Income and Consumption over the Business Cycle: Evidence from Matched Administrative Data*

Calogero Brancatelli[1] Roman Inderst[2]

This version: September 2021

Abstract

This paper revisits the effects of income changes on consumption of private households by focusing on a commonly disregarded and yet sizeable component of household expenditures: consumption of food and non-food consumer packaged goods. We exploit a new data source from the Netherlands that combines on the level of individual households administrative data from tax records with household scanner data, thus minimizing measurement error for both expenditures and the key explanatory variable, household disposable income. Even after controlling for differences in needs and for consumption volume, we document significant variation in expenditures and thereby reveal substantial scope for potential savings. Still, even though the Netherlands experienced a recession and a subsequent recovery in the analysed period from 2011 to 2018, we find only an economically small relationship with income, which is also not higher for households with low income or low liquidity. Despite remaining small in magnitude, we document inter alia a much higher coefficient for single households. We can exclude various potentially confounding effects as we show that retailers practice national pricing and as we control for sample composition and potential substitution between in-house and out-of-house consumption.

Keywords: income effects, consumer-packaged goods, administrative data
JEL Classification: D12, E21, E32, M30

---

* This paper is the result of a collaboration with the AiMark foundation and Centraal Bureau voor de Statistiek (CBS) Netherlands. We thank both for their continuous support. We are also grateful for comments received from Thomas Otter and Adrian Fritzsche. The authors benefited from financial support from a long-term research grant by the Think Forward Initiative (TFI). During the study implementation, Brancatelli worked for Goethe University Frankfurt and is now working for the European Central Bank. The opinions expressed in this paper are our own and do not necessarily reflect the views of the European Central Bank.

[1] Goethe University Frankfurt and European Central Bank, brancatelli@econ.uni-frankfurt.de

[2] Goethe University Frankfurt, inderst@finance.uni-frankfurt.de

I Introduction

Consumption of (fast-moving) consumer packaged goods (CPGs) represents a sizable fraction of all but the richest households' consumption and disposable income.[3] With the growing importance of national branded products in the middle of the last century, market research companies have developed panels in which consumers can record all of their respective purchases. With regard to CPGs, the quality of information available to the industry through such sources exceeds that available from public statistical sources.[4] Brand manufactures as well as marketing scholars and industrial economists typically make use of such information to understand substitution patterns within a given product category. It is only more recently that scholars have been exploiting these particularly rich and accurate data to answer questions concerning the aggregate economy, as we discuss below.

In the aggregate, an immediate question is how total expenditures for such CPGs vary with income, notably over the business cycle. Answering this question poses various data challenges that we address by matching household scanner data with administrative income data, by drawing on variation from the Netherlands' double-dip recession following the financial crisis, as well as by generating controls for out-of-house-consumption and for changes in supply. Our key finding is that there is only an economically very small relationship between household income and this part of consumption expenditures. In our key panel regressions, we find that a 10% increase in disposable household income is related to only a 0.17% increase in CPG consumption. While some (staple) food products may be considered necessities, our finding seems still surprising given the large differences in actual expenditures across households, which prevail even after controlling for households' consumption needs and thus suggest sizeable savings possibilities. Coefficients remain economically small even when we include lags, thereby possibly capturing longer-term changes to income, or when we restrict consideration to households with low income or low overall financial wealth (liquidity). However, while remaining small, the key coefficient becomes five times larger when considering single households only.

Our results are notably not downward biased by measurement errors as we match households' true expenditures from GfK's household scanner panel[5] with total disposable income from administrative data, compiled inter alia for tax purposes and collected by the Dutch statistical office (Centraal Bureau voor de Statistiek, CBS). Such matching was only possible after obtaining individual consent from over 5,000 households. Thus, our data allow us to measure both expenditures and income on the level of the overall household, rather than individual household members (such as the main earner). Through CBS we have also access to income data for all Dutch households, which allows us to make our analysis representative

[3] The unweighted average over our full sample of matched household-year observations amounts to 9.95%.

[4] In Germany, the leading marketing company, GfK, has initiated in 1957 its panel of then 1.000 households in direct cooperation with one of the leading brand manufacturers.

[5] Researchers from notably the U.S. may be more familiar with Nielsen's household (or Homescan) panel data. In many other countries, the panels provided by GfK or the sister organization Kantar furnish the same service. Notably in the Netherlands, Nielsen does not have a household panel, so that GfK's panel represents the industry standard.

(using respective weights). Such matching of household scanner data with administrative (income) data constitutes a major data innovation and is shared with Brancatelli, Fritzsche, Inderst, and Otter (2020), where the focus lies, however, on the substitution between national brands and private-label products.

As CPG consumption consists mainly of food products, we need to consider the possibility of substitution with out-of-house consumption. For instance, higher income may induce households to consume more food in restaurants, which could, in principle, even lead to a negative association between CPG expenditures and income. For this reason, we construct and include variables that capture the nutritional intake of households' food expenditures. Still, the measured coefficient remains economically small. In contrast, the measured coefficient could overstate households' responsiveness if, for instance, reduced expenditures during the downturn were largely due to a reduction in prices or to an expansion of cheaper stores. While such concerns are mitigated by national pricing and assortment strategies of retailers in the Netherlands, which we confirm empirically, and the relative uniformity of retailers' presence, we also include supply-side controls using data on store openings and including a measure of changes in district income, to which retailers may respond (e.g., with such openings). We also conduct all analyses in (first) differences, thereby reducing biases from spurious correlations. These controls have, however, no significant impact on the anyway economically small coefficient. Its small size is also not due to a lack of variation in income, as our considered time period from 2011 to 2018 covers the "second dip" (2011 to 2013) of the Netherlands' double-dip recession following the financial crisis as well as the subsequent recovery. Over this time period, GDP growth dropped by 3.6 percentage points from the fourth quarter in 2011 to the third quarter in 2013 and unemployment rose by 3.3 percentage points. Over the considered time period, one third of income changes were negative.

For countries with a similar socioeconomic background as the Netherlands, differences in household (per capita) CPG expenditures may thus be largely a result of differences in taste and, consequently, one may not expect households to significantly adjust these expenditures when facing changes in income.[6] In other words, expenditures for CPGs thus represent a particularly stable component of total consumption.[7]

[6] Of course, there are also other reasons for why other components of consumption, such as durable (or slow-moving consumer) goods, may respond more strongly, e.g., as such purchases can be delayed. Parker, Souleles, Johnson, and McClelland (2013) also document very small effects for food and non-durable items in response to a windfall income shock. In marketing research, Dutt and Padmanabhan (2011), using aggregate data in a cross-country analysis of currency crises, show how consumer reaction in slow-moving goods categories is indeed much stronger.

[7] However, we would not claim that these findings also apply to countries with different institutional backgrounds. In fact, in many aspects, the Netherlands are comparable to other Continental European countries, but less so to the US or UK. This applies to the mitigation of changes in income and employment during recessions, as well as to a more egalitarian distribution of income. For example, looking at the World Bank Development Indicators from 2016, the GINI coefficient of the Netherlands, Germany and France were 28.2, 31.9 and 31.9 respectively. In comparison, the Gini coefficient of the UK in that period was 34.8 and the US coefficient was 41.4, indicating substantially higher concentration of income in the population of those countries.

Changes in households' total consumption are often measured by surveys, based on diaries or interviews[8], rather than by recording actual expenditures, as we do in this paper, albeit only for one component of total consumption. We acknowledge that in such surveys also hypothetical questions can be asked, which allow disentangling responses to anticipated vs. non-anticipated as well as temporary vs. permanent changes in income.[9] Although the fact that observed income changes relate to the business cycle swing should mitigate concerns about anticipation in part, a formal distinction of separate effects is not possible in our setting. However, rather than testing a particular hypothesis, our aim is to quantify the relationship between income and expenditures on the household level across the business cycle.[10] That said, within these limits we still try to distinguish between temporary and longer-term changes by including lags.[11]

As noted above, household scanner data have been widely used in marketing research.[12] Only recently, such data have been used more widely in macroeconomics. Kaplan and Shulhofer-Wohl (2017), Jaravel (2019), and Redding and Weinstein (2020) all use household scanner data to measure inflation and cost of living. Alcott, Lockwood, and Taubinsky (2019) study the implications of sin taxes. Various papers have used such data to analyse nutritional intake and its possible determinants (e.g., Alcott et al. (2019), Oster (2018)). As noted above, for the Netherlands we also build up a database that provides nutritional information at the level of individual products.

The remainder of this paper is structured as follows: In section II, we introduce our dataset. In section III, we present stylised facts about CPG consumption and income. In section IV, we undertake our key panel analysis. We conclude in section V. A detailed appendix is provided in section VI.

II Data

We use a novel dataset introduced in Brancatelli et al. (2020) matching household scanner data from the Dutch GfK consumer panel with administrative data on wealth, income and sociodemographic characteristics, as provided by CBS.[13] Matching the datasets was possible as more than 5,000 households

---

[8] Cf. on the use of these two data sources in the US e.g. Attanasio, Battistin, and Ichimura (2004) or Heathcote, Perri, and Violante (2010).

[9] See, for instance, Jappelli and Pistaferri (2014) or Christelis, Georgarakos, Jappelli, Pistaferri, and Van Rooij (2019).

[10] The majority of in particular negative income changes should thus result from the business cycle, rather than, say, transition to retirement, which would typically represent both a permanent and largely anticipated change in income. Our sample, for which we had to obtain individual permissions to match data (see below), does not contain sufficient observations to decompose income changes accordingly.

[11] For instance, to test the "permanent income hypothesis", researchers rely on specific events such as the federal government shutdown in 2013 (e.g. Baker and Yanellis, 2017) or tax rebates in 2018 (e.g. Misra and Surico, 2014).

[12] Dubé, Hitsch, and Rossi (2018) and Biswas, Chintagunta, and Dhar (2019) relate self-reported income data from Nielsen to households' choices between private labels and different store formats. In Brancatelli et al. (2020) we relate households' private-label share to CBS income from administrative data.

[13] Household scanner data from GfK consumer panel are provided by the AiMark foundation.

individually consented.[14] Between 2011 and 2018[15] we are thus able to observe on the level of individual households both their purchases and actual, rather than only reported, information regarding income and wealth. Our level of aggregation is the household, rather than individuals, which is appropriate both for the considered expenditures as well as for income.[16] We enrich these data with a panel of retailer outlets and their geographical location (geo coordinates), thus providing additional supply-side information. Moreover, we complement these data with information on nutritional values in order to control for potential changes in out-of-house consumption as well as to estimate the potential for savings. In the following subsections, we describe each dataset in detail.

II.i Household scanner data

We use household scanner data from the GfK consumer panel, which is the main provider of such information in the Netherlands. It covers household purchases for a wide range of CPGs, both food and non-food.[17] To register purchases, households selected for the GfK consumer panel are equipped with an electronic home scanning device. Our unit of observation is thus a single product on the barcode level bought by an individual household on a certain day at a specific retail chain. For our analysis, we aggregate the data on a household-year level.

Matching with CBS data is only possible for those households which explicitly consented to the matching of their data.[18] This restricts our unbalanced panel from originally 11,041 to 6,151 distinct households in 2018. Of these, we observe 89% over the entire period. Also, note that we are still able to draw on the full Dutch population from CBS and we use this information both to analyse the representativeness of our sample and to construct regression weights.

We conduct the following data selection. We neglect households for which we observe less than six months of purchase data or less than three shopping trips on average per month. Next, we exclude transactions with unreasonably high or low prices by removing transactions with prices four times higher than the median or less than a quarter of the median of the same product.[19] Overall, we lose less than 1% of total expenditures.

[14] Matching, based on household characteristics provided directly by GfK, took place in the secure environment of CBS and was undertaken by its statistical staff. All analyses with matched data took place in the secure environment of CBS.

[15] The chosen time period is limited by data availability from CBS.

[16] When using household scanner data alone, typically reported income instead only relates to the main earner.

[17] This includes fresh, frozen, refrigerated food, alcoholic and non-alcoholic beverages, health and beauty products as well as pet needs, cleaning, detergent products, and general merchandise at a wide variety of retail outlets.

[18] In the European Union in May 2018 the General Data Protection Regulation (GDPR) came into force, which severely limits the creation, storage, and use of personal data. Due to the GDPR GfK is obliged to ask for the consent of each household. In 2018 all households, which were in the panel at the time, had the opportunity to voluntarily agree to their household scanner data being linked to CBS data.

[19] We identify the same product via barcodes and calculate median prices across households and months for each year.

To capture full years of transactions for each household, thereby avoiding seasonal biases as our level of observation are household-years, we exclude panel members that entered the panel after January of a particular year. We thereby lose 6.42% of households.

II.ii CBS data

The Dutch Centraal Bureau voor de Statistiek (CBS), also referred to as Statistics Netherlands, is a governmental institution that collects statistical information in the Netherlands. It provides microdata to authorised institutions under strict conditions of confidentiality.[20] These include, in our case, households' sociodemographic information as well as information on income and wealth. We use CBS information on the number of household members (at the same address), the number of children, the number of male or female persons, and the number of household members receiving an income.[21] Changes in these variables may affect the required consumption basket and thereby expenditures, which is why we account for the composition of a household through the inclusion of these variables in our regression analysis below.

In the descriptive parts of this paper, we adjust income, wealth, and consumption figures for household size, by dividing the respective values by the following OECD scale: $1 + 0.7(n-1) + 0.3k$, where n is the number of adults and k the number of children in the household. For instance, if a single-person household grows by the addition of an adult, the headcount increases from 1 to 1.7. Unless indicated otherwise, we express all Euro values in constant 2018 terms by multiplying the nominal Euro values with the corresponding inflation rate of the respective year vis-à-vis 2018 using the Dutch Harmonised Index of Consumer Prices (HICP).

Our key explanatory variable refers to disposable household income. CBS provides researchers with highly accurate income data from Dutch tax authorities. As some components of income, such as that from self-employment, are only available on a yearly basis and as this generally applies to net income after taxes, which is correctly calculated only at the end of the year, we use yearly income for our analysis. Disposable (household) income is defined as gross household income minus deductions, consisting of transfers and tax payments of all household members.[22]

---

[20] CBS microdata data have been used previously in the literature, for instance in in De Meijer, O'Donnell, Koopmanschap, and Van Doorslaer (2013), Lammers, Bloemen, and Hochguertel (2013), Raymond, Mohnen, Palm, and Van Der Loeff (2010) and Vellekoop and Wiederholt (2019).

[21] CBS provides information on the year of birth and gender of each individual household member. We compare key sociodemographic variables of CBS and GfK data in Appendix VI.i.

[22] Gross household income comprises labour income, business income, capital income, social transfers, insurance benefits, allowances, and alimony from ex-spouses of all household members. Deducted transfer payments comprise income transfers, such as alimony to ex-spouses, as well as mandatory social transfers for social assistance and national insurance benefits. Taxes deducted from gross household income consist of all income tax and a wealth tax paid within a household.

CBS also provides detailed information about households' financial wealth. Financial assets consist of deposit and current account balances and the market value of stocks and bonds as per 1st January of the following year. Dutch households need to declare this information in their annual income tax declaration in order to derive the base for the income tax on return on assets. [23] As we explain below, the use of financial wealth data allows us to analyse whether households with smaller savings react more strongly to income changes as often suggested by the literature.

Through CBS we also have access to the granular data on the income of all Dutch households. We use this both to check for the representativeness of our sample and to perform a weighted analysis. We comment below on the representativeness and, in this context, also provide descriptive statistics for sociodemographics, including income and financial wealth, both for our sample and the population (cf. also Table 13 in Appendix VI.v).

II.iii Supplementary data

For our analysis, we are able to draw on two additional data sets. First, data obtained from Distrifood Dynamics provide us with quarterly observations of all retailer outlets including addresses. Linking households to the respective geo coordinates allows us to measure distances between retailer outlets and homes.[24] As discussed below, this provides an additional control for possible changes in the proximity of notably cheaper discounters.

Second, we enrich our household scanner data with detailed information on nutritional values and ingredients of food products, making use of product barcodes and product group classifications available. Most products reported in the GfK data are provided with a standardised Global Trade Item Number (GTIN), which identifies products uniquely on a global level. To obtain information on the level of GTINs, we combine three sources of data: web scraped data from a large Dutch retailer, web scraped data from a large German provider of food information as well as data obtained from a cooperation with Atrify GmbH (formerly known as *1Worldsync*), a provider of a global product information platform for retailers and producers. In terms of value, this covers already 39% of the entire food expenditures in our sample. In addition, we impute non-matched records using nutritional data from the Dutch Food Composition database, provided by the National Institute for Public Health and the Environment. This database contains detailed records of nutritional data on the level of more than 2,000 food products and product clusters. Overall, we thereby cover 99.8% of the entire food expenditures made by households. As we combine

[23] Whenever no such information is provided in the income tax declaration, CBS uses reported data from the respective financial institutions to impute these values individually. As of 2016, this applies also to foreign financial assets that are located in countries which are part of the OECD Common Reporting Standard adopted by the European Commission.

[24] Distances are calculated as the shortest line between the midpoint of the 6-digit zip code of household residency and store location. The resulting imprecision is assumed to be below 1km on average as there are overall 459,438 6-digit postal codes in the Netherlands, each yielding approx. 37 inhabitants.

different data sources, we provide robustness analyses for different compositions. We report details about the construction of our nutritional database in Appendix VI.iv.

II.iv Representativeness of matched data

Via CBS, we have access to household-level information for the whole Dutch population and can compare this to our sample. In Table 1, we first provide detailed information on our main explanatory variable of interest, disposable income, which we have deflated and standardised for household size. Table 1 shows that mean, median and key percentiles of our income variable in our sample are similar to those obtained from the population as a whole. Notably, when looking at similarities along the distribution, we find that the largest differences prevail at the 10th and 90th percentiles. In particular, this may reflect well-known problems in the under-representativeness of commercial panels at the tails of the income distribution. To alleviate concerns that this may bias the effects of interest in our analysis, we exploit the access to CBS registry data to construct sampling weights based on the procedure described in Brancatelli et al. (2020). For our main regression and robustness checks, we provide separate results for when we include those weights.[25]

In Table 13 (Appendix VI.v), we provide in addition key statistical descriptives for all relevant sociodemographic variables, including income and financial wealth, and for both endpoints of our considered period from 2011 to 2018. Once more, we report figures both for our sample and for the whole population. Overall, this confirms that our sample is largely comparable to the whole population along these dimensions. With the noted under-representativeness of commercial panels at the tails of the income distribution, the representativeness originates from the endeavour of commercial panel providers to keep the panel representative over time and as we have no reason to suspect that there are systematic selection biases for the consenting sample.

[25] Specifically, we construct year-specific adjustment weights using post-stratification. We follow Bethlehem and Biffignandi (2011) and start by stratifying our regression samples into strata using disposable income from the CBS household income registry INHATAB. For each year, we first split the population into income deciles to create 10 strata of equal size. Next, we compute for each stratum a correction weight as the ratio between the population percentage of households over the sample percentage of households. Finally, we multiply the correction weight with the inverse probability of each household of being in the sample. Since the true (and household-specific) probability of being sampled is unknown and since we assume that households within the stratum are homogeneous with regard to their inclusion probability, we assign an equal inclusion probability to each household within the stratum. We impute the stratum-specific inclusion probability with the ratio of households within the stratum over the total number of households in the population.

Table 1: Summary statistics of disposable income in sample and population

	Disposable income p.c.	
	2011	2018
Mean - S.	23,379.33	24,037.68
Mean- Pop.	23,532.35	24,852.26
P10 - S.	14,736.30	14,683.00
P10 - Pop.	12,164.27	12,513.91
P25 - S.	17,660.22	17,795.22
P25 - Pop.	15,925.80	16,552.00
P50 - S.	21,993.73	22,571.00
P50 - Pop.	21,001.24	22,349.00
P75 - S.	27,269.93	28,410.00
P75 - Pop.	27,768.59	29,635.88
P90 - S.	33,986.74	35,008.50
P90 - Pop.	36,471.79	38,164.17
N - S.	3,110	5,451
N - Pop.	7,544,704	7,973,137

Notes: Table 1 shows summary statistics for disposable income per standardized household member. Values are expressed in constant 2018 Euro. Household size is standardised using an OECD scale defined as 1 + 0.7(n- 1) + 0.3k, where n is the number of adults and k the number of children in the household.

III Stylised facts: CPG expenditures and income

III.i CPG expenditures and CPG retailing in the Netherlands

The considered products fall into two main segments: food and non-food items. The latter includes, for instance, detergents and cleaning products, as well as pet needs. Food items comprise beverages as well as various types of fresh or refrigerated food. Substitution of recorded purchases with out-of-home consumption and variation of it over time should be less of an issue for non-food items. Our data allow us to decompose total CPGs into a food and non-food component. At various occasions, we subsequently distinguish between these components. Overall, food consumption represents 82.1% of total CPG consumption expenditures.

With respect to food items, while we have only information on purchases intended primarily for consumption at home, we control for possible out-of-house substitution by including as a control variable measures of the nutritional content of the purchased basket. Figure 1 shows that household expenditures for CPGs represent a sizable component of total disposable income, namely between 9.5% and 11.5% over time.

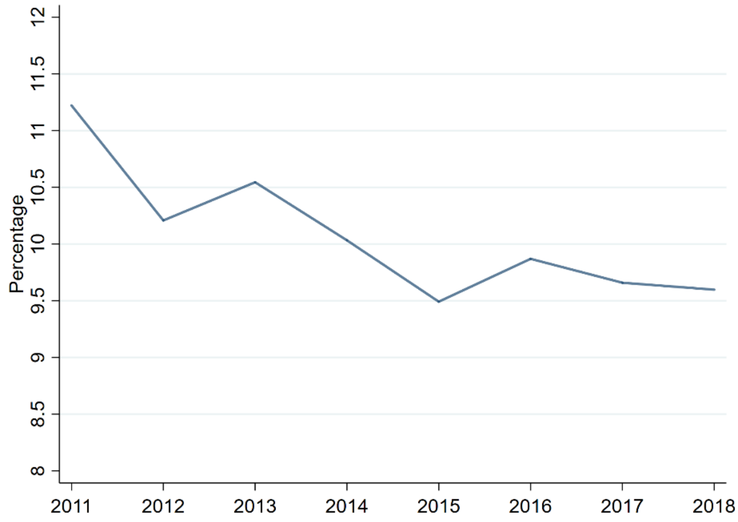

Figure 1: Average household CPG expenditure share in terms of disposable household income

Notes: Figure 1 shows the average share of CPG expenditures across households expressed as percentage of total CPG consumption over nominal disposable household income per year.

Next, we provide some background on retailers. This is important both to understand the scope of potential savings (through purchasing a similar basket of goods at lower price) and for how we subsequently control for potential different and varying access to the respective assortments. Retailers in the Netherlands fall broadly into two categories: (full-line) supermarket chains and hard discounters. Hard discounters have a much smaller assortment, consisting mainly or even almost exclusively of store brands, and consequently, on average, lower prices. The purchase of store brands instead of national brands provides more generally a source of potential savings. According to industry reports, even at the largest supermarket chain, Albert Heijn, private labels are, on average, 31 % to 55 % cheaper than comparable national brands (IPLC, 2016). In a cross-country comparison (IRI, 2016), the ratio of private label to national brand prices was 73.8 % in the Netherlands, compared to 88.5 % in the US. The granular nature of our data allows us to confirm this. Therefore, we have looked into categories that are both important to households and that allow to clearly identify the units of purchases. Across the whole sample of observed purchases, the price ratio of private labels to brands equalled 86.79 % for coffee, 67.75 % for ketchup, and 66.16 % for washing detergents. Together with retailers' national strategies, which we confirm next, we can thus document that households have large savings potentials. This provides an important background for the subsequently documented considerable differences in households' (per capita) expenditures.

We know from industry experts that retail chains in the Netherlands practice national pricing. We acknowledge that this is markedly different to observed practices in other countries. One explanation for this is the small size of the Netherlands combined with the relative homogeneity of living conditions. We provide confirming evidence for national pricing in Appendix VI.iii.[26] Potential differences in expenditures across households are thus not explained by store-specific retailer pricing. National pricing also alleviates concerns that retailers' local price adjustments, e.g., in regions that were more or less affected by the recession, could confound the measured coefficients.

III.ii Variation in disposable income

Over the considered time period, the Netherlands underwent the second "dip" of the so-called double-dip recession in the aftermath of the financial crisis (and the consequent "Great Recession"). Figure 2 depicts the percentage changes of GDP and household income. The figure highlights that both GDP and household income dropped between 2011 and 2013 and increased subsequently. This macroeconomic development provides a natural source of exogenous variation.

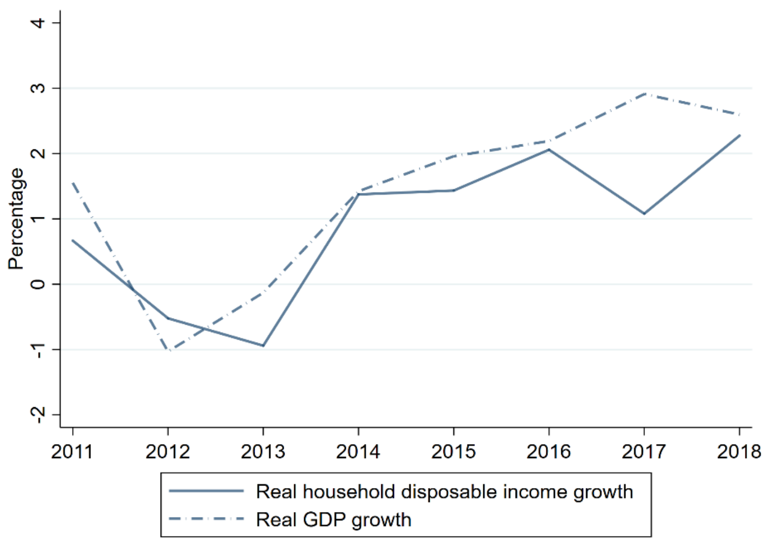

Figure 2: Real GDP and disposable household income growth

Notes: Figure 2 shows annual growth rates of real GDP and real disposable household income in the Netherlands.

Source: World Bank World Development Indicators and OECD National Accounts Data.

We next depict variation of per-household income in our estimation sample. Such variation has two sources: first, adjustments in any of the recorded income components, as described in section II.ii, and second, changes in the household's composition. While in our regression analysis below we account for household

[26] Information from industry experts also suggests that, apart from regional specialities, retailers follow a national assortment strategy. Using our granular data, we have also checked for such regional differences in assortments in Brancatelli et al. (2020). Recalling that store brands are typically even considerably cheaper than national brands, we calculate for each retailer in each year and each region the Euro share of store brands. We do not find significant variations within a specific retailer.

composition through the inclusion of household size and members characteristics, for the descriptive analysis in this subsection we calculate disposable income per (adjusted) household member, using the OECD scale introduced above.[27]

Table 2: Main percentiles of annual percentage changes in household disposable income

Variable	p5	p10	p25	p50	p75	p90	p95	N
Income (%)	-19.73	-12.11	-3.35	0.77	5.92	16.66	28.06	25,338

Notes: This table shows the main percentiles of annual disposable income growth for the sample used in the baseline regression. Underlying values are expressed in constant 2018 Euro. Household size is standardised using an OECD scale (defined as 1 + 0.7(n- 1) + 0.3k, where n is the number of adults and k the number of children.

Table 2 shows the main percentiles of the income growth distribution across households and time. Observed income increases by 0.77% per year at the median. Figure 6 in Appendix VI.iv depicts the distribution of changes over time, showing that also on a year-by-year basis, there is considerable variation in the data. The year-by-year split also uncovers the macroeconomic shifts showed above, as the fraction of positive changes increases over time.

III.iii Cross-sectional variation of CPG expenditures

Our main analysis exploits the panel structure of our dataset and thus within-household variation in income, including various controls. In this section, however, we first analyse cross-sectional variation in expenditures. This also allows us to explore the potential for savings.

Table 3: Summary statistics of CPG, food and non-food consumption levels

Variable	p5	p10	p25	p50	p75	p90	p95	N
CPG consumption level (in 100 EUR)	8.36	10.60	14.68	19.68	25.47	32.39	37.66	30,593
Food consumption level (in 100 EUR)	7.06	8.99	12.28	16.21	20.70	25.82	29.51	30,593
Non-food consumption level (in 100 EUR)	0.61	0.89	1.60	2.85	4.87	7.96	10.84	30,593

Notes: Table 3 shows the main percentiles of CPG, food and non-food consumption levels for the sample used in the baseline regression. All values are expressed in constant 2018 Euro. Household size is standardised using an OECD scale (defined as 1 + 0.7(n- 1) + 0.3k, where n is the number of adults and k the number of children.

[27] We note, however, that by the nature of our data from CBS, changes to disposable income also take into account potential gains and losses in terms of applicable tax brackets.

Table 3 depicts the distribution of CPG (food and non-food) spending levels pooled across all household-year observations, where, again, household size has been standardised using the OECD scale and values are expressed in constant 2018 Euro. The median CPG consumption level per year and standardised household member is 1,968 Euro. However, the distribution shows that the median masks substantial heterogeneity. Remarkably, total CPG consumption per standardised household member at the 90th percentile is three times larger than consumption at the 10th percentile. As we discuss in more detail below, these figures suggest a large potential for savings, since there are also no substantial supply-side obstacles to realise such savings (given national pricing and the relatively uniform availability of stores). Recall that we also identified various potential sources for such a savings potential, such as shopping at hard discounters or purchasing store brands. Given the documented national strategies of retailers, these alternatives are available to all households (cf. also our subsequent controls).

To strengthen this conclusion, we next analyse the relationship in a multivariate setting. For this purpose, we estimate a set of OLS regressions of different expenditure measures on a set of explanatory variables that proxy for household needs: the number household members, the number of females, children and income earners. We report the results in Table 4.

Table 4: Potential savings regression

	(1) All CPG	(2) All CPG	(3) Food	(4) Food	(5) Non-food
HH members	977.548***	-76.647**	867.859***	-8.489	109.620***
	(52.26)	(35.37)	(42.87)	(26.69)	(16.73)
Fem. HH members	-184.252***	27.146	-180.951***	-4.475	-3.287
	(42.30)	(28.52)	(33.99)	(21.61)	(13.75)
HH members <18	-508.934***	-174.360***	-399.814***	-112.646***	-109.083***
	(52.00)	(33.82)	(43.14)	(25.57)	(15.67)
HH income earners	30.368	60.580**	41.487	70.252***	-11.137
	(44.38)	(28.82)	(36.30)	(22.02)	(13.65)
Carbohydrates		-0.032***		-0.027***	
		(0.00)		(0.00)	
Sugars		0.013***		0.009***	
		(0.00)		(0.00)	
Fats		-0.049***		-0.046***	
		(0.00)		(0.00)	
Calories		0.007***		0.006***	
		(0.00)		(0.00)	
Constant	1682.904***	924.644***	1223.919***	594.980***	458.902***
	(55.85)	(38.47)	(44.91)	(28.16)	(19.14)
Observations	30,566	30,566	30,566	30,566	30566
R^2	0.222	0.641	0.278	0.710	0.018
Adjusted R^2	0.222	0.641	0.278	0.710	0.018

Notes: Table 4 presents estimation results from a pooled OLS regression of different consumption measures on sociodemographic characteristics and nutritional intake. The dependent variables are either CPG consumption (columns 1-2), food consumption (columns 3-4) or non-food consumption (columns 5-6). Expenditure levels are expressed in constant 2018 prices. Clustered standard errors are displayed in parenthesis. Significance levels are defined as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

For non-food expenditures, key sociodemographic variables explain only 1.5% of total variation. This is larger for food, where, as reported in column (3) around 30% of variation is explained by sociodemographic variables that proxy for household needs. This still leaves 70% of expenditure variation unexplained. As we consider the full basket, rather than only a certain category, we cannot control for the number of purchased items of a given product or some other measure of volume. For food, we use instead the purchased intake of calories, carbohydrates, sugars and fats, thereby controlling for the volume of purchased food. The (adjusted) R^2 increases indeed to 71%, leaving still almost 30% of variation unexplained.

Before moving on, we note that when we repeat the analysis in Table 4 by also including our subsequently introduced supply-side controls (see Section IV), results remain virtually unchanged..

III.iv Cross-sectional variation of CPG expenditures with disposable income

Still in a cross-sectional analysis, we next explore the association between spending on CPG consumption and disposable income. Figure 3 through Figure 5 report this relationship across all household-year observations.

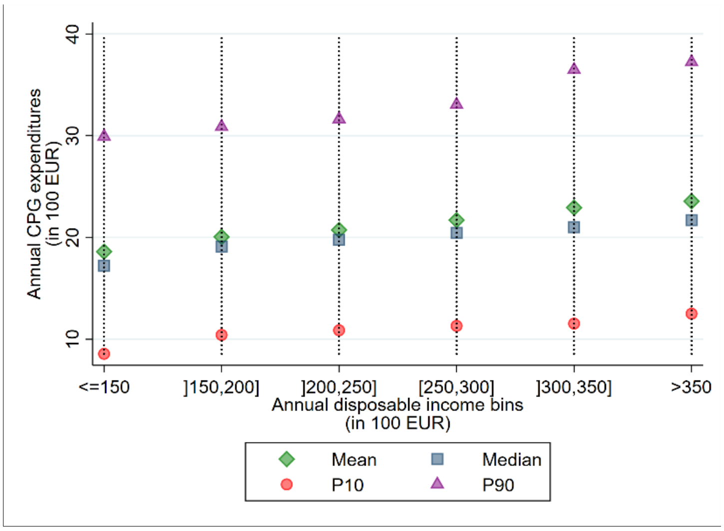

Figure 3: Total CPG consumption by disposable income bin

Notes: Figure 3 shows mean, median, p10, and p90 of CPG consumption by bins of disposable household income across all household-year observations. Consumption expenditures and income are expressed in constant 2018 Euro. Household size is standardised using an OECD scale defined as 1 + 0.7(n- 1) + 0.3k, where n is the number of adults and k the number of children.

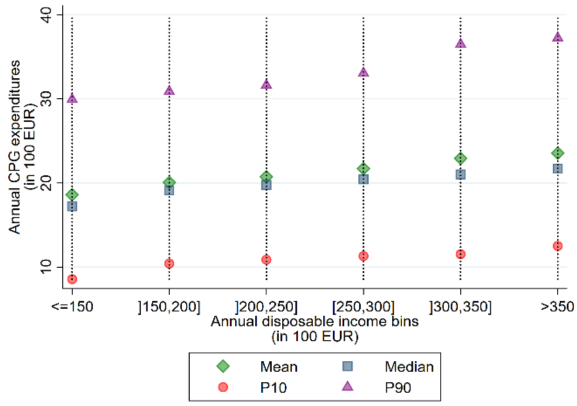

Figure 4: Food consumption by disposable income bin

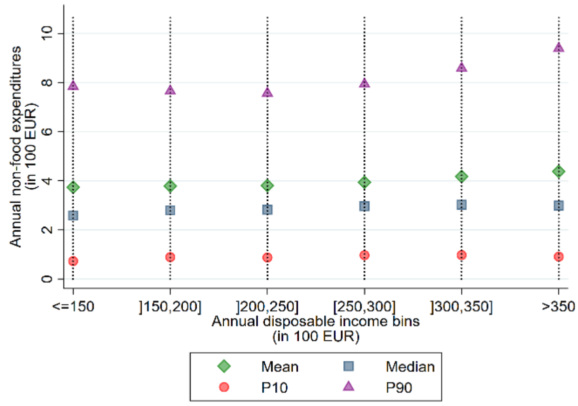

Figure 5: Non-food consumption by disposable income bin

Notes: Figure 4 and Figure 5 show mean, median, p10, and p90 of food and non-food consumption by bins of disposable household income across all household-year observations. Consumption expenditures and income are expressed in constant 2018 Euro. Household size is standardised using an OECD scale defined as 1 + 0.7(n- 1) + 0.3k, where n is the number of adults and k the number of children.

In line with findings in the previous section, the figures confirm large variation in per-capita expenditures, which persist also within a given income bin. Precisely, for each considered income bin, the 10th percentile exhibits total CPG expenditures of around 1,000 Euro, while the standardised per-capita expenditures for the 90th percentile is always three times as large. Median standardised per-capita expenditure exhibits instead a much lower variation across income bins: At the lowest yearly income bin (up to 15,000 Euro) it equals 1,723 Euro and at the highest income bin (over 35,000 Euro) 2,171 Euro.

In Appendix VI.v, we explore this (largely cross-sectional) relationship between income and expenditures also in a multivariate setting as follows. We estimate a pooled OLS regression of CPG expenditures on income derived from our baseline model introduced below, which includes also various sociodemographic and supply-side control variables (as discussed subsequently). The estimation results are summarised in Table 14 in the appendix. For the various specifications, we find that a change of 1,000 Euro in household income is associated with an increase in CPG consumption between 10 and 30 Euro, all else being equal. We consider this to be a fairly low figure, noting that CPG expenditures make up around 10% of total income across all households. If this ratio remained constant across income categories, the figure should be in the order of 100 Euro (instead of ranging between 10 and 30 Euro). As we explore in the subsequent section, the coefficient still drops considerably when we consider only within-household variation in income.

IV Within-household estimation of the relationship between income and CPG consumption changes

IV.i Empirical strategy

We now exploit the panel structure of our dataset. The respective coefficient from the pooled regression, which we reported above, clearly cannot be interpreted causally, e.g., as individuals with particular (unobserved) preferences may both strive for higher income and products that are more expensive. Relying now on within-household variation, our research question is how changes in income affect total CPG expenditures, notably in the context of business cycle fluctuations, such as the Dutch double-dip recession. For this purpose, we estimate a fixed effects regression of CPG consumption on disposable household income while considering various time-varying supply- and demand-side controls. In our baseline analysis we estimate the following model:

$$\Delta log(C_{ht}) = \beta\, \Delta log(I_{ht}) + \gamma \Delta \boldsymbol{X}_{ht} + \phi \Delta log(I_{rt}) + \boldsymbol{\alpha}_h + \boldsymbol{\psi}_t + \epsilon_{ht}\,.$$

Here, $\Delta log(C_{ht})$ denotes the first difference of the logarithm of CPG consumption of household h in year t. The term $\Delta log(I_{ht})$ captures the first difference of the logarithm of disposable income of household h in year t. By taking first differences, we also alleviate concerns about potential spurious correlation due to a common long-run trend of consumption and income, though we use throughout deflated values (cf. Baker (2018) for a similar approach in household finance).

The coefficient β represents our main coefficient of interest. In our main regression, income is thereby only considered contemporaneously and may capture both permanent and transitory changes in income.[28] Unfortunately, the limited number of observations does not allow to distinguish between permanent changes, respectively changes that are perceived to be permanent. For instance, in the respective downturn the recorded negative income changes represent a composition of changes that notably younger workers may perceive as relatively short-lived, while job losses close to retirement age should be perceived as a permanent reduction in income.

While we are mainly concerned with quantifying the respective relationship over the business cycle, rather than testing a particular theory, we still wish to account for the possibility that households may at first not react to a change in income that was (wrongly) perceived to be only short-lived.[29] We re-estimate our model including the first lag of (yearly) income.[30] Note that our focus on the business cycle may also somewhat

[28] See Dubé et al. (2018) for a similar discussion.

[29] In fact, economic theory would suggest that, in the absence of liquidity constraints, households should hardly adjust consumption due to short-term fluctuations in income. Still, empirical literature has documented excess sensitivity of households even after transitory shocks (cf. Jappelli and Pistaferri, 2010).

[30] Considering lagged and contemporaneous estimates jointly when applying first differences has been frequently adopted in other strands of literature, e.g. when estimating cost pass-through (cf. Nakamura and Zerom 2010).

alleviate concerns that the coefficient could be attenuated as households had already anticipated the change in income (e.g., by planning to reduce working hours or transiting in retirement).

The vector $\Delta \boldsymbol{X}_{ht}$ in our model collects a set of time-varying household socio-demographics in first differences that are derived from CBS data. The variables comprise household size, calculated as the number of officially registered persons in the household, number of children, number of females, and number of income earners in the household, measured accurately through the matched administrative data.

To control for the nationwide price level in each year, all Euro values are expressed in constant 2018 prices using the Dutch HICP. We also include year fixed effects. Recall now that in the Netherlands, retailers practice national (pricing and assortment) strategies, which alleviates concerns that retailers may react to regional economic changes through targeted price cuts, thereby confounding our coefficient of interest. Still, given notably considerable price differences between hard discounters and supermarket chains, there may be concerns that retailers react, for instance, by opening or closing outlets. To address this, we proceed as follows. In our main specification, we include a generic control for such possible (local) responses through a measure of local income changes. For this we calculate for each household h and each year t the log of yearly median household income in each household's region (district) r, captured in differences by $\Delta log(I_{rt})$.[31] To capture potential dynamics of such a response, we estimate the baseline model also including the lag of the log of district income and discuss results below. Based on geo coordinates of retailer and household locations available in our dataset, in a robustness analysis we also include the number of supermarket and discounter stores in the radius of 5km of each household.

All models are estimated accounting for household and year fixed effects, collected in vectors $\boldsymbol{\alpha}_h$ and $\boldsymbol{\psi}_t$ respectively, together with a constant. We estimate our baseline models both with and without sampling weights, constructed directly from administrative CBS, to address concerns about potential selection bias.

As discussed above, to the extent that households substitute between in-house and out-of-house consumption of food, variation of CPG expenditures alone may however provide a misleading picture of the impact of the business cycle on individual consumption. Instead, the respective would then comprise also such a compensating effect. For food we therefore add as an additional control nutritional intake of households, adding the log of the following variables to the vector $\Delta \boldsymbol{X}_{ht}$: carbohydrates (in g) per year, sugars (in g) per year, fats (in g) per year and calories (in Kcal) per year.

[31] We define *district income* as median disposable yearly household income within a district. A district is a geographical subunit of a municipality in which usually a certain form of land use or building predominates, for instance, an industrial area or a residential area. The exact boundaries of a district are set by municipalities. In 2018, there were 1,818 districts in which any of our panelists were living with an average of 3,993 households per district.

IV.ii Main estimation results

In this subsection, we present the results of our baseline fixed effects regressions. Table 5 documents the findings using six different specifications with varying supply- and demand-side controls as well as weights.

Table 5: Fixed effects regression - CPG consumption

	(1) ΔLog(CPG)	(2) ΔLog(CPG)	(3) ΔLog(CPG)	(4) ΔLog(CPG)	(5) ΔLog(CPG)	(6) ΔLog(CPG)
ΔLog(HH Income)	0.041***	0.037***	0.018**	0.017*	0.025***	0.013***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.00)
ΔL1. Log(HH Income)					0.012*	
					(0.01)	
ΔHH members			0.032***	0.030***	0.027***	-0.006
			(0.01)	(0.01)	(0.01)	(0.00)
ΔFem. HH members			-0.008	-0.007	-0.005	0.011**
			(0.01)	(0.01)	(0.01)	(0.01)
ΔHH members <18			0.015**	0.022***	0.017**	0.008**
			(0.01)	(0.01)	(0.01)	(0.00)
ΔHH income earners			0.001	0.003	0.003	-0.001
			(0.00)	(0.01)	(0.01)	(0.00)
ΔLog(Carbohydrates)						-0.360***
						(0.04)
ΔLog(Sugars)						0.110***
						(0.01)
ΔLog(Fats)						-0.384***
						(0.04)
ΔLog(Calories)						1.411***
						(0.07)
ΔLog(District income)			0.030	0.017	0.049	0.022
			(0.04)	(0.04)	(0.04)	(0.02)
D.2013			-0.002	0.001		-0.000
			(0.01)	(0.01)		(0.00)
D.2014			-0.004	-0.002	0.005	0.014***
			(0.01)	(0.01)	(0.01)	(0.00)
D.2015			-0.009*	-0.009	-0.001	0.016***
			(0.01)	(0.01)	(0.01)	(0.00)
D.2016			0.019***	0.020***	0.027***	0.023***
			(0.01)	(0.01)	(0.01)	(0.00)
D.2017			0.009*	0.013**	0.018***	0.024***
			(0.01)	(0.01)	(0.01)	(0.00)
D.2018			0.014***	0.016***	0.018***	0.024***
			(0.01)	(0.01)	(0.01)	(0.00)
Constant	-0.018***	-0.020***	-0.023***	-0.026***	-0.029***	-0.009***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Observations	25,309	25,309	25,309	25,309	20,585	25,309
R^2	0.003	0.003	0.010	0.011	0.012	0.646
Adjusted R^2	0.003	0.003	0.010	0.010	0.011	0.646

Notes: Table 5 presents the estimation results from the baseline fixed effects regression. The dependent variable is the first difference of the log of total CPG expenditures per year. Column (1) presents results including only disposable income. Column (3) adds supply and demand side variables to the model from column (1). Column (5) includes the lag of disposable income. Column (6) adds nutritional variables to the model from column (3). Columns (2) and (4) replicate the regressions in columns (1) and (3) respectively using sampling weights. Clustered standard errors are displayed in parenthesis. Significance levels are defined as follows: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Columns 3 contains our key result using the baseline model with the basic set of supply and demand variables. We find that a 10% increase in disposable income is associated with a 0.18% increase in CPG consumption.[32] Using weights (column 4) the relationship is slightly reduced to 0.17%. While in both cases the coefficient is statistically significant at the 5% and 10% level respectively, it is economically rather small.

Our finding suggests that households hardly adjusted CPG consumption in response to income variation as experienced during the recession and the subsequent boom. These findings hold despite large savings potential, as documented above.

In our estimates, the supply-side control (district income) turns out to be statistically insignificant. We find that this holds also when including both a contemporaneous and the first lag of log median district income, leaving our main coefficient virtually unaffected. Our analysis is also robust to the inclusion of the number of supermarket and discounter stores in the radius of 5km of each household. The respective results for these checks are omitted for the sake of brevity.

We now discuss the additional reported regressions. The results in column (5) show that the income lag of first order is highly significant but, compared to the contemporaneous effect, smaller in magnitude. The estimation suggests that for a 10% change of disposable income the associated cumulative shift in CPG consumption is 0.35%, i.e., almost twice as larger as the contemporaneous effect alone. We interpret these results to be consistent with the view that households may adjust consumption also with some delay to an initial income change, for instance due to a change in perceptions or due to revised expectations about future income. Still, also the cumulative effect remains economically small, given that, as a benchmark, the coefficient needs to be one so that CPG expenditures would account for a constant fraction of income.

Column (6) extends the baseline model with a set of variables measuring the nutritional intake of in-house food consumption.[33] As the coefficient changes only modestly and notably does not increase, there seems to be little associated change in in-house versus out-of-house consumption.

Following the literature, we now explore separately positive and negative income changes. For this we proceed as follows.[34] First, we define an income change by calculating the year-on-year percentage change of deflated disposable income for each household. Next, we divide the sample by separating negative rates from those being positive or zero. Overall, 36% of all measured income changes are negative. Previous

---

[32] To obtain the exact average effect for a 10%-change, we transform the elasticity in the following way: $100 * (e^{0.018*\log(1.1)}$-1). Accordingly, the average effect for a 25% change in income would be 0.40%.

[33] Notice that, as we explain more in depth in Appendix VI.i, the nutritional data used for the estimation in column (6) prioritises data sources on the GTIN-level and imputes nutritional values with more aggregated data from the NEVO database. As a robustness check, we also re-estimate the regression using an alternative version of our database in which we prioritise NEVO data over GTIN-level data to accomplish a more stable composition over time. In the results not shown here, we find that results are not affected. Note also that, with regard to reported coefficients, as the different nutritional values, as well as sociodemographics, all relate to consumption needs and purchasing volume, this can generate substantial multicollinearity.

[34] Usually the literature analysing positive and negative income shocks can rely on survey data, in which households state whether they have been subject to a positive or negative income change from the previous to the present year. Given that we observe actual (administrative) income data, we cannot rely on such pre-classified data.

studies suggest that households react more strongly to negative changes (Christelis et al., 2019). This is confirmed in our results. As reported in Table 6 for our main regression, for negative income changes the coefficient of 0.037 is substantially larger than the effect for positive income changes (coefficient of 0.024).

The literature has discussed various rationales for such a difference, such as loss aversion or a precautionary savings motive, next to imperfect access to credit markets (e.g., Bunn, Le Roux, Reinold, & Surico, 2018). We cannot disentangle these different hypotheses. However, we note that our subsequent analysis of potential heterogeneity across households does not reveal differences for low income or low liquidity households.

Table 6: Fixed effects regression – negative vs. positive income changes

	Panel A: Negative income changes	Panel B: Positive income changes
	(1) ΔLog(CPG)	(2) ΔLog(CPG)
ΔLog(HH Income)	0.037**	0.024*
	(0.02)	(0.01)
ΔHH members	0.040***	0.021
	(0.01)	(0.01)
ΔFem. HH members	-0.008	-0.003
	(0.01)	(0.02)
ΔHH members <18	0.021	0.015
	(0.01)	(0.01)
ΔHH income earners	0.005	-0.002
	(0.01)	(0.01)
ΔLog(District income)	0.002	0.058
	(0.05)	(0.07)
D.2013	-0.009	-0.008
	(0.01)	(0.01)
D.2014	-0.004	-0.017**
	(0.01)	(0.01)
D.2015	-0.014*	-0.023***
	(0.01)	(0.01)
D.2016	0.012	0.009
	(0.01)	(0.01)
D.2017	0.007	0.000
	(0.01)	(0.01)
D.2018	0.005	0.003
	(0.01)	(0.01)
Constant	-0.018***	-0.009
	(0.01)	(0.01)
Observations	10,981	14,328
R^2	0.018	0.008
Adjusted R^2	0.017	0.007

Notes: Table 6 shows the results from estimations of the unweighted baseline fixed effects regression using a sample split for positive and negative income changes. All estimations use the same supply and demand side controls employed in column (3) of Table 5. Column (1) (Panel A) reports the results for negative income changes. Column (2) (Panel B) reports the results for positive income changes. Clustered standard errors are displayed in parenthesis. Significance levels are defined as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

IV.iii Heterogeneity across households

We now analyse whether the estimated coefficient is sensitive to the composition of households. For this we first analyse whether the overall low effect masks differences between households with respect to their level of income or liquidity. Such differences in the reaction to income are predicted by household finance theory and have been confirmed empirically (e.g., Jappelli and Pistaferri (2010) and recently Baker (2018)).

Table 7: Fixed effects regression – low income and low liquidity households

	Panel A: Baseline results	Panel B: Low income	Panel C: Low financial wealth
	(1) ΔLog(CPG)	(2) ΔLog(CPG)	(3) ΔLog(CPG)
ΔLog(HH Income)	0.018**	0.021***	0.017**
	(0.01)	(0.01)	(0.01)
ΔLog(HH Income) x D.LHHI		-0.010	
		(0.02)	
ΔLog(HH Income) x D.LHHFW			0.004
			(0.02)
ΔHH members	0.032***	0.032***	0.032***
	(0.01)	(0.01)	(0.01)
ΔFem. HH members	-0.008	-0.008	-0.008
	(0.01)	(0.01)	(0.01)
ΔHH members <18	0.015**	0.015**	0.015**
	(0.01)	(0.01)	(0.01)
ΔHH income earners	0.001	0.001	0.001
	(0.00)	(0.00)	(0.00)
ΔLog(District income)	0.030	0.030	0.030
	(0.04)	(0.04)	(0.04)
D.2013	-0.002	-0.002	-0.002
	(0.01)	(0.01)	(0.01)
D.2014	-0.004	-0.004	-0.004
	(0.01)	(0.01)	(0.01)
D.2015	-0.009*	-0.009*	-0.009*
	(0.01)	(0.01)	(0.01)
D.2016	0.019***	0.019***	0.019***
	(0.01)	(0.01)	(0.01)
D.2017	0.009*	0.009*	0.009*
	(0.01)	(0.01)	(0.01)
D.2018	0.014***	0.014***	0.014***
	(0.01)	(0.01)	(0.01)
Constant	-0.023***	-0.023***	-0.023***
	(0.00)	(0.00)	(0.00)
Observations	25,309	25,309	25,309
R^2	0.010	0.010	0.010
Adjusted R^2	0.010	0.010	0.010

Notes: Table 7 presents the results from estimations of the unweighted baseline fixed effects regression adding interaction terms for households with low income and low liquidity. All estimations use the same supply and demand side controls employed in column (3) of Table 5. The baseline results are repeated in Column (1) (Panel A). Column (2) (Panel B) reports the results for low income households. In this regression we interact household disposable income with a dummy (D.LHHI) that is 1 if a household is in the lowest quartile of the standardized and deflated household disposable income distribution and 0 else. Column (3) reports the results for low liquidity households. Accordingly, in this regression we interact household disposable income with a dummy (D.LHHFW) that is 1 if a household is in the lowest quartile of the standardized and deflated financial wealth distribution and 0 else. Column (3) reports the results for low liquidity households. Clustered standard errors are displayed in parenthesis. Significance levels are defined as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

We note first that from columns (2) and (4) in Table 5 results are robust to the inclusion of weights, so that a low coefficient is not due to, for instance, oversampling households with high income and potentially low response. We now make use of our CBS information on households' financial assets. Precisely, we proceed as follows. We classify a household as having low income if its deflated standardised income falls into the lowest quartile. Likewise we define a low-liquidity household if this holds with respect to standardised financial assets. Table 7 reports estimation results when adding respective interaction terms of household disposable income with low-income and low-liquidity dummies respectively. While the main coefficient of interest is robust and statistically significant throughout specifications, the interacting terms turn out to be economically and statistically insignificant, suggesting that low-income and low-liquidity households do not react more strongly than the average in our setting.

Household reactions are potentially masked by income changes that stem from changes in life circumstances. For instance, a household member can decide to reduce working hours in response to the birth of a child, which would lower household income but potentially increase expenditures in the same period. While in our regression we carefully control for the composition and the size of a household, we cannot exclude that such more complex interactions are not fully accounted for by our econometric model. We now restrict the estimation sample to single-member households. A household is defined to be a single-member household and thus included in the sample as long as it consists of only one household member at a given point in time.

Table 8 reports respective results for different specifications. Specifically, the coefficient of 0.097 in column (2) is more than five times as large as in the baseline estimate (coefficient of 0.018). The respective coefficients, including for lagged income (0.040 in column (3) compared to 0.027 in the baseline specification), are substantially larger also in all other specifications, albeit they remain economically small. This result may be of separate interest for other researchers using household panel data. It should be noted that our data provide us with a precise measure of total household income. Hence, the difference between the original regression and the regression focusing only on single-member households does not stem from a possible failure to capture total household income in the former analysis (as would result when only main earner income was reported). We note, however, that the larger effect could also be due to preferences or life circumstances of single-member households that differ systematically from those encountered in larger households.

Table 8: Fixed effects regression – CPG consumption of single households

	(1) ΔLog(CPG)	(2) ΔLog(CPG)	(3) ΔLog(CPG)	(4) ΔLog(CPG)
ΔLog(HH Income)	0.098***	0.097***	0.114***	0.031***
	(0.02)	(0.02)	(0.02)	(0.01)
ΔL1. Log(HH Income)			0.040**	
			(0.02)	
ΔLog(Carbohydrates)				-0.329***
				(0.05)
ΔLog(Sugars)				0.107***
				(0.02)
ΔLog(Fats)				-0.330***
				(0.04)
ΔLog(Calories)				1.251***
				(0.09)
ΔLog(District income)		-0.078	-0.106	-0.025
		(0.07)	(0.08)	(0.04)
D.2013		-0.006		0.004
		(0.01)		(0.01)
D.2014		-0.002	0.000	0.018**
		(0.01)	(0.01)	(0.01)
D.2015		-0.010	-0.012	0.020***
		(0.01)	(0.01)	(0.01)
D.2016		0.030***	0.034***	0.034***
		(0.01)	(0.01)	(0.01)
D.2017		0.012	0.013	0.026***
		(0.01)	(0.01)	(0.01)
D.2018		0.020*	0.016	0.025***
		(0.01)	(0.01)	(0.01)
Constant	-0.013***	-0.020**	-0.019**	-0.015**
	(0.00)	(0.01)	(0.01)	(0.01)
Observations	6,187	6,187	5,048	6,187
R^2	0.011	0.017	0.020	0.553
Adjusted R^2	0.011	0.016	0.018	0.552

Notes: Table 8 presents the estimation results from the unweighted baseline fixed effects regression using a sample split for single-member households. Households are defined to be single-member households if a household consists of only one household member. The dependent variable is the first difference of the log of total CPG expenditures per year. Column (1) presents results including only disposable income (cf. column (1) in Table 5). Column (2) adds supply and demand side variables to the model from column (1) (cf. column (3) in Table 5). Column (3) includes the lag of disposable income (cf. column (5) in Table 5). Column (4) adds nutritional variables to the model from column (3) (cf. column (6) in Table 5). Clustered standard errors are displayed in parenthesis. Significance levels are defined as follows: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

IV.iv Heterogeneity across categories

In a final analysis, we now make use of the granularity of our expenditure data, which basically record every CPG purchase per household. As we noted in the introduction, this granularity is the main reason why practitioners and researchers in marketing and industrial organization use such data, as it allows, for instance, to measure substitution patterns between brands in a given category. In this paper, however, we are interested in harnessing household panel data analyzing changes in aggregate expenditures for this part of total consumption. While a product- or category-level analysis is thus not warranted, we still consider separately the two already discussed key components of CPG expenditures: food and non-food.[35] Such an analysis could be of interest for at least the following reasons.

[35] We note that the consideration of heterogeneity across food- and non-food expenditures is shared with Brancatelli et al. (2020), despite the differences in focus of both papers outlined in the introduction.

Despite the obvious heterogeneity of purchases, food items could be considered as more essential, which would suggest that its relation to income is lower than that of non-food items. In fact, running separately our key regression (column 3 in Table 5), we see in Table 9 that for non-food consumption (column 3) the coefficient is substantially higher than for food consumption (column 2), being almost twice as large.[36] A separate analysis of different categories should also be of interest to researchers that make use of such data for notably the analysis of household-specific inflation (cf. the Introduction). To the extent that our results lend themselves to a causal interpretation, households' potentially different responses to macroeconomic fluctuations, e.g., their reduction of expenditures, should be an integral part of the overall picture.

Table 9: Fixed Effects regression - food and non-food consumption

	(1) ΔLog(CPG)	(2) ΔLog(Food)	(3) ΔLog(Non-food)
ΔLog(HH Income)	0.018**	0.016**	0.036**
	(0.01)	(0.01)	(0.02)
ΔHH members	0.032***	0.034***	0.016
	(0.01)	(0.01)	(0.02)
ΔFem. HH members.	-0.008	-0.012	0.008
	(0.01)	(0.01)	(0.02)
ΔHH members <18	0.015**	0.017***	0.008
	(0.01)	(0.01)	(0.01)
ΔHH income earners	0.001	-0.001	0.007
	(0.00)	(0.00)	(0.01)
ΔLog(District income)	0.030	0.033	0.006
	(0.04)	(0.04)	(0.08)
D.2013	-0.002	0.000	-0.022*
	(0.01)	(0.01)	(0.01)
D.2014	-0.004	-0.008	-0.004
	(0.01)	(0.01)	(0.01)
D.2015	-0.009*	-0.014***	0.005
	(0.01)	(0.01)	(0.01)
D.2016	0.019***	0.018***	0.006
	(0.01)	(0.01)	(0.01)
D.2017	0.009*	0.012**	-0.013
	(0.01)	(0.01)	(0.01)
D.2018	0.014***	0.011**	0.012
	(0.01)	(0.01)	(0.01)
Constant	-0.023***	-0.015***	-0.056***
	(0.00)	(0.00)	(0.01)
Observations	25,309	25,309	25,305
R^2	0.010	0.011	0.003
Adjusted R^2	0.010	0.010	0.002

Notes: Table 9 presents the estimation results from the unweighted baseline fixed effects regression with different dependent variables. All estimations use the same supply and demand side controls employed in column (3) of Table 5. Column (1) repeats the results using total CPG expenditures per year as dependent variable. Columns (2) and (3) show the results using food and non-food consumption respectively. Clustered standard errors are displayed in parenthesis. Significance levels are defined as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

[36] As food constitutes 82% of total expenditures in our sample, this category dominates also in the aggregate regression (column 1).

V Conclusion

This paper builds on a data innovation that consist in the linkage of household scanner data with administrative income data, both for the same unit of observation, namely the individual household. Our key finding is an economically small relationship between household income and expenditures for CPGs. As our unique data match allows us to minimise measurement error for our two key variables, we can exclude such concerns as explanation for our estimated small coefficient.

As we conduct our analysis over a relatively large swing of the business cycle in the Netherlands following the financial crisis, we are confident that this small magnitude is also not due to a lack of variation in income. We also report that over one third of income changes are negative. We also include statistical weights, built from income data for the whole Dutch population, and we conduct a separate analysis for households with low income and low financial wealth, which does, however, not change our results. Importantly, the small association between income and expenditures applies despite considerable differences in households' expenditures, even after controlling for differences in needs. We confirmed that these differences in expenditures indeed reflect large savings potentials and we also documented that in the Netherlands these are not driven by differences in households' access to outlets with lower prices or different assortment. In sum, we can thus rule out many potential reasons, such as measurement errors, lack of representativeness of our sample, or lack of savings potential, for the economically small size of our coefficient of interest. Finally, we can also exclude that our results are affected by potential substitution between in-house and out-of-house consumption as we control for households' nutritional intake by complementing our data with detailed information on nutritional values and ingredients of food products, which leaves our coefficient of interest virtually unaffected.

Thus, our key conclusion is that in economies like that of the Netherlands, households only marginally adjust their expenditures CPGs in response to changes in income (across the business cycle). Throughout the paper, however, we carefully interpret our results as we cannot distinguish between anticipated and non-anticipated in income changes. As we discussed, the measurement of income changes across a pronounced down- and upwards swing in the business cycle may however also somewhat attenuate these concerns.

Results for low-income and low-liquidity households confirm our major conclusion, in that the measured coefficient is not higher. The income coefficient is, however, five orders of magnitude higher for single-member households (albeit still economically small). We conjecture that this may be due to the fact that for other households there may be confounding effects related to, for instance, changes in household composition. As the precision of our measure of household income is independent of household composition, other than if, for instance, only the income of the main earner was reported, this finding may be of interest also for conducting other analyses. Our split along food and non-food components of CPG consumption may finally be of relevance to research that uses CPG (household scan) data for other purposes, as it documents potentially different responses of households to macroeconomic changes, depending on the respective items.

Currently, a limitation of the dataset that we have constructed with the help of AiMark and CBS is the restricted number of consenting participants, though the fraction of consenting households in 2018 was at 75%. As all new members of the Dutch GfK panel are asked to agree to such data matching, this will continuously enlarge the sample. Given the increasing interest of economists to work with scanner data also outside traditional applications in marketing, we hope that our research more generally suggests how such data can be fruitfully expanded, in this case by matching with administrative data.

References

Allcott, H., Diamond, R., Dubé, J. P., Handbury, J., Rahkovsky, I., & Schnell, M. (2019). Food deserts and the causes of nutritional inequality. *The Quarterly Journal of Economics*, *134*(4), 1793-1844.

Allcott, H., Lockwood, B. B., & Taubinsky, D. (2019). Regressive sin taxes, with an application to the optimal soda tax. *The Quarterly Journal of Economics*, *134*(3), 1557-1626.

Attanasio, O., Battistin, E., & Ichimura, H. (2004). *What really happened to consumption inequality in the US?* (No. w10338). National Bureau of Economic Research.

Baker, S. R. (2018). Debt and the response to household income shocks: Validation and application of linked financial account data. *Journal of Political Economy*, *126*(4), 1504-1557.

Baker, S. R., & Yannelis, C. (2017). Income changes and consumption: Evidence from the 2013 federal government shutdown. *Review of Economic Dynamics*, *23*, 99-124.

Bethlehem, J., & Biffignandi, S. (2011). *Handbook of web surveys* . Wiley, New York.

Biswas, S., Chintagunta, P. K., & Dhar, S. K. (2019). Income and Wealth Effects on Consumer Packaged Goods Purchases.

Brancatelli, C., Fritzsche, A., Inderst, R. & Otter, T. (2020). Measuring Income and Wealth Effects on Private-Label Demand with Matched Administrative Data.

Bunn, P., Le Roux, J., Reinold, K., & Surico, P. (2018). The consumption response to positive and negative income shocks. *Journal of Monetary Economics*, *96*, 1-15.

Christelis, D., Georgarakos, D., Jappelli, T., Pistaferri, L., & Van Rooij, M. (2019). Asymmetric consumption effects of transitory income shocks. *The Economic Journal*, *129*(622), 2322-2341.

Dekimpe, M. G., & Deleersnyder, B. (2018). Business cycle research in marketing: A review and research agenda. *Journal of the Academy of Marketing Science*, *46*(1), 31-58.

De Meijer, C., O'Donnell, O., Koopmanschap, M., & Van Doorslaer, E. (2013). Health expenditure growth: Looking beyond the average through decomposition of the full distribution. *Journal of Health Economics*, *32*(1), 88-105.

Dubé, J. P., Hitsch, G. J., & Rossi, P. E. (2018). Income and Wealth Effects on Private-Label Demand: Evidence from the Great Recession. *Marketing Science*, *37*(1), 22-53.

Dutt, P., & Padmanabhan, V. (2011). Crisis and consumption smoothing. *Marketing Science*, *30*(3), 491-512.

Heathcote, J., Perri, F., & Violante, G. L. (2010). Unequal we stand: An empirical analysis of economic inequality in the United States, 1967–2006. *Review of Economic dynamics*, *13*(1), 15-51.

IPLC. (2016). *Retaining consumers tempted by the discount model* (Research Report).

IRI. (2016). *Special Report: Private Label in Western Economies* (Research Report).

Jappelli, T., & Pistaferri, L. (2010). The Consumption Response to Income Changes. *Annual Review of. Econonomics*, *2*, 479-506.

Jappelli, T., & Pistaferri, L. (2014). Fiscal policy and MPC heterogeneity. *American Economic Journal: Macroeconomics*, *6*(4), 107-36.

Jaravel, X. (2019). The unequal gains from product innovations: Evidence from the US retail sector. *The Quarterly Journal of Economics*, *134*(2), 715-783.

Kaplan, G., & Schulhofer-Wohl, S. (2017). Inflation at the household level. *Journal of Monetary Economics*, *91*, 19-38.

Lammers, M., Bloemen, H., & Hochguertel, S. (2013). Job search requirements for older unemployed: Transitions to employment, early retirement and disability benefits. *European Economic Review*, *58*(2), 31-57.

Misra, K., & Surico, P. (2014). Consumption, income changes, and heterogeneity: Evidence from two fiscal stimulus programs. *American Economic Journal: Macroeconomics*, *6*(4), 84-106.

Nakamura, E., & Zerom, D. (2010). Accounting for incomplete pass-through. *The Review of Economic Studies*, *77*(3), 1192-1230.

Oster, E. (2018). Diabetes and diet: Purchasing behavior change in response to health information. *American Economic Journal: Applied Economics*, *10*(4), 308-48.

Parker, J. A., Souleles, N. S., Johnson, D. S., & McClelland, R. (2013). Consumer spending and the economic stimulus payments of 2008. *American Economic Review*, *103*(6), 2530-53.

Raymond, W., Mohnen, P., Palm, F., & Van Der Loeff, S. S. (2010). Persistence of innovation in Dutch manufacturing: Is it spurious?. *The Review of Economics and Statistics*, *92*(3), 495-504.

Redding, S. J., & Weinstein, D. E. (2020). Measuring aggregate price indices with taste shocks: Theory and evidence for CES preferences. *The Quarterly Journal of Economics*, *135*(1), 503-560.

Vellekoop, N., & Wiederholt, M. (2019). Inflation expectations and choices of households.

VI Appendix

VI.i Nutritional data

In this appendix, we describe the construction of our nutritional database. The process comprises essentially two steps: first, the data collection and second the matching of GfK products vis-à-vis the nutritional records, followed by the aggregation of nutritional data on the level of households per year. In the following, we describe each step more in depth.

To collect the data, we proceed as follows. First, we collect nutritional data on the level of individual GTINs from three sources: web scraped data from a large Dutch retailer, web scraped data from a large German online platform for food information, and finally data obtained from a cooperation with Atrify GmbH (formerly known as *1Worldsync*), a provider of a global product information platform for retailers and producers. In all three cases, the GTIN is the identifying variable and thus allows us to conduct a one-for-one match with the GfK product database. Overall, among these sources we collect information on more than 600,000 unique GTINs across all product categories.[37]

Next, we complement these data with information from the Dutch Food Composition database (NEVO), provided by the Dutch National Institute for Public Health and the Environment (NIPHE).[38] The NEVO data are identified on the level of so-called "product codes", which reflect single products or product clusters and are usually derived from a variety of sources, including publicly available food composition databases or research on particular products in the market.[39] The database is updated in intervals of between two and three years. To construct our database, we use the version published in 2016.[40] Overall, the database contains 2,389 product clusters ("product codes"), each with a unique description. All product codes can be grouped into 23 more aggregated product groups ("Product groep"). For each data record obtained through either approach, we observe a battery of different nutritional values standardised per 100g/ml, including calories, fat and fatty acids, total sugars, carbohydrates, vitamins, and minerals.

Due to the different levels of granularity and precision of the data, the matching between GfK products and the nutritional database evolves sequentially, granting priority to more precise data sources.

[37] Notice that coverage of GTINs overlaps across all sources. We address the deduplication of such cases below. Notably, not all barcodes in the GfK data represent a formal GTIN that would identify a barcode across different data sources. Instead, some barcodes are introduced by GfK artificially to classify unpackaged fresh food products which usually do not have a GTIN and which households scan to record these products. Overall, we find 40,294 such entries in the GfK barcode data, which represents 4% of all barcodes in the data.

[38] For details, see https://www.rivm.nl/en/dutch-food-composition-database.

[39] The NIPHE provides detailed information on the sources for each product on the NEVO website. Nutritional values are usually derived using analytical results from other organisations, research institutes and universities as well as data from foreign food composition tables, information from manufacturers and recipes.

[40] As this paper is being written, the 2019 version is being released and the version used in our paper is only available on request.

Specifically, we first conduct individual GTIN-level matching of GfK products vis-à-vis GTIN-level data in the nutritional database.[41] In terms of Euro value, this part of the matching covers already 39% of the entire food expenditures in our sample. Next, we impute non-matched records using nutritional data from the NEVO database.[42] Overall, with this matching procedure, we manage to match 90% of the entire food expenditure distribution. Finally, we impute the remaining unmatched GfK records by matching all GfK food product subcategories to product codes or even product groep descriptions from NEVO.[43]

To alleviate concerns that the sequential matching procedure induces artificial variation from its composition, we also construct an alternative database, in which we prioritise the category-matching with NEVO data. While this alternative database may identify nutritional intake less precisely, the stable composition allows us to rule out artificial over-time variation and to therefore validate our results using the more precise matching.

In sum, the final database (in both versions) covers 99.8% of all GfK food expenditures. Summary statistics for expenditure coverage rates on the level of households are reported in Table 10 for the main database.

Table 10: Summary statistics of household coverage rates of nutrition data

Mean	SD	P1	P10	P25	P50	P75	P90	P99
99.84	0.88	97.12	99.73	99.95	100.00	100.00	100.00	100.00

Notes: Table 10 shows summary statistics of household coverage rates of nutrition data with regard to GfK household scanner data. The figures are expressed as percentages of food expenditures covered.

VI.ii Sociodemographic characteristics: GfK data vs. CBS data

In this appendix, we report key summary statistics of panellists in GfK data compared to only those who consented to the matching with CBS.

[41] For this, we first deduplicate all GTIN-level data in our nutritional database by prioritising data from Atrify over web-scraped data in case multiple entries from different sources are available for a given GTIN. When we match each GTIN-level entry one-for-one to the GfK data, we run a series of plausibility checks to ensure that the matched GTINs records are not mistakenly different products due to reassignments of GTINs over time. Normally, a producer or retailer introduces a new GTIN whenever a product changes. In rare cases, it may happen that a GTIN is then re-assigned to a different product over time. Since all barcodes in the GfK data are overwritten by the latest record, we cannot track the GTIN-product relationship over time.

[42] This step mostly evolves manually. To rule out inconsistencies and mistakes in the matching procedure, we run a series of plausibility checks on the matched data and, in addition, carefully validate the matching procedure independently in the following way. We draw multiple random samples of matched data and classify erroneous matches. For each draw, we calculate an error rate. We find that on average the error rates for five subsamples are 10%.

[43] In total, we find more than 300 food product subcategories in the GfK data.

Table 11: Summary statistics of panellists in GfK compared to consenting households in 2011

	Panel A: Age		Panel B: Household size		Panel C: Social class		Panel D: Income class	
	All	Consent	All	Consent	All	Consent	All	Consent
Mean	6.94	7.40	2.49	2.40	4.53	4.61	14.35	14.72
SD	3.02	2.53	1.24	1.20	1.64	1.61	6.25	6.00
P5	1.00	1.00	1.00	1.00	1.00	1.00	3.00	3.00
P25	6.00	6.00	2.00	2.00	4.00	4.00	9.00	11.00
P50	7.00	8.00	2.00	2.00	5.00	5.00	16.00	16.00
P75	9.00	9.00	4.00	3.00	6.00	6.00	20.00	20.00
P90	11.00	11.00	4.00	4.00	6.00	6.00	21.00	21.00
P95	11.00	11.00	5.00	5.00	6.00	6.00	22.00	22.00

Notes: Table 11 shows key summary statistics for the year 2011 for sociodemographic variables obtained from GfK data for both all GfK households and households consenting to the matching with CBS. Panel A shows summary statistics for age, panel B shows summary statistics for household size, panel C focusses on social class, and panel D on income of the household head. The table is based on appendix table 14 in Brancatelli et al. (2020).

VI.iii National pricing

In this appendix, we describe how retailers in the Netherlands practice national pricing. To illustrate that national pricing existed throughout the entire time period for Dutch retailers, we analyse price variation across regions for six major retailers in the Netherlands.

For this purpose, we remove products with prices that are fourfold higher than the median or smaller than a quarter of the median price. Furthermore, we also disregard barcodes artificially introduced by GfK to classify unpackaged fresh food products, which usually do not have a GTIN. From this set of products, we identify and keep the 500 most popular products per retailer per year and separately estimate a price regression for each retailer and year from 2011 to 2018. In our regression, the dependent variable is the log of a household-shopping-trip-specific price of any given product. The independent variables comprise a set of retailer-specific barcode dummies, monthly time fixed effects and three regional fixed effects.[44] Throughout the analysis, we use North Netherlands as the basis. Overall, we estimate price regressions for six retailers and eight years[45] and extract the set of regional fixed effects for each regression.

Summary statistics are represented in Table 12, showing that across all retailers and years the fixed effects are considerably small, with more than 90% being smaller than 1%.

[44] Household scanner data do not contain geo information about individual shopping trips. Therefore, we rely on household location (postcode) to construct regional fixed effects. Regions are derived from the official NUTS 1 codes of the Netherlands: North Netherlands, East Netherlands, West Netherlands, South Netherlands.

[45] In sum, we estimate 44 regression specifications: Until 2014, we estimate the regression for six major retailers. Since two major retailers merged in 2014, we only estimate the regression for five retailers from 2015 onward.

Table 12: Summary statistics of regional fixed effects across all years and retailers

	Mean	SD	Skew	P5	P10	P25	Median	P75	P90	P95	N
Regional FE	0.0017	0.0042	1.1826	-0.0031	-0.0022	-0.0009	0.0008	0.0031	0.0078	0.010	132

Notes: Table 12 depicts summary statistics of regional fixed effects from independent yearly regressions for six major retailers based on the household scanner data of the GfK consumer panel between 2011 and 2018. The table is based on table 4 in Brancatelli et al. (2020).

We finally note that even if a given retailer practised regional price differentiation not captured by this approach, this would not risk confounding our interpretation of the income coefficient as long as price differences do not change over time.

VI.iv Additional graphs

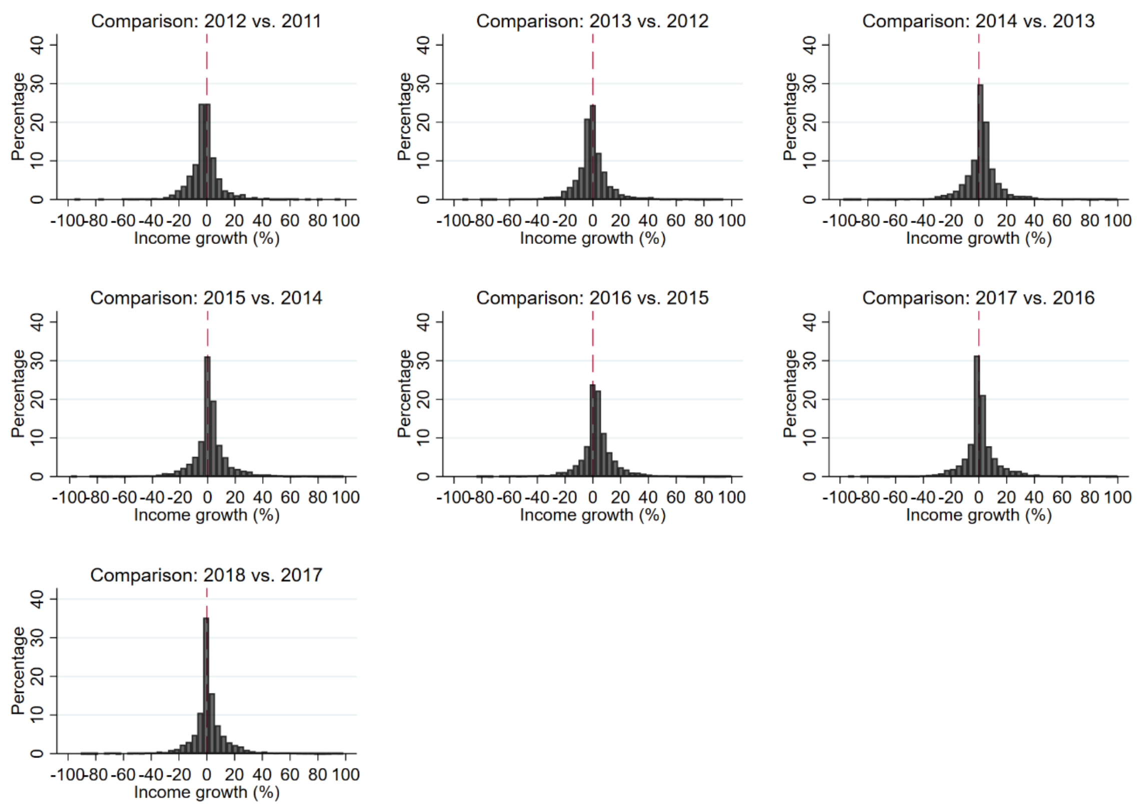

Figure 6: Disposable income changes

Notes: This figure shows the distribution of disposable income growth rates over time. Household size is standardised using an OECD scale defined as 1 + 0.7(n- 1) + 0.3k, where n is the number of adults and k the number of children. Income is expressed in constant 2018 Euro. The figure is truncated at values below -100% and above +100%.

VI.v Additional Tables

Table 13: Summary statistics of key socio-demographic variables

year	label	m - S.	m - Pop.	p10 - S.	p10 - Pop.	p25 - S.	p25 - Pop.	p50 - S.	p50 - Pop.	p75 - S.	p75 - Pop.	p90 - S.	p90 - Pop.
2011	Disposable income p.c.	23,379.33	23,532.35	14,736.30	12,164.27	17,660.22	15,925.80	21,993.73	21,001.24	27,269.93	27,768.59	33,986.74	36,471.79
2018	Disposable income p.c.	24,037.68	24,852.26	14,683.00	12,513.91	17,795.22	16,552.00	22,571.00	22,349.00	28,410.00	29,635.88	35,008.50	38,164.17
2011	Financial assets p.c.	29,624.24	37,910.28	1,060.96	344.06	3,333.90	1,943.71	11,589.33	9,032.25	29,338.02	26,900.76	66,394.65	76,083.19
2018	Financial assets p.c.	27,996.16	36,221.45	715.00	346.15	2,677.65	1,790.77	9,814.00	8,731.18	25,378.06	26,476.00	58,580.00	70,196.00
2011	Household members	2.46	2.18	1.00	1.00	2.00	1.00	2.00	2.00	3.00	3.00	4.00	4.00
2018	Household members	2.40	2.12	1.00	1.00	1.00	1.00	2.00	2.00	3.00	3.00	4.00	4.00
2011	Number of children	0.49	0.43	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	2.00	2.00
2018	Number of children	0.49	0.39	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	2.00	2.00
2011	Number of females	1.26	1.10	1.00	0.00	1.00	1.00	1.00	1.00	2.00	1.00	2.00	2.00
2018	Number of females	1.25	1.07	1.00	0.00	1.00	1.00	1.00	1.00	2.00	1.00	2.00	2.00
2011	Number of inc. earners	1.94	1.72	1.00	1.00	1.00	1.00	2.00	2.00	2.00	2.00	3.00	3.00
2018	Number of inc. earners	1.92	1.73	1.00	1.00	1.00	1.00	2.00	2.00	2.00	2.00	3.00	3.00

Notes: Table 13 shows summary statistics for both the analysis sample and the population for the following socio-demographic variables for the years 2011 and 2018: disposable income per standardised household member, financial assets per standardised household member, the number of household members, the number of females, the number of children and the number of income earners in a household. Columns abbreviated with S. show the corresponding statistics for the sample and columns abbreviated with Pop. Show corresponding population figures Note that household size for disposable income and financial assets is standardised using an OECD scale (defined as 1 + 0.7(n- 1) + 0.3k, where n is the number of adults and k the number of children. Euro values are expressed in constant 2018 Euro.

Table 14: Pooled OLS regression on the determinants of CPG consumption

	(1) CPG	(2) CPG	(3) CPG	(4) CPG
HH Income	0.028***	0.030***	0.009***	0.010***
	(0.00)	(0.00)	(0.00)	(0.00)
HH members			873.313***	886.568***
			(55.16)	(55.51)
Fem. HH members.			-167.582***	-179.588***
			(42.04)	(42.33)
HH members <18			-443.430***	-432.556***
			(52.76)	(53.42)
HH income earners			-18.937	-13.150
			(45.76)	(47.16)
District income			0.010***	0.011***
			(0.00)	(0.00)
D.2012			-304.052***	-286.333***
			(26.43)	(27.51)
D.2013			-402.257***	-378.903***
			(29.55)	(30.56)
D.2014			-498.934***	-480.447***
			(31.46)	(33.26)
D.2015			-598.597***	-580.521***
			(33.38)	(35.34)
D.2016			-604.205***	-592.607***
			(35.38)	(37.14)
D.2017			-653.015***	-639.071***
			(37.31)	(38.52)
D.2018			-688.228***	-666.369***
			(38.94)	(40.16)
Constant	2474.231***	2280.293***	1809.208***	1612.391***
	(123.35)	(113.60)	(121.05)	(121.37)
Observations	30,566	30,566	30,566	30,566
R^2	0.113	0.161	0.240	0.294
Adjusted R^2	0.113	0.160	0.240	0.293

Notes: Table 14 presents the estimation results from the baseline regression using OLS. The dependent variable is total CPG expenditures per year. Column (1) presents results including only disposable income. Column (3) adds supply and demand side variables to the model from column (1). Columns (2) and (4) replicate the regressions in columns (1) and (3) respectively using sampling weights. Clustered standard errors are displayed in parenthesis. Significance levels are defined as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 15: Fixed effects regression – food and non-food consumption of single households

	(1) ΔLog(CPG)	(2) ΔLog(Food)	(3) ΔLog(Non-food)
ΔLog(HH Income)	0.097***	0.095***	0.114***
	(0.02)	(0.02)	(0.04)
ΔLog(District income)	-0.078	-0.076	-0.040
	(0.07)	(0.07)	(0.18)
D.2013	-0.006	-0.011	0.022
	(0.01)	(0.01)	(0.03)
D.2014	-0.002	-0.010	0.025
	(0.01)	(0.01)	(0.02)
D.2015	-0.010	-0.021*	0.043
	(0.01)	(0.01)	(0.03)
D.2016	0.030***	0.025**	0.032
	(0.01)	(0.01)	(0.03)
D.2017	0.012	0.015	0.004
	(0.01)	(0.01)	(0.03)
D.2018	0.020*	0.015	0.023
	(0.01)	(0.01)	(0.03)
Constant	-0.020**	-0.007	-0.076***
	(0.01)	(0.01)	(0.02)
Observations	6,187	6,187	6,183
R^2	0.017	0.018	0.004
Adjusted R^2	0.016	0.017	0.003

Notes: Table 15 shows the results from estimations of the unweighted baseline fixed effects regression using a sample split for single-member households. Households are defined to be single-member households if a household consists of only one household member. All estimations use the same supply and demand side controls employed in column (3) of Table 5. Column (1) shows the results using the first difference of the log of CPG consumption as dependent variable. Column (2) shows the results using the first difference of the log of food consumption as dependent variable. Column (3) shows the results using the first difference of the log of non-food consumption as dependent variable. Clustered standard errors are displayed in parenthesis. Significance levels are defined as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.