

KU LEUVEN

DEPARTMENT OF ECONOMICS

A new year, a new you? Heterogeneity and self- control in food purchases

Laurens CHERCHYE, Bram DE ROCK, Rachel GRIFFITH, Martin O'CONNELL, Kate SMITH and Frederic VERMEULEN

FACULTY OF ECONOMICS AND BUSINESS



DISCUSSION PAPER SERIES DPS17.19

DECEMBER 2017

A new year, a new you?

Heterogeneity and self-control in food purchases

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

December 8, 2017

Abstract

We document considerable within-person (over time) variation in diet quality that is not fully explained by responses to fluctuations in the economic environment. We propose a two-selves model that provides a structural interpretation to this variation, in which food choices are a compromise between a healthy and an unhealthy self, each with well-behaved preferences. We show that the data are consistent with this model using revealed preference methods. The extent of self-control problems is higher among younger and lower income consumers, though this is overstated if we do not control for responses to fluctuations in the economic environment. Our results are intuitively related to stated attitudes on self-control.

Keywords: Two-selves model, self-control, revealed preferences

JEL classification: C14, D12, D90, I12.

Acknowledgements: The authors gratefully acknowledge financial support from the European Research Council (ERC) under ERC-2015-AdG-694822 and ERC-2013-CoG-614221, the Economic and Social Research Council (ESRC) under the Centre for the Microeconomic Analysis of Public Policy (CPP), grant number RES-544-28-0001, the British Academy under pf160093, the FWO (Research Fund-Flanders) and the FNRS (Research Fund-Wallonia). 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. We thank participants at the 2017 Annual Meeting of ASSET in Algiers, the 2017 (Ce)² Workshop in Warsaw, and seminar participants at the Institute for Fiscal Studies, HEC Montréal and LISER for their comments. All errors and omissions remained the responsibility of the authors.

*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.

“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

A growing theoretical literature posits that people sometimes make decisions that they later regret. The presence of self-control problems opens the possibility that public policy may be able to improve welfare by increasing the cost of consumption today to reflect unanticipated future costs, or by facilitating the use of commitment devices (O’Donoghue and Rabin (2003), Gruber and Koszegi (2004), O’Donoghue and Rabin (2006) and Bernheim (2009)). A leading example of a setting in which self-control problems may play a role is in food consumption. Evidence from the experimental literature (for instance, Read and Van Leeuwen (1998) and Gilbert et al. (2002)) and the existence of a multi-billion dollar diet industry (Cutler et al. (2003)) attest to this.

However, there is limited direct evidence on self-control problems from observational consumption data. An important reason for this is that to provide convincing evidence requires that we account for the considerable heterogeneity in individuals’ food preferences that leads both to wide cross-sectional variation in food choices and to variation in how individuals respond to changes in the economic environment. In this paper we provide empirical evidence on the existence, size and variation in self-control problems in food choice, while controlling for heterogeneity in preferences and in responses to changes in the economic environment.

We make two contributions to the literature. First, we document substantial within-person variation in diet quality over time. Considerable attention has been paid to the variation in diet quality across people; we show that variation within people is of a similar magnitude to the cross-sectional variation. The patterns of within-person variation are suggestive of self-control problems, with individuals exhibiting periods of significant deterioration in the nutritional quality of their food purchases followed by periodic “resets” (e.g. around New Years), where diet quality improves substantially for a limited time and then deteriorates again.

Second, we develop an empirically tractable framework that allows us to account for the heterogeneous food choice behavior of individuals and that provides a structural interpretation of the within-person time-series fluctuations in diet quality. To do this we bring together insights from the behavioral literature on

multi-selves models with the literature on collective household behavior. We model individual food purchase behavior as a compromise between two selves, a healthy and an unhealthy self, each with stable preferences over two separate baskets of food and drink products. The first basket is associated with a healthy lifestyle and contains products such as fruits, vegetables and whole grains. The second basket contains products such as soda, crisps and confectionery. The allocation of products between the two selves is determined endogenously and is consumer specific. Our framework encompasses standard models of consumer purchase behavior. This means that it incorporates the impact of variation in the relative prices of foods and in food budgets, in addition to psychological factors related to temptation and self-control, to explain the observed purchase behavior.

The two selves bargain over the food budget, with bargaining power that can change over time. We model the two selves as entering into a bargaining process for which we make only the minimal assumption that it results in a Pareto optimal outcome. Pareto efficiency is a weak requirement that underlies most existing multi-selves models. We make use of the sharing rule concept, often used to quantify the bargaining power of individuals in collective models, to be able to interpret variation in the diet quality of individuals' food baskets over time as due to variation in the bargaining weights of the healthy and unhealthy selves. For an individual who exhibits self-control, either from an intrinsic psychological trait, or from external commitment devices, we expect the bargaining weights of the two selves to remain broadly stable throughout the year. In contrast, an individual who lacks self-control may experience fluctuations in the bargaining weight of the unhealthy self, reflecting temporary resoluteness followed by periods where they succumb to temptation.

This definition of self-control is related to that used in the theoretical literature. We focus on increases in the influence of the unhealthy self in decision making as indicating failure to exert self-control. We also describe variation in the mean sharing rule across individuals; however, without making further assumptions, it is not possible to identify whether this reflects heterogeneity in (stable) preferences across individuals, or variation in the average influence that the unhealthy self has over the food budget. A strength of our focus on within-consumer deviations in choices over time is that we avoid confounding self-control problems with heterogeneity in preferences, something that we would do if we were to use cross-consumer comparisons to make inference about the extent of self-control issues (see also Ameriks et al. (2007) and Bucciol (2012)).

Our empirical analysis exploits longitudinal data on the grocery purchases of a sample of British individuals. Revealed preference methods suggest that the two-selves model does a good job at explaining variation in the data, and a better job than a single-self model, in which each individual is characterized by a single, stable utility function. We recover the sharing rule, which describes the share of the food budget that is allocated to the healthy self. We show that the level of this sharing rule, which represents the healthy self's relative bargaining power, is correlated with stated preferences that support our interpretation of consumers' behavior. For example, individuals with a higher mean sharing rule are more likely to state that they try to eat a healthy diet and are less likely to state that they have tastes for processed foods.

We show that around 20% of within-person variation in the sharing rule is explained by responses to changes in prices or food budgets. We also show that seasonality in preferences does not explain the remaining variation in the sharing rule. We investigate the relationship between our measure of self-control (residual variation in the sharing rule after controlling for fluctuations in prices and food budgets) with age and income. Several papers have posited that self-control may be causally related to poverty (see, e.g., Mani et al. (2013), Haushofer and Fehr (2014) and Bernheim et al. (2015)). We find evidence consistent with this hypothesis: individuals with lower income experience greater variation in their spending on healthy foods, even after controlling for their responses to changes in prices and food budgets. We find a similar relationship with age: younger people have more variation in their residual sharing rule than older people. This is consistent with findings from the savings literature (see, e.g., Ameriks et al. (2007) and Bucciol (2012)). In addition, we find that self-control issues are correlated with stated preferences. For example, individuals who state that they regularly make a shopping list and commit to buying the same brands also exhibit smaller fluctuations. On the other hand, individuals who state that they often spend money without thinking, or spend more on their credit card than they should, have larger fluctuations in their sharing rule.

Our work relates to several important literatures. Firstly, it relates to the literature on multiple-selves models, which dates back to Strotz (1955) and Peleg and Yaari (1973) and which remains very popular in the theoretical literature (see, e.g., Gul and Pesendorfer (2001), Kalai et al. (2002), Gul and Pesendorfer (2004), Fudenberg and Levine (2006), Manzini and Mariotti (2007), de Clippel and Eliaz (2012), Manzini and Mariotti (2015), Ambrus and Rozen (2015) and Manzini and Mariotti (2016)). In this paper our aim is to study the choices that individu-

als make and to consider what they can tell us about the extent of self-control problems. Many of these multi-selves models deliver observationally equivalent predictions about choice behavior. They do have importantly different implications for welfare, but this is not our focus in this paper. The main intuition that we take away from these multi-selves models can be distilled down to capturing consumer decision making as a process in which two selves each have an influence on decision-making (Manzini and Mariotti (2007) and Spiegler (2011)).

There is an extensive psychological literature on the use of self-regulation and behavior modification (see references and discussion in Rabin (1998), DellaVigna (2009), Dolan et al. (2012) and Bernheim et al. (2015)). The literature has found relatively little evidence of the effective use of such commitment devices (see DellaVigna and Malmendier (2006), DellaVigna (2009), Giné et al. (2010), Bryan et al. (2010), Garon et al. (2015) and Bernheim et al. (2016)). An extension of the multi-selves model formalizes the notion of costly self-control (see Gul and Pesendorfer (2001) and Fudenberg and Levine (2006)).

Our paper also closely relates to the collective approach to choice behavior, which is popular in the household economics literature (see Chiappori (1988), Chiappori (1992), Browning and Chiappori (1998), Chiappori and Ekeland (2009), Dunbar et al. (2013) and Browning et al. (2013)). The structure of the collective model turns out to be well suited to bring the multiple-selves model to data. We consider an individual with a healthy and an unhealthy self characterized by (different) rational preferences and that enter into a bargaining process that results in a Pareto optimal outcome. Chiappori (1988) shows that the sharing rule – in our set up this is the share of the budget allocated to healthy goods – is a direct indication of the bargaining power of the healthy self. Therefore, we are able to remain agnostic about the specific interaction between the selves (as long as it results in a Pareto optimal outcome), which means that our approach encompasses most of the cited multiple-selves models, and we are still able to interpret variation in the sharing rule as indicative of variation in the bargaining power of the healthy self.

We check whether the theoretical implications of our two-selves model are satisfied for our data by making use of revealed preference methods in the tradition of Samuelson (1938), Afriat (1967) and Varian (1982). Our empirical application builds on the work of Cherchye, De Rock, and Vermeulen (2007, 2011), who developed revealed preference methods to analyze collective choice behavior. These methods provide a useful tool to evaluate the consistency of our two-selves model with the data. Revealed preference methods typically place restrictions on the time

variation in preferences, but are very flexible regarding the assumptions about the functional form of demand. The methods allow us to check behavioral consistency with the two-selves model for each individual separately, and thus avoid assuming that individuals are characterized by homogeneous preferences.

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 by individuals across time. We show that this variation is heterogeneous across individuals and provide descriptive evidence that is indicative of temptation and self-control problems. In Section 3 we develop a model of behavior that captures temptation by modeling an individual's decision making as reflecting the tension between a healthy and an unhealthy self. Importantly, we allow for preference heterogeneity across individuals. In Section 4 we show that individual purchase behavior can be rationalized by the two-selves model. We analyze how the sharing rule varies with prices, food budgets and what it suggests about self-control and temptation. A final section concludes and discusses some avenues for future work.


2 Food purchasing behavior

Considerable attention has been paid to variation in diet quality across individuals and the impact this has had on rising obesity and diet-related disease (Cutler et al. (2003), Finkelstein and Zuckerman (2008), Bleich et al. (2008), Baum and Ruhm (2009) and Cutler and Lleras-Muney (2010)). Cross-sectional heterogeneity in diet quality has been shown to be associated with important inequalities in health outcomes; for instance, Cutler and Lleras-Muney (2010) show that obesity rates are half as high among the better educated. Much less attention, however, has been paid to the variation in diet quality within-person over time. We emphasize the importance of this type of variation by studying the decisions that individuals make over food purchases and the nutritional characteristics of their food basket through time. This is a setting in which the idea that some consumers may suffer from self-control problems is particularly salient (O'Donoghue and Rabin (2000), Gul and Pesendorfer (2001) and Downs et al. (2009)), and where poor decisions have potentially important consequences for well-being and welfare (Cawley (2000), Bhattacharya and Sood (2011) and Finkelstein et al. (2013)).

People's tendency to make (and fail to keep) New Year's resolutions to lead more healthy lifestyles indicates the potential importance of within-person variation in diet quality over time. Figure 2.1 shows data from Google Trends for both the US and UK; panel (a) shows the time trends in google searches for the

word “diet” and panel (b) shows the trends for searches for “healthy foods”. Both searches, in both countries, show spikes in January and a steady decline as the year progresses.

Figure 2.1: *Google searches*




Notes: Data from Google Trends for the search term “diet” (left hand panel) for the US and UK, collected on 6th January 2017. Data from Google Trends for the search term “healthy food” (right hand panel) for the US and UK, collected on 18th April 2017. Numbers represent search interest relative to the highest point on the chart for the given region and time. A value of 100 is the peak popularity for the term; a score of 50 means the term was half as popular as the peak.

This pattern is reflected in the composition of food products in the shopping baskets of individuals in the UK. We demonstrate this in Figure 2.2, which shows how the share of calories from “healthy foods” varies over time on each day between 2005 and 2012 (where we use food as shorthand for food and non-alcoholic drinks). The graph shows a similar pattern to the Google Trends data – a spike in healthiness in January, followed by some decline, plateauing around the middle of the year, then further decline until the end of the year.

The data that we use in this graph, and in the analysis below, is for a sample of 3,645 single individuals in the UK from the Kantar Worldpanel, each of whom we observe for at least 24 months. The data record all grocery purchases at the transaction level made and brought into the home by these individuals. The data 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. We focus on single-person households to avoid confounding self-control problems with intra-household allocation issues. We describe the data in more detail in Appendix A.1.

Figure 2.2: *Share of calories from healthy foods over 2005–2011*



Notes: We calculate the percentage of calories from healthy foods purchased on each shopping trip for a sample of 3,645 individuals. The definition of healthy foods is given in the text. The figure shows the mean of this variable across individuals on each day over 2005–2011.

We use the nutrient profile score (NPS) to rank foods by their “healthiness”. The NPS is an index developed by nutritionists, and used by the UK Government, that converts the multidimensional nutrient profile of a food product into a single score. A lower score means the product is more healthy.¹ Products with scores below 4 are deemed “healthy” – this threshold is used by the UK Government to restrict advertising of unhealthy products to children.² In this section, we use the government’s cutoff of 4 to class foods as healthy and unhealthy. In Section 4 we allow for the fact that individuals might have heterogeneous views about what counts as healthy by endogenising the classification of foods into the healthy and unhealthy sets. More details on the NPS are provided in Appendix A.2.

Figure 2.2 shows a decline in the calories purchased from healthy foods of over 10%, on average, over the calendar year. This is a sizable change. To provide some intuition for the magnitude of this change, if we considered the average shopping basket, this change would be approximately equivalent to cutting the calories from


¹The NPS gives positive points for saturated fat, sodium, sugar and calories and negative points for fibers, proteins and fruit, vegetables and nuts. The NPS ranges from -15 to 40; 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.

²For drinks the NPS scale is slightly different; we use the government classification that drinks with a score below 1 are categorized as healthy.

chocolate, sweets, and cake in half, and doubling calories from fruit.³ The variation within-person is of a similar magnitude to the variation we see across people (see Appendix A.3).

Figure 2.2 depicts an important aggregate trend in diet quality, with a general decline over the course of the year. However, it masks a great deal of heterogeneity in within-person variation. January is the healthiest month for only around one-third of individuals; the healthiest month for the remaining two-thirds is roughly evenly spread over the rest of the year, with October, November and December being the months that are least often the healthiest. Similarly, roughly one-third of individuals purchase their least healthy grocery basket in December, but the remaining two-thirds buy their least healthy basket in a different month. Individuals also vary considerably in the frequency with which they experience a decline, with some experiencing many dips over the year, and others experiencing only a few.

Figure 2.3: *Share of calories from healthy foods around significant dates*



Notes: We calculate the percentage of calories from healthy foods for each shopping trip. The dotted line shows the mean of this (across people) on each day leading up to and immediately after Easter (left hand panel) and the individual's birthday (right hand panel) for all years 2005-2011.


Easter and birthdays are two other times of the year that are also associated with systematic changes in purchases of healthy foods – see Figure 2.3. In the run up to both Easter and to individuals' own birthdays the share of healthy foods purchased tends to decline, gradually over several days in the former case and more starkly a few days beforehand in the latter case. In both cases the share of

³Chocolate, sweets, and cake together account for an average of 9% of total calories, cutting calories from this source by half would reduce the share of calories from unhealthy foods by around 4.5 percentage points; fruit accounts for around 4% of calories, doubling this would increase the share of calories from healthy foods by around 4 percentage points.

calories from unhealthy products recovers towards the pre-event level immediately following the occasion. These spikes are often contemporaneous with spikes in food expenditure; in exploring what drives fluctuations in diet quality over time it is important to control for price and income effects.

In Figure 2.4 we summarize the overall variation in both the average share of healthy foods purchased across individuals and the within-individual variation over time. Panel (a) shows a histogram of the mean share of the calories that individuals purchase as healthy food. Variation in diet quality across individuals is considerable; 5% of individuals purchase more than 70% of their calories in healthy foods, while at the other extreme 5% purchase less than 35% from healthy foods. On average, people purchase just over half (53%) of calories from healthy foods. This variation is likely to reflect both preference heterogeneity (e.g., it is strongly correlated with individuals' stated preference that they prefer a healthy diet) and differences in the economic environment they face (e.g., differences in prices faced or food budgets). It may also reflect persistent differences in individuals' propensity to yield to temptation.

Figure 2.4: *Heterogeneity in purchases of healthy foods*



Notes: For each person-year-month we calculate the percentage of calories from healthy food. Panel (a) shows the distribution of each person's mean of this variable (calculated across year-months in which they are in the sample). Panel (b) shows the distribution of each person's standard deviation (across year-months) of this variable.

Panel (b) shows a histogram of the within-person standard deviation in the share of calories from healthy food. The average standard deviation is around 10 percentage points; it is less than 6 percentage points for 5% of individuals and more than 14 percentage points for 5% of individuals. Variation in diet quality over time may be driven by temptation and self-control problems. However, it could also reflect rational responses to changes in the economic environment. If the relative

prices of healthy and unhealthy foods varied over time, or if the expenditure elasticity of healthy and unhealthy foods departs from unity and food budgets fluctuate, then we might expect the share of calories from healthy food to vary over time. This motivates the need to develop a model of food choice that we can take to data and that can accommodate the considerable heterogeneity in preferences that are a clear feature of the data.

3 Theory

In this section we present a two-selves model that incorporates rich preference heterogeneity and that allows us to capture the possibility that individuals exhibit self-control problems. In this model, individual choice is driven by the influence of a healthy self and an unhealthy self, that are both characterized by their own stable and well-behaved preferences. We make use of a revealed preference methodology (Samuelson (1938), Afriat (1967) and Varian (1982)) that allows us to investigate whether the observed purchasing behavior can be rationalized by a two-selves model. This method checks consistency of purchasing behavior with the two-selves model for each individual separately, which fully accounts for across individual preference heterogeneity.

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 and drinks. 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 4.1, 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 also individual specific, which does not pose a problem since the application of our model is individual specific. 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. In our empirical application prices vary by region, but for notational simplicity we omit that here. 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 which may not be stable over time. Note that an important feature of our model is that it incorporates elements of non-standard decision making, without abandoning the assumption of rational choice behavior altogether.

Intuitively, a more resolute individual is one whose bargaining process is stable over time. Heterogeneity in preferences (and other characteristics, such as metabolism, and lifestyle choices, such as levels of exercise) might mean that an individual prefers more or less healthy foods. They may regularly consume a diet that gives more or less weight to the healthy self, but this does not change much over time. By contrast, an individual who suffers from self-control problems will, from time to time, be tempted by their unhealthy self, and so experience greater fluctuations in the bargaining process. Note that our approach, which focuses on within-person variation in bargaining, does not rule out that consistently buying an unhealthy diet may reflect systematic self-control problems, however, we cannot separately identify this from the individual simply having strong preferences for unhealthy goods.

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 (see, e.g., Chiappori (1988), Browning and Chiappori (1998) and Chiappori and Ekeland (2009)). The sharing rule distributes the food budget, x_{it} , of individual i 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 . One important question is how the selves' utility functions are aggregated to determine the observed purchase behavior $(\mathbf{q}_{it}^h, \mathbf{q}_{it}^l)$ of individual i . Chiappori (1988) demonstrated that the sharing rule, as defined above, is compatible with the assumption that both selves choose Pareto efficient allocations, in the setting with stable rational preferences over the food bundles \mathbf{q}_{it}^h and \mathbf{q}_{it}^l . The sharing rule is then a direct indication of the bargaining power of both selves.⁴

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 μ_{it} equals one, then the individual behaves according to the healthy self's preferences, while if μ_{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 and drinks 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 rational choice model without any such constraints. In the empirical analysis we show that our two-selves model outperforms the standard model and thus, by implication, the rational choice model with strong separability.

However, 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 expenditures 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 selves' preferences or the budget constraint.

⁴Our 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, Demuyne, and De Rock (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 expenditures in total food expenditure, $\eta_{it} = x_{it}^h/x_{it}$. There is a one-to-one relation between the Pareto weight μ_{it} and the sharing rule. Hence, under some regularity conditions, 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.

To summarize the two-selves model, consider an individual who has the New Year's resolution to eat more healthy food. In the first few months of the year this individual's food purchasing behavior is driven, to a large extent, by the preferences of the healthy self. The relatively high Pareto weight μ_{it} that is attached to the healthy self would be reflected by a high share of the healthy food items in the food budget, as determined by the individual-specific sharing rule $\eta_i(\mathbf{p}_t^h, \mathbf{p}_t^l, x_{it}, \mathbf{z}_{it})$. As the year progresses her resolution deteriorates and she is increasingly tempted, reflected by a change in \mathbf{z}_{it} . This decline in her resolution is associated with, all else equal, a decrease over time of the Pareto weight and an increase of the share of unhealthy food items in the food budget. Of course, variation in prices and budgets might also lead to variation in the share of the budget spent on healthy foods. In the empirical analysis of the model we isolate variation in the sharing rule that is driven by \mathbf{z}_{it} and link this to self-control problems. We first discuss the testable implications of the two-selves model.

3.3 Testable implications

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, De Rock, and Vermeulen (2007, 2011) characterize similar conditions in a revealed preference setting *à la* Afriat (1967) and Varian (1982).

We assume that both selves of individual i have 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 $\mathbf{p}_t^{j'} \mathbf{q}_i^j \leq \mathbf{p}_t^j \mathbf{q}_{it}^j$.*

Cherchye, De Rock, and Vermeulen (2011) show that both sets $S^{ih} = \{(\mathbf{p}_t^h; \mathbf{q}_{it}^h), t = 1, \dots, T_i\}$ and $S^{il} = \{(\mathbf{p}_t^l; \mathbf{q}_{it}^l), t = 1, \dots, T_i\}$ have to 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:*

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

Intuitively, for a given data set S^{ij} , the Afriat Inequalities allow us to obtain an explicit construction of the utility level and the marginal utility of income associated with each observation t : they define a utility level U_t^{ij} and a marginal utility of income λ_t^{ij} (associated with the observed budget x_{it}^j for each observed bundle \mathbf{q}_{it}^j). Note that the Afriat Inequalities are linear in the unknowns U_t^{ij} and λ_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.

Finally, note that this approach can be separately applied to each individual i , so 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.

3.4 Afriat Index

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} or 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. In revealed preference applications this is widely used 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 is equivalent 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. We refer to Choi et al. (2014) for additional discussion on the interpretation of the Afriat Index as a measure for the degree of data consistency with rationalizable behavior.⁵

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 outlays. 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.

3.5 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.

⁵These 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 empirically equivalent.

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 π_{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 ν_{git} . To define each ν_{git} we first draw a random number from the uniform distribution between 0 and 1. If this number is greater than π_{gi} , then we set ν_{git} equal to zero. If it is less than π_{gi} then ν_{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.

4 Empirical implementation

In this section we empirically implement the model described above. We present evidence on how well the two-selves model rationalizes the data. We recover variation in the sharing rule that is not due to changes in the economic environment, and describe how this varies with individuals' characteristics.


4.1 Fit of the model

To implement the methods described in Section 3 we aggregate products to the level of 85 goods based on their nutritional characteristics.⁶ For each good we construct a price index that is a weighted average of the product level region-month specific prices, where the weights reflect the quantity share of products – see Appendix B.1 for details.

As described in Section 3.2, the two-selves model requires us to partition the goods into those purchased by the healthy and unhealthy self. To do this, we make use of the NPS, which summarizes the nutrient profile of a food product in a single continuous measure (see Appendix A.2). This allows us to rank goods by how healthy they are. In Section 2, we used the government’s cutoff of 4 (or 1 for drinks) to group foods into healthy and unhealthy sets. However, 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. 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 Table 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. Figure 4.1 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 that best rationalizes their food purchases.

⁶In our data 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.

Figure 4.1: *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.

For the selected partitioning of the goods, Figure 4.2(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 Section 3.5. 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 4.2(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.

Figure 4.2: *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 3.4. To compute the power of the test we construct Afriat indices for the random draws for each individual, as described in Section 3.5. 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.

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 higher for around two thirds of the individuals. 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). This leads us to conclude that our two-selves model provides a good fit of the data. Note 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. Empirically if we consider this model, we find that the extra structure decreases even further the goodness-of-fit, and thus the Afriat indices, of the single-self model.

4.2 Recovering the sharing rule

The two-selves model gives a structural interpretation to the share of spending on healthy food via the sharing rule; in particular, in our set up the share of spending on healthy foods is a 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}_{tr}^h , total food spending, x_{it} , and other factors that affect the bargaining weight, \mathbf{z}_{it} :

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

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

The function η_i is a consumer-specific nonparametric function and its arguments include 85 prices, the food budget and the vector \mathbf{z}_{it} , which includes potentially unobservable variables that affect the sharing rule. In order to estimate the portion of variation in η_i that is driven by variation in prices and food budgets, and the portion that is due to other factors, we impose the following restrictions. First, we assume that the unobservable information is one dimensional and separable, so that we can write $\eta_{it} = g_i(\mathbf{p}_{rt}^h, \mathbf{p}_{rt}^l, x_{it}) + z_{it}$, where g_i contains the observable part of the sharing rule. Second, we assume that the two price vectors, which together contain 85 separate price series, can be approximated by four price indices: one for very healthy foods (those with an NPS less than 0), Π_{irt}^{vh} ; one for healthy foods (those with an NPS between 0 and 10, but classed by the individual as healthy), Π_{irt}^h ; one for unhealthy foods (those with an NPS between 0 and 10 and classed by the individual as unhealthy), Π_{irt}^l ; and one for very unhealthy foods (those with an NPS greater than 10), Π_{irt}^{vl} . 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 quantity purchased on each of the four sets of foods – very healthy, healthy, unhealthy, very unhealthy) – see Appendix B.1 for further details.

Under these assumptions we get:

$$\eta_{it} = g_i(\Pi_{irt}^{vh}, \Pi_{irt}^h, \Pi_{irt}^l, \Pi_{irt}^{vl}, x_{it}) + z_{it}.$$

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

$$\eta_{it} = \alpha_i + \beta_i^h \ln \left(\widetilde{\frac{\Pi_{irt}^h}{\Pi_{irt}^{vh}}} \right) + \beta_i^l \ln \left(\widetilde{\frac{\Pi_{irt}^l}{\Pi_{irt}^{vh}}} \right) + \beta_i^{vl} \ln \left(\widetilde{\frac{\Pi_{irt}^{vl}}{\Pi_{irt}^{vh}}} \right) + \gamma_i \left(\widetilde{\ln r x_{it}} \right) + z_{it} \quad (4.1)$$

where $\widetilde{\ln(\cdot)}$ denotes that we normalize each variable by subtracting the individual specific mean⁷ and,

$$\ln r x_{it} \equiv \ln x_{it} - [\bar{\omega}_i^{vh} \ln(\Pi_{it}^{vh}) + \bar{\omega}_i^h \ln(\Pi_{it}^h) + \bar{\omega}_i^l \ln(\Pi_{it}^l) + \bar{\omega}_i^{vl} \ln(\Pi_{it}^{vl})],$$

⁷For instance, $\ln \left(\widetilde{\frac{\Pi_{irt}^h}{\Pi_{irt}^{vh}}} \right) = \ln \left(\frac{\Pi_{irt}^h}{\Pi_{irt}^{vh}} \right) - \frac{1}{T_i} \sum_t \ln \left(\frac{\Pi_{irt}^h}{\Pi_{irt}^{vh}} \right)$.

is the log of real food expenditure.⁸ 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 α_i as the individual's preference for healthy food when they face their average food budget and average relative prices.

Estimation

We estimate equation (4.1) for each individual using OLS. (In Section 4.5, we show that the results below are robust with respect to instrumenting log real food expenditure). The R^2 is a measure of the amount of the variation in the sharing rule that is explained by variation in the economic environment (relative prices and food budgets). The mean R^2 across individuals is 0.21 – on average, around 20% of the within-person variation in the sharing rule is explained by variation in the economic environment. For one-quarter of individuals, more than 27% of variation in their sharing rule is explained by economic variables; for one-quarter of individuals, less than 10% is explained by variation in prices and food budgets.

Estimates of the coefficients in equation (4.1) are summarized in Appendix B.2. For each individual, we test the joint significance of the coefficients on relative prices. For around 33% of individuals, the relative price terms are jointly significant at the 5% level, and for around 43% of individuals, the relative price terms are jointly significant at the 10% level. For 25% of individuals, the coefficient on the log real food expenditure term is statistically different from zero; for 80% of these individuals the coefficient is negative, meaning that increases in the food budget lower the bargaining weight of the healthy self, while for the remaining 20%, the coefficient is positive, meaning an increase in the food budget raises the healthy self bargaining weight.

For each individual, we calculate the elasticity of the sharing rule with respect to the four price indices $(\varepsilon_i^{vh}, \varepsilon_i^h, \varepsilon_i^l, \varepsilon_i^{vl})$ and real food expenditure, ε_i^{rx} ; these describe the percentage change in the sharing rule allocation to the healthy self from a 1% increase in each variable. Table 4.1 shows the distribution of these elasticities across individuals. Theory does not give a clear prediction on the sign of these elasticities. At the median, increases in the price of very healthy and healthy foods are associated with increases in the bargaining weight of the healthy self. However, at the median, an increase in the price of unhealthy foods increase the bargaining

⁸This is the log of nominal food expenditure deflated with a food price index (the weights in the index, $(\bar{\omega}_i^{vh}, \bar{\omega}_i^h, \bar{\omega}_i^l, \bar{\omega}_i^{vl})$ are the individual's mean share of spending on each of the four sets of goods).

weight of the healthy self, whereas increases in the price of very unhealthy foods act to reduce the bargaining weight. For the majority of individuals, increases in food budgets lower the bargaining weight of the healthy self. Table 4.1 makes clear, however, that there is a lot of heterogeneity across individuals, and for some people the size of the effects are the opposite of the median effects.

Table 4.1: *Elasticities of the sharing rule with respect to prices and food expenditure*


<i>Elasticity of sharing rule w.r.t:</i>	Mean	Percentile		
		25 th	50 th	75 th
Real food expenditure, ε_i^{rx}	-0.05	-0.15	-0.05	0.04
Price of very healthy foods, ε_i^{vh}	0.20	-0.85	0.12	1.20
Price of healthy foods, ε_i^h	0.35	-1.10	0.39	1.88
Price of unhealthy foods, ε_i^l	0.28	-1.12	0.31	1.73
Price of very unhealthy foods, ε_i^{vl}	-0.61	-2.42	-0.61	1.20

Notes: For each individual we compute the elasticity of the sharing rule with respect to the four price indices and real food expenditure evaluated at the individual's mean sharing rule, $\bar{\eta}_i$. The elasticity of the sharing rule with respect to the price indices for healthy, unhealthy, and very unhealthy goods are given by: $\varepsilon_i^g = \frac{1}{\bar{\eta}_i}(\beta_i^g - \gamma_i \bar{w}_i^g)$, for $g = h, l, vl$. The elasticity of the sharing rule with respect to the price index for very healthy food is given by: $\varepsilon_i^{vh} = -\frac{1}{\bar{\eta}_i}(\beta_i^h + \beta_i^l + \beta_i^{vl} + \gamma_i \bar{w}_i^{vh})$. The elasticity of the sharing rule with respect to the food budget is given by: $\varepsilon_i^{rx} = \frac{\gamma_i}{\bar{\eta}_i}$. The table shows the mean and 25, 50, 75 percentiles of the distribution across individuals.

4.3 Variation in preferences for healthy food

The parameter estimate $\hat{\alpha}_i$ is consumer i 's "mean" sharing rule, evaluated at their mean real food expenditure and when the prices of the healthy, unhealthy and very unhealthy goods (relative to the very healthy goods) are equal to the average that the individual faces. It therefore captures the consumer's average propensity to choose healthy relative to unhealthy foods. This in part reflects preferences, but may also reflect the consumer's average tendency to succumb to temptation. The median value of $\hat{\alpha}_i$ is 52%, the 25th percentile is 40% and the 75th percentile is 61%. Figure 4.3 shows the distribution across individuals.


Figure 4.3: *Distribution of $\hat{\alpha}_i$ across individuals*



Notes: The figure shows the histogram of $\hat{\alpha}_i$ from estimating equation 4.1.

Figure 4.4(a) shows the mean sharing rule for different age groups. Older people spend a greater share of their food budget on healthy foods than younger people. Socioeconomic differences in health are well documented. For instance, Cutler and Lleras-Muney (2010) show that obesity rates are half as high among the better educated. They estimate that income, health insurance and family background can explain about 30% of the relationship between education and health behaviors. We use the average grocery expenditure of each individual – i.e. expenditure on all fast moving consumer goods (food, alcohol, household supplies, toiletries etc.) – over the course of their time in the sample (at least two years) as a proxy for income. In the Appendix B.3 we show that this is strongly correlated with income using data from the Living Costs and Food Survey. Henceforth, we refer to the distribution of average grocery expenditure as the income distribution. Figure 4.4(b) shows the mean sharing rule for quintiles of this distribution. People in the top quintile have a sharing rule that implies they spend 2 percentage points more of their food budget on healthy foods than people in the bottom quintile. In terms of magnitude, this difference would be roughly equivalent to reducing by half spending on chocolate and confectionery (which has an average expenditure share of 4.6%) and doubling spending on dark green vegetables (which has an average expenditure share of 1.1%).

Figure 4.4: Variation in the mean sharing rule, $\hat{\alpha}_i$, with age and income



Notes: Age group is defined based on the individual’s median age while in the sample. For each individual we calculate their mean total spending on fast moving consumer goods (food, alcohol, household supplies, toiletries etc.) across the period they are in the sample. We use this as a measure of individual income. Quintiles are based on the distribution of this variable across individuals. 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.

In Appendix B.3 we summarize differences in the mean estimated sharing rule across some other demographic groups. 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 than non-smokers. We also show that individuals' preferences for healthy food accord with their stated preferences: the mean sharing rule is increasing in stated preferences for healthy food and declining in stated preferences for processed food. These correlations are reassuring. In addition, we find that individuals with lower stated self-control have a lower mean sharing rule; this suggests differences in the mean sharing rule may partly reflect permanent differences in self-control across individuals.

4.4 Variation in residual sharing rule

In our model, changes in the economic variables (prices and food budgets) have a twofold effect: (i) the standard effect on demand (captured by the within self utility functions) and (ii) an impact on the bargaining process, which implies a change in the respective shares going to healthy and unhealthy foods and drinks. In Section 4.2 we summarize the impact of economic variables on the bargaining process. In what follows, we focus on the residual sharing rule z_{it} . Recall that z_{it} measures an individual's deviation from the mean sharing rule after accounting for the impact of variation in prices and food budgets. In Section 4.5 we show robustness of our results to removing seasonality in preferences, along with a number of other possible concerns.


Variation over time

For each individual, the residual sharing rule, z_{it} , 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 or food expenditures. In Figure 4.5 we summarize the average (across individuals) variation in \hat{z}_{it} 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 expenditure allocated to healthy foods, η_{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 economic variables (prices and food budgets) and unobservables (z_{it}).

The figure shows that, on average, the residual sharing rule declines from January to March, increases from March to July, and then deteriorates from July onwards. This pattern is similar to the share of expenditure on healthy goods (which reflects both variation driven by the observed and unobserved parts of the sharing rule). The decline in the last quarter of the year is larger for the share

of spending on healthy food than for the residual sharing rule. This means that individuals' responses to the changes they experience in the economic environment (prices and food budgets) accounts for some of the decline in diet quality over the year, but not all of it. Indeed, more than half of the decline is accounted for by variation in the unobserved part of the sharing rule. Individuals see the bargaining weight of their unhealthy self rise over the second half of the year, for reasons unrelated to variation in prices or food budgets.

Figure 4.5: *Deviations from the mean sharing rule over the calendar year*



Notes: Black dots show the mean sharing rule, η_{it} , in each calendar month, pooled across years. Red triangles show the mean residual sharing rule, \hat{z}_{it} , which removes the effect of variation in prices and food budgets, 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.


Variation across individuals

Figure 4.5 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. We construct the standard deviation of the residual sharing rule for each individual i over t : $\hat{\sigma}_i = \text{sd}(\hat{z}_{it})$. This measures variation in the individual's sharing rule that is not driven by changes in prices and food budgets, which we interpret as indicating self-control problems. In Appendix B.4 we show that the size of the deviations in the residual sharing rule are

correlated with individuals' stated attitudes. Individuals who have lower stated self-control (i.e. state that they spend money without thinking, or spend more on their credit card than they should) have larger fluctuations in their 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.

We also construct the standard deviation of the sharing rule for each individual i over t : $\tilde{\sigma}_i = \text{sd}(\eta_{it})$. This measures the total variation in the individual's sharing rule, driven by both changes in the economic environment and other factors. Figure 4.6 compares the distributions of $\tilde{\sigma}_i$ and $\hat{\sigma}_i$ across individuals. The distribution of $\hat{\sigma}_i$ is to the left of the distribution of $\tilde{\sigma}_i$. The shift shows the extent to which accounting for individuals' responses to variation in prices and their food budgets 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 7.9 percentage points.

Figure 4.6: *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 and food budgets.

Figure 4.7(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. This can be interpreted as the extent of variation explained by the economic environment, and the size of the error from interpreting the entire sharing rule (without removing the effect of the economic environment) as a measure of self-control problems. Together these graphs show two things. First, there is an age gradient in our measure of self-control. Young people, on average, have more unexplained variation in the bargaining weight of their healthy self, indicating they suffer more from self-control problems than older people. This is consistent with previous findings (see Ameriks et al. (2007), Bucciol (2012), Dohmen et al. (2017)). Second, failing to account for the effects of the economic environment would lead to an over-statement of this gradient. For individuals aged below 40, variation in the economic environment is, on average, responsible for 1.25 percentage points of the standard deviation in their sharing rule; for individuals aged over 70, 0.95 percentage points is explained by responses to the economic environment.


Figure 4.8 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. Part of this is driven by the residual sharing rule (panel (a)) and is therefore reflective of self-control problems. Part of it is driven by responses to the economic environment (panel (b)), which must be properly controlled for to avoid over-stating the gradient in self-control problems.

Despite this, it is still 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 larger than the cross-sectional difference in the average sharing rule between the bottom to top quintile (which is roughly 2 percentage points – see Figure 4.4(b)). This 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 lifecycle are less likely to smooth their food consumption over their Social Security pay periods.

Figure 4.7: *Variation in $\hat{\sigma}_i$ with age*

(a) Standard deviation of residual sharing rule




(b) Effect of responses to changes in prices and food expenditure




Notes: Age group is defined based on the individual's median age while in the sample. The top panel shows the mean standard deviation in the residual sharing rule across individuals in each age group. 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 bottom panel shows the mean of this difference across individuals in age group. 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 4.8: *Variation in $\hat{\sigma}_i$ with income*

(a) Standard deviation of residual sharing rule



(b) Effect of responses to changes in prices and food budgets



Notes: For each individual we calculate their 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. Quintiles are based on the distribution of this variable across individuals. The top panel shows the mean standard deviation in the residual sharing rule across individuals in each quintile. 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 bottom panel shows the mean of this difference across individuals in each quintile. 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' response to changes in their economic environment leads to an overestimate of the gradient of self-control 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 indi-

viduals 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%). This illustrates the importance of distinguishing between cross-sectional and within-individual comparisons: Figure 4.4(b) shows that higher income individuals, on average, allocate a higher share of their food budgets to healthy foods, but our within-individual estimates show that, on average, increases in food budgets *reduce* the bargaining weight of the healthy self. These patterns are consistent with results in the literature that there is a socioeconomic gradient in diet quality and that there are reductions in diet quality during periods when individuals experience positive income shocks (e.g. Ruhm (2000)).

4.5 Robustness

We check the robustness of our results to three things: (i) seasonal variation in preferences, (ii) instrumenting food expenditure in equation (4.1), and (iii) using only the subset of individuals for whom the two-selves model rationalizes their observed behavior better than the single-self model.

Seasonal preferences

In Section 4.4 we interpret residual variation in the sharing rule as reflecting self-control. However, it may be the case that individuals have seasonal variation in their preferences due to other time varying factors that do not represent or affect self-control, for instance, the weather. We investigate the robustness of our results to this possibility by estimating an alternative specification of equation 4.1, in which we allow for quarterly effects, τ_{it} . In particular, we define $z_{it} = \tau_{it} + z_{it}^T$ and construct $\widehat{\sigma}_i^T = \text{sd}(\widehat{z}_{it}^T)$ as the residual sharing rule (net of seasonal effects).

We repeat our analysis in Section 4.4 with $\widehat{\sigma}_i^T$ – see Appendix B.5. Removing quarterly variation in preferences reduces the standard deviation of the residual sharing rule, but only by a relatively small amount – the mean across individuals falls from 7.9 to 7.4. This means that there is still substantial variation in the sharing rule over time that is unaccounted for by either responses to fluctuations in prices and food budgets, and any seasonality (at the quarterly level) in preferences.

The age and income gradients in the residual sharing rule remain, after taking out seasonal effects.⁹

Instrumenting food expenditure

In Section 4.4 we partial out the effect of variation in food prices and budgets. However, individuals' responses to the changing economic environment may partly reflect lapses in self-control. For instance, for some individuals, times in which they spend a lot on food may also be times in which they are more inclined to succumb to temptation (which implies a low value of z_{it}). 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 to be uncorrelated with individuals' self-control problems with respect to unhealthy foods. Our instrument set includes a set of prices from the consumer price index (CPI). These 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). The instruments also include individual monthly spending on non-food items (cleaning products, toiletries, cosmetics). 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 extent of their propensity to succumb to the temptation associated with unhealthy foods.¹⁰

We present the results from this alternative IV approach in Appendix B.5. 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 the IV. This is consistent with the idea that individuals might spend more when they are more tempted. However, the relationships with age and income, and the importance of the economic environment, are qualitatively similar.

⁹We see similar results when we take out common month effects.

¹⁰Pooling 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.

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. The two-selves model performs better for individuals who are less likely to state that they make shopping lists or commit to buying the same things.

In our main results, above, we study variation in the sharing rule for all individuals. We check whether our results are robust to using only the two-thirds of individuals for whom the two-selves model better rationalizes the data than the single-self model. We present these results in Appendix B.5. When we look just at the subset of individuals for whom the two-self model fits better, we continue to see larger fluctuations in the sharing rule for younger and lower income individuals.

5 Summary and concluding comments

We propose a two-selves model in which food choices are a compromise between a healthy and an unhealthy self. We use data on a sample of single person households and nonparametric revealed preference methods to show that the model is consistent with the data. This model allows us to put a structural interpretation on the data; specifically, the sharing rule, or share of spending on healthy foods, can be interpreted as the bargaining weight of the healthy self. Within-person variation in the sharing rule (reflecting variation in diet quality) is significant and of a similar magnitude to variation across individuals.

We show that these fluctuations are not fully explained by the responses to the economic environment. We recover the “residual” sharing rule, which removes individuals' responses to changes in the economic environment, and interpret within-person variation in the residual sharing rule as indicative of self-control problems. We find that the extent of self-control problems is larger for younger and lower income individuals. If we fail to take into account how individuals respond to changes in the economic environment we would overestimate both the extent of self-control and the gradient of self-control with age and income.

The existence of self-control problems suggests that there may be scope for public policy to improve welfare in this setting. However, we are cautious about

making statements pertaining to welfare using our current framework. As discussed in the introduction, many of the existing multi-selves models deliver observationally equivalent predictions about choice behavior, which are captured by our revealed preference conditions, but may have substantially different welfare implications. Given our focus in the current study, we deliberately remain agnostic about the specific interaction between the selves (except from assuming Pareto efficient outcomes). But our findings do suggest that a more detailed investigation of the particular interaction mechanisms at play constitutes a potentially important avenue for follow-up research.

Finally, our study opens up a new application of the collective model of choice behavior, which has become the workhorse model in the household economics literature. Our results show that individuals' food purchase behavior can be modeled as the outcome of a collective interaction process between two selves. This shows that the large and well-established toolkit of empirical methods to analyze collective household behavior can also be used to analyse individuals' choice behavior in terms of multiple selves.

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.
- Ambrus, A. and K. Rozen (2015). Rationalising Choice with Multi-self Models. *Economic Journal* 125(585), 1136–1156.
- 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.
- 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.
- Becker, G. S. (1962). Irrational Behavior and Economic Theory. *Journal of Political Economy* 70(1), 1–13.
- Bernheim, B. D. (2009). Behavioral Welfare Economics. *Journal of the European Economic Association* 7(2-3), 267–319.
- Bernheim, B. D., J. Meer, and N. K. Navarro (2016). Do Consumers Exploit Commitment Opportunities? Evidence from Natural Experiments Involving

- Liquor Consumption. *American Economic Journal: Economic Policy* 8(4), 41–69.
- 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.
- Bronars, S. G. (1987). The Power of Nonparametric Tests of Preference Maximization. *Econometrica* 55(3), 693–698.
- 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.
- Cawley, J. (2000). Body Weight and Women’s Labor Market Outcomes. Working Paper 7841, National Bureau of Economic Research.
- 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. Demuyne, 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.
- 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.
- de Clippel, G. and K. Eliaz (2012). Reason-based choice: A bargaining rationale for the attraction and compromise effects. *Theoretical Economics* 7(1), 125–162.

- 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.
- Dohmen, T., A. Falk, B. H. H. Golsteyn, D. Huffman, and U. Sunde (2017, October). Risk Attitudes Across The Life Course. *The Economic Journal* 127(605), F95–F116.
- Dolan, P., M. Hallsworth, D. Halpern, D. King, R. Metcalfe, and I. Vlaev (2012). Influencing behaviour: The mindspace way. *Journal of Economic Psychology* 33(1), 264–277.
- Downs, J. S., G. Loewenstein, J. Wisdom, and B. J. S. Downs (2009). Strategies for Promoting Healthier Food Choices. *American Economic Review: Papers & Proceedings* 99(2), 159–164.
- 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. 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.
- Garon, J.-D., A. Masse, and P.-C. Michaud (2015, November). Health club attendance, expectations and self-control. *Journal of Economic Behavior & Organization* 119(Supplement C), 364–374.
- 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.
- Giné, X., D. Karlan, and J. Zinman (2010). Put Your Money Where Your Butt Is: A Commitment Contract for Smoking Cessation. *American Economic Journal: Applied Economics* 2(4), 213–235.
- Gruber, J. and B. Koszegi (2004). Tax incidence when individuals are time-inconsistent: The case of cigarette excise taxes. *Journal of Public Economics* 88(9-10), 1959–1987.
- 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.
- Hagenaars, A. J., K. De Vos, M. Asghar Zaidi, and others (1994). Poverty statistics in the late 1980s: Research based on micro-data.
- Haushofer, J. and E. Fehr (2014). On the psychology of poverty. *Science* 344(6186), 862–867.

- Kalai, G., A. Rubinstein, and R. Spiegel (2002). Rationalizing choice functions by multiple rationales. *Econometrica* 70(6), 2481–2488.
- Mani, A., S. Mullainathan, E. Shafir, and J. Zhao (2013). Poverty Impedes Cognitive Function. *Science* 341(6149), 976–980.
- Manzini, P. and M. Mariotti (2007). Sequentially rationalizable choice. *American Economic Review* 97(5), 1824–1839.
- Manzini, P. and M. Mariotti (2015). State dependent choice. *Social Choice and Welfare* 45(2), 239–268.
- Manzini, P. and M. Mariotti (2016). Dual Random Utility Maximisation. Discussion Paper 1605, University of St Andrews.
- 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.
- O’Donoghue, T. and M. Rabin (2000, April). The economics of immediate gratification. *Journal of Behavioral Decision Making* 13(2), 233–250.
- O’Donoghue, T. and M. Rabin (2003). Studying optimal paternalism, illustrated by a model of sin taxes. *American Economic Review* 93(2), 186–191.
- O’Donoghue, T. and M. Rabin (2006). Optimal sin taxes. *Journal of Public Economics* 90(10-11), 1825–1849.
- 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.
- 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. 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.
- Ruhm, C. J. (2000). Are recessions good for your health? *Quarterly Journal of Economics* 115(2), 617–650.
- Samuelson, P. A. (1938). A Note on the Pure Theory of Consumer’s Behaviour. *Economica* 5(17), 61–71.
- Spiegler, R. (2011). *Bounded Rationality and Industrial Organization*. Oxford University Press.
- Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies* 23(3), 165–180.
- Varian, H. R. (1982). The nonparametric approach to demand analysis. *Econometrica* 50(4), 945–973.

APPENDIX

A new year, a new you? Heterogeneity and self-control in food purchases

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

December 8, 2017

A Data and measurement

A.1 Kantar Worldpanel

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. For each transaction, the data contain information on exact transaction prices, where the product was purchased, whether it was on promotion, and the nutritional composition of the product. The data also contain demographic information on the panel of households.

We focus on a sample of 3,645 single individuals over the period 2005 to 2011. Table A.1 compares selected demographics in our data with single individuals in the nationally representative Living Costs and Food Survey. 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 temptation 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.1: *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 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), 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.2: *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.2. 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.

A.3 Within and between variance

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 A.3 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 cross sectional variation. In the main analysis in the paper we use only the monthly aggregation.

Table A.3: *Variation in diet quality*

<i>% calories from healthy food</i>	(1)	(2)
	Purchases aggregated to:	
	Week	Month
Mean	51.02	48.34
Standard deviation	22.64	13.77
within people	20.50	10.00
between person	9.60	9.46

Notes: In column (1) we calculate the % of calories from healthy foods across shopping trips in each person-year-week, in column (2) we calculate the % calories from healthy food across shopping trips in each person-year-month. The table shows the decomposition of the standard deviation of these variables into the within individual versus between-person variation.

B Empirical implementation

B.1 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 .

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 . 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 are 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 are given by

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

In Table B.1 we describe the price variation and expenditure shares of each good.

Very healthy, healthy, unhealthy and very unhealthy food price indices

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

$$\Pi_{it}^{vh} = \sum_{g \in VH_i} w_{ig} P_{grt} \quad (\text{B.2})$$

$$\Pi_{it}^h = \sum_{g \in H_i} w_{ig} P_{grt} \quad (\text{B.3})$$

$$\Pi_{it}^l = \sum_{g \in L_i} w_{ig} P_{grt} \quad (\text{B.4})$$

$$\Pi_{it}^{vl} = \sum_{g \in VL_i} w_{ig} P_{grt} \quad (\text{B.5})$$

Table B.1: *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 are 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.


where VH_i, H_i, L_i, VL_i are the sets of goods in the very healthy, healthy, unhealthy and very unhealthy sets, respectively, and:

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

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

B.2 Coefficient estimates

Figure B.1: *Joint significance of price coefficients*



Notes: The figure shows the distribution of p-values for tests of the joint significance of the coefficients on the relative prices of healthy, unhealthy, and very unhealthy (relative to the price of very healthy foods).

Figure B.2: *Distributions of estimated coefficients on:*




Notes: The figures show the distribution of estimated coefficients in equation 4.1. The lighter shaded areas represent parts of the distribution for which the estimated coefficients are not significantly different from zero (at the 5% level), the darker areas show the the coefficients that are significantly different from zero.

B.3 Variation in mean sharing rule

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 B.3 shows that there is a strong relationship between households' annual equivalized income and equivalized weekly grocery spending.

Figure B.3: *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)).

Relationship between mean sharing rule and demographics

Table B.2: *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]
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.

Relationship between mean sharing rule and stated preferences

Kantar Wordpanel 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 lack of self-control. The questions we use are listed in Table B.3. Questions change from year to year; no questions were asked in 2008 and 2011. When an individual answers the same question across years, we take his/her median response. We use 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 lack of self-control.

Table B.4 shows the relationship between individuals' mean sharing rule and whether they have above or below median (across individuals) preferences for the five factors.

Table B.4: *Variation in mean sharing rule by stated preferences*

(1)	(2) Mean	(3) Difference	(4) 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]
Above median tendency to buy on promotion	50.3		
Below median tendency to buy on promotion	50.9	0.6	[-0.4, 1.6]
Above median shopping commitment	50.8		
Below median shopping commitment	50.3	-0.5	[-1.5, 0.5]
Above median stated self-control	52.4		
Below median stated self-control	50.0	-2.4	[-3.7, -1.0]

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.

Table B.3: 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>Self-control</i>						
I stick to a set routine for doing the housework	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	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 scale: 1 "Strongly disagree", 2 "Disagree", 3 "Neither agree nor disagree (or missing)", 4 "Agree", 5 "Strongly Agree". † for these questions we reverse the ordering of the answer scale so 1 is "Strongly agree" etc.

B.4 Variation in self-control

Relationship between self-control and demographics

Table B.5: *Variation in $\hat{\sigma}_i$ by demographic characteristics*

	(1) Mean	(2) Difference	(3) 95% CI for diff.
Female	8.07		
Male	8.16	0.09	[-0.06, 0.24]
Age less than 40	8.87		
Age 40-65	8.15	-0.72	[-0.95, -0.49]
Age over 65	7.79	-1.08	[-1.31, -0.84]
High skilled	8.23		
Medium skilled	8.09	-0.14	[-0.37, 0.09]
Low skilled	8.08	-0.16	[-0.39, 0.08]
Full time work	8.40		
Part time work	7.92	-0.48	[-0.76, -0.19]
Not working	8.31	-0.09	[-0.34, 0.17]
Retired	7.80	-0.60	[-0.78, -0.42]
Non smoker	8.07		
Smoker	8.21	0.14	[-0.04, 0.31]
Vegetarian	8.25		
Non vegetarian	8.09	-0.16	[-0.39, 0.07]
Above median income	7.07		
Below median income	8.49	1.43	[1.29, 1.56]

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.

Relationship between self-control and stated preferences

Table B.6 summarize the relationships between variation in the residual sharing rule and individuals' stated preferences. These are intuitive. Individuals that have a stated preference for a healthy diet, a stated preference for buying the same things, and stated higher self-control tend to exhibit less variation, while those that have a preferences for processed food and state that they typically buy things on promotion tend to have more variation.

Table B.6: *Variation in $\hat{\sigma}_i$ by stated preferences*

	(1) Mean	(2) Difference	(3) 95% CI for diff.
Above median preferences for healthy food	8.03		
Below median preferences for healthy food	8.22	0.19	[0.04, 0.34]
Above median preferences for processed food	8.27		
Below median preferences for processed food	7.96	-0.32	[-0.47, -0.16]
Above median tendency to buy on promotion	8.38		
Below median tendency to buy on promotion	7.82	-0.56	[-0.71, -0.41]
Above median shopping commitment	7.94		
Below median shopping commitment	8.30	0.36	[0.21, 0.51]
Above median stated self-control	7.53		
Below median stated self-control	7.99	0.46	[0.25, 0.66]


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.

B.5 Robustness

Seasonal preferences

In order to allow for the fact that individuals may have seasonal preferences that are driven by factors unrelated to self-control, we estimate a variant of equation 4.1 in which we allow for quarterly effects. Figure B.4 shows that there is still substantial variation in the residual sharing rule, even after taking out quarterly effects. Figures B.5 and B.6 are analogous to Figures 4.7 and 4.8. They show that the age and income gradients remain, even after controlling for seasonal preferences.


Figure B.4: *Standard deviation of the sharing rule and the residual sharing rule (taking out heterogeneous quarterly effects)*




Notes: The white bars show the distribution of individuals' standard deviation in the sharing rule, σ_i . The grey bars show the distribution of individuals' standard deviation in the residual sharing rule, $\hat{\sigma}_i^r$, which takes out the responses to changes in prices and total food budgets and quarterly variation in preferences.

Figure B.5: Variation in $\hat{\sigma}_i^T$ with age

(a) Standard deviation of residual sharing rule




(b) Effect of responses to changes in prices and food expenditure




Notes: Age group is defined based on the individual's median age while in the sample. The top panel shows the mean standard deviation in the residual sharing rule across individuals in age group (taking out seasonal effects). For each individual we calculate the difference in the standard deviation in the sharing rule and the standard deviation in the residual sharing rule, $\sigma_{\tilde{\sigma}_i} - \hat{\sigma}_i^T$; the bottom panel shows the mean of this difference across individuals in age group. 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 B.6: *Variation in $\hat{\sigma}_i^r$ with income*

(a) Standard deviation of residual sharing rule



(b) Effect of responses to changes in prices and food expenditure




Notes: For each individual we calculate their 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. Quintiles are based on the distribution of this variable across individuals. The top panel shows the mean standard deviation in the residual sharing rule across individuals in each quintile (taking out seasonal effects). For each individual we calculate the difference in the standard deviation in the sharing rule and the standard deviation in the residual sharing rule, $\sigma_i - \hat{\sigma}_i^r$; the bottom panel shows the mean of this difference across individuals in each quintile. 95% confidence intervals are shown. Standard errors reflect uncertainty over the sample of individuals, but not over the time series for each individual.

Instrumenting food expenditure

We estimate equation (4.1) using an IV approach, in which we instrument for the expenditure term using variables that are likely to drive total spending on food,

while being uncorrelated with individuals' self-control problems with respect to unhealthy foods. We use $\hat{\sigma}_i^{IV}$ to denote individual i 's estimated standard deviation in their residual sharing rule based on these estimates. More details of this are provided in the main paper.


Figure B.7: *Standard deviation of the sharing rule and the residual sharing rule (based on IV estimates)*




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, $\hat{\sigma}_i^{IV}$, which takes out the responses to changes in prices and total food budgets.

Figure B.8: Variation in $\hat{\sigma}_i^{IV}$ with age

(a) Standard deviation of residual sharing rule




(b) Effect of responses to changes in prices and food expenditure




Notes: Age group is defined based on the individual's median age while in the sample. The top panel shows the mean standard deviation in the residual sharing rule across individuals in age group. For each individual we calculate the difference in the standard deviation in the sharing rule and the standard deviation in the residual sharing rule, $\tilde{\sigma}_i - \hat{\sigma}_i^{IV}$; the bottom panel shows the mean of this difference across individuals in age group. 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 B.9: Variation in $\hat{\sigma}_i^{IV}$ with income

(a) Standard deviation of residual sharing rule



(b) Effect of responses to changes in prices and food expenditure




Notes: For each individual we calculate their 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. Quintiles are based on the distribution of this variable across individuals. The top panel shows the mean standard deviation in the residual sharing rule across individuals in each quintile. For each individual we calculate the difference in the standard deviation in the sharing rule and the standard deviation in the residual sharing rule, $\tilde{\sigma}_i - \hat{\sigma}_i^{IV}$; the bottom panel shows the mean of this difference across individuals in each quintile. 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.

Individuals for whom the two-self model fits better

We repeat our analysis in Section 4.4 using only the subset of individuals (2470 out of a total 3645) for whom the two-self model fits better than the single-self model.


Figure B.10: *Standard deviation of the sharing rule and the residual sharing rule (based subset of 2470 individuals)*




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 and total food budgets. The distributions are drawn for the subset of individuals (2470) for whom the two-self model fits better than the single-self model.

Figure B.11: Variation in $\hat{\sigma}_i$ with age (based subset of 2470 individuals)

(a) Standard deviation of residual sharing rule




(b) Effect of responses to changes in prices and food expenditure




Notes: Age group is defined based on the individual's median age while in the sample. The top panel shows the mean standard deviation in the residual sharing rule across individuals in age group. 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 bottom panel shows the mean of this difference across individuals in age group. The estimates are for the subset of individuals (2470) for whom the two-self model fits better than the single-self model. 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 B.12: Variation in $\hat{\sigma}_i$ with income (based subset of 2470 individuals)

(a) Standard deviation of residual sharing rule



(b) Effect of responses to changes in prices and food expenditure



Notes: For each individual we calculate their 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. Quintiles are based on the distribution of this variable across individuals. The top panel shows the mean standard deviation in the residual sharing rule across individuals in each quintile. 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 bottom panel shows the mean of this difference across individuals in each quintile. The estimates are for the subset of individuals (2470) for whom the two-self model fits better than the single-self model. 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.

