.4 we show that our two-selves model outperforms the standard single-self model and thus, by implication, the single-self model with strong separability.

Under the assumption of bargaining between the two selves, the Pareto weight generally depends on the food prices  $\mathbf{p}_t^h$  and  $\mathbf{p}_t^l$ , and on the food expenditure  $x_{it}$ . The Pareto weight may also depend on other factors, captured by  $\mathbf{z}_{it}$ , that have an impact on the bargaining weight between the two selves, but that do not affect the budget constraint. In our empirical application we include in  $\mathbf{z}_{it}$  the advertising of unhealthy foods, and weather variables.

---

<sup>20</sup>Our setting is also compatible with a situation in which the healthy and unhealthy selves behave noncooperatively. The intuition behind this result is that free-riding behavior is excluded by default, since the healthy food items are exclusively consumed by the healthy self, while the less healthy food items are exclusively consumed by the unhealthy self (see Cherchye et al. (2011)). This result implies that, in principle, in this setting any cooperative behavior can be represented as noncooperative behavior (and vice versa). We focus on the sharing rule interpretation.

The sharing rule gives the relative share of the healthy self's food expenditure in total food expenditure,  $\eta_{it} = x_{it}^h/x_{it}$ . There is a one-to-one relation between the Pareto weight  $\mu_{it}$  and the sharing rule for our collective consumption model; for given prices and food budget, the share of spending on healthy food is a monotone transformation of the bargaining weight of the healthy self. Hence, the sharing rule is a function of the same variables that affect the Pareto weight. We can therefore write:

$$\eta_{it} = \eta_i(\mathbf{p}_t^h, \mathbf{p}_t^l, x_{it}, \mathbf{z}_{it}). \quad (3.2)$$

In the next section we make extensive use of this sharing rule to describe the variation in the bargaining weights of the two selves over time. We show how much variation in the sharing rule is driven by changes in prices, budgets and other observable factors, such as advertising and weather, and how much variation remains unexplained by these observables.

### 3.3 Testable implications

#### Exact rationalizability

Browning and Chiappori (1998) and Chiappori and Ekeland (2009) characterize the testable implications of the collective model, which has a similar structure to our two-selves model. Cherchye et al. (2007, 2011) characterize similar conditions in a revealed preference setting *à la* Afriat (1967) and Varian (1982).

Pareto efficiency requires that the behavior of both selves of individual  $i$  can be rationalized in terms of stable and rational preferences:

**Definition 1** (Stable and rational preferences). *Let  $S^{ij} = \{(\mathbf{p}_t^j; \mathbf{q}_{it}^j), t = 1, \dots, T_i\}$  be a set of observations of self  $j$ , where  $j = h, l$ . Self  $j$ 's behavior is rationalizable if there exists a non-satiated utility function  $u^{ij}$  such that, for each observation  $t$ , we have  $u^{ij}(\mathbf{q}_{it}^j) \geq u^{ij}(\mathbf{q}_i^j)$  for all  $\mathbf{q}_i^j$  such that  $\mathbf{p}_i^{j'} \mathbf{q}_i^j \leq \mathbf{p}_i^{j'} \mathbf{q}_{it}^j$ .*

Cherchye et al. (2011) that the sets  $S^{ih} = \{(\mathbf{p}_t^h; \mathbf{q}_{it}^h), t = 1, \dots, T_i\}$  and  $S^{il} = \{(\mathbf{p}_t^l; \mathbf{q}_{it}^l), t = 1, \dots, T_i\}$  are rationalizable if and only if they satisfy a series of Afriat-style inequalities, which we refer to as the *Afriat Condition*.

**Definition 2** (Afriat Condition). *Let  $S^{ij} = \{(\mathbf{p}_t^j; \mathbf{q}_{it}^j), t = 1, \dots, T_i\}$  be a set of observations of self  $j$ , where  $j = h, l$ . The set  $S^{ij}$  satisfies the Afriat Condition if there exist numbers  $U_t^{ij}, \lambda_t^{ij} \in \mathbb{R}_{++}$  that meet, for all observations  $s$  and  $t$ , the Afriat Inequalities:*

$$U_s^{ij} - U_t^{ij} \leq \lambda_t^{ij} \mathbf{p}_t^{j'} (\mathbf{q}_{is}^j - \mathbf{q}_{it}^j).$$

These conditions do not impose explicit structure on the bargaining process, besides assuming Pareto optimal outcomes. In particular, the impact of the  $\mathbf{z}_{it}$  variables is implicitly modeled via the sharing of the budget. This implies that in the empirical application we do not need an exhaustive list of variables that influence the bargaining process in order to implement our testable implications.

The Afriat Inequalities are linear in the unknowns  $U_t^{ij}$  and  $\lambda_t^{ij}$ . Thus, we can use standard linear programming techniques to verify rationalizability of self  $j$ 's behavior for a given individual  $i$  (corresponding to the data set  $S^{ij}$ ). Checking behavioral consistency with the two-selves model requires verifying the Afriat Condition for each self separately. We can also apply this approach separately to each individual  $i$ , so it does not impose homogeneity of preferences across consumers. This is an important feature in light of the empirical evidence showing considerable variation across consumers in their food purchasing behavior.

### Extended Afriat Condition

The Afriat Condition provides a pass/fail test of stable rational preferences: either the data satisfy the condition or they do not. It is therefore useful to measure how close the observed behavior is to exact rationalizability in the case that one or both of the sets  $S^{ih}$  and  $S^{il}$  violates the Afriat Condition for a given individual  $i$ . For this purpose, we use (a two-selves, weighted, version of) the Afriat Index (Afriat (1973)). This index measures the fraction by which observed expenditures must be decreased for the data to be rationalized by the model. The Afriat Index is often used in revealed preference applications to assess the goodness-of-fit of a rationalizability condition such as the one in Definition 2.

We first define the Afriat Index for a given self  $j$  of individual  $i$ . To do so, we make use of the modified concept, the Extended Afriat Condition, which is defined for  $0 \leq e_i \leq 1$ .

**Definition 3** (Extended Afriat Condition). *Let  $S^{ij} = \{(\mathbf{p}_t^j; \mathbf{q}_{it}^j), t = 1, \dots, T_i\}$  be a set of observations of self  $j$ , where  $j = h, l$ . The set  $S^{ij}$  satisfies the Extended Afriat Condition if there exist numbers  $U_t^{ij}, \lambda_t^{ij} \in \mathbb{R}_{++}$  that meet, for all observations  $s$  and  $t$ , the Extended Afriat Inequalities:*

$$U_s^{ij} - U_t^{ij} \leq \lambda_t^{ij} \mathbf{p}_t^{j'} (\mathbf{q}_{is}^j - e_i * \mathbf{q}_{it}^j).$$

The exact Afriat Condition in Definition 2 corresponds to  $e_i = 1$ . Generally, lower values of  $e_i$  imply weaker rationalizability restrictions. For a given data set  $S^{ij}$ , the Afriat Index equals the largest value of  $e_i$  such that  $S^{ij}$  satisfies the

Extended Afriat Condition. It measures how close the observed behavior is to exactly rationalizable behavior. Choi et al. (2014) provide further discussion of the Afriat Index as a measure of the degree of data consistency with rationalizable behavior.<sup>21</sup>

For our two-selves model, we can define a separate Afriat Index for the healthy and the unhealthy self of each individual  $i$ . From this, we construct a weighted Afriat Index as the weighted average of these self-specific indices, by setting the weights equal to the shares in the total food expenditures (over all observations) of the respective selves. This weighting accounts for the “importance” of each self in the individual’s total food budget. The interpretation of the resulting weighted index is that the smaller the index, the more the observed food purchasing behavior deviates from the behavior induced by the two-selves model.

### 3.4 Implementation

#### Set up

We empirically implement the model described above and present evidence on how well the two-selves model rationalizes the data. Recall that the data are collected at the transaction level using handheld scanners in the home. We observe transactions at the UPC (barcode) level; there are several hundred thousand UPCs. Many of these are the same product available in different pack sizes, formats and, in some cases, flavors. There are 113,025 distinct products recorded as being purchased over the seven year period that we consider. Dealing with this number of products is not tractable. We aggregate the data to the individual-year-month-good level. We aggregate products to the level of 85 goods based on their nutritional characteristics; see Appendix A.2. We construct a price index for each good that is a weighted average of the product level region-month specific prices, where the weights reflect the quantity share of products; see Appendix A.3. To summarize, for each individual, we use data on the prices and quantities purchased of 85 goods in  $T_i \in [24, 84]$  months.

The two-selves model requires us to partition the goods into those purchased by the healthy and unhealthy self. Without further restrictions, there are  $1.9 \times 10^{25}$  possible ways to divide the 85 goods into two mutually exclusive sets. To simplify


---

<sup>21</sup>These authors refer to the Afriat Index as “Critical Cost Efficiency Index” (CCEI). In the revealed preference literature, the two denominations are used interchangeably. In defining their CCEI, Choi et al. (2014) start from the generalized axiom of revealed preference (GARP) condition for rationalizable consumer behavior, whereas we start from the Afriat Condition in Definition 2. As shown by Varian (1982), the two conditions for rationalizability are equivalent.



the problem, we make use of the NPS, which was developed by nutritionists to convert the multi-dimensional nutrient profile of a product into a single score (we discuss this in Section 2, and provide details in Appendix A.2). This allows us to rank goods by how healthy they are. There are a number of reasons why people might differ in terms of what is, and what is perceived to be, “healthy” and “unhealthy”. We therefore use the observed purchases of individuals to determine the cutoff endogenously.

Figure 3.1: *Cutoffs between healthy and less healthy foods*



Notes: For each individual the selected cutoff is the one that gives the best fit of the two-selves model with the data, i.e., the highest weighted Afriat Index. The goods are ranked from more to less healthy; the red dashed line shows the government’s cutoff (which corresponds to an NPS of 4). Drinks have a government specified cutoff of 1: we classify drinks with an NPS less than 1 as belonging to the healthy set of goods for all individuals, and drinks with an NPS greater than 1 as belonging to the unhealthy set of goods for all individuals.

We classify the 34 goods that have an average NPS less than 0 as preferred only by the healthy self, and the 24 goods with an average NPS of more than 10 (or more than 1 for drinks) as preferred only by the unhealthy self (see again Appendix A.2). We consider the 27 goods with an average NPS between 0 and 10 as uncertain and potentially belonging to either the healthy or the unhealthy set. For each individual we empirically identify the cutoff within these 27 goods as follows. We compute the Afriat Index evaluated at each possible cutoff and choose the classification that corresponds to the highest index value, i.e. that best rationalizes the observed purchase patterns. Figure 3.1 shows the share of individuals for which the good listed on the horizontal axis is the cutoff between healthy and unhealthy. The goods are ordered in decreasing healthiness from left to right. The red dashed line shows the location of the government’s cutoff of 4.

There is considerable variation across individuals in the cutoff between healthy and unhealthy goods that best rationalizes their food purchases. In Section 4.3 we show the robustness of our results to using the government’s cutoff of 4.

### **The two-selves model fits the data**

For the selected partitioning of the goods, Figure 3.2(a) shows the distribution of the weighted Afriat index for the two-selves model. Almost 20% of individuals have observed purchase behavior that is exactly rationalizable by the two-selves model. The Afriat indices for the remaining individuals are very high, indicating that only small perturbations (1% on average) of the budget are needed to ensure purchase behavior is rationalized by the two-selves model.

To compute a measure of the power of the revealed preference test we construct Afriat indices for random draws from budget sets for each individual, as described in Appendix C. We calculate the proportion of random draws that have Afriat indices greater than the true Afriat index computed with the data. This can be interpreted as the probability that the true Afriat index is below that implied by random behavior. Figure 3.2(b) shows the distribution of the probabilities – they are concentrated around zero, indicating that the test has sufficient power to discriminate between observed and random behavior.

A natural alternative to the two-selves model is a single-self model, in which an individual has a single stable utility function defined over all 85 goods. The pass rate of the two-selves model is almost twice as high as the pass rate of the single-self model, and the Afriat index for the two-selves model is strictly higher for around two thirds of the individuals.<sup>22</sup> The Kolomogorov-Smirnov test statistically confirms the difference between the distribution of the Afriat indices: the test statistic is 0.131, which is higher than the threshold for a 1% test (which is 0.04). It is worth noting that the two-selves model and the single-self model are not nested (see, for example, Section 3.B.2 in Chiappori (1988)). Our framework nests a single-self model with a strong separability structure related to the healthy and unhealthy goods. If we consider this model empirically, we find that the extra structure decreases even further the goodness-of-fit, and thus the Afriat indices, of the single-self model.


We conclude that our two-selves model provides a good fit of the data. We cannot, of course, rule out that alternative models would also explain the observed

---

<sup>22</sup>Our empirical analysis of the single-self model applies suitably adapted versions of Definitions 2 and 3 to the individual specific data  $S^i$  (comprising both healthy and unhealthy food expenditures).

food purchasing behavior of individuals. However, the appeal of our model is that it introduces time-varying preferences in a parsimonious way, while retaining an economic interpretation. In the next section, we investigate the drivers of variation in the sharing rule, or bargaining weight of the unhealthy self vis-à-vis the healthy self, over time.

Figure 3.2: *Model fit and power of the two-selves model*



*Notes: The construction of the weighted Afriat index for the two-selves model is described in Section 3.3. To compute the power of the test we construct Afriat indices for the random draws for each individual, as described in Appendix C. For each individual we calculate the proportion of random draws that have Afriat indices greater than the Afriat index for the observed value. The figure shows the distribution across individuals.*

## 4 Quantifying the determinants of variation in food choices

In this section we quantify the extent to which standard economic variables (prices, budgets), along with other advertising and weather, drive variation in diet quality through their influence on the sharing rule. We describe differences in the mean sharing rule across individuals, and we recover residual variation in the sharing rule (after conditioning on these observables) in order to study within-individual variation in diet quality. We show that variation is higher for poorer and younger households. We analyze the behavior of our full sample of individuals and in Section 4.3 we show robustness of our results to restricting our analysis to the subsample of individuals for whom the two-selves model outperforms the single-self model, as well as a number of other potential concerns.

## 4.1 Estimating the sharing rule

In Section 2 we describe variation in the share of calories from healthy foods. However, in order to understand the determinants of the variation in diet quality, we use the model set out in Section 3. The model generates predictions in terms of expenditure shares (rather than calorie or quantity shares), and the results from Chiappori (1988) allow us to interpret the share of spending on healthy food as a monotonic transformation of the bargaining weight of the healthy self.

The sharing rule depends on the vector of prices for healthy and unhealthy foods,  $\mathbf{p}_{rt}^h$  and  $\mathbf{p}_{rt}^l$ , total food spending,  $x_{it}$ , and a vector of other factors,  $\mathbf{z}_{it}$ , that might affect the bargaining weight:

$$\eta_{it} = \eta_i(\mathbf{p}_{rt}^h, \mathbf{p}_{rt}^l, x_{it}, \mathbf{z}_{it}).$$

Notice that here we make explicit the dependency of prices on region,  $r$ , which takes values: {north, central, south} of the UK.

The function  $\eta_i$  is a consumer-specific nonparametric function and its arguments include 85 prices, the food budget and the vector  $\mathbf{z}_{it}$ , which includes both observable and unobservable variables that affect the sharing rule. Two potentially important observable factors are advertising and weather. Firms may advertise unhealthy products more intensely in some time periods than in others. Weather might also affect food choices over time, for example, if people are more likely to purchase soft drinks in hotter weather.

In order to recover the variation in  $\eta_i$  that is driven by variation in prices and food budgets, and variation that is due to other factors, we impose the following restrictions. First, we separate the vector  $\mathbf{z}_{it}$  into a component with observable variables,  $\tilde{\mathbf{z}}_{it}$ , and an unobservable scalar,  $\epsilon_{it}$ , which we assume is additively separable. The observable variables include monthly advertising expenditure on unhealthy foods (confectionery, snacks, soft drinks, prepared and convenience foods), which we obtain from the AC Nielsen Advertising Digest – a data set that records all monthly food and drink advertising expenditure at the brand level; see Appendix A.4. The vector  $\tilde{\mathbf{z}}_{it}$  also includes three variables that capture the weather conditions in individual  $i$ 's local area at time period  $t$  – the minimum temperature in that month, the maximum temperature, and the amount of rain; see Appendix A.5 for details.

Second, we assume that the two price vectors, which together contain 85 separate price series, can be approximated by two price indices: one for healthy foods (those with an NPS below the individual's cutoff),  $\Pi_{irt}^h$ ; and one for unhealthy

foods (those with an NPS above the individual’s cutoff),  $\Pi_{irt}^l$ . These price indices are weighted averages of the prices for the goods which comprise each set. The weights are equal to the individual’s mean quantity share of each good (out of the quantity purchased of each of the two sets of foods – healthy, unhealthy); see Appendix A.3 for details. In Section 4.3, we show robustness of our results to approximating the 85 prices series with a larger set of price indices.

Under these assumptions we get:

$$\eta_{it} = g_i(\Pi_{irt}^h, \Pi_{irt}^l, x_{it}, \tilde{\mathbf{z}}_{it}) + \epsilon_{it}.$$

We approximate  $g_i$  with an expression that is linear in the log of the two price indices, the log of a deflated expenditure term, and the weather and advertising variables, and that has individual specific coefficients. The approximation of the sharing rule that we estimate is:

$$\eta_{it} = \alpha_i + \beta_i \ln \left( \widetilde{\frac{\Pi_{irt}^l}{\Pi_{irt}^h}} \right) + \gamma_i \left( \widetilde{\ln rx_{it}} \right) + \boldsymbol{\theta}'_i \tilde{\mathbf{z}}_{it} + \epsilon_{it}, \quad (4.1)$$

where  $\widetilde{\ln(\cdot)}$  denotes that we normalize each variable by subtracting the individual specific mean,<sup>23</sup> and:

$$\widetilde{\ln rx_{it}} \equiv \ln x_{it} - [\bar{\eta}_i \ln(\Pi_{it}^h) + (1 - \bar{\eta}_i) \ln(\Pi_{it}^l)]$$

is the log of real food expenditure.<sup>24</sup> Deflating by the price index for food ensures that the sharing rule is homogeneous of degree zero in prices and expenditure. Normalizing the variables allows us to interpret  $\alpha_i$  as the individual’s preference for healthy food when she faces her average food budget, average relative prices, and average other observed factors (advertising and weather conditions).

We estimate equation (4.1) for each individual using OLS. In Section 4.3, we show that the results are robust to instrumenting log real food expenditure. The parameter estimate  $\hat{\alpha}_i$  is consumer  $i$ ’s “mean” sharing rule – i.e. her sharing rule evaluated at her mean real food expenditure, mean relative prices and mean values for the other observed factors (advertising and weather conditions). It therefore captures the consumer’s average propensity to choose healthy relative to unhealthy foods (i.e. the average share of their budget allocated to healthy foods).

---

<sup>23</sup>For instance,  $\ln \left( \widetilde{\frac{\Pi_{irt}^l}{\Pi_{irt}^h}} \right) = \ln \left( \frac{\Pi_{irt}^l}{\Pi_{irt}^h} \right) - \frac{1}{T_i} \sum_t \ln \left( \frac{\Pi_{irt}^l}{\Pi_{irt}^h} \right)$ .

<sup>24</sup>This is the log of nominal food expenditure deflated with a food price index (the weights in the index,  $(\bar{\eta}_i, 1 - \bar{\eta}_i)$  are the individual’s mean share on healthy and unhealthy foods).

The median value of  $\hat{\alpha}_i$  is 52%, the 25th percentile is 41% and the 75th percentile is 62%.

In Appendix E we summarize differences in the mean estimated sharing rule across demographic groups. Older and higher income individuals spend a greater proportion of their food budget on healthy foods. This lines up with the literature, in which socioeconomic differences in health are well documented (e.g. Cutler and Lleras-Muney (2010)). On average, men have lower mean values of the sharing rule than women – they allocate 1.9 percentage points less of their food budget to healthy foods than women. People not in work tend to spend a lower share on healthy foods than people in work – with an average sharing rule that allocates 2.3 percentage points less to healthy foods, and smokers allocate 3.5 percentage points less to healthy food than non-smokers. We show that this pattern also holds if we use an alternative measure of socioeconomic status, based on occupation. We also show that overweight and obese individuals spend less of their budget on healthy foods than normal weight individuals.

Kantar Wordpanel asks participants a selection of questions to gauge their attitude to a variety of lifestyle factors, including preferences for healthy and processed food. (see Appendix A.6 for details). Table 4.1 shows the correlation of two of these stated preferences with  $\hat{\alpha}_i$ . Individuals’ estimated preferences for healthy food accord with their stated preferences: the mean sharing rule,  $\hat{\alpha}_i$ , is increasing in stated preferences for healthy food and declining in stated preferences for processed food. In Appendix D we show that these correlations are robust to conditioning on gender, age, socioeconomic status, employment status, BMI, whether the individual is a smoker or vegetarian, and income.

Table 4.1: *Variation in mean sharing rule by stated preferences*

(1)	(2)	(3)	(4)
	Mean	Difference	95% CI for diff.
Above median preferences for healthy food	52.7		
Below median preferences for healthy food	48.2	-4.5	[-5.4, -3.5]
Above median preferences for processed food	48.9		
Below median preferences for processed food	52.3	3.4	[2.5, 4.4]

*Notes: The numbers in column (3) are the difference in means from the first row in each group. Confidence intervals reflect uncertainty arising from the sample of individuals, but not over the time series variation for each individual.*

## 4.2 Within-individual variation in the sharing rule

To show the importance of observables in explaining within-individual variation in the sharing rule, we compute several partial  $R^2$  indices. One of these measures the variation in the sharing rule that is explained by variation in the economic environment (relative prices and food budgets). The others measure the variation in the sharing rule that is explained by variation in the weather or advertising. The total  $R^2$  measures the amount of the variation in the sharing rule that is explained by variation in the economic environment (relative prices and food budgets), and other factors such as advertising and weather. Table 4.2 summarizes these partial  $R^2$ s.

Table 4.2: *Partial  $R^2$ s from regression of  $\eta_{it}$*

Regressors:	Mean	25th	50th	75th
Price and budget	0.12	0.03	0.09	0.17
Advertising	0.05	0.01	0.02	0.07
Weather	0.10	0.04	0.07	0.13
Price, budget, advertising, weather	0.23	0.13	0.21	0.30


*Notes: Table reports moments of the distribution of partial  $R^2$  of 3645 separate individual level regressions of 4.1.*

The mean partial  $R^2$  with respect to prices and food budgets across individuals is 0.12 – on average, around 12% of the within-individual variation in the sharing rule is explained by variation in price and food budgets. The mean total  $R^2$  across individuals is 0.23 – on average, around 23% of within-individual variation in the sharing rule is explained by prices, food budgets, advertising and weather. This illustrates that individuals’ response to variation in these variables does explain a portion of the variation in the sharing rule, but a considerable fraction of the fluctuations in the share of spending on healthy foods over the year remains unexplained.

### Variation over time

In Figure 4.1 we summarize the average (across individuals) variation in  $\hat{\epsilon}_{it}$  (i.e. the residual sharing rule) over the year. We plot the deviation in the mean relative to January (pooled over years). We also show the mean (relative to January) of the observed share of spending on healthy foods,  $\eta_{it}$ , over the year, which captures the average variation in the share of spending allocated to the healthy self that is driven both by observable and unobservable variables.

Figure 4.1: *Deviations from the mean sharing rule over the calendar year*



Notes: Black dots show the mean sharing rule,  $\eta_{it}$ , in each calendar month, pooled across years. Red squares show the mean residual sharing rule,  $\hat{\epsilon}_{it}$ , which removes the effect of variation in prices, food budgets, advertising, and weather, in each calendar month, pooled across years. Both are measured relative to the mean values in January. 95% confidence intervals are shown. Confidence intervals reflect uncertainty arising from the sample of individuals, but not over the time series variation for each individual.

Figure 4.1 shows that the observed share of expenditure on healthy goods declines from January to March, increases from March to July, and then deteriorates from July onward. Once we control for prices, food budgets, advertising and weather variation, the residual sharing rule does not rebound in the middle of the year. This in part reflects that people tend to switch to eating more healthily when the weather improves, and therefore controlling for this switching means that the unexplained part of the sharing rule becomes flatter through the middle of the year. However, there is still a clear downward trend over the calendar year.

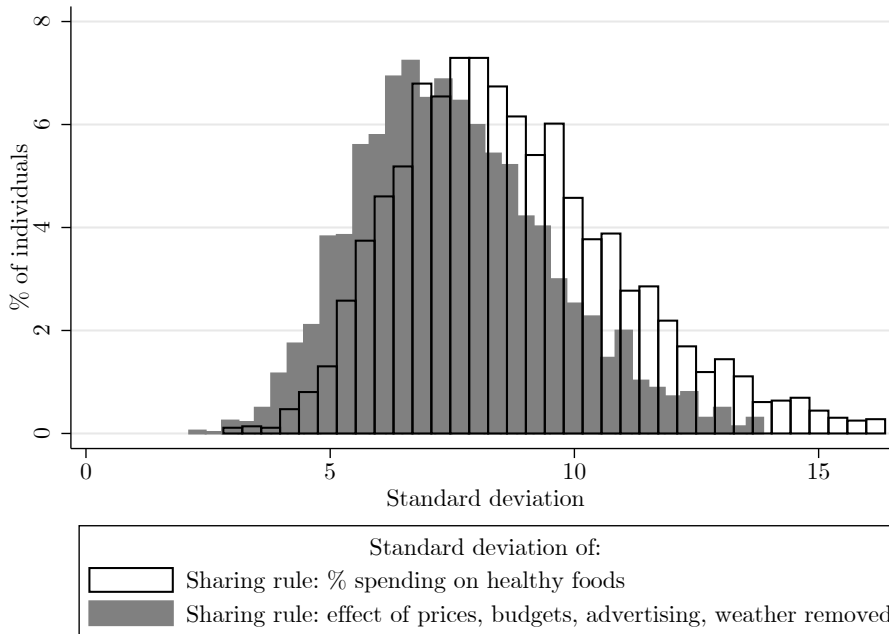
Figure 4.1 shows that, on average, preferences of the unhealthy self become more important in determining food choice over the calendar year. However, it masks a great deal of variation across individuals. Although the majority of people have diets that deteriorate over the calendar year, the magnitude of this decline, and when the decline starts to occur, varies across individuals.

We construct the standard deviation of the residual sharing rule for each individual  $i$  over  $t$ :  $\hat{\sigma}_i = \text{sd}(\hat{\epsilon}_{it})$ ; this captures variation over time in the sharing rule (and hence the bargaining weight) around the mean that is not driven by that individual's responses to changes in food prices, food budgets, weather, or advertising. We also construct the standard deviation of the sharing rule for each individual  $i$  over  $t$ :  $\tilde{\sigma}_i = \text{sd}(\eta_{it})$ , which measures the total variation in the in-



dividual’s sharing rule, driven by both unexplained changes and changes in the economic environment, weather, and advertising.

Figure 4.2: *Standard deviation of the sharing rule and the residual sharing rule*



*Notes: The white bars show the distribution of individuals’ standard deviation in the sharing rule,  $\tilde{\sigma}_i$ . The grey bars show the distribution of individuals’ standard deviation in the residual sharing rule,  $\hat{\sigma}_i$ , which takes out the responses to changes in prices, food budgets, advertising and weather.*

Figure 4.2 compares the distributions of  $\tilde{\sigma}_i$  and  $\hat{\sigma}_i$  across individuals. The leftward shift of the distribution shows the extent to which accounting for individuals’ responses to variation in prices, food budgets, advertising and weather explains variation in their spending on healthy food over time. The average standard deviation of the sharing rule,  $\tilde{\sigma}_i$ , is 8.8 percentage points, but of the residual sharing rule,  $\hat{\sigma}_i$ , is 7.6 percentage points. This means that a substantial proportion of the within-individual variation in the sharing rule is unexplained by prices, budgets, advertising or weather, even allowing for individuals’ heterogeneous responses to changes in these variables.


### Variation with demographics

We consider how the within-individual variation in the sharing rule differs with age and income. We focus on these demographics because they are the focus of policy intervention. For example, the UK government restricts advertising of

unhealthy foods targeted at young people, and the motivation for the Soft Drinks Industry Levy, a tax levied on drinks products with higher sugar content, is to reduce sugar consumption, and ultimately obesity, in young people. In the US many policies, such as Supplemental Nutrition Assistance Program, are aimed at improving the diets of low income households. Policies that aim to improve the nutrition of lower income and younger people have been shown to have long run impacts (e.g. Hoynes et al. (2016)). In addition, Mani et al. (2013) show that poverty directly impedes cognitive function, thereby leading to poor choices.

Figure 4.3: *Variation in  $\hat{\sigma}_i$  with age*

(a) Standard deviation of residual sharing rule (b) Effect of responses to changes in prices, food expenditure, advertising and weather



Notes: Age is the individual's median age while in the sample. The left panel shows the mean standard deviation in the residual sharing rule across individuals at each age. For each individual we calculate the difference in the standard deviation in the sharing rule and the standard deviation in the residual sharing rule,  $\tilde{\sigma}_i - \hat{\sigma}_i$ ; the right panel shows the conditional mean of this difference across ages. 95% confidence intervals are shown. Confidence intervals reflect uncertainty arising from the sample of individuals, but not over the time series variation for each individual.


Figure 4.3(a) shows how the standard deviation in the residual sharing rule,  $\hat{\sigma}_i$ , varies across age groups. Panel (b) shows how the difference in variation in the total and residual sharing rule,  $(\tilde{\sigma}_i - \hat{\sigma}_i)$ , varies with age. Together these graphs show two things. First, there is an age gradient. Young people, on average, have more unexplained variation in the bargaining weight of their healthy self. Second, failing to account for the effects of the economic environment, advertising and weather conditions would lead to an over-statement of this gradient. For individuals aged below 40, variation in these observables is, on average, responsible for 1.45 percentage points of the standard deviation in their sharing rule; for individuals aged over 70, 1.05 percentage points is explained by responses to these variables. One possible explanation for why the quality of older individuals' diets does not react as strongly to variation in prices and incomes is because they have

more time to shop around and more scope to engage in home production due to a lower opportunity cost of time relative to younger individuals (Aguiar and Hurst (2007, 2013)).

Figure 4.4 shows a similar pattern for the relationship between  $\hat{\sigma}_i$ ,  $\tilde{\sigma}_i - \hat{\sigma}_i$  and the income distribution. People with low income exhibit more variation in their sharing rule. This is partly driven by their responses to changes in observable factors – see panel (b). However, it is also the case that lower income individuals have more variation in their residual sharing rule. This difference is meaningful: individuals in the bottom quintile have a standard deviation in their residual sharing rule that is more than 2 percentage points larger than individuals in the top quintile. This difference is comparable to the cross-sectional difference in the average sharing rule between the bottom and top quintiles, which is roughly 2 percentage points.

Figure 4.4: *Variation in  $\hat{\sigma}_i$  with income*

(a) Standard deviation of residual sharing rule (b) Effect of responses to changes in prices, food expenditure, advertising and weather



*Notes: For all individuals we calculate the mean total spending on fast moving consumer goods (food, alcohol, household supplies, toiletries etc.) across the period they are in the sample, which we use as a proxy for income. In the Appendix A.7 we provide further evidence on our measure of income. Percentiles are based on the distribution of this variable across individuals. The left panel shows the conditional mean standard deviation in the residual sharing rule across income percentiles. For each individual we calculate the difference in the standard deviation in the sharing rule and the standard deviation in the residual sharing rule,  $\tilde{\sigma}_i - \hat{\sigma}_i$ ; the right panel shows the conditional mean of this difference across income percentiles. 95% confidence intervals are shown. Confidence intervals reflect uncertainty arising from the sample of individuals, but not over the time series variation for each individual.*

Failure to account for individuals' responses to changes in prices, budgets, advertising and weather leads to an overestimate of the gradient with age and income. This is primarily due to differences in the fluctuations in individuals' food budgets and how they respond to them. Younger and lower income individuals have larger fluctuations in their food budgets: the standard deviation of logged

real food expenditure is 30% higher for individuals aged under 40 compared with individuals aged over 70, and it is 70% higher for individuals in the bottom income quintile compared with individuals in the top. Across the income gradient, this is amplified by the fact that real food expenditure statistically significantly affects the sharing rule for more individuals in the bottom quintile (32%) than the top quintile (18%).

### 4.3 Robustness

In this section we show robustness to our specification choices. Table 4.3 shows that for a range of specifications (details below) there remains substantial within-individual time-series fluctuations in diet quality, and that these fluctuations are larger for younger and lower income individuals.

Table 4.3: *Robustness*

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Specification	Mean $\hat{\sigma}_i$	$\hat{\sigma}_i$ by age group			$\hat{\sigma}_i$ by income quint.		
		<40	70+	Ratio	1st	5th	Ratio
Baseline	7.62	8.43	7.36	1.15	9.06	6.53	1.39
Controls for four price indices	7.22	7.96	7.01	1.14	8.60	6.15	1.40
Instrumenting food expenditure	7.77	8.55	7.52	1.14	9.23	6.66	1.39
Exogenous cutoff	7.90	8.68	7.60	1.14	9.45	6.69	1.41
Individuals: two-self model is better	7.60	8.40	7.29	1.15	9.11	6.51	1.40

*Notes:* Column (1) lists the different robustness specifications for which we estimate equation (4.1) and construct  $\hat{\sigma}_i$  as the individual's standard deviation in the residual from each regression. Column (2) shows the mean  $\hat{\sigma}_i$  across individuals. Columns (3) and (4) show the mean  $\hat{\sigma}_i$  for individuals aged under 40 and over 70, respectively, and column (5) shows the ratio of the two means. Columns (6) and (7) show the mean  $\hat{\sigma}_i$  for individuals in the lowest and highest income quintiles, respectively, and column (8) shows the ratio of the two means.

#### Four price indices

One possible concern is that, in our estimation of the sharing rule, we use two price indices to capture variation in the prices of healthy and unhealthy foods over time. It is not possible to include the full vector of 85 price indices in equation (4.1) while allowing for individual heterogeneity in price effects. However, we do consider a more flexible specification in which we assume that the movements in the price vector can be approximated by four price indices: one for very healthy foods, one for healthy foods, one for unhealthy foods, and one for very unhealthy foods.<sup>25</sup>

<sup>25</sup>Very healthy foods are those with an NPS less than 0,  $\Pi_{irt}^{vh}$ ; healthy foods are those with an NPS between 0 and 10, but classed by the individual as healthy,  $\Pi_{irt}^h$ ; unhealthy foods are

Including additional price indices does reduce the variation in the residual sharing rule, but only slightly – the average standard deviation falls from 7.6 percentage points to 7.2 percentage points, and Table 4.3 shows that the correlations with age and income are robust to including more controls for price variation.

### **Instrumenting food expenditure**

In Section 4.1 we partial out the effect of variation in food prices and budgets. However, it may be the case that unobserved factors that increase the bargaining weight of either of the two selves also lead to changes in total food budgets. To deal with this possibility, which could affect our interpretation of the residual sharing rule, we instrument for the expenditure term in equation (4.1) using variables that are likely to drive total spending on food, but are less plausible shifters of the bargaining weight. Our instrument set includes a set of prices from the consumer price index (CPI) and individual monthly spending on non-food items (cleaning products, toiletries, cosmetics).<sup>26</sup> We expect non-food spending and the relative price of food and non-food to influence individuals' allocation of their total budget between food and other commodities, but not to be correlated with the bargaining weights of the two selves.<sup>27</sup> The results are very similar across the OLS and IV specifications. Slightly more of the observed variation in the sharing rule is explained by prices and budgets in the OLS specification than in the IV specification (the mean standard deviation in the residual sharing rule is 7.6 percentage points in the OLS specification, compared with 7.8 percentage points in the IV specification). However, the relationships with age and income are qualitatively similar; see Table 4.3.

### **Exogenous cutoff between healthy and unhealthy foods**

In our main specification we allow the cutoff between healthy and unhealthy foods to be determined endogenously by individuals' food purchasing patterns. This

---

those with an NPS between 0 and 10 and classed by the individual as unhealthy,  $\Pi_{irt}^l$ ; and very unhealthy foods are those with an NPS greater than 10,  $\Pi_{irt}^v$ . These price indices are weighted averages of the prices for the goods which comprise each set, and are constructed in an analogous way to the two price indices described in Appendix A.3.

<sup>26</sup>The prices consist of the all-items CPI, which captures the general price level in the economy, and the CPI component indices for the set of non-housing goods (food, alcohol and tobacco, furniture and equipment, health care, transport, communications, recreation, education, restaurants and hotels, other goods and services).

<sup>27</sup>Pooling in the first stage across individuals results in an F-statistic for a test of the joint significance of the instruments of over 700. Estimating the first stage individual-by-individual results in lower F-statistics, and, for some individuals, weak instruments.

allows for the fact that there may be heterogeneity in the perception of what is “healthy”. An alternative is to use the government’s cutoff: foods with an NPS below 4 (below 1 for drinks) are classified as healthy, while those above are classified as unhealthy. There is slightly more variation in the residual sharing rule when the exogenous cutoff is used (7.9 percentage points compared with 7.6 percentage points in the baseline), but this is not statistically different. Table 4.3 shows that the age and income gradients are robust to using the exogenous cutoff.

### **Individuals for whom the two-self model fits better**

We check the consistency of individuals’ behavior with both the two-selves model and a natural alternative, the single-self model. For two-thirds of individuals the two-selves model is better able to rationalize their behavior than the single-self model. The two-selves model performs better (relative to the single-self model) for younger individuals; there is no significant relationship with income. Table 4.3 shows that when we look just at this subset of individuals with a better fit of the two-selves model, we also continue to see larger fluctuations in the sharing rule for younger and lower income individuals.

## **5 Discussion**

In this final section, we discuss various alternative explanations for the residual variation in the sharing rule.

### **5.1 Holidays and birthdays**


Christmas, Easter and birthdays are times at which individuals’ diet quality deteriorates and then again improves (see also Dai et al. (2014)). We estimate a variant of the sharing rule regression (equation (4.1)) including dummy variables for Christmas, Easter and the individual’s birthday. These celebrations explain a small fraction of the residual sharing rule – the average standard deviation in the residual sharing rule falls from 7.6 percentage points to 7.2 percentage points when we control for holidays and birthdays; there remains substantial unexplained variation in diet quality over the year. We find that the correlations with age and income are robust to taking out the variation due to holidays and birthdays.

## 5.2 Food outside the home

Our analysis focuses on grocery purchases, which constitute 85% of calories purchased (calculated using the Living Costs and Food Survey, described below), and are therefore key to understanding the nutritional implications of people’s food choices. However, there may be concern that variation in the nutritional quality of people’s at home food purchases is partly offset by changes in their purchases made for consumption outside of the home.

We show that the aggregate patterns presented in Section 2 are neither an artifact of the Kantar Worldpanel data, nor are affected by food consumed outside the home. We use data from the Living Costs and Food Survey (LCFS), which is the main UK expenditure survey and similar to the CEX in the US. It is a repeated cross-sectional two-week expenditure diary and records information on the quantities of foods consumed in and outside of the home. The fact that it is a repeated cross-section means that we cannot look at within-individual time series variation, but we can check whether the aggregate decline in diet quality over the calendar year holds when we look at food in and out of the home. Figure 5.1 shows that there is clear decline in the share of calories from healthy foods from January through to December. In addition, we also verified that the share of calories purchased from food out is stable over the calendar year, which suggests that substitution to food out over the year is unlikely to be confounding our results.

Figure 5.1: *Calorie share of healthy food over the calendar year for food in and out, LCFS 2005-11*



*Notes: We use data on the food and drink purchases (consumed in and out of the home) made by 42,062 households over 2005-11 recorded in the Living Costs and Food Survey. We construct the share of calories purchased from healthy foods (foods with an NPS below 4, and drinks with an NPS below 1); the line shows the mean across households and years.*

### 5.3 Does within-individual variation reflect impulsiveness?

Our results show that within-individual fluctuations in diet quality over time are considerable, and are only partly explained by observable factors such as prices, food budgets, advertising and weather changes. The two-selves model provides a natural way to interpret this variation as the outcome of a process in which two selves each have an influence on decision making.

We investigate the relationship between the standard deviations of individuals' residual sharing rule,  $\hat{\sigma}_i$  (which captures the magnitude of an individual's fluctuation in the share of food allocated to the healthy self that are not explained by prices, food budgets, advertising or weather), and individuals' stated attitudes. We use responses to the stated preference and attitude questions from the Kantar Worldpanel on: the tendency to buy things on offer, shopping commitment, and stated impulsiveness (see Appendix A.6 for details). Table 5.1 shows the correlation between these variables and  $\hat{\sigma}_i$ . Individuals who have higher stated impulsiveness (i.e. state that they spend money without thinking, or spend more on their credit card than they should) have larger fluctuations in their residual sharing rule over the year. Individuals who state that they commit to buying the same brands experience smaller deviations, while those who say that they are more spontaneous and, for example, are influenced by promotions, experience larger deviations in their residual sharing rule. In Appendix D we show that these correlations are robust to conditioning on gender, age, socioeconomic status, employment status, BMI, whether the individual is a smoker or vegetarian and income. We also show that individuals with lower stated impulsiveness have a lower mean sharing rule i.e. spending less on healthy foods, on average.

Lack of self-control can manifest itself in many ways; for example, it may lead people to consistently eat unhealthy foods, or it may lead to larger deviations from average behaviour as individuals try to commit themselves to healthy diets, and then succumb to temptation. We cannot separately identify whether average differences in food choices across individuals reflect lack of self-control or preferences heterogeneity. However, by looking within individuals across time, we can measure the extent to which people deviate from their long-run behaviour. Thus, one interpretation of the bargaining process between the selves is that it captures an individual's ability to act in accordance with her long-run preferences, and that larger fluctuations reflect, at least in part, a lack of self control.

With this interpretation in mind, we return to the demographic correlations described in Section 4. Figure 4.3 indicates that younger people have larger fluctu-



tuations in their residual sharing rule than older people; this is consistent with the findings in Ameriks et al. (2007) that show that the young suffer more from self-control problems than older people. Figure 4.4 accords with evidence that low income people are more susceptible to self-control problems. Indeed, a number of papers point to low income being causally related to self-control problems. For example, Haushofer and Fehr (2014) and Mani et al. (2013) suggest that the stress and cognitive loads of being in poverty means people are more likely to make unwise decisions and underweight the future. Bernheim et al. (2015) argue that poverty can perpetuate itself by undermining the capacity for self-control: low initial wealth precludes self-control, and hence asset accumulation, creating a poverty trap. Banerjee and Mullainathan (2010) take an alternative approach by assuming that “temptation goods” are inferior goods, which leads to a similar conclusion that self-control problems give rise to asset traps. Mastrobuoni and Weinberg (2009) find that retired individuals who have accumulated lower savings over their life cycle are less likely to smooth their food consumption over their Social Security pay periods.

Table 5.1: *Variation in  $\hat{\sigma}_i$  by stated preferences*

	(1) Mean	(2) Difference	(3) 95% CI for diff.
Above median tendency to buy on promotion	7.86		
Below median tendency to buy on promotion	7.39	-0.46	[-0.61, -0.32]
Above median shopping commitment	7.50		
Below median shopping commitment	7.78	0.27	[0.13, 0.41]
Above median stated impulsiveness	7.73		
Below median stated impulsiveness	7.41	-0.32	[-0.52, -0.12]

*Notes: The numbers in column (3) are the difference in means from the first row in each group. Confidence intervals for  $\hat{\sigma}_i$  reflect uncertainty arising from the sample of individuals, but not over the time series variation for each individual.*

## 6 Summary and concluding comments

We show that within-individual across time fluctuations in food purchasing are large, and have first-order implications for diet quality. The variation follows a clear within-year pattern, where diet deteriorates over the course of the calendar year, and resets each January. We provide empirical evidence using observational data on the entire grocery basket that is consistent with other findings in the

behavioral literature. We also document substantial heterogeneity across people, both in their average purchasing patterns and the extent to which their food choices vary over time.

We use nonparametric revealed preference methods to show that the data can be rationalized by a model in which food choices are a compromise between a healthy and an unhealthy self. This model outperforms the standard single-self model, and provides an economic interpretation to variation in the data. Specifically, the sharing rule, or share of spending on healthy foods, can be interpreted as the bargaining weight of the healthy self.

We show that the fluctuations in the sharing rule across time are not fully explained by responses to the economic environment or other factors, such as advertising and the weather. We recover the “residual” sharing rule, which removes individuals’ responses to changes in these other factors. We explore whether the within-individual variation in the residual sharing rule is indicative of self-control problems. Using data on individuals’ stated impulsiveness, we provide suggestive evidence that sharing rule fluctuations are larger for individuals with higher impulsiveness. We also find that variation in the residual sharing rule is larger for younger and lower income individuals.

Our results are informative for governments struggling to deal with diet related disease. There remains a challenge in how to design policy that encourages individuals to behave as they do at their healthiest times. One potentially fruitful avenue for further research is the possible welfare enhancing role of commitment devices, which can provide a way for people to achieve their long-run preferred outcome and avoid self-control problems (see Bryan et al. (2010) for a survey). It is unclear whether the market will provide such commitment devices (Gottlieb (2008)), and government policy can compensate for the inability of the market to efficiently provide commitment devices. For example, “sin taxes”, which raise the price of tempting goods, can increase welfare if enough people are time inconsistent (O’Donoghue and Rabin (2003)). Determining the optimal sin tax level entails trading off the welfare gain to those suffering from internalities against the welfare loss of those that do not suffer from internalities but who face higher prices (Griffith et al. (2017)). Our results suggest this trade-off may also apply within individual over time.

## References

- Afriat, S. N. (1967). The construction of utility functions from expenditure data. *International Economic Review* 8(1), 67–77.
- Afriat, S. N. (1973). On a System of Inequalities in Demand Analysis: An Extension of the Classical Method. *International Economic Review* 14(2), 460–472.
- Aguiar, M. and E. Hurst (2007). Life-cycle prices and production. *American Economic Review* 97(5), 1533–1559.
- Aguiar, M. and E. Hurst (2013). Deconstructing Life Cycle Expenditure. *Journal of Political Economy* 121(3), 437–492.
- Allcott, H., R. Diamond, and J.-P. Dubé (2017, December). The geography of poverty and nutrition: Food deserts and food choices across the United States. *NBER Working Paper 24094*.
- Ameriks, J., A. Caplin, J. Leahy, and T. Tyler (2007). Measuring Self-Control Problems. *American Economic Review* 97(3).
- Arambepola, C., P. Scarborough, and M. Rayner (2008). Validating a nutrient profile model. *Public Health Nutrition* 11(04), 371–378.
- Ariely, D. and K. Wertenbroch (2002, May). Procrastination, Deadlines, and Performance: Self-Control by Precommitment. *Psychological Science* 13(3), 219–224.
- Ascherio, A., M. Katan, J. Stampfer, and W. Willett (1999). Trans fatty acids and coronary heart disease. *New England Journal of Medicine* 340, 1–11.
- Ashraf, N., D. Karlan, and W. Yin (2006, May). Tying Odysseus to the Mast: Evidence From a Commitment Savings Product in the Philippines. *The Quarterly Journal of Economics* 121(2), 635–672.
- Banerjee, A. and S. Mullainathan (2010). The Shape of Temptation: Implications for the Economic Lives of the Poor. Working Paper 15973, National Bureau of Economic Research.
- Baum, C. L. and C. J. Ruhm (2009). Age, socioeconomic status and obesity growth. *Journal of Health Economics* 28(3), 635–648.
- Bénabou, R. and M. Pycia (2002, November). Dynamic inconsistency and self-control: A planner–doer interpretation. *Economics Letters* 77(3), 419–424.
- Bernheim, B. D., D. Ray, and Ş. Yeltekin (2015). Poverty and Self-Control. *Econometrica* 83(5), 1877–1911.
- Bhattacharya, J. and N. Sood (2011). Who pays for obesity? *Journal of Economic Perspectives* 25(1), 139–158.
- Bleich, S., D. Cutler, C. Murray, and A. Adams (2008). Why is the developed world so obese? *Annual Review of Public Health* 29, 273–295.
- Broda, C. and D. E. Weinstein (2010, June). Product Creation and Destruction: Evidence and Price Implications. *American Economic Review* 100(3), 691–723.
- Browning, M. and P.-A. Chiappori (1998). Efficient intra-household allocations: A general characterization and empirical tests. *Econometrica* 66(6), 1241–1278.
- Browning, M., P.-A. Chiappori, and A. Lewbel (2013). Estimating Consumption Economies of Scale, Adult Equivalence Scales, and Household Bargaining Power. *Review of Economic Studies* 80(4), 1267–1303.

- Bryan, G., D. Karlan, and S. Nelson (2010). Commitment Devices. *Annual Review of Economics* 2(1), 671–698.
- Buccioli, A. (2012). Measuring self-control problems: A structural estimation. *Journal of the European Economic Association* 10(5), 1084–1115.
- Cherchye, L., B. De Rock, and F. Vermeulen (2007). The collective model of household consumption: A nonparametric characterization. *Econometrica* 75(2), 553–574.
- Cherchye, L., B. De Rock, and F. Vermeulen (2011). The revealed preference approach to collective consumption behaviour: Testing and sharing rule recovery. *Review of Economic Studies* 78(1), 176–198.
- Cherchye, L., T. Demuynck, and B. De Rock (2011). Revealed Preference Analysis of Non-Cooperative Household Consumption. *Economic Journal* 121(555), 1073–1096.
- Chiappori, P.-A. (1988). Rational household labor supply. *Econometrica* 56(1), 63–90.
- Chiappori, P.-A. (1992). Collective labor supply and welfare. *Journal of Political Economy* 100(3), 437–467.
- Chiappori, P.-A. and I. Ekeland (2009). The microeconomics of efficient group behavior: Identification. *Econometrica* 77(3), 763–799.
- Choi, S., S. Kariv, W. Müller, and D. Silverman (2014). Who Is (More) Rational? *American Economic Review* 104(6), 1518–1550.
- Currie, J. (2009). Healthy, Wealthy, and Wise: Socioeconomic Status, Poor Health in Childhood, and Human Capital Development. *Journal of Economic Literature* 47(1), 87–122.
- Currie, J., S. DellaVigna, E. Moretti, and V. Pathania (2010, August). The Effect of Fast Food Restaurants on Obesity and Weight Gain. *American Economic Journal: Economic Policy* 2(3), 32–63.
- Cutler, D., E. Glaeser, and J. Shapiro (2003). Why Have Americans Become More Obese? *Journal of Economic Perspectives* 17(3), 93–118.
- Cutler, D. M. and A. Lleras-Muney (2010). Understanding differences in health behaviors by education. *Journal of Health Economics* 29(1), 1–28.
- Dai, H., K. L. Milkman, and J. Riis (2014, October). The Fresh Start Effect: Temporal Landmarks Motivate Aspirational Behavior. *Management Science* 60(10), 2563–2582.
- Dai, H., K. L. Milkman, and J. Riis (2015). Put your imperfections behind you: Temporal landmarks spur goal initiation when they signal new beginnings. *Psychological science* 26(12), 1927–1936.
- DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic Literature* 47(2), 315–372.
- DellaVigna, S. and U. Malmendier (2006). Paying not to go to the gym. *American Economic Review* 96(3), 694–719.
- Dubois, P., R. Griffith, and A. Nevo (2014, March). Do Prices and Attributes Explain International Differences in Food Purchases? *American Economic Review* 104(3), 832–867.

- Dunbar, G. R., A. Lewbel, and K. Pendakur (2013). Children’s Resources in Collective Households: Identification, Estimation, and an Application to Child Poverty in Malawi. *American Economic Review* 103(1), 438–471.
- Finkelstein, A., E. F. P. Luttmer, and M. J. Notowidigdo (2013). What good is wealth without health? The effect of health on the marginal utility of consumption. *Journal of the European Economic Association* 11(SUPPL. 1), 221–258.
- Finkelstein, E. A., C. J. Ruhm, and K. M. Kosa (2005). Economic Causes and Consequences of Obesity. *Annual Review of Public Health* 26(1), 239–257.
- Finkelstein, E. A. and L. Zuckerman (2008). *The Fattening of America: How The Economy Makes Us Fat, If It Matters and What To Do About It*. New Jersey, USA: Wiley.
- Fudenberg, D. and D. K. Levine (2006). A dual-self model of impulse control. *American Economic Review* 96(5), 1449–1476.
- Gilbert, D. T., M. J. Gill, and T. D. Wilson (2002). The Future Is Now: Temporal Correction in Affective Forecasting. *Organizational Behavior and Human Decision Processes* 88(1), 430–444.
- Gottlieb, D. (2008). Competition over Time-Inconsistent Consumers. *Journal of Public Economic Theory* 10(4).
- Griffith, R., M. O’Connell, and K. Smith (2017). Corrective Taxation and Externalities from Food Consumption. *CESifo Economic Studies*.
- Gul, F. and W. Pesendorfer (2001). Temptation and Self-Control. *Econometrica* 69(6), 1403–1435.
- Gul, F. and W. Pesendorfer (2004). Self-Control and the Theory of Consumption. *Econometrica* 72(1), 119–158.
- Haushofer, J. and E. Fehr (2014). On the psychology of poverty. *Science* 344(6186), 862–867.
- Hausman, J. and E. Leibtag (2007). Consumer benefits from increased competition in shopping outlets: Measuring the effect of Wal-Mart. *Journal of Applied Econometrics* 22(7), 1157–1177.
- Hinnosaar, M. (2016, August). Time inconsistency and alcohol sales restrictions. *European Economic Review* 87, 108–131.
- Hoynes, H., D. W. Schanzenbach, and D. Almond (2016, April). Long-Run Impacts of Childhood Access to the Safety Net. *American Economic Review* 106(4), 903–934.
- Kaplan, G., J. Albrecht, G. Alessandria, F. Alvarez, K. Burdett, J. Campbell, G. Ellison, S. Ellison, D. Fitzgerald, G. Gaballo, M. Gentzkow, J. Moraggonzales, E. Hurst, N. Jaimovich, B. Jovanovic, S. Seiler, R. Shimer, R. Wright, G. Kaplan, and F. Hall (2016). Relative Price Dispersion: Evidence and Theory. Technical report.
- Kaplan, G. and G. Menzio (2015). The morphology of price dispersion. *International Economic Review* 56(4), 1165–1206.
- Kaplan, G. and S. Schulhofer-Wohl (2017, November). Inflation at the household level. *Journal of Monetary Economics* 91, 19–38.
- Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics*, 443–477.

- Mani, A., S. Mullainathan, E. Shafir, and J. Zhao (2013). Poverty Impedes Cognitive Function. *Science* 341(6149), 976–980.
- Mastrobuoni, G. and M. Weinberg (2009, August). Heterogeneity in Intra-monthly Consumption Patterns, Self-Control, and Savings at Retirement. *American Economic Journal: Economic Policy* 1(2), 163–189.
- Milkman, K. L., T. Rogers, and M. H. Bazerman (2010). I’ll have the ice cream soon and the vegetables later: A study of online grocery purchases and order lead time. *Marketing Letters* 21(1), 17–35.
- O’Donoghue, T. and M. Rabin (1999a). Doing it now or later. *American Economic Review* 89(1), 103–124.
- O’Donoghue, T. and M. Rabin (1999b). Incentives for procrastinators. *Quarterly Journal of Economics* 114(3), 769–816.
- O’Donoghue, T. and M. Rabin (2003). Studying optimal paternalism, illustrated by a model of sin taxes. *American Economic Review* 93(2), 186–191.
- Oster, S. M. and F. M. S. Morton (2005). Behavioral biases meet the market: The case of magazine subscription prices. *The BE Journal of Economic Analysis & Policy* 5(1).
- Peleg, B. and M. E. Yaari (1973). On the Existence of a Consistent Course of Action when Tastes are Changing. *Review of Economic Studies* 40(3), 391.
- Public Health England (2017, March). Sugar reduction: Achieving the 20% - GOV.UK. <https://www.gov.uk/government/publications/sugar-reduction-achieving-the-20>.
- Rabin, M. (1998). Psychology and Economics. *Journal of Economic Literature* 36(1), 11–46.
- Rayner, M., P. Scarborough, A. Boxer, and L. Stockley (2005). *Nutrient Profiles: Development of Final Model. Final Report for the Food Standards Agency*. Oxford: British Heart Foundation Health Promotion Research Group, Department of Public Health, University of Oxford.
- Rayner, M., P. Scarborough, and T. Lobstein (2009). *The UK Ofcom Nutrient Profiling Model: Defining Healthy and Unhealthy Foods and Drinks for TV Advertising to Children*.
- Read, D., G. Loewenstein, and S. Kalyanaraman (1999, December). Mixing virtue and vice: Combining the immediacy effect and the diversification heuristic. *Journal of Behavioral Decision Making* 12(4), 257–273.
- Read, D. and B. Van Leeuwen (1998). Predicting hunger: The effects of appetite and delay on choice. *Organizational behavior and human decision processes* 76(2), 189–205.
- Sadoff, S., A. S. Samek, and C. Sprenger (2015). Dynamic inconsistency in food choice: Experimental evidence from a food desert. *mimeo*.
- Samuelson, P. A. (1938). A Note on the Pure Theory of Consumer’s Behaviour. *Economica* 5(17), 61–71.
- Shapiro, J. M. (2005, February). Is there a daily discount rate? Evidence from the food stamp nutrition cycle. *Journal of Public Economics* 89(2-3), 303–325.
- Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies* 23(3), 165–180.

- Thaler, R. H. and H. M. Shefrin (1981, April). An Economic Theory of Self-Control. *Journal of Political Economy* 89(2), 392–406.
- Trope, Y. and A. Fishbach (2000, October). Counteractive self-control in overcoming temptation. *Journal of personality and social psychology* 79(4), 493–506.
- US Department of Agriculture (2015). *Dietary Guidelines for Americans 2015-2020*, Volume 8.
- Varian, H. R. (1982). The nonparametric approach to demand analysis. *Econometrica* 50(4), 945–973.
- World Health Organization (2015). *Healthy Diet*.

# APPENDIX

## FOR ONLINE PUBLICATION

A new year, a new you?

A two-selves model of within-individual  
variation in food purchases

Laurens Cherchye, Bram De Rock, Rachel Griffith,  
Martin O’Connell, Kate Smith and Frederic Vermeulen  
September, 2019

## A Data and measurement

### A.1 Purchase data

The Kantar Worldpanel records all grocery purchases (food, household supplies, toiletries, etc.) made and brought into the home by a representative panel of roughly 25,000 households. The households record (using handheld scanners in the home) all grocery purchases made and brought into the home. The data are at the transaction (barcode or UPC) level and include foods and drinks, as well as household goods such as cleaning supplies, toiletries, etc. We know the exact products purchased, the price paid for them, and we have information on the nutritional characteristics of each product.

Demographics are collected via a telephone survey for the main shopper in each household in each year. We use the demographic variables: gender, age, socioeconomic status, employment status, body mass index (BMI), whether or not the individual is a smoker, and whether or not the individual is a vegetarian - see Table A.1. For demographic variables that change over time (e.g., age) we take the median value for the time the individual is in the sample.

Socioeconomic status is based on the occupation of the individual - see <http://www.nrs.co.uk/nrs-print/lifestyle-and-classification-data/social-grade/> for details of the classification. Individuals in social grades A and B are classified as “high skilled”, individuals in social grades C1 and C2 are classified as “medium skilled”, and individuals in social grades D and E are classified as “low skilled”.



Table A.1: *Demographics*

	Mean
Female	0.58
Male	0.42
Age less than 40	0.13
Age 40-65	0.49
Age over 65	0.38
High skilled	0.13
Medium skilled	0.46
Low skilled	0.41
Full time work	0.33
Part time work	0.10
Not working	0.13
Retired	0.44
Normal weight	0.32
Overweight	0.31
Obese	0.37
Non smoker	0.78
Smoker	0.22
Vegetarian	0.11
Not vegetarian	0.89

*Notes: Demographic information is collected each year. If an individual's demographic information changes over time (e.g. age), we take the median value for the time the individual is in the sample. We construct dummy variables, listed in the first column, and the second column shows the mean of each variable across individuals.*

We focus on a sample of 3,645 single individuals over the period 2005 to 2011 (see the main paper for a description). Table A.2 compares selected demographics in our data with single individuals in the nationally representative Living Costs and Food Survey, which is a sample of repeated cross-sections. We drop months in which the individual is recorded not making any purchases for longer than 14 days. This removes periods on which individuals are on holiday. We also require that individuals are present for at least 24 (not necessarily consecutive) year-months.

There are a number of advantages to these data: they cover all grocery purchases brought into the home; they contain information at the transaction level, including price, date and nutrient information at the UPC (or barcode) level; the data are longitudinal. One drawback is that the data do not contain information on consumption; our results relate to diet quality as it affects purchase decisions.

The data also do not include information on sharing within the households, we therefore focus on individuals to avoid issues of intra-household allocation of food.

Table A.2: *Demographics of single person households*

	Kantar 2005-2011		LCFS 2005-2011	
	Mean	95% C.I	Mean	95% C.I
Male	41.6	[40.0, 43.2]	43.3	[42.4, 44.2]
Age	62.7	[62.3, 63.2]	59.6	[59.3, 59.9]
SES: Highly skilled	16.0	[14.9, 17.2]	14.9	[13.9, 15.9]
SES: Semi-skilled	50.1	[48.4, 51.7]	54.6	[53.2, 55.9]
SES: Unskilled	33.8	[32.2, 35.4]	30.5	[29.2, 31.8]

*Notes: There are 3,645 individuals in the Kantar Worldpanel sample. There are 11,628 individuals in the Living Costs and Food Survey (LCFS) sample. Male, highly skilled, semi-skilled and unskilled are dummy variables equal to 1 if the individual belongs to that group. Socioeconomic status (SES) groupings of highly skilled, semi-skilled and unskilled individuals are based on occupation.*

## A.2 Nutrient Data and the Nutrient Profile Score (NPS)

### Nutrient Data

Data on the nutritional composition of foods is collected by Kantar directly from the information available on the products' packaging. Information is provided on the energy, saturated fat, sugar, sodium, fibre and protein content of the product. We take a quantity-weighted average of nutrients at the UPC level for each of the 85 food groups – listed in Table A.4. Table A.3 lists the mean, minimum and maximum of each nutrient across the 85 food groups.

Table A.3: *Nutrients across food groups*

	Mean	Min	Max
Energy (kj/100g)	843.8	1.5	3524.7
Saturated fat (g/100g)	3.6	0.0	40.6
Sugar (g/100g)	9.0	0.0	81.4
Sodium (mg/100g)	0.2	0.0	1.0
Fibre (g/100g)	1.8	0.0	10.3
Protein (g/100g)	7.7	0.0	25.6

*Notes: Nutrients are collected from back-of-pack information for each UPC. We take a quantity-weighted average of nutrients within each food group. The table shows the mean, min, and max of nutrients across the 85 food groups.*

## Nutrient Profile Score

A number of different measures and data sources are used by nutritionists, researchers and policy makers to measure the range of foods and nutrients that an individual purchases or consumes. Measuring nutritional quality is complicated; for example, how does a product that is high in sugar and low in fat compare to one that is high in fat but low in sugar. We use the nutrient profiling score (NPS), which converts the multidimensional nutrient profile of a food product into a single dimensional score (Rayner et al. (2005), Arambepola et al. (2008) and Rayner et al. (2009)). A higher score means that the product is less healthy. Specifically, products get points based on the amount of each nutrient they contain; 1 point is given for each 335kJ per 100g, for each 1g of saturated fat per 100g, for each 4.5g of sugar per 100g, and for each 90mg of sodium per 100g. Each gram of fiber reduces the score by 1 point. Products also get scores based on their fruit and protein content. Protein enters the score only if the score omitting protein is below a threshold of 11 points. In theory, a product can score a maximum of 40 points, and a minimum of -15. The UK Food Standard Agency classifies a food product with a score of 4 points or more (and a drink with a score of 1 point or more) as “less healthy”, and these products are not allowed to be advertised during TV programmes mainly watched by children.

Table A.4: Goods: nutrient profile scores

Good	NPS=	“Unhealthy” points					– “Healthy” points			
		Energy	Saturates	Sugar	Sodium	Total	FVN	Fibre	Protein	Total
Pulses	-10.1	2.7	0.2	0.1	0.8	3.8	5.0	4.0	4.9	13.9
Frozen veg	-9.4	0.3	0.2	0.1	0.3	0.9	5.0	3.4	1.9	10.3
Dark green veg	-8.7	0.0	0.0	0.0	0.0	0.1	5.0	2.4	1.4	8.7
Potatoes	-7.0	0.3	0.0	0.0	0.0	0.3	5.0	1.2	1.1	7.3
Onions	-7.0	0.0	0.0	0.0	0.0	0.0	5.0	1.0	1.0	7.0
Other fresh veg	-6.9	0.2	0.1	0.4	0.1	0.7	5.0	1.9	0.7	7.6
Canned veg	-6.5	1.0	0.4	0.3	1.5	3.1	5.0	2.7	1.9	9.6
Orange veg	-6.3	0.0	0.0	0.9	0.0	1.0	5.0	2.2	0.1	7.2
Other salad	-5.7	0.0	0.2	0.1	0.1	0.4	5.0	0.6	0.5	6.1
Dark green salad	-5.5	0.0	0.0	0.0	0.0	0.0	5.0	0.5	0.1	5.6
Tomatoes	-5.4	0.0	0.0	0.1	0.2	0.4	5.0	0.6	0.1	5.7
Beans	-5.1	2.4	0.6	0.4	0.4	3.8	0.0	4.6	4.3	8.9
Citrus	-5.1	0.0	0.0	0.9	0.0	0.9	5.0	1.0	0.0	6.0
Other fresh fruit	-4.8	0.0	0.0	1.4	0.0	1.4	5.0	1.2	0.0	6.2
Tropical fruit	-4.3	0.0	0.1	1.8	0.0	2.0	5.0	1.3	0.0	6.3
Apples	-4.0	0.0	0.0	2.0	0.0	2.0	5.0	1.0	0.0	6.0
Oatmeal	-4.0	4.0	1.0	0.5	0.2	5.7	0.0	4.9	4.8	9.7
Berries	-3.9	0.1	0.1	1.6	0.1	1.8	5.0	0.7	0.0	5.8
Canned fruit	-3.4	0.1	0.0	1.9	0.1	2.1	5.0	0.5	0.0	5.5
Wholegrains	-3.3	2.7	0.1	0.2	3.5	6.5	0.0	4.9	4.8	9.7
Fruit juice	-3.2	0.0	0.0	1.7	0.1	1.8	5.0	0.0	0.0	5.0
Pasta	-2.8	3.3	0.2	0.0	0.5	4.0	0.0	2.4	4.5	6.8
Flour	-2.7	3.9	0.1	0.0	1.3	5.3	0.0	3.1	4.8	8.0
Grapes	-2.1	0.0	0.0	3.0	0.0	3.0	5.0	0.0	0.0	5.0
Dried fruit	-2.0	1.3	0.4	4.2	0.1	6.0	5.0	2.5	0.6	8.1
Whole chicken	-1.9	1.1	0.4	0.0	1.5	3.0	0.0	0.1	4.8	4.9
Bananas	-1.5	1.0	0.0	3.5	0.0	4.5	5.0	1.0	0.0	6.0
Other poultry	-1.3	1.3	1.1	0.0	1.4	3.7	0.0	0.0	5.0	5.0
Chicken pieces	-0.9	1.4	0.9	0.0	1.7	4.0	0.0	0.0	4.9	4.9
Other grains	-0.8	2.6	0.3	0.1	2.9	5.9	0.0	2.8	3.9	6.7
Carbonated diet drinks	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Non-carbonated diet drinks	0.0	0.0	0.0	0.1	0.0	0.1	0.0	0.0	0.0	0.0
Water	0.0	0.0	0.0	0.1	0.0	0.1	0.0	0.0	0.0	0.0
Tea	0.3	0.1	0.1	0.1	0.0	0.3	0.0	0.0	0.0	0.0
Eggs	0.1	1.0	2.8	0.0	0.9	4.7	0.0	0.0	4.7	4.7
Bread	0.3	2.4	0.4	0.2	4.3	7.3	0.0	3.0	4.1	7.1
Milk	0.4	0.0	1.0	0.9	0.5	2.4	0.0	0.0	1.9	1.9
Frozen fish	0.4	1.7	0.9	0.0	3.3	6.0	0.0	0.7	4.9	5.6
Frozen meals	0.5	1.7	0.9	0.0	1.7	4.3	0.0	2.0	1.9	3.8
Canned fish	0.7	1.2	0.7	0.0	3.6	5.6	0.0	0.0	4.8	4.9
Soup	0.8	0.5	0.4	0.6	3.1	4.7	0.0	2.4	1.5	3.9
Pork	0.9	1.6	3.6	0.0	0.4	5.6	0.0	0.0	4.7	4.7
Rice	1.2	2.5	0.4	0.0	1.7	4.6	0.0	0.9	2.5	3.3
High quality beef	1.3	1.6	3.7	0.0	0.7	5.9	0.0	0.0	4.6	4.6
Fresh fish	1.4	1.3	1.1	0.0	3.3	5.7	0.0	0.2	4.1	4.3
Muesli	1.5	4.0	0.9	3.8	0.9	9.6	0.0	5.0	3.1	8.0
Meat appetizers	1.6	2.0	1.5	0.1	3.3	7.0	0.0	1.0	4.4	5.4
Yogurt	1.7	0.6	1.0	2.0	0.4	4.0	0.0	0.2	2.1	2.3
Lamb	1.7	1.5	4.4	0.0	0.6	6.4	0.0	0.0	4.7	4.7
Healthier pizza	2.0	2.2	2.5	0.2	3.5	8.4	0.0	2.0	4.4	6.4
Ready meals	2.7	1.3	2.1	0.1	2.9	6.4	0.0	0.9	2.8	3.7
Dips	5.0	2.4	2.7	0.3	3.6	8.9	0.0	1.3	2.6	3.9
Less healthy pizza	6.9	2.8	4.0	0.1	4.3	11.2	0.0	1.9	2.4	4.3
Low quality beef	7.7	2.5	6.4	0.0	0.9	9.8	0.0	0.0	2.1	2.1
Buns and scones	8.0	3.1	2.0	3.1	3.3	11.5	0.0	2.0	1.5	3.5
Fatty pork	8.0	2.5	6.9	0.0	0.7	10.1	0.0	0.0	2.1	2.1
Breakfast cereal	8.1	4.1	0.9	4.7	3.1	12.9	0.0	3.7	1.1	4.8
Condiments	8.2	1.3	0.9	1.9	5.1	9.2	0.0	0.8	0.2	1.0
Chilled desserts	8.4	2.1	3.6	3.2	0.9	9.8	0.0	0.6	0.8	1.5
Chips and snacks	8.6	2.3	4.4	0.1	4.8	11.6	0.0	1.8	1.2	3.0
Fatty lamb	9.9	2.5	8.1	0.0	0.6	11.3	0.0	0.0	1.4	1.4
Non-carbonated non-diet drinks	1.5	0.1	0.0	0.9	0.6	1.5	0.0	0.0	0.0	0.0
Carbonated non-diet drinks	1.6	0.0	0.0	1.6	0.0	1.6	0.0	0.0	0.0	0.0
Coffee	2.8	1.1	1.5	1.8	0.5	4.8	0.0	0.9	1.1	2.1
Snacks	10.5	5.8	3.4	0.5	6.1	15.9	0.9	3.5	1.0	5.4
Pastry and pies	10.8	3.0	6.6	0.0	3.6	13.2	0.0	1.5	0.9	2.5
Frozen desserts	10.9	2.1	5.2	3.9	0.5	11.8	0.0	0.5	0.3	0.9
Milkshake mixes	11.0	3.3	0.3	8.9	0.3	12.8	0.0	1.8	0.0	1.8
Bacon and ham	11.1	1.7	3.1	0.0	8.1	12.9	0.0	0.1	1.8	1.8
Jams and syrups	12.2	3.5	1.7	7.7	0.5	13.4	0.0	1.0	0.1	1.2
Cream	13.3	3.5	9.6	0.1	0.2	13.4	0.0	0.0	0.1	0.1
Other dairy	13.5	3.2	5.9	3.4	1.9	14.4	0.0	0.1	0.8	0.9
Sausage	13.8	2.8	6.4	0.0	6.0	15.2	0.0	1.0	0.3	1.3
Croissants and waffles	14.1	4.7	6.9	1.4	3.9	16.8	0.0	2.2	0.6	2.8
Baking ingredients	14.2	3.7	0.5	8.8	1.5	14.6	0.0	0.4	0.0	0.4
Solid cheese	14.6	2.4	7.4	0.3	5.6	15.7	0.0	0.1	0.9	1.0
Cake	14.7	4.2	4.8	5.8	1.8	16.6	0.0	1.5	0.5	1.9
Pate and other deli	16.4	2.9	7.3	0.1	7.1	17.4	0.0	0.9	0.2	1.1
Margarine	18.1	9.9	8.1	0.0	0.1	18.1	0.0	0.0	0.0	0.0
Hot chocolate	18.2	4.0	3.9	8.7	5.1	21.7	0.0	3.5	0.0	3.6
Biscuits	18.5	5.3	7.4	5.5	3.0	21.2	0.0	2.5	0.2	2.7
Chocolate and confectionery	20.9	5.1	6.8	9.4	0.6	21.8	0.0	0.9	0.1	1.0
Cream cheese	21.2	4.5	9.9	0.0	6.8	21.2	0.0	0.0	0.0	0.0
Oils	21.3	5.9	9.5	0.0	6.0	21.4	0.0	0.0	0.0	0.0
Solid fats	22.6	8.5	10.0	0.0	4.1	22.6	0.0	0.0	0.0	0.0

We aggregate individual food UPCs (barcodes), of which there are around several hundred thousand, into 85 goods, listed in Table A.4. The table shows the NPS of products within each good, and the components that constitute the NPS. For the empirical implementation we need to split goods into a “healthy” and “unhealthy” set. The 34 goods with an average NPS less than 0 are deemed to be preferred only by the healthy self (i.e., “always healthy” items). The 24 goods with an average NPS of more than 10 (or more than 1 for drinks) are deemed to be preferred only by the unhealthy self (i.e., “always unhealthy” items). We consider the 27 goods with an average NPS of between 0 and 10 as uncertain and potentially belonging to either the healthy or the unhealthy category for each individual. Section 3.4 explains how we partition the goods into those purchased by the healthy and unhealthy self for our sample of individuals.

### A.3 Prices

#### Prices of goods

We model demand at the level of 85 goods. We observe prices at the UPC (barcode) level; there are several hundred thousand of these. Many of these are the same product available in different pack sizes, formats and, in some cases, flavors. There are 113,025 distinct products (defined by brand). We aggregate these to the good level based on their nutritional characteristics.

Let  $b$  index 113,025 products,  $g$  index 85 goods,  $r$  index three geographic regions (north, central, south),  $t$  index month. The set of products in each good is  $\mathcal{B}_g$ . The set of individuals living in region  $r$  is  $\mathcal{R}_r$ .

The variable  $p_{brt}$  denotes the mean (quantity weighted across different pack sizes and formats) price per kilogram of product  $b$  in region  $r$  in month  $t$ . The variable  $q_{ib}$  denotes the total (over time) quantity of product  $b$  purchased by consumer  $i$ . The price of foods in good  $g$ , in region  $r$  in month  $t$  is measured as the weighted average of the prices of products in that good, where the weights are the mean (over time and consumers) quantity shares of consumers in that region.

Define the share of product  $b$  in region  $r$  as  $w_{br} = \frac{\sum_{i \in \mathcal{R}_r} q_{ib}}{\sum_{i \in \mathcal{R}_r} \sum_{b' \in \mathcal{B}_g} q_{ib'}}$ .

The price of foods in good  $g$  is given by

$$P_{grt} = \sum_{b \in \mathcal{B}_g} w_{br} p_{brt}, \quad (\text{A.1})$$

In Table A.5 we describe the price variation and expenditure share of each good.

Table A.5: *Goods: budget shares and relative prices*

Good	Budget share	Price variation	
		Min.	Max.
Pulses	0.02	0.75	1.26
Frozen veg	0.69	0.89	1.10
Dark green veg	1.09	0.80	1.52
Potatoes	1.52	0.79	1.26
Onions	0.14	0.72	1.25
Other fresh veg	1.95	0.81	1.27
Canned veg	0.81	0.86	1.23
Orange veg	0.46	0.76	1.35
Other salad	1.07	0.69	1.37
Dark green salad	0.87	0.82	1.20
Tomatoes	1.39	0.78	1.32
Citrus	1.04	0.76	1.32
Beans	0.02	0.81	1.18
Other fresh fruit	1.44	0.80	1.29
Tropical fruit	0.18	0.82	1.37
Apples	0.97	0.85	1.33
Oatmeal	0.31	0.85	1.24
Berries	0.20	0.69	1.46
Canned fruit	0.34	0.83	1.21
Wholegrains	1.30	0.82	1.14
Fruit juice	1.74	0.79	1.22
Pasta	0.31	0.83	1.21
Flour	0.15	0.78	1.32
Grapes	0.85	0.70	1.57
Dried fruit	1.77	0.65	1.65
Whole chicken	1.68	0.81	1.15
Bananas	1.10	0.81	1.23
Other poultry	0.48	0.80	1.38
Chicken pieces	0.78	0.81	1.18
Other grains	0.05	0.87	1.13
Carbonated diet drinks	0.81	0.79	1.24
Non-carbonated diet drinks	0.11	0.83	1.18
Water	0.52	0.91	1.14
Tea	1.02	0.85	1.22
Eggs	1.15	0.72	1.25
Bread	3.37	0.86	1.16
Milk	5.34	0.83	1.16
Frozen fish	1.17	0.93	1.09
Frozen meals	1.33	0.92	1.10
Canned fish	0.85	0.76	1.38
Soup	1.91	0.85	1.20
Pork	0.77	0.86	1.19
Rice	0.45	0.78	1.27
High quality beef	1.31	0.83	1.27
Fresh fish	1.81	0.81	1.25
Muesli	0.24	0.91	1.17
Meat appetizers	0.15	0.91	1.13
Yogurt	2.80	0.89	1.14
Lamb	0.58	0.83	1.29
Healthier pizza	0.11	0.92	1.11
Ready meals	6.98	0.91	1.13
Dips	0.44	0.95	1.06
Less healthy pizza	0.66	0.94	1.12
Low quality beef	0.76	0.78	1.31
Buns and scones	0.38	0.85	1.31
Fatty pork	0.39	0.85	1.17
Breakfast cereal	1.19	0.93	1.18
Condiments	2.21	0.89	1.15
Chilled desserts	2.02	0.88	1.15
Chips and snacks	0.20	0.84	1.18
Fatty lamb	0.29	0.83	1.38
Non-carbonated non-diet drinks	0.71	0.91	1.19
Carbonated non-diet drinks	0.83	0.84	1.29
Coffee	1.73	0.92	1.16
Snacks	2.30	0.90	1.19
Pastry and pies	1.86	0.88	1.15
Frozen desserts	1.53	0.93	1.09
Milkshake mixes	0.02	0.89	1.16
Bacon and ham	4.99	0.83	1.17
Jams and syrups	0.68	0.84	1.20
Cream	0.43	0.79	1.25
Other dairy	0.18	0.83	1.23
Sausage	1.07	0.88	1.16
Croissants and waffles	0.15	0.87	1.19
Baking ingredients	0.84	0.82	1.22
Solid cheese	0.86	0.85	1.18
Cake	2.21	0.88	1.18
Pate and other deli	0.18	0.90	1.18
Margarine	0.39	0.75	1.21
Hot chocolate	0.33	0.91	1.11
Biscuits	3.52	0.88	1.19
Chocolate and confectionery	4.59	0.93	1.13
Cream cheese	2.51	0.86	1.18
Oils	1.07	0.83	1.31
Solid fats	0.97	0.81	1.34

Notes: Column (1) shows the share of total spending (across individuals and year-months) on each good. The construction of prices for each good in each region-year-month is described in the text. For each good in each region-year-month, we divide price over the average price for the good across all regions and year-months. Columns (2) and (3) show the minimum and maximum values of this for each good.

## Healthy and unhealthy food price indices

The price indices for healthy and unhealthy sets of foods that we use in Section 4 are defined as:

$$\Pi_{it}^h = \sum_{g \in H_i} w_{ig} P_{grt} \quad (\text{A.2})$$

$$\Pi_{it}^l = \sum_{g \in L_i} w_{ig} P_{grt} \quad (\text{A.3})$$

where  $H_i, L_i$  are the sets of goods in the healthy and unhealthy sets, respectively, and:

$$w_{ig} = \frac{\sum_{b \in \mathcal{B}_g} q_{ib}}{\sum_{g' \in X_i} \sum_{b \in \mathcal{B}_{g'}} q_{ib}}, \text{ where } X = H, L$$

i.e. the quantity share that each good  $g$  constitutes within the set of healthy,  $H_i$ , and unhealthy,  $L_i$ , foods for individual  $i$  across all time periods.

## A.4 Advertising

We use data from the AC Nielsen Digest of Advertising. These data record monthly advertising expenditure by all firms at the brand level in the UK across the main media types, including television, online, radio, press. We use information on advertising in the categories “Confectionery & snacks”, “Drinks”, “Prepared and convenience food” and class these as advertising of unhealthy food and drinks products. We sum all advertising expenditure of all brands in these categories across all media formats to construct a monthly advertising expenditure variable,  $\tilde{a}_t$ .

It is standard in the literature to consider that past advertising affects current demand (e.g. Erdem et al. (2008)); we therefore use the monthly advertising expenditure flow variable to construct a stock measure,  $a_t$ :

$$a_t = \sum_{n=0}^{t-t_0} \delta^n \tilde{a}_{t-n}$$

where  $\delta$  is a depreciation factor applied to advertising expenditure; we assume  $\delta = 0.8$ .

## A.5 Weather

We use weather data provided by the UK Met Office (<https://www.metoffice.gov.uk/public/weather/climate-historic/#?tab=climateHistoric>). We use data on the monthly minimum temperature, maximum temperature and monthly total rainfall

collected at 37 weather stations across the UK. We match households to their nearest weather station based on their geographic location.

## A.6 Stated preferences

Kantar Worldpanel asks participants a selection of questions to gauge their attitude to a variety of lifestyle factors. We use a subset of these to construct a measure of individuals' preferences for healthy food, processed food, tendency to buy things on offer, shopping commitment, and impulsiveness. The questions we use are listed in Table A.6. Questions change from year to year; no questions were asked in 2008 and 2011.

Responses to the questions are recorded using a Likert scale. We code the responses as follows: 1 "Strongly disagree", 2 "Disagree", 3 "Neither agree nor disagree (or missing)", 4 "Agree", 5 "Strongly Agree". When an individual answers the same question across years, we take her median response.

We take an unweighted average of responses to the questions within each group to construct individuals' stated preferences for healthy food, processed food, tendency to buy things on offer, shopping commitment, and impulsiveness. On the basis of these aggregated preferences, we divide individuals into groups: above and below the median for each set of stated preferences. Figure A.1 shows the distribution of individuals' stated preferences in each of these 5 areas, and also shows the median in each case.

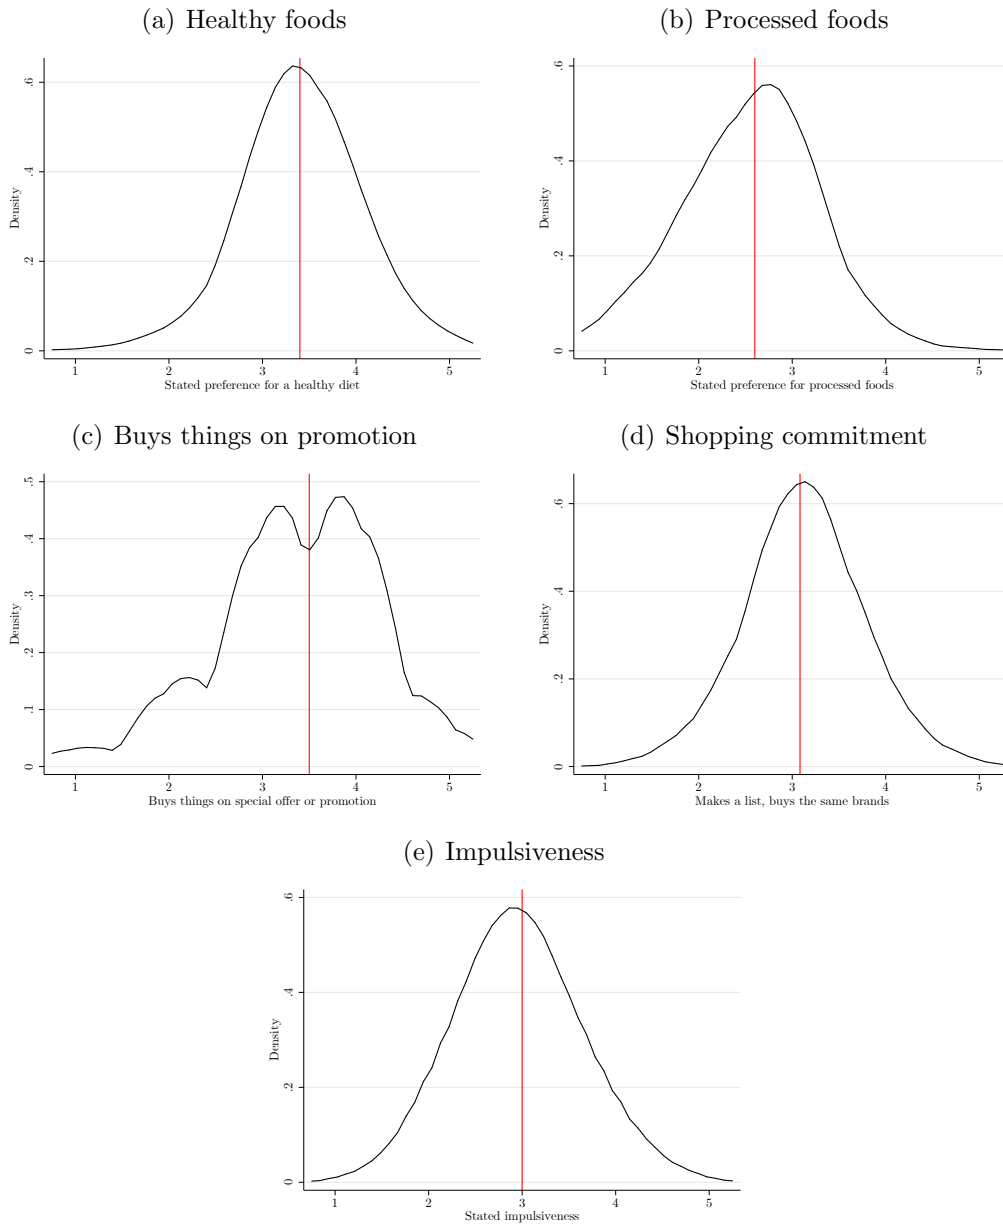


Table A.6: Stated preferences

Question:	Collected in:					
	2005	2006	2007	2009	2010	2012
<i>Preferences for healthy foods</i>						
I try to buy a healthy range of foods these days	Yes	Yes	Yes	No	No	No
I look out for the 'healthy' products	Yes	Yes	Yes	No	No	No
I prefer to buy low salt products if they're available	Yes	Yes	Yes	No	No	No
My diet is very important to me	Yes	Yes	Yes	Yes	Yes	Yes
I make sure I eat well-balanced meals	Yes	Yes	Yes	No	No	No
I read the list of ingredients on the pack before buying	Yes	Yes	Yes	No	No	No
I try to lead a healthy lifestyle	No	No	No	Yes	Yes	Yes
I am actively trying to manage my cholesterol level	No	No	No	Yes	Yes	Yes
The nutritional labelling on food and drink products has an effect on what I buy	No	No	No	Yes	Yes	Yes
I restrict how much sugary food I eat	No	No	No	Yes	Yes	Yes
<i>Preferences for processed foods</i>						
Snacks are one of my favourite kind of foods to eat	Yes	No	No	No	No	No
I tend to eat when I am bored	Yes	Yes	Yes	No	No	No
I'm not really worried about eating healthily	Yes	Yes	Yes	No	No	No
I rely heavily on convenience products to make cooking simple and quick	Yes	Yes	Yes	Yes	Yes	Yes
I'm a busy person so often eat on the run	Yes	No	No	No	No	No
I often buy take-away meals to eat at home	No	No	No	Yes	Yes	Yes
<i>Tendency to buy things on promotion</i>						
I often buy things just because I see them on the shelf	Yes	Yes	Yes	No	No	No
I will buy a brand I don't normally buy if it's on offer	Yes	Yes	Yes	Yes	Yes	Yes
I shop around to take advantage of special offers	No	No	No	Yes	Yes	Yes
<i>Shopping commitment</i>						
Once I find a brand I like I tend to stick to it	Yes	Yes	Yes	Yes	Yes	Yes
I decide which brands to buy before I go shopping	Yes	No	No	No	No	No
It is important to me which brand I buy	Yes	No	No	No	No	No
I make a shopping list before I go out and stick to it	No	Yes	Yes	Yes	Yes	Yes
I plan my shopping trips so I know I can be as efficient as possible	No	No	Yes	No	No	No
<i>Impulsiveness</i>						
I stick to a set routine for doing the housework <sup>†</sup>	Yes	No	No	No	No	No
With a credit card I sometimes spend more than I should	Yes	No	No	No	No	No
I make a shopping list before I go out and stick to it <sup>†</sup>	Yes	No	No	No	No	No
I tend to spend money without thinking	Yes	No	No	No	No	No
I spend more money in the supermarket than I intend to	Yes	No	No	No	No	No

Notes: Each question is answered on a 5-point Likert scale: 1 "Strongly disagree", 2 "Disagree", 3 "Neither agree nor disagree (or missing)", 4 "Agree", 5 "Strongly Agree". <sup>†</sup> for these questions we reverse the ordering of the answer scale so 1 is "Strongly agree" etc.

Figure A.1: *Distribution of stated preference variables*




*Notes: The questions that we use to construct each set of preferences are listed in Table A.6. We take an unweighted average of responses to questions within each set. The responses are measured on a Likert scale: 1 “Strongly disagree”, 2 “Disagree”, 3 “Neither agree nor disagree (or missing)”, 4 “Agree”, 5 “Strongly Agree”. The red vertical line shows the median stated preference in each case; we divide individuals into two groups (for each stated preference measure) on the basis of whether their stated preference is above or below the median.*

## A.7 Total expenditure as a proxy for income

We use total grocery expenditure to proxy for household income. The Living Costs and Food Survey (LCFS) is an expenditure survey that collects data on spending for a repeated cross-section of households (in contrast to the Kantar data, which has a panel structure). It also contains information on household income. Figure A.2 shows that there is a strong relationship between households' annual equivalized income and equivalized weekly grocery spending.

Figure A.2: *Relationship between household income and grocery expenditure*




Notes: Figure drawn using data on 4937 households in the Living Costs and Food Survey 2011. The horizontal axis shows logged equivalised annual income of the household, and the vertical axis shows equivalised weekly grocery expenditure. Figure trims the 5th and 95th percentiles of the logged equivalised annual income distribution. We equivalise using the standard OECD modified equivalence scale (see Hagenaars et al. (1994)).

## B Variation in diet quality, all households

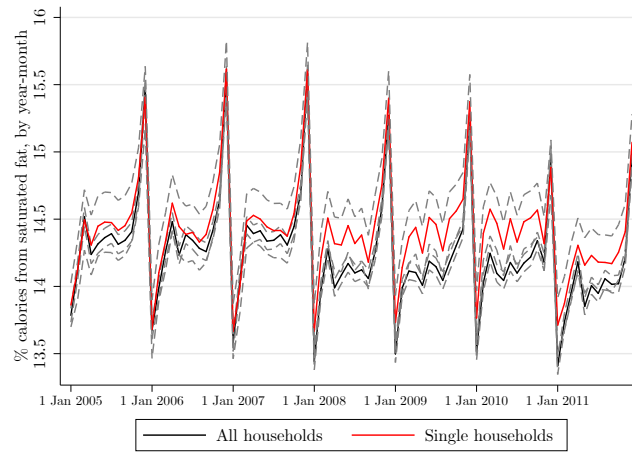
In our main analysis we focus on a sample of single households, to avoid the confounding implications of intra-household allocation. Here we show that that broad patterns of declining diet quality across the calendar year can be seen in the full, nationally representative Kantar Worldpanel. Figure B.1 shows the mean share of calories from added sugar, saturated fat, and protein for all households and single households only, over year-months, 2005-2011.

Figure B.1: *Diet quality, 2005–2011*


(a) % calories from added sugar



(b) % calories from saturated fat



(c) % calories from protein



Notes: The figure shows the percentage of calories from added sugar, saturated fat and protein purchased in each year-month for a sample of 51,969 households over 2005–2011. The red lines show the equivalent for the sample of single households we study in the main paper. The grey lines show the 95% confidence interval.

## C Power

A power analysis evaluates the probability of detecting an alternative hypothesis to the model under study. Bronars (1987) first defined a procedure to assess the power for revealed preference conditions of rationalizability. His alternative hypothesis was based on the notion of irrational behavior of Becker (1962), which states that individuals randomly choose consumption bundles that exhaust the available budget. Our power assessment adapts Bronars' procedure for the two-selves model.

An important requirement for our power evaluation method relates to the prevalence of zero expenditures in the data. In fact, without explicit correction, randomly drawing bundles from a budget constraint obtains a zero probability of simulating zero consumption of a certain item. Such a simulation would therefore never match reality if zero expenditures are present. To take the observed zero expenditures into account, we calculate the proportion  $\pi_{gi}$  of strictly positive expenditures for each individual  $i$  for each good  $g$  over the  $T_i$  observations. We subsequently draw individual-specific random bundles defined by  $\nu_{git}$ . To define each  $\nu_{git}$  we first draw a random number from the uniform distribution between 0 and 1. If this number is greater than  $\pi_{gi}$ , then we set  $\nu_{git}$  equal to zero. If it is less than  $\pi_{gi}$  then  $\nu_{git}$  is the result of a new drawing from the uniform distribution (between 0 and 1). The (random) budget share of good  $g$  for individual  $i$  in observation  $t$  is defined as  $w_{git} \equiv (\nu_{git} / \sum_i \nu_{git})$ . The random quantity bundle for individual  $i$  in observation  $t$  is obtained by multiplying this budget share  $w_{git}$  by the observed expenditure level  $x_{it}$  and dividing the outcome by the corresponding components of the price vector  $\mathbf{p}_t$ .

For each individual  $i$  and each observation  $t$  we use this procedure to construct 200 random consumption bundles. This defines 200 series of  $T_i$  random consumption bundles. The advantage of this procedure is that it results in an expected proportion of zero expenditures that complies with the observed proportion. If an individual has no expenditures on a particular good across all observations  $t$ , then it is never randomly allocated a consumption bundle with strictly positive expenditures on that good. The randomly constructed consumption bundles can be used to evaluate the power of the rationalizability conditions for our two-selves model. For each individual  $i$  (characterized by  $T_i$  observed price-income regimes), we compute the proportion of random draws with Afriat indices above the true Afriat index computed with the data. This captures the probability that this true index is below the Afriat index associated with random behavior.

## D Correlation with stated preferences

### D.1 Relationship between mean sharing rule and stated preferences

Table D.1 shows the relationship between individuals' mean sharing rule and individuals' stated preferences, conditioning on demographic controls (listed in Table A.1).

Table D.1: *Variation in mean sharing rule by stated preferences*

(1)	(2)	(3)	(4)
	Mean	Difference	95% CI for diff.
<i>With demographic controls</i>			
Above median preferences for healthy food	58.7		
Below median preferences for healthy food	55.2	-3.4	[-4.5, -2.4]
Above median preferences for processed food	56.4		
Below median preferences for processed food	58.9	2.5	[1.5, 3.5]
Above median tendency to buy on promotion	57.6		
Below median tendency to buy on promotion	57.8	0.2	[-0.8, 1.2]
Above median shopping commitment	57.9		
Below median shopping commitment	57.5	-0.4	[-1.4, 0.6]
Above median stated impulsiveness	56.4		
Below median stated impulsiveness	58.1	1.7	[0.3, 3.1]

*Notes: The numbers in column (3) are the difference in means from the first row in each group. Confidence intervals reflect uncertainty arising from the sample of individuals, but not over the time series variation for each individual. The table is analogous to Table 4.1, but conditions on demographic controls. The demographic controls included are the variables listed in Table A.1.*

### D.2 Relationship between variation in residual sharing rule and stated preferences

Table D.2 summarizes the relationships between variation in the residual sharing rule and individuals' stated preferences, conditioning on demographic controls (listed in Table A.1).

Table D.2: *Variation in  $\hat{\sigma}_i$  by stated preferences*

	(1) Mean	(2) Difference	(3) 95% CI for diff.
<i>With demographic controls</i>			
Above median preferences for healthy food	7.40		
Below median preferences for healthy food	7.40	0.00	[-0.14, 0.14]
Above median preferences for processed food	7.50		
Below median preferences for processed food	7.30	-0.19	[-0.34, -0.05]
Above median tendency to buy on promotion	7.48		
Below median tendency to buy on promotion	7.21	-0.26	[-0.41, -0.12]
Above median shopping commitment	7.30		
Below median shopping commitment	7.45	0.15	[0.01, 0.29]
Above median stated impulsiveness	7.95		
Below median stated impulsiveness	7.67	-0.28	[-0.47, -0.09]

*Notes: The numbers in column (3) are the difference in means from the first row in each group. Confidence intervals for  $\hat{\sigma}_i$  reflect uncertainty arising from the sample of individuals, but not over the time series variation for each individual. The table is analogous to Table 5.1, but conditions on demographic controls. The demographic controls included are the variables listed in Table A.1.*

## E Variation with demographics

### E.1 Relationship between mean sharing rule and demographics

Table E.1 describes the relationship between individuals' mean sharing rule and individuals' demographics.

Table E.1: *Variation in mean sharing rule by demographic characteristics*

(1)	(2)	(3)	(4)
	Mean	Difference	95% CI for diff.
Female	51.3		
Male	49.5	-1.9	[-2.8, -0.9]
Age less than 40	49.6		
Age 40-65	49.9	0.3	[-1.2, 1.8]
Age over 65	51.7	2.1	[0.6, 3.6]
High skilled	52.2		
Medium skilled	51.2	-1.0	[-2.5, 0.5]
Low skilled	49.4	-2.8	[-4.3, -1.3]
Full time work	50.9		
Part time work	50.9	-0.0	[-1.8, 1.8]
Not working	48.6	-2.3	[-3.9, -0.7]
Retired	51.4	0.5	[-0.7, 1.6]
Normal weight	51.6		
Overweight	50.9	-0.7	[-1.9, 0.5]
Obese	49.4	-2.3	[-3.4, -1.1]
Non smoker	51.3		
Smoker	47.8	-3.5	[-4.6, -2.4]
Vegetarian	50.9		
Non vegetarian	50.5	-0.3	[-1.8, 1.2]
Above median income	51.1		
Below median income	50.0	-1.1	[-2.0, -0.2]

*Notes: The numbers in column (3) are the difference in means from the first row in each group. Confidence intervals for  $\hat{\alpha}_i$  reflect uncertainty arising from the sample of individuals, but not over the time series variation for each individual.*

## E.2 Relationship between variation in residual sharing rule and demographics

Table E.2 summarizes the relationships between variation in the residual sharing rule and individuals' demographics.



Table E.2: Variation in  $\hat{\sigma}_i$  by demographic characteristics

	(1) Mean	(2) Difference	(3) 95% CI for diff.
Female	7.59		
Male	7.67	0.08	[-0.06, 0.22]
Age less than 40	8.32		
Age 40-65	7.65	-0.68	[-0.89, -0.46]
Age over 65	7.35	-0.98	[-1.20, -0.76]
High skilled	7.74		
Medium skilled	7.60	-0.14	[-0.36, 0.07]
Low skilled	7.60	-0.14	[-0.36, 0.08]
Full time work	7.89		
Part time work	7.42	-0.47	[-0.74, -0.21]
Not working	7.80	-0.09	[-0.32, 0.15]
Retired	7.34	-0.55	[-0.71, -0.38]
Normal weight	7.67		
Overweight	7.68	0.01	[-0.16, 0.18]
Obese	7.53	-0.14	[-0.30, 0.03]
Non smoker	7.58		
Smoker	7.75	0.17	[0.00, 0.33]
Vegetarian	7.74		
Non vegetarian	7.60	-0.14	[-0.35, 0.08]
Above median income	6.92		
Below median income	8.32	1.39	[1.26, 1.52]

Notes: The numbers in column (3) are the difference in means from the first row in each group. Confidence intervals for  $\hat{\sigma}_i$  reflect uncertainty arising from the sample of individuals, but not over the time series variation for each individual.

## References for Online Appendix

- Arambepola, C., P. Scarborough, and M. Rayner (2008). Validating a nutrient profile model. *Public Health Nutrition* 11(04), 371–378.
- Becker, G. S. (1962). Irrational Behavior and Economic Theory. *Journal of Political Economy* 70(1), 1–13.
- Bronars, S. G. (1987). The Power of Nonparametric Tests of Preference Maximization. *Econometrica* 55(3), 693–698.

- Erdem, T., M. Keane, and B. Sun (2008). The impact of advertising on consumer price sensitivity in experience goods markets. *Quantitative Marketing and Economics* 6(2), 139–176.
- Hagenaars, A. J., K. De Vos, M. Asghar Zaidi, and others (1994). Poverty statistics in the late 1980s: Research based on micro-data.
- Rayner, M., P. Scarborough, A. Boxer, and L. Stockley (2005). *Nutrient Profiles: Development of Final Model. Final Report for the Food Standards Agency*. Oxford: British Heart Foundation Health Promotion Research Group, Department of Public Health, University of Oxford.
- Rayner, M., P. Scarborough, and T. Lobstein (2009). *The UK Ofcom Nutrient Profiling Model: Defining Healthy and Unhealthy Foods and Drinks for TV Advertising to Children*.