Online Appendix to *Unequal expenditure switching: Evidence from Switzerland*

Raphael Auer, Ariel Burstein, Sarah Lein, and Jonathan Vogel

[Not for publication]

Appendix A Data appendix

A.1 Processing the Homescan data

Households record if a purchase occurs within Switzerland or at a retailer abroad. We drop all transactions that occur abroad. Throughout the analysis, we focus on prices including the local VAT.

For expositional purposes, to examine the period around the January 2015 appreciation we shift the data of all transactions by 15 days, so that the appreciation coincides with the change in the calendar year. For example, what is referred to as 2015 (or the first quarter of 2015) includes the actual calendar dates January 15, 2015 through January 14, 2016 (January 15, 2015 through April 14, 2015).

Participating households manually enter data on their transactions. We remove potential errors in the data using a two-step procedure. First, for each transaction we calculate the unweighted average log price across all other transactions of the same product. We then identify all transactions with a price level exactly equal to 1 and, within this set of transactions, drop any transaction for which the absolute value of the log average price excluding this transaction is greater than 2; we do this because it appears that some transactions are accidentally coded as having a price of 1. Second, on the remaining sample, for each transaction we re-calculate the unweighted average log price across all other transactions in the same product and drop each transaction for which the absolute value of the log price minus the log average price excluding this transaction within the product is greater than 2. These transactions may correspond to instances in which quantity and price have been switched. This two-step procedure drops very few transactions: e.g., 273 in 2014 and 585 in 2015.

Whereas EANs are generally product-specific rather than retailer-specific, a block of numbers—all EANs starting with the number 2, termed "in-store" EANs—is reserved for assignment by the retailer. In-store EANs have a variety of uses. They can be assigned by the retail chain, for example if a specific good is sold exclusively by the respective retail chain. However, they can also be assigned at the outlet level, for example when applying a discount to food approaching its expiration date. The same in-store EAN could be used for different products across the different outlets of a retail chain. In-store EANs are thus dropped, unless we can find a product description on codecheck.info that allows us to uniquely map the in-store EAN to a product and its origin.

In the raw data, an observation is a transaction. A transaction is defined by the combination of the household identifier, EAN code, quantity purchased, price paid (net of good-specific discounts due to e.g. coupons), date of the shopping trip, and the name of the

Table A1: Homescan data summary statistics in 2014

	All	Known origin		
		All	Imported	Domestic
Number of products	77,176	8,409	4,084	4,325
Expenditures	118.1	41.9	11.3	30.6
Transactions	254.6	110.4	27.7	82.7

Notes: The sample is all purchases made within Switzerland in 2014 across all households in the Homescan data. The first column includes all purchases made within Switzerland in 2014, the second column includes all such purchases for which the production location of the good is known, and the third and fourth columns decompose the second column into imported and domestically produced purchases. Number of Products is the number of distinct barcode products that are sold within each sample. Expenditures and Transactions are total expenditures (in hundreds of thousands of CHF).

Table A2: Household summary statistics by Homescan income bin in 2014

Income bin	0-35k	35-50k	50-70k	70-90k	90-110k	110-160k	>160k	Total
Median income	15,069	45,410	55,566	76,005	96,569	128,035	257,259	
No. of households	398	554	733	739	391	458	29	3,302
Avg household size	1.7	2.1	2.5	2.9	3.1	3.2	3.8	2.6
Share with kids	7	8	13	17	20	20	24	14
Share elderly HH	22	21	13	9	5	3	0	12
Share higher education	12	15	17	24	33	53	45	17
Median expenditure	735	935	1,043	1,252	1,246	1,292	1,617	1,270

Notes: Household characteristics by income bin in the Homescan data (for our sample of households with positive expenditure in 2014 on products with known production locations). Share higher education is the share of household main earners who have university or college degrees. Share with kids is the share of HHs with at least one child under the age of 10. Share elderly HH is the share of HHs in which all members are over the age of 70. Each HH's total pre-tax annual income is constructed using the relationship between HH characteristics and the level of total household pre-tax annual income in FORS. Median income reports the median value within each Homescan income bin.

retailer. In an abuse of terminology, we redefine a transaction as follows. We aggregate all purchases within a particular household identifier, EAN code, date of the shopping trip, and the name of the retailer into one. To do so, we construct the price of this "transaction" as the unweighted average of prices across transactions in the raw data.

We restrict our sample to households with positive expenditures inside Switzerland in 2014 on products with known import status; this yields a sample of 3,302 households.[1]

A.2 Household pre-tax income

Overview. The Homescan data includes a comprehensive set of household socioeconomic characteristics, as reported in Table A2. However, a household's total pre-tax annual income is reported only in seven bins. We construct a more granular measure of household pre-tax income by using information from two supplementary data sets, the Swiss Household Panel compiled by the Swiss Centre of Expertise in the Social Sciences (henceforth

[1]We construct the first column of Table A1 including all households with positive expenditures in 2014 without restricting to those with positive expenditures on products with known import status.

FORS) and data from the Swiss Federal Tax Administration (2014) (henceforth SFTA). Our approach is to estimate the relationship between household characteristics and total pre-tax income in 2014 in the FORS data and to use this relationship—augmented by the SFTA data—to predict the level of household income for all households in the Homescan sample.[2] We also predict the change in each household's income between 2014–15 following a similar procedure (using the panel structure of FORS).[3]

FORS data. FORS surveys household members regarding their total annual net income in CHF at the time of the survey. The sum of all household members' net income is defined as the sum of labor earnings, asset income, private transfers, public transfers, and social security pensions, all net of taxes.[4] From the data, we calculate household-specific income for calendar years and the socioeconomic characteristics of the household's main earner (which we observe in the Homescan data). Last, we use weights that adjust for non-responses to the household questionnaire in the FORS survey. The population FORS is sampling from is representative, but the response rates differ by socioeconomic characteristics, so FORS has developed weights to adjust for these differences in response rates, which we employ; see Kuhn (2018) for a description.

We adjust the FORS survey waves to correspond to calendar years. FORS is conducted once each year, but the surveying takes place from September to February, with e.g. the 2013 survey wave being sampled from 09/2013 to 02/2014 and the 2014 survey wave being sampled from 09/2014 to 02/2015. The survey includes the date each household was interviewed on, and we thus allocate incomes to calendar years rather than survey waves. We may observe two surveys per calendar year for a household when a household is surveyed between January and February in one wave and between September and December in the following wave. In such cases, we use only the later survey. For the year 2014, the resulting data set contains information on the socioeconomic characteristics of 6,658 households interviewed during January, February, September, October, November, and December 2014.

SFTA data. We additionally use data on the distribution of annual taxable household in-

[2]In practice, the predicted level of household income falls within the relevant Homescan income bin for each household.

[3]When regressing changes of income on household characteristics, to address potential measurement error in income in the FORS data, the 2014 income bins in FORS are instrumented with bins corresponding to average income during the period 2013–16. We also remove outliers of income changes. Finally, we do not use the SFTA data in predicting changes in income.

[4]There are two types of surveys sent to each household. One is a questionnaire for the household as a whole. The other one includes individual questionnaires for each member of the household. FORS judges the individual responses for income to be more reliable, and we thus use the income measure that is summed over individual incomes. FORS conducts manual checks when the individual responses and the household responses are very inconsistent. See Kuhn (2018) for further explanations.

come of natural persons in 2014—from Swiss Federal Tax Administration (2014) (henceforth SFTA)—to measure the median household income level associated with each of the Homescan income bins. Taxable household income ("Steuerpflichtiges Einkommen") is equal to total pre-tax household labor income minus social security contributions and other tax deductions. The SFTA records the number of Swiss households for each 10,000 CHF income step (and steps of 100,000 for incomes above 200,000 CHF). This is the official data for the distribution of pre-tax household income in Switzerland.

We use the SFTA data to obtain the best possible measure of median income within each of the seven Nielsen income bins. In doing so, we split the 30,000–40,000 step in the SFTA data (which includes 30,000–35,000 in the lowest Nielsen income bin and 35,000–40,000 in the next Nielsen income bin) and allocate the number of households equally to the 0–35,000 and 35,000–50,000 CHF brackets in the Homescan data. The resulting median income levels within each Homescan bracket are 15,0000, 45,000, 55,000, 75,000, 95,000, 125,000, and 250,000 CHF.

Specifics. We use these datasets as follows. First, using the FORS data, we project the log of household pre-tax income on the following characteristics: an indicator variable for each of the income bins in the Homescan data, an indicator variable for the household's canton of residence, the education of the household's main earner, the number of household members, the number of household members 17 and under, and the number of household members 70 years old and above.[5] Second, we then predict household income for our Homescan households using these coefficient estimates and the information on a household's socioeconomic characteristics included in the Nielsen database, but replacing the income-bin fixed effects estimated in FORS with the median income in the SFTA data associated with each of the Homescan income bins.[6]

We consider two robustness exercises for predicting household income. In one, we allocate each household in a given Nielsen income bin to a common income level, equal to that of the median household's income in that income range in the SFTA data. In the other, we use the above approach, but leveraging only FORS data (we do not replace the estimated income fixed effects using SFTA data).

[5]The FORS data provides information on the canton of residence. Cantons are more aggregated geographies than two-digit zip codes. However, in some instances two-digit zip codes do not map uniquely to cantons. Of the 76 two-digit zip codes in the Homescan data, we can map all but 29 into a unique canton. Of these 29 two-digit zip codes, we map 22 into 2 cantons and 7 into 3 cantons. In these cases, we allocate the respective canton fixed effects to two-digit zip codes weighing equally the respective fixed effects. In the FORS data, we observe the number of household members up to and including age 17 and the number 70 and above, whereas in Homescan we observe the number of household members under age 10 and the number above age 70.

[6]When we predict changes in income in the Homescan data, we use only the estimates from the FORS data because the SFTA data does not provide a household panel.

Table A3: Relationship between household income and expenditure share on products with known import status in 2014

	(1)	(2)	(3)
log(Income)	0.25	-0.00	0.07
	[0.24]	[0.26]	[0.25]
Observations	3302	3302	3302
Control size		X	X
All controls			X

Notes: Estimation of equation (A1), replacing the dependent variable with the share of household expenditure on products with known import status. Column 2 controls for household size. Column 3 additionally controls for an indicator for whether there is a child under 10 and an indicator if everyone in the HH is older than 70. Robust standard errors are clustered by income quantile (of which there are fifty) and observations are weighted by the product of the number of households in each quantile $\times$ the household's share of expenditure in 2014 within its income quantile. *p<.1; **p<.05; ***p<.01

Footnote 13. In the paper (footnote 13) we state that household income is not significantly correlated with the household's share of expenditure in 2014 on products for which we do not observe import status. To document this fact, we estimate equation (A1), but replace the dependent variable with the share of household expenditure on products with known import status. Table A3 displays the results.

A.3 Details of the Swiss Federal Statistical Office (SFSO) data

In our analysis, we require budget shares across three sectors by income group, inflation rates by income group, and import shares by income group within each of our three sectors. We construct these using three data sets provided by the SFSO. In these data sets, products are defined at a much more disaggregated level than at our sector level. Here, we describe how we concord the three data sets provided by the SFSO and how we construct these variables for the five income groups within the SFSO data.

The first data set, the Swiss Household Budget Survey (HBS), includes information on consumption expenditures by income group and consumption category.[7] The HBS is collected by the SFSO via a rotating and non-overlapping survey, randomly sampled throughout the year from the SFSO's register of all Swiss households. Around 250 households participate each month and record consumption expenditures during the following month for 296 HBS consumption categories. The latter include both goods and services, in categories such as "rice", "pasta", or "tickets for public transport." The survey also collects data on households' socioeconomic characteristics, including income. The SFSO publishes HBS category-specific expenditure shares averaged over a three-year horizon for each of five income groups. The expenditure share data we use in our analysis covers the years

[7] See Swiss Federal Statistical Office (2014) and Swiss Federal Statistical Office (2013) for a detailed description. One purpose of the survey is to calculate the category weights underlying the consumer price index.

2012–14.[8]

The second data source is the disaggregated data underlying the Swiss CPI, which is also published by the SFSO and described in Swiss Federal Statistical Office (2016). It includes price indices for 217 disaggregated CPI consumption categories. The data includes annual price index levels, from which we calculate the category-specific annual inflation rate. We use the data from the 2016 release, which includes the annual rate of inflation for the years 2013–16. Finally, we also use data from the SFSO that reports import shares per CPI consumption category. These import shares are collected periodically via firm surveys. They are used by the SFSO to publish an inflation rate for imported consumer goods.

We concord the HBS expenditure categories with the CPI expenditure categories. Many CPI expenditure categories are identical to the ones from the HBS data. However, not all categories are identical in the two data sets. Therefore, we rely on coarser categories to concord the HBS and CPI schemes.[9] The resulting concordance includes 187 consumption categories.[10]

To compute *(i)* inflation rates by income group and *(ii)* import shares by income group within each of our three broad sectors, we use the expenditure shares by income group across the 187 consumption categories as an income-group-specific weight. We construct the inflation rate by income group in each year as the income-group-specific weighted average of inflation rates across the 187 consumption categories (using the 2012-2014 expenditure shares). We construct the import share in each of our three aggregate sectors for each income group as the income-group-specific weighted average of the import shares of each of the 187 consumption categories within the relevant aggregate sector. Hence, variation across income groups in aggregate inflation rates and in import shares within each of our three aggregated sectors arises exclusively from differences across income groups in expenditure shares across the 187 consumption categories (inflation rates and import shares are assumed to be identical across income groups within each of the 187 consumption categories).

When aggregating from the 187 consumption categories into our three broad sectors—groceries, non-grocery goods, and services—we divide goods as follows. Groceries include all food and beverages at home as well as additional products that are included in the Homescan data, such as "cleaning articles", or "soaps and foam baths." Non-groceries includes all other goods categories.

---

[8]Due to data sparsity, the SFSO does not publish expenditure shares for all income group-category combinations. We impute missing income group-category expenditure shares by the overall expenditure share for the category.

[9]When using coarser HBS categories, we sum the expenditures of the HBS categories we aggregate. When using coarser CPI categories, we use the CPI weights to aggregate the CPI categories.

[10]This concordance is available in the replication material.

A.4 Comparing grocery import shares across Homescan and SFSO

Table 1 shows that the aggregate import share for groceries is substantially higher in the SFSO sample (37.9%) than in the 2014 Homescan data (26.9%); we reproduce these numbers in column 1 of Table A4. Here, we show that more disaggregated import shares—at the product category level—are broadly similar in the SFSO and Homescan data. The difference in the aggregate import share is mostly due to expenditures in the SFSO data being concentrated in sectors with high import shares, particularly in goods other than food and non-alcoholic beverages (that is, alcohol, tobacco, and non-food grocery items).

To compare import shares at a disaggregated level, we concord Homescan "product-groups" with SFSO "product names," resulting in a data set of 44 common categories. We then separately calculate Homescan and SFSO expenditure and import shares for these categories using the 2014 Homescan micro data and the SFSO data underlying Table 1.

In this sample of matched categories (which does not comprise the entire sample), the weighted import share is 37.9% in the SFSO (the SFSO import share reported in column 3 of Table A4) and 26.2% in the Homescan data (the Homescan import share reported in column 4 of Table A4). Hence, we obtain the same discrepancy between Homescan and SFSO grocery import shares in our matched data set as in the full data set.

Using this matched sample, we now provide evidence that the discrepancy between these grocery import shares is driven by differences in expenditure shares across categories rather than differences in import shares within them. As a first exercise, we calculate the correlation coefficient between these category-level import shares. This correlation is 0.63 (significant at the 0.1 percent level). As a second exercise, we construct the unweighted import share across categories in the SFSO data and in the Homescan data. These shares, reported in the second column of Table A4, are substantially more similar, 43.5% in the SFSO data and 39.4% in the Homescan data. Column 3 reports weighted average import shares across categories in the SFSO and Homescan data using expenditure shares from the SFSO data and column 4 replicates this exercise using expenditure shares from the Homescan data. Each of these columns reports shares that are more similar than using different expenditure weights (i.e., compared to the first column).

These differences in expenditure shares are largely accounted for by expenditures on items other than food and non-alcoholic beverages. The latter category has a high import share (56.4% in the SFSO data). While these categories account for 31.2% of expenditures in the SFSO data, they represent only 21.3% of expenditures in the Homescan data. Such differences may reflect that the SFSO adjusts expenditure shares for tobacco or that the Homescan sample captures food and beverage expenditures better than non-food grocery expenditures (e.g., medicines, household equipment, cosmetics, personal care appliances),

Table A4: Import shares in groceries using SFSO and Homescan data

	All categories	Matched categories		
	Own weights	Unweighted	SFSO weights	Homescan weights
SFSO import shares	37.9	43.5	37.9	26.4
Homescan import shares	26.9	39.4	31.0	26.2

Notes: Column 1 displays the grocery import share in SFSO and in Homescan data. Column 2–4 use a subset of each data set (44 common categories) that can be matched. Column 2 displays the unweighted average import share across common categories. Columns 3 and 4 display the weighted average import share across common categories weighted using expenditure weights from SFSO data in column 3 and from Homescan data in column 4.

which also tend to be purchased in non-grocery retail outlets.

Appendix B Empirical appendix

B.1 Details on stylized facts

Here we provide additional details, tables, and figures associated with our stylized facts presented in Section 2.2.

SF 1 Part 2 (SFSO and Homescan): Import shares within groceries before the 2015 CHF appreciation were not strongly correlated with income.

In Section 2.2 we show that import shares are higher among higher-income households in the 2012–14 SFSO data. On the other hand, we state that import shares within groceries are not strongly correlated with income.

This is evident in the SFSO data from Table 1. We next show that import shares are also not significantly correlated with income in the product-level Homescan data. To do so, we estimate

$$100 \times \frac{X_{hM}}{X_{hM} + X_{hD}} = \alpha + \beta \log(Income_h) + [\zeta' K_h] + \varepsilon_h \tag{A1}$$

where X_{hM} and X_{hD} are household h's expenditure on imports and domestic goods in 2014, $\log(Income_h)$ is the logarithm of household h's income in 2014, and K_h is a vector of household controls. Robust standard errors are clustered by income quantiles (of which there are fifty) and observations are weighted by the product of the number of households in each income quantile times the household's share of expenditure in 2014 within its quantile. The coefficient β identifies the difference in import shares in 2014 between higher- and lower-income households. Table A5 displays the results, which are insignificantly different from zero whether or not we control for additional household characteristics.

SF 2 (Homescan): The import share increased following the 2015 CHF appreciation.

Table A5: Household income and import shares in Homescan in 2014

	(1)	(2)	(3)
log(Income)	-0.06	0.50	0.42
	[0.46]	[0.51]	[0.52]
Observations	3302	3302	3302
Control size		X	X
All controls			X

Notes: Estimation of equation (A1). Column 2 controls for household size. Column 3 additionally controls for an indicator for whether there is a child under 10 and an indicator if everyone in the HH is older than 70. Robust standard errors are clustered by income quantile (of which there are fifty) and observations are weighted by the product of the number of households in each quantile $\times$ the household's share of expenditure in 2014 within its income quantile. *p<.1; **p<.05; ***p<.01

The aggregate import share increased from 26.9% to 27.5% between 2014 and 2015. To show that this rise occurred within individual households—rather than from a change in the composition of expenditures across households—we estimate the following regression

$$100 \times \frac{X_{hMt}}{X_{hMt} + X_{hDt}} = \alpha + \mathbb{FE}_h + \sum_{y \neq 2014} \beta_t \mathbb{I}_{y=t} + \varepsilon_{ht} \tag{A2}$$

where X_{hMt} and X_{hDt} are expenditures on imports and domestic goods for household h in year t, $\mathbb{FE}_h$ is a household fixed effect that controls for systematic differences across households in import shares, and $\mathbb{I}_{y=t}$ is an indicator that equals one if $y = t$. Robust standard errors are clustered by income quantiles (of which there are fifty) and observations are weighted by the product of the number of households in each income quantile times the household's share of expenditure in 2014 within its quantile. The coefficients β_t identify the change within households in the share of expenditures on imports between year t and 2014.

Figure A1 displays our estimated year fixed effects, β_t, together with their associated 95% confidence intervals when estimating regression (A2) separately for each of twelve horizons, where we define horizon j as the first j months in year t and in 2014; our annual regressions are equivalent to horizon 12. Over the full year, there are no economically or statistically significant differences between 2013 and 2014. On the other hand, within households the import share was higher in 2015 than it was in 2014—the increase in the import share in 2015 is largely stable over all twelve horizons—and this persists through 2016.

SF 3 (Homescan): Import shares increased less for higher-income households following the 2015 CHF appreciation.

Figure A1: Plotting β_t by time horizon from equation (A2)

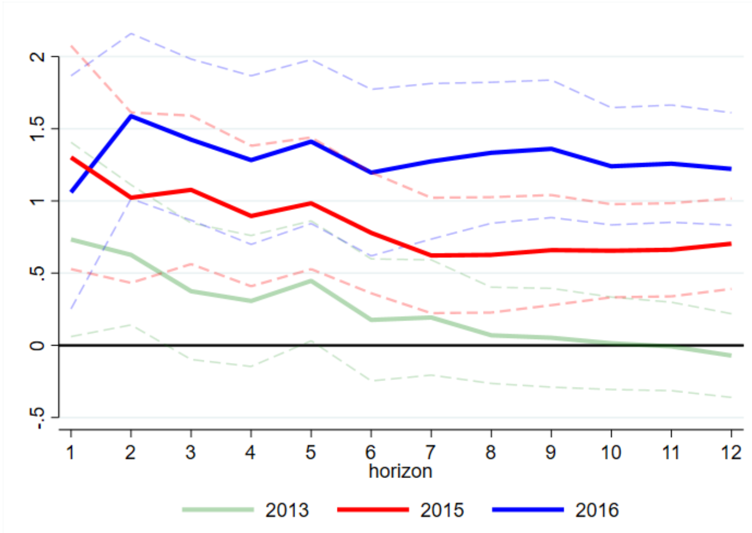

Notes: Estimation of equation (A2) separately by horizon (for horizons 1–12), showing estimated coefficients, β_y, and associated 95% CIs. Robust standard errors are clustered by income quantiles (of which there are fifty) and observations are weighted by the product of the number of households in each income quantile times the household's share of expenditure in 2014 within its quantile.

Table 2 reports results from estimating the following household level regression

$$100 \times \frac{X_{hMt}}{X_{hMt} + X_{hDt}} = \mathbb{FE}_t + \mathbb{FE}_h + \sum_{y \neq 2014} \mathbb{I}_{y=t} \left[\beta_t Inc_h + [\zeta_t' K_h] \right] + \varepsilon_{ht} \tag{A3}$$

where $\mathbb{FE}_h$ and $\mathbb{FE}_t$ are household and time fixed effects that soak up any systematic differences in import shares across households or years, $\mathbb{I}_{y=t}$ is an indicator that equals one if $y = t$, K_h is a vector of household controls, and Inc_h is a measure of household h's income in 2014.[11] The coefficient β_t identifies the difference-in-difference—between year t and 2014 and between higher relative to lower-income households—in the log of imports relative to domestic purchases.

SF 4 (Homescan): The price of imported relative to domestic goods fell following the 2015 CHF appreciation. Neither import nor domestic price changes varied systematically with household income.

We measure the monthly log price of each barcode product as the average of log prices across all transactions, weighing transactions by expenditures within the relevant month. The average change in log prices—relative to December 2014—within the set of domestic goods and, separately, the set of imports is identified estimating the following regression

[11] The additional controls that are interacted with year are: household size, an indicator for whether there is a child under 10, and an indicator if everyone in the HH is older than 70.

Figure A2: Price changes and household income

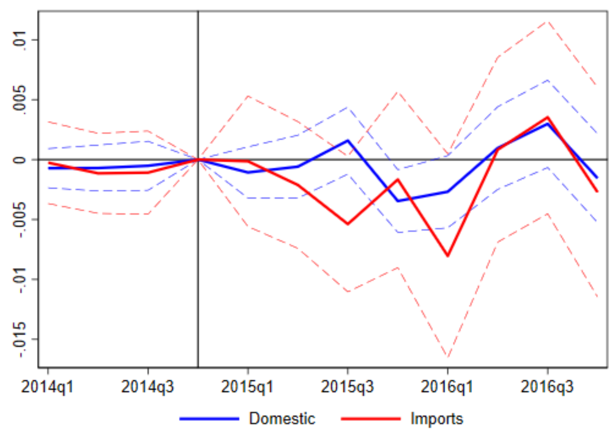

Notes: Estimation of (A5) displaying estimated coefficient β_q and associated 95% confidence interval for each quarter q. Coefficients for imported and domestic goods are estimated separately. Robust standard errors are clustered by product and observations are weighted by 2014 expenditures by group j on product i.

separately for domestic and imported goods,

$$\log p_{im} = \alpha + \mathbb{FE}_i + \sum_{m' \neq \text{Dec 2014}} \mathbb{I}_{m'=m} \beta_m + \varepsilon_{im} \tag{A4}$$

where i indexes product and m indexes month. We weigh each observation by total expenditure on that product in 2014. The coefficient β_m identifies the average difference in product prices—separately for imported and domestic goods—between month m and December 2014. Figure 1 in the Introduction displays our results with robust standard errors clustered at the product level. Before the 2015 appreciation, import prices and domestic prices moved together. Following the appreciation, import prices fell by approximately 2.1% relative to domestic prices (averaging the change between December 2014 and each month in 2015).

Did prices paid change differentially for households with different incomes? Separately on the sample of imported goods and domestic goods, we estimate

$$\log p_{ihq} = \alpha + \alpha_{ih} + \alpha_q + \sum_{y \neq 2014Q4} \mathbb{I}_{y=q} \beta_q \log(Income_h) + \varepsilon_{ihq} \tag{A5}$$

where p_{ihq} is the level of the price of product i paid by household aggregation h (defined as the 50 income quantiles) in quarter q.[12] We measure this price as the geometric weighted average product price across transactions within hq, weighing by expenditures in the cur-

[12]We aggregate up from months in (A4) to quarters in (A5) given the finer disaggregation across incomes in (A5).

Table A6: A lack of systematic price variation across space

	(1) Geneva&Valais	(2) Neuchatel	(3) Berne	(4) Basel	(5) Aarau	(6) Central CH	(7) Grisons	(8) Eastern CH
Region FEs	0.005***	0.002*	0.003***	0.000	-0.000	-0.005***	-0.001	-0.002*
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)

Notes: Estimation of equation (A6). Observations weighted by expenditure in 2014. Clustered by product. The one-digit zip code containing Zurich (which is the most populous) is omitted.

rent quarter. We weigh observations in (A5) by 2014 expenditures by household aggregation h on product i and cluster standard errors by product. The coefficient β_q identifies the difference-in-difference—between quarter q and the fourth quarter of 2014 and between higher- relative to lower-income households—in the average log price.

Results for the differences-in-differences coefficients, β_q, are displayed visually in Figure A2. As indicated in the figure, point estimates are economically small and statistically insignificantly different from zero. Changes in prices paid at the individual product level surrounding the 2015 appreciation do not differ systematically across incomes.

A related observation is that average price *levels* do not vary much across regions in Switzerland in 2014. To document this fact, we estimate

$$\log p_{ij} = \alpha + \mathbb{FE}_i + \mathbb{FE}_j + \varepsilon_{ij} \tag{A6}$$

where $\log p_{ij}$ is the weighted average log price for domestic purchases in 2014 of product i within one-digit-zip code j, $\mathbb{FE}_j$ is a one-digit zip-code specific fixed effect, and $\mathbb{FE}_i$ is a product-specific fixed effect. We weigh observations by expenditure in 2014 and cluster by product.

Table A6 displays our estimated one-digit-zip code fixed effects. The omitted fixed effect is for the most populous one-digit zip (which contains Zurich). There are at most tiny systematic differences in average prices across regions (conditioning on the range of offered products), with the greatest difference from Zurich being half of one log point.

SF 5 (Homescan): The dispersion of price changes across goods was greater between 2014–15 than between 2013–14 and than between 2015–16, especially within the set of imported goods.

For each barcode product we calculate an expenditure-weighted average log price across all transactions for each year t. For each product and each t in 2013, 2014, and 2015 we then calculate the change in log price between t and $t+1$. For each t, we restrict our sample either to imported (column 1), domestic (column 2), or all goods (column 3). To reduce the role of abnormally large price changes, we drop products with year-to-year price ra-

Table A7: Standard deviation of log price changes across barcode products

	Balanced all years			Expend. weight 2014		
	Imports	Domestic	All	Imports	Domestic	All
2013–14	0.054	0.042	0.046	0.053	0.045	0.048
2014–15	0.074	0.048	0.057	0.074	0.048	0.057
2015–16	0.061	0.041	0.047	0.062	0.042	0.048

Notes: Expenditure-weighted standard deviation of annual log price change across barcode products for imported goods, domestic goods, and all goods. The left panel includes only products purchased in all years 2013, 2014, 2015, and 2016. The right panel additionally uses common weights (given by 2014 expenditures) across years.

tios above 3 or below 1/3. We then construct the weighted standard deviation of log price changes for each t, weighting by expenditures in year t. Results are shown in Table 3. Restricting the sample to the set of products that were purchased in all three years, or additionally imposing common weights (given by 2014 expenditures) across t leaves results broadly unchanged, as shown in Table A7.

B.2 Additional details for robustness of estimation of η_s

In this section, we describe a range of additional robustness and sensitivity exercises focusing on our second approach to estimating η_s. In the first exercise, we consider an alternative estimation approach—using cross-sectional income elasticities—that allows us to relax the restriction that the good-specific component of income elasticities is common across households in the initial period. In the second set of exercises, we vary specific baseline choices and show that our baseline point estimate is robust. Third, we show that our results are robust if we do not infer household income using household characteristics beyond Homescan income bin or if we drop high- or low-income households from our estimation (both in approaches 1 and 2). In the final set of exercises, we demonstrate that our results are robust to incorporating spatial variation in both expenditures and prices.

Using cross-sectional income elasticities. In our baseline approach, we estimate differences in Hicksian elasticities without the need to first estimate income elasticities in the cross section. This approach leverages restriction (15), which imposes that the good-specific component of income elasticities is common across households in the initial period. In this sensitivity (mentioned in footnotes 24 and 30), we consider an alternative estimation approach that allows us to relax this restriction. This alternative approach involves first estimating cross-sectional income elasticities. We apply this alternative procedure in Approach 1 and obtain very similar results.

Define

$$\kappa_{hi} \equiv \left(\frac{\partial \log e_h}{\partial \log u_h}\right)^{-1} \times \left(\gamma_i - \frac{\partial \eta_s}{\partial \log u_h} \log p_{hit_0}\right)$$

where all derivatives in this section are evaluated at t_0. In our baseline procedure, we assume $\kappa_{hi} = \kappa_i + \kappa_{hs}$, so that the income elasticity for good i at t_0 can be expressed as the sum of κ_i, which is common for all households, and a household-sector-specific term. Here we drop this restriction. Equation (18) becomes

$$d \log b_{hit} = \mathbb{FE}_{it} + \mathbb{FE}^M_{hst} + \kappa_{hi} d \log \left(\frac{I_{ht}}{P_{ht}}\right) - \eta_s \log(I_{ht_0}) d \log p_{hit} + \iota_{hit}. \tag{A7}$$

In our modified procedure, we first estimate κ_{hi} from the cross section and then estimate η_s.

We implement this procedure in Approach 1, where we only need to estimate a single income elasticity, that of imports relative to domestic goods: $\kappa_h^{MD} \equiv \kappa_{hM} - \kappa_{hD}$.[13] Given an estimate of κ_h^{MD}, we obtain η_s from a modified version of equation (19):

$$d \log \left(\frac{b_{hMt}}{b_{hDt}}\right) - \kappa_h^{MD} d \log \left(\frac{I_{ht}}{P_{ht}}\right) = \alpha - \eta_s \log(I_{ht_0}) d \log \left(\frac{p_{Mt}}{p_{Dt}}\right) + \iota_{ht} \tag{A8}$$

To estimate κ_h^{MD}, we estimate a standard Engel-curve regression for the share of imports relative to domestic goods by household in $t_0 = 2014$,

$$\log \left(\frac{b_{hMt_0}}{b_{hDt_0}}\right) = \beta_0 + \beta_1 \log(I_{ht_0}) + \beta_2 \log(I_{ht_0})^2 + X_h' \gamma + \iota_{ht_0} \tag{A9}$$

where we have imposed that $\kappa_h^{MD} = \beta_1 + 2 \times \beta_2 \log(I_{ht_0})$. Tastes for imports relative to domestic goods at t_0 must be uncorrelated with income conditional on other controls, X_h. We do not require this assumption in our baseline approach.

If we impose $\beta_2 = 0$, we are back to our baseline assumption that the good-specific-component of income elasticities is common across households. However, even in this case, the approach differs from our baseline approach, as we estimate κ_h^{MD} using cross-sectional rather than time-series variation. In the case of $\beta_2 = 0$, regression (A9) is very similar to regression (A1), where the dependent variable is $\frac{b_{hMt_0}}{b_{hMt_0} + b_{hDt_0}}$ rather than $\log\left(\frac{b_{hMt_0}}{b_{hDt_0}}\right)$.

Table A8 displays estimates of η_s obtained using this procedure. Columns 1–3 contain results imposing $\beta_2 = 0$ varying the set of controls and columns 4–6 contain results without this restriction. Across specifications, estimates are very similar to our baseline. Standard errors should be interpreted with caution since we do not adjust for the fact that

[13] In Approach 2, we would have to estimate thousands of income elasticities at the barcode product level.

Table A8: Robustness of Approach 1: Using cross-sectional Engel curves

	Linear			Quadratic		
	(1)	(2)	(3)	(4)	(5)	(6)
$\log(I_{ht_0}) d \log(p_{Mt}/p_{Dt})$	2.17***	2.16***	2.17***	2.09***	2.06***	2.08***
	[0.54]	[0.54]	[0.54]	[0.54]	[0.54]	[0.54]
Observations	2901	2901	2901	2901	2901	2901
Control size		X	X		X	X
All controls			X			X

Notes: Results of estimating equation (A8) using estimates of κ_h^{MD} estimated using equation (A9) (in 2014 data) under the assumption that $\beta_2 = 0$ in columns 1-3 and without this assumption in columns 4–6. Columns 1 and 4 include no controls. Columns 2 and 5 include household size controls. Columns 3 and 6 additionally include an indicator for whether there is a child under 10 and an indicator for whether everyone in the HH is older than 70. The regression (A8) is clustered and weighted as in the baseline of Approach 1. The regression (A9) is unweighted. Standard errors in this table do not correct for the fact that the dependent variable depends on an estimated coefficient. *p<.1; **p<.05; ***p<.01

Table A9: Robustness of Approach 2: Varying baseline choices

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\log(I_{ht_0}) \times d \log p_{it}$	1.93**	1.80**	1.87*	2.03*	1.61*	1.89**	1.87*
	[0.87]	[0.73]	[0.97]	[1.12]	[0.94]	[0.89]	[0.99]
Observations	95,325	95,325	95,325	120,889	97,366	92,383	95,325
Baseline	X						
Winsorize 5%		X					
No winsorizing			X				
Unbalanced sample				X			
Sample >20 border prices					X		
Sample >32 border prices						X	
Prices rel to 14Q4							X
K-P F Stat (first stage)	13.1	13.1	13.1	16.1	8.7	13.9	12.1

Notes: Column 1 replicates our baseline 2SLS estimate of η_s in column 3 of Table 5. Columns 2–7 each vary one choice in our baseline specification. Column 2 winsorizes at the 5th percentile whereas column 3 does not winsorize at all. Column 4 drops the sample restriction that a product is only included if it was purchased at least once per month in the year and a half before and after the CHF appreciation. Column 5 (column 6) includes products in border groups with more than 20 (more than 32) border price observations. Column 7 defines $dlogp_{it}$ as the log price change between 2015 and the fourth quarter of 2014. *p<.1; **p<.05; ***p<.01

the left-hand side of equation (A8) depends on a coefficient estimated in the first step of the procedure.

Varying baseline choices. Column 1 of Table A9 displays our baseline 2SLS estimate and the remaining columns display results from various robustness exercises. In our baseline we winsorize changes in log expenditures at the first percentile (in the top and bottom tails). In columns 2 and 3 we instead winsorize at the 5th percentile and not at all. Our baseline sample only includes products if they were purchased at least once per month in the year-and-a-half before and after the CHF appreciation. In column 4 we drop this sample restriction. Our baseline sample only includes products in border groups for which there are more than 28 border price observations in 2014. In columns 5 and 6 we include additional border groups (those with more than 20 border price observations) and fewer

Table A10: Robustness of Approach 2: Varying baseline choices

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\log(I_{ht_0}) \times d\log p_{it}$	1.93**	3.55	2.62**	2.27**	1.97***	1.62*	1.83**	2.35*
	[0.87]	[2.23]	[1.10]	[1.01]	[0.62]	[0.94]	[0.85]	[1.33]
Observations	95,325	43,559	67,179	82,995	116,930	95,325	95,325	95,325
Baseline	X							
Horizon 3m		X						
Horizon 6m			X					
Horizon 9m				X				
Percent change					X			
Omit $d\log(I_{ht}/P_{ht})$						X		
All inv. currencies							X	
HH size interaction								X
K-P F Stat (first stage)	13.1	7.6	8.6	11.4	12.8	12.8	12.7	14.6

Notes: Column 1 replicates our baseline 2SLS estimate of η_s in column 3 of Table 5. Columns 2–7 each vary one choice in our baseline specification. Columns 2–4 use price and expenditure changes measured over the first 3, 6, and 9 months of 2014 and 2015. Column 5 replaces log changes in expenditures and in prices with percent changes. Column 6 omits the covariate $d\log(I_{ht}/P_{ht})$ from the regression. Column 7 uses an alternative instrument using the share of non-CHF invoiced border prices, including all currencies. Column 8 includes a control for household-size interacted with the change in product price, instrumented using a version of our baseline instrument replacing the log of household income with household size. In column 8, the reported F statistic is the SW F on $\log(I_{ht_0})d\log p_{it}$. The unreported SW F stat on the household-size interaction is over 14. *p<.1; **p<.05; ***p<.01

border groups (those with more than 32 border price observations). In our baseline, we use price changes and expenditure changes defined using the full years of 2014 and 2015. In column 7 we use retail price changes between the fourth quarter of 2014 and the first quarter of 2015 as calculated in Auer et al. (2021) and changes in expenditures over the full years of 2014 and 2015.[14] Each of these choices has little effect on either first-stage or second-stage results.

Column 1 of Table A10 again displays our baseline 2SLS estimate and the remaining columns display results from additional robustness exercises. In columns 2, 3, and 4 we use price changes and expenditure changes measured over the first 3, 6, and 9 months of 2014 and 2015. In all cases, changes in real income are still measured over the full year given data availability. Results remain largely stable across these specifications; the elasticity difference is larger when estimated using changes in expenditures and prices over the first 3 months, but it is not precisely estimated.

In our baseline, we use log changes in prices and in expenditure shares. This approach drops all observations for which initial (i.e. 2014) or terminal (i.e. 2015) expenditures are zero. In column 5, we replace log changes in expenditures and in prices with percent changes. This alternative approach keeps any observation for which consumption in 2014 is positive (as long as any household in any income group consumes the product in 2015). Our main result is largely unchanged. In our baseline, we control for changes in real in-

[14]For each product, we first calculate average retail prices by region, retailer, and month, then average these across regions and retailers by month, and finally average monthly prices by quarter.

come. If we omit this covariate, our estimated difference in elasticities falls; see column 6. Our baseline instrument uses the share of imported goods in each border group that are denominated in EUR out of all goods denominated in either EUR or CHF. If we instead use the share of non-CHF invoiced border prices including all currencies, results are largely unchanged as shown in column 7. Another concern is that household income is correlated with household size and that households of different sizes have different elasticities. In column 8 we control for the interaction between household size and the log product price change and instrument for this interaction using a version of our baseline instrument in which we replace log income with household size. The SW F stats for both endogenous variables are above 14 and our main result is largely unchanged.

In our baseline, we two-way cluster standard errors at the level of household income and, separately, the interaction between import status and the value of the share of imported goods denominated in EUR in the corresponding border group. Here, we report how the first-stage F statistic and the second-stage standard error vary with these choices. If we two-way cluster standard errors at the level of household income and, separately, the barcode product, the first-stage F statistic is approximately 25. If we two-way cluster standard errors at the level of household income and, separately, the interaction between import status and the border group (rather than by the share of imported goods denominated in EUR in the border group, which makes a difference because 7 of the 35 border groups have a common EUR invoicing share equal to zero), the first-stage F statistic is approximately 6. If we one-way cluster standard errors at the level of the triple interaction between import status, the share of products denominated in EUR in the border group, and household income, the first-stage F statistic is well over 100. If we one-way cluster standard errors at the level of the interaction between import status and the share of products denominated in EUR in the border group, the first-stage F statistic is largely unchanged. In all cases, the second-stage standard error is very similar to its value in our baseline.

Alternative measures of household income. In our baseline in Approaches 1 and 2, we infer household income and changes in income combining Homescan information on household characteristics, the Swiss Household Panel (FORS), and the Swiss Federal Tax Administration (2014) (SFTA). Here, we replicate our baseline estimation of Approaches 1 and 2 using two alternative methodologies to measure household income.

In the first alternative methodology to measure household income, we assign each household in a given Nielsen income bin a common income level equal to the median level of income associated with that income bin in the SFTA data (as described in Section A.2). Because we do not use FORS to infer household income, we similarly do not use it to infer changes in household income; hence, we omit the covariate measuring changes in real income from both approaches in this robustness. Finally, because we only have 7 income bins in these exercises, we do not two-way cluster including income; instead, we one-way cluster. Assigning households to these incomes leaves our baseline estimate in Approach 1 largely unchanged and slightly increases the estimate in Approach 2, as shown in the left panel of Table A11.

In the second alternative methodology to measure household income, we predict household income almost exactly as in our baseline. However, we do not augment the prediction from FORS data using any information from the SFTA data; instead, we keep the estimated income-bin fixed effects estimated in FORS data. Assigning households to these incomes increases our estimates of η_s slightly, as shown in the right panel of Table A11.

Dropping specific income groups. Are the specific income groups driving the variation that identifies differences in elasticities particularly high- or low-income households? In Tables A12 and A13 we replicate our baseline estimation of Approaches 1 and 2, respectively, dropping either all households in the lowest Homescan income group, the two lowest Homescan income groups, the highest Homescan income group, or the two highest Homescan income groups (out of the seven income groups). While Approach 1 is estimated at the household level, in Approach 2 we combine households into 50 aggregates. Hence, in Approach 2, rather than dropping individual households and then reconstructing 50 new aggregates across the remaining households, we instead start from the same 50 aggregates and drop the minimal number of these such that we drop all households in the relevant Nielsen income bins.

Across the eight cases (two approaches and dropping four distinct sets of households), we obtain a positive coefficient. This coefficient is similar to our baseline estimates in all cases but one (dropping the two lowest income groups in Approach 1, where we lose almost a third of our observations). Our estimates, however, are less precise, especially when we drop the lowest income groups. We conclude that the negative relationship between incomes and price elasticities is not driven by either high- or low-income households; although for precision, low-income households play an important role.

Incorporating spatial variation. In our baseline we did not incorporate geography at all. We aggregated households by 2014 income alone and, therefore, used common price changes within each individual product across household aggregates.

Here, we show that further disaggregating our household groups by both geography and income leaves our results largely unchanged. Column 1 of Table A14 replicates our baseline 2SLS result from column 3 of Table 5. In the remaining columns in Table A14 we disaggregate households both across 50 income quantiles (as before) and across each of 9 one-digit zip codes in Switzerland; our regression specification incorporates correspond-

Table A11: Robustness of Approaches 1 and 2 to inferring household income

	Without using FORS data		Without using SFTA data	
	(1) Approach 1	(2) Approach 2	(3) Approach 1	(4) Approach 2
$\log(I_{ht_0}) d \log(p_{Mt}/p_{Dt})$	2.14^{***}		2.60^{***}	
	[0.54]		[0.72]	
$\log(I_{ht_0}) \times d \log p_{it}$		2.29^{***}		2.27^{**}
		[0.72]		[1.01]
Observations	2,901	19,881	2,901	95,325
K-P F Stat (first stage)		11.2		12.6

Notes: We replicate our baseline in Approaches 1 and 2, inferring household incomes differently. In columns 1 and 2 we assign a common value of household income across all households in each of the 7 Homescan income bins equal to the median income in that bin, as described in the text. In both columns we omit the covariate measuring changes in real income and do not include income in our clustering. In columns 3 and 4 we assign income without using data from SFTA, as described in the text. *p<.1; **p<.05; ***p<.01

Table A12: Robustness of Approach 1 dropping household income ranges

	(1) Drop lowest	(2) Drop 2 lowest	(3) Drop 2 highest	(4) Drop highest
$\log(I_{ht_0}) d \log(p_{Mt}/p_{Dt})$	1.58	0.72	2.52^{***}	2.52^{***}
	[1.29]	[1.55]	[0.42]	[0.46]
Observations	2569	2085	2460	2872

Notes: Each column of this table replicates column 1 of Table 4 while omitting a subset of the estimation sample. Column 1 drops all households in the lowest Homescan income bin whereas column 2 additionally drops the second-lowest income bin. Column 4 drops all households in the highest Homescan income bin whereas column 3 additionally drops the second-highest income bin. *p<.1; **p<.05; ***p<.01

Table A13: Robustness of Approach 2 dropping household income ranges

	(1) Drop lowest	(2) Drop 2 lowest	(3) Drop 2 highest	(4) Drop highest
$\log(I_{ht_0}) \times d \log p_{it}$	2.60	2.29	1.55	1.69^{**}
	[1.84]	[1.73]	[1.21]	[0.80]
Observations	83,897	69,301	79,007	93,126
K-P F Stat (first stage)	12.0	12.3	13.4	13.2

Notes: Each column of this table replicates column 3 of Table 5 (the baseline 2SLS estimate in Approach 2) while omitting a subset of the estimation sample. Column 1 drops the minimum number of the 50 household aggregates such that all households in the lowest Homescan income bin (out of 7) are excluded, whereas column 2 drops additional aggregates to exclude all households in the second-lowest income bin. Column 4 drops the minimum number of the 50 household aggregates such that all households in the highest Homescan income bin are excluded, whereas column 3 drops additional aggregates to exclude all households in the second-highest income bin. *p<.1; **p<.05; ***p<.01

Table A14: Robustness of Approach 2: Incorporating spatial variation

	(1)	(2)	(3)
$\log(I_{ht_0}) \times d\log p_{it}$	1.930**	2.170***	
	[0.867]	[0.663]	
$\log(I_{ht_0}) \times d\log p_{hit}$			1.542***
			[0.572]
Observations	95,325	134,596	134,596
Baseline	X		
Spatial variation: outcome		X	X
Spatial variation: price			X
K-P F Stat (fist stage)	13.1	12.4	18.5

Notes: Columns 1 replicates our baseline 2SLS estimate of η_s in column 3 of Table 5 in which an observation is a product $\times$ household income quantile (of which there are fifty). In columns 2 and 3 we further disaggregate households by one-digit zip code and in column 3 we measure product-specific price changes separately across each one-digit zip code. In Columns 2 and 3 we two-way cluster by the interaction between import status and the share of imported goods that are denominated in EUR and, separately, the household aggregation (income quantile $\times$ one-digit zip code). *p<.1; **p<.05; ***p<.01

ingly more disaggregated household fixed effects, where h is now the interaction between the income quantile and zip code. Column 2 displays the results of estimating the baseline specification—continuing to use a common price change within each good—using this more disaggregated data; first- and second-stage results are largely unchanged. In column 3, we additionally use price changes measured separately within each of the 9 one-digit zip codes. Incorporating price variation across regions leads to a modest attenuation in our baseline estimate of η_s (from -1.93 to -1.54) and our instrument remains strong.

Finally, we describe an alternative instrument leveraging spatial price variation, a Hausman instrument interacted with household income. Using this instrument, we find much smaller differences in elasticities across incomes. We also show that this Hausman instrument may be endogenous in our particular Swiss setting (where there is little price variation across space).

In a first step, we omit our cost-shock instrument and use an alternative: the interaction between a Hausman instrument and initial log income. Specifically, for households in a particular income quantile $h \in \{1, ..., 50\}$ living in a particular one-digit zip code $j \in \{1, ..., 9\}$, we instrument for the interaction between the income of quantile h and the product-specific price change in one-digit zip code j using the income of quantile h and the product-specific price change measured outside of j. The instrument is very strong, with an F statistic of over 250. The very strong first stage can be understood by the fact that there is very little variation in regional prices of individual products set by the major national retailers in Switzerland. This also explains why this specification yields very similar estimates to the baseline OLS using common national price changes displayed in column 1 of Table 5. In particular, the second-stage coefficient of interest, $\eta_s = 0.093$, is over an order of magnitude smaller than our baseline 2SLS estimate.

The exclusion restriction when using a Hausman instrument—without interacting with income—is that there are no product-specific demand shocks at the national level that are correlated with price changes whereas the exclusion restriction when using a cost-shock instrument is that the cost shock is uncorrelated with demand shocks. Given that we are over-identified—with two instruments and one endogenous variable—we can use Hansen's (1982) J test, an over-identification test of all instruments: the joint null hypothesis is that all instruments are valid. Estimating (18) using both instruments, we obtain a Hansen J statistic of 5.739 and a Chi-sq p value of 0.0166, thus rejecting the null hypothesis that both instruments are exogenous. Given that cost-based instruments are the gold-standard in demand estimation—or 'textbook instrumental variables' as Nevo (2000) refers to them—one conclusion might be that the Hausman-based instrument is endogenous in our setting. Of course, even if the Hausman-based instrument is endogenous in our setting, that does not imply endogeneity in other contexts.

B.3 Estimating $\bar{\eta}_s$

Neither of the two approaches in Section 4 identify the intercept $\bar{\eta}_s$ defined in equation (16). However, under stronger assumptions they can be adjusted to do so.

In our first approach in Section 4.2 using equation (19), if we assume that the average import demand shifter ν_{it} is zero between 2014 and 2015, then $\bar{\eta}_s$ is identified from the constant α as $\bar{\eta}_s = 1 - \alpha / \left(d \log \left(p_{Mt}/p_{Dt}\right)\right)$. Given $d \log \left(p_{Mt}/p_{Dt}\right) = -0.0216$ and the constant displayed in column 1 of Table 4, we obtain $\bar{\eta}_s \approx 26.6$. Together with our estimate of $\eta_s = -2.189$ from this approach, this implies that the initial elasticity of substitution is 4.92 for a household with income of 20,000 CHF and that this elasticity remains positive for all household incomes below approximately $190,000$ CHF.

In our second approach in Section 4.3 we cannot recover $\bar{\eta}_s$ without moving the average product-specific demand shock ν_{it} to the residual. In this case, rather than re-estimate η_s under a stronger exclusion restriction, we subtract the estimated price interaction from both the left- and right-hand sides of equation (18) and then instrument for the log change in product price using our cost shifter. In our baseline we obtain $\bar{\eta}_s = 20.87$. In combination with the baseline estimate of $\eta_s = -1.930$, the initial elasticity of substitution for a household with income of 20,000 in 2014 is 1.76 and this elasticity remains positive for all household incomes below approximately 50,000 CHF.

The *levels* of initial elasticities of substitution (e.g., 4.92 and 1.76 in approaches 1 and 2 for a household with income of 20,000) are much less stable than the implied *differences* across household incomes across approaches (e.g., 2.40 and 2.12 in approaches 1 and 2 comparing across households with income differences of a factor of three).[15]

Appendix C Theoretical appendix

We use a particular formulation of the non-homothetic CES preferences presented in Fally (2022). Given the consumption bundle c_{ht} and preference parameters ζ_{ht} for household h at time t, utility u is implicitly given by

$$f_h\left(u\right)^{\frac{\rho-1}{\rho}} = \sum_s \left(\zeta_{hst} u^{\gamma_s}\right)^{\frac{1}{\rho}} \left(c_{hst}\right)^{\frac{\rho-1}{\rho}}, \tag{A10}$$

where

$$c_{hst} = \left(\sum_{i \in \mathcal{I}(s)} \left(\zeta_{hit} u^{\gamma_i}\right)^{\frac{1}{\eta_s(u)}} \left(c_{hit}\right)^{\frac{\eta_s(u)-1}{\eta_s(u)}}\right)^{\frac{\eta_s(u)}{\eta_s(u)-1}}, \tag{A11}$$

$f_h(\cdot) > 0$ and $\rho, \eta_s(\cdot) \in [0,1) \cup (1,\infty)$. These preferences reduce to nested homothetic CES if, for example, $\eta_s(u)$ is independent of u, $\gamma_i = \gamma_s = 0$, and $f_h'(u) > 0$. The household chooses $\{c_{hit}\}$ to maximize u subject to the budget constraint $I_{ht} = \sum_i p_{it} c_{hit}$. The expenditure function associated with these preferences is given by (3). The maximum utility achieved by household h at time t is $v_h(p_{ht}, I_{ht}; \zeta_{ht}) \equiv u_{ht}$ where $e\left(p_{ht}, u_{ht}; \zeta_{ht}\right) = I_{ht}$. We discuss below conditions that ensure that the expenditure function is monotonic in u.

Deriving equation (13). Log-linearizing $I_{ht} = e_h\left(\mathbf{p}_{ht}, u_{ht}; \zeta_{ht}\right)$ at t_0 yields

$$d\log I_{ht} = \frac{\partial \log e_h}{\partial \log u_h} d\log u_{ht} + \sum_i b_{hit_0} d\log p_{hit} + \bar{\varepsilon}_{ht},$$

where $\bar{\varepsilon}_{ht} \equiv \sum_i \frac{\partial \log e_h}{\partial \zeta_{hi}} d\zeta_{hit}$ and derivatives are evaluated at t_0. Solving for $d\log u_{ht}$ yields

$$d\log u_{ht} = \left(\frac{\partial \log e_h}{\partial \log u_h}\right)^{-1} \times \left(d\log I_{ht} - \sum_i b_{hit_0} d\log p_{hit} - \bar{\varepsilon}_{ht}\right) \tag{A12}$$

This is equation (13) in the text.

[15] In addition to instability of the estimated levels across approaches, each estimate has its own confidence interval. In the first approach, the estimated value of $\bar{\eta}$ is highly sensitive to the estimated constant. A one standard deviation change in the regression constant (0.129), moves the level of $\bar{\eta}$ by $5.97 \approx 0.129/0.0216$. In the second approach, we do not report standard errors because it is not straightforward to do so with a dependent variable that depends on previous estimates, two-way clustering, and a large set of fixed effects.

Deriving equation (17). Substituting equation (A12) into equation (12) yields

$$d\log b_{hit} = \left(\frac{\partial \log e_h}{\partial \log u_h}\right)^{-1} \times \left(\gamma_i - \frac{\partial \eta_s}{\partial \log u_h}\log p_{hit_0}\right)\left(d\log I_{ht} - \sum_i b_{hit_0} d\log p_{hit} - \bar{\varepsilon}_{ht}\right)$$
$$+ d\log \zeta_{hit} + (1-\eta_{hst_0})d\log p_{hit} + \psi_{hst}$$

The previous expression and assumption (15) yield

$$d\log b_{hit} = (\kappa_i + \kappa_{hs})\left(d\log I_{ht} - \sum_i b_{hit_0} d\log p_{hit} - \bar{\varepsilon}_{ht}\right)$$
$$+ d\log \zeta_{hit} + (1-\eta_{hst_0})d\log p_{hit} + \psi_{hst}$$

Note that the only i-specific term multiplying changes in real income is κ_i. This implies that household h's income elasticity for good i in sector s in the initial period can be expressed as the sum of a good-specific and a household-sector specific component. The previous expression and assumption (16) yield

$$d\log b_{hit} =(\kappa_i + \kappa_{hs})\left(d\log I_{ht} - \sum_i b_{hit_0} d\log p_{hit} - \bar{\varepsilon}_{ht}\right)$$
$$+ d\log \zeta_{hit} + (1-\bar{\eta}_s - \eta_s \log I_{ht_0})d\log p_{hit} + \psi_{hst}$$

The previous expression is equation (17) given the definitions $\nu_{hit} \equiv d\log\zeta_{hit} - \kappa_i\bar{\varepsilon}_{ht}$ and $\widetilde{\psi}_{hst} \equiv \psi_{hst} + \kappa_{hs}\left(d\log(I_{ht}/P_{ht}) - \bar{\varepsilon}_{ht}\right)$. The demand shifter ν_{hit} combines the taste shifter for good i, $d\log\zeta_{hit}$, and the change in utility due to taste shifters, $\bar{\varepsilon}_{ht}$ interacted with the utility elasticity κ_i.

Assumptions (15) and (16). We consider a cardinalization of the utility function that satisfies two properties. First, the elasticity of substitution η is log-linearly related to u_{ht},

$$\eta_{hst} \equiv \tilde{\bar{\eta}}_s + \tilde{\eta}_s \log(u_{ht}). \tag{A13}$$

If $\tilde{\eta}_s < 0$, then a household that attains a higher indifference curve is less price sensitive in sector s. In combination with the assumption that initial prices of individual goods within s are given by $\log p_{hit_0} = \log p_{it_0} + \log p_{hst_0}$ we obtain

$$\frac{\partial \eta_s}{\partial \log u_h}\log p_{hit_0} = \tilde{\eta}_s\left(\log p_{it_0} + \log p_{hst_0}\right)$$

The second property of our utility function is that the elasticity of the expenditure function with respect to u_{ht} in the initial period is common across households. To achieve this outcome, we assume that $f_h(\cdot)$ introduced in (3) is

$$f_h(x) = a_0 x^{a_1} \left[\sum_s \zeta_{hst} x^{\gamma_s} \left(P_{hs}(x) \right)^{1-\rho} \right]^{\frac{1}{\rho-1}} \tag{A14}$$

with $a_0 > 0$ and $a_1 > 0$ and where

$$P_{hs}(x) = \left(\sum_{i \in \mathcal{I}(s)} \zeta_{hit_0} x^{\gamma_i} \left(p_{hit_0} \right)^{1-\eta_s(x)} \right)^{\frac{1}{1-\eta_s(x)}} \tag{A15}$$

In this case, $e_h \left(\mathbf{p}_{ht_0}, u_{ht_0}; \zeta_{ht_0} \right) = I_{ht_0} = a_0 \times u_{ht_0}^{a_1}$ and $\partial \log e_h / \partial \log u_h = a_1$ when evaluated at t_0. These cardinalization assumptions imply equation (16), where $\bar{\eta}_s \equiv \tilde{\bar{\eta}}_s - a_1^{-1} \tilde{\eta}_s \log(a_0)$ and $\eta_s \equiv a_1^{-1} \tilde{\eta}_s$, and also imply equation (15), where $\kappa_i \equiv a_1^{-1} \gamma_i - \eta_s \log p_{it_0}$ and $\kappa_{hs} \equiv -\eta_s \log p_{hst_0}$.

Monotonicity of the expenditure function. For any constant u, the shape of the indifference curves implied by the non-homothetic utility function (A11) is the same as under homothetic CES. Similarly, for any given u, the shape of the expenditure function (3) and corresponding Hicksian demand under non-homothetic CES is the same as under homothetic CES. In order for our utility function to be well-defined there must be a unique solution for u in equations (A10)–(A11). In order for our expenditure function to be well-defined, there must be a unique u that solves $e(p, u; \zeta) = I$, and the expenditure must be increasing in u to ensure budget exhaustion.

We examine these properties first analytically—applying results in Fally (2022)—and then numerically. We focus on the empirically relevant case in which the elasticity of substitution is decreasing in u, in a specification with a single sector (or, equivalently, all sectors are symmetric). In this case, the utility function (A10) is

$$f(u)^{\frac{\eta(u)-1}{\eta(u)}} = \sum_i \left(\zeta_i u^{\gamma_i} \right)^{\frac{1}{\eta(u)}} c_i^{\frac{\eta(u)-1}{\eta(u)}}$$

where we have dropped household and time sub-indices, $\zeta_i \geq 0$ for all i and $\sum_i \zeta_i = 1$. To use the notation of Fally (2022), define $G_i(u) \equiv f(u) \left(\zeta_i u^{\gamma_i} \right)^{\frac{1}{1-\eta(u)}}$, and re-express the utility function as

$$1 = \sum_i \left(c_i / G_i(u) \right)^{\frac{\eta(u)-1}{\eta(u)}} \tag{A16}$$

the expenditure function as

$$e\left(p,u;\zeta\right)=\left(\sum_{i}\left(G_{i}(u)p_{i}\right)^{1-\eta(u)}\right)^{\frac{1}{1-\eta(u)}}, \tag{A17}$$

and demand for good i as

$$\frac{p_{i}c_{i}}{I}=\left(\frac{G_{i}(u)p_{i}}{I}\right)^{1-\eta(u)} \tag{A18}$$

with $\sum_{i}\left(\frac{G_{i}(u)p_{i}}{I}\right)^{1-\eta(u)}=1$.

Proposition 4 in Fally (2022) states that a sufficient condition for the demand system (A18) with $\eta'(u)<0$ to be integrable is

$$K(u)\equiv\sum_{i}\exp\left(\frac{(\eta(u)-1)^{2}}{\eta'(u)}\frac{G_{i}'(u)}{G_{i}(u)}\right)<1. \tag{A19}$$

The proof of Proposition 4 in Fally (2022) shows that if condition (A19) is satisfied, then there is a unique solution u in (A16) and u in $e\left(p,u;\zeta\right)=I$, and that around each of those values of u the expenditure function is increasing in u.

We prove that (A19) is satisfied under our functional form assumption $\eta(u)=\bar{\eta}+\eta\log(u)$ with $\bar{\eta}\neq1$, $\eta<0$, and $f(u)=\left(u^{k_{1}}\right)^{\frac{1}{1-\eta(u)}}$. In this case, $G_{i}(u)=\left(\zeta_{i}u^{\tilde{\gamma}_{i}}\right)^{\frac{1}{1-\eta(u)}}$, where $\tilde{\gamma}_{i}\equiv\gamma_{i}+k_{1}$.[16] Hence,

$$\frac{G_{i}'}{G_{i}}=\log\left(\zeta_{i}u^{\tilde{\gamma}_{i}}\right)+\frac{(1-\eta(u))}{\eta'(u)}\frac{\tilde{\gamma}_{i}}{u}$$

Combining the previous expression with the definition of $K(u)$ yields

$$K(u)\equiv\sum_{i}\zeta_{i}u^{\tilde{\gamma}_{i}}\exp\left(\frac{\tilde{\gamma}_{i}}{u}\frac{1-\eta(u)}{\eta'(u)}\right)$$

Using the functional form $\eta(u)=\bar{\eta}+\eta\log(u)$, the previous expression implies

$$K(u)=K=\sum_{i}\zeta_{i}\exp\left[\tilde{\gamma}_{i}\left(\frac{1-\bar{\eta}}{\eta}\right)\right]$$

Since $\sum_{i}\zeta_{i}=1$, K is a weighted average of $\exp(x_{i})$ for $x_{i}\equiv\tilde{\gamma}_{i}(1-\bar{\eta})/\eta$. If $\bar{\eta}>1$, then $(1-\bar{\eta})/\eta>0$ and $\exp(x_{i})<1$ for all i if $\tilde{\gamma}_{i}<0$ for all i. Hence, if $k_{1}<-\max_{i}\{\gamma_{i}\}$

[16] As in Fally (2022), we do not consider the case of $\eta(u)=1$. To maintain $\eta(u)>1$, we could assume $\eta(u)=\max\{\delta,\eta+\eta_{1}\log(u)\}$ for some $\delta>1$. Here, we do not make this assumption and simply show that (A19) holds in a neighborhood of any u for which $\eta(u)>1$.

then condition (A19) is satisfied. If $\bar{\eta} < 1$, then $(1-\bar{\eta})/\eta < 0$ and $\exp(x_i) < 1$ for all i if $\tilde{\gamma}_i > 0$ for all i. Hence, if $k_1 > -\min_i\{\gamma_i\}$ then condition (A19) is satisfied. For any $\bar{\eta} \neq 1$, condition (A19) can always be ensured to hold since the level of k_1 and γ_i are not pinned down by observable choices (which only depend on differences in γ_i) and do not affect changes in welfare.

The functional form $f(u) = (u^{k_1})^{\frac{1}{1-\eta(u)}}$ used in the previous result differs from assumption (A14) used in deriving the estimation equation (which gives $(\partial \log e_h)/(\partial \log u_h) = a_1$ at t_0 prices). In order to check whether the expenditure function is increasing in u under (A14) away from t_0 prices, we resort to numerical simulations. We consider a range of incomes I from 15,000 to 250,000 CHF and elasticities of substitution as a function of income $\eta(I) = 3 - 2 \times \log(I/250,000)$. We consider 10 goods and draw random utility elasticities $\gamma_i \sim U(0,2)$, initial prices $p_i \sim U(0,1)$, and initial taste shifters $\zeta \sim U(0,1)$; we then renormalize to satisfy $\sum \zeta_i = 1$. We set $a_0 = 1$ and $a_1 = 1,000$. For small deviations in prices relative to their t_0 levels, the expenditure function is approximately equal to $a_0 u^{a_1}$. To allow for larger price changes, we draw price changes from a log-normal distribution with mean zero and standard deviation 0.3. Across a large number (4,280,000) random simulations, only 108 (or 0.0025%) contain a non-increasing portion of the expenditure function (across a large range of utilities). As with quadratic or translog utility, in these cases one must restrict the space of feasible choices or prices to ensure that we are in the monotonic region of the expenditure function.

Appendix D Additional quantitative results

D.1 Sensitivity analysis of results from Section 5.1

Here we present the additional results described briefly in Section 5.1.

In the right-hand panel of Table A15 we display results imposing common expenditure shares across households, using the expenditure share calculated across all households. Whereas the first-order effects are, obviously, now identical across households, the second-order effects are little changed from our baseline.

In Table A16 we display the full non-linear effect of price changes for alternative levels of $\bar{\eta}_s$—so that the elasticity for households with income of 120,000 CHF ranges between 1.5 and 5—while holding the differences in elasticities across households fixed. Greater substitution generates larger declines in the welfare-relevant price index; however, differences between income groups are not very sensitive even for the large range of $\bar{\eta}_s$ considered.

Table A15: Welfare-relevant grocery price changes: Additional results I

Annual income	2013–14 Heterogeneous elasticities 1st-order	Switching	Exact	2014–15 Common exp. shares 1st-order	Switching	Exact
1: 20,000 elasticity 6.6	1.2	-0.6	0.4	-1.2	-0.9	-2.2
2: 60,000 elasticity 4.4	1.1	-0.4	0.7	-1.2	-0.6	-1.8
3: 120,000 elasticity 3.0	1.0	-0.2	0.7	-1.2	-0.3	-1.5

Notes: The left panel replicates the left panel of Table 7, but using 2013–14 changes. The right panel replicates the left panel of Table 7, but imposing common expenditure shares across HHs (calculated across all HHs).

Table A16: 2014–15 Exact welfare-relevant grocery price changes: Additional results II

Annual income	Varying high-income elasticity ($\eta_{High,s}$) $\eta_{High,s} = 1.5$	$\eta_{High,s} = 3$	$\eta_{High,s} = 5$
1: 20,000 elasticity $\eta_{High,s} + 3.6$	-1.9	-2.2	-2.7
2: 60,000 elasticity $\eta_{High,s} + 1.4$	-1.5	-1.7	-2.1
3: 120,000 elasticity $\eta_{High,s}$	-1.3	-1.6	-2.0

Notes: Column 2 exactly replicates column 3 of the left panel of Table 7. Columns 1 and 3 display results for alternative values of the elasticity of substitution for the highest-income household.

D.2 Sensitivity analysis of results from Section 5.2

Here we present the additional results described briefly in Section 5.2.

First, in response to import price declines (compared to increases displayed in Table 8), the first-order and expenditure-switching effects push welfare of higher- relative to lower-income households in opposite directions. High income households benefit more from the first-order effect because they have higher initial import shares. On the other hand, low-income households benefit more from the expenditure-switching effect because they have higher price elasticities. If we assume $\sigma_j = 0$, which mitigates the expenditure switching effect, then the first channel dominates for small import price declines and the second channel dominates for larger import price declines. If we set $\sigma_j > 0$, then lower-income households gain slightly more in response to the 2.2% import price decline. This is because the observed increase in the variance of price changes in 2014–15 is sufficiently strong to make the expenditure-switching effect dominate. Table A17 displays the results.

Second, in our baseline we choose $\eta_s = -2$. Table A18 reports results in which we use $\eta_s = -1.5$, which is at the lower end of our estimates. We maintain the assumption that the elasticity of substitution for the highest-income household equals 3, which pins down $\bar{\eta}$. As expected, the importance of heterogeneous elasticities for shaping the unequal welfare implications of foreign prices is smaller.

Third, in our baseline we choose $\bar{\eta}_s$ so that the lowest initial elasticity of substitution (that for the highest-income household with income of $120,000$ CHF) is equal to 3. Tables

Table A17: Import price declines

Annual income	Import price shock -2.2	-10	-20	-40	-2.2
	$\sigma = 0$				$\sigma > 0$
1: 20,000 elasticity 6.6	0.47	2.4	5.4	13.2	0.73
2: 60,000 elasticity 4.4	0.53	2.6	5.6	12.7	0.70
3: 120,000 elasticity 3.0	0.59	2.8	5.8	12.6	0.69
% difference in CV btw					
income groups 2 and 1	13	9	4	-4	-5
income groups 3 and 1	25	17	9	-4	-6

Notes: This table replicates the exercise in Table 8 but studying import price declines. We omit the contribution of heterogeneous ηs because the first-order and higher-order effects move in opposite directions.

Table A18: Smaller differences in elasticities of substitution

Annual income	Import price shock +2.2	+10	+20	+40	+1000	+2.2
	$\sigma = 0$					$\sigma > 0$
1: 20,000 elasticity 5.7	-0.4	-1.9	-3.3	-5.1	-6.6	-0.2
2: 60,000 elasticity 4.0	-0.5	-2.3	-4.2	-7.2	-12.4	-0.4
3: 120,000 elasticity 3.0	-0.6	-2.6	-5.0	-9.1	-22.0	-0.5
% difference in CV btw						
income groups 2 and 1	16	20	27	41	86	58
income groups 3 and 1	30	38	49	77	232	104
Contribution of heterogeneous ηs						
income groups 2 and 1	6	23	37	55	76	62
income groups 3 and 1	5	20	34	52	82	59

Notes: This table replicates the exercise in Table 8 but imposing $\eta_s = 1.5$ rather than $\eta_s = 2$, while maintaining that the lowest elasticity of substitution (that for the highest-income household with income of $120,000$ CHF), η_{hst_0}, is equal to 3.

Table A19: Elasticity of substitution of high-income group = 1.5

	Import price shock					
	+2.2	+10	+20	+40	+1000	+2.2
Annual income			$\sigma = 0$			$\sigma > 0$
1: 20,000 elasticity 5.1	-0.4	-1.9	-3.4	-5.5	-7.6	-0.3
2: 60,000 elasticity 2.9	-0.5	-2.3	-4.5	-8.1	-20.0	-0.4
3: 120,000 elasticity 1.5	-0.6	-2.7	-5.3	-10.4	-87.3	-0.6
% difference in CV btw						
income groups 2 and 1	16	22	30	48	163	63
income groups 3 and 1	30	40	54	89	1047	112
Contribution of heterogeneous ηs						
income groups 2 and 1	8	28	44	62	87	67
income groups 3 and 1	7	25	41	60	96	65

Notes: This table replicates the exercise in Table 8 but imposing the lowest elasticity of substitution (that for the highest-income household with income of $120,000$ CHF), η_{hst_0}, is equal to 1.5 rather than 3.

A19 and A20 report results in which we use an elasticity of substitution for the highest-income household equal to 1.5 and 5, respectively. Lower levels of price elasticities imply much larger welfare losses for every income group. However, except for the movement to autarky experiment, the percentage difference in CV between income groups and the contribution of heterogeneous elasticities are not very sensitive to the level of the elasticities keeping the elasticity difference between income groups unchanged.

Fourth, in our baseline we choose $\rho = 0.99$ so that expenditure shares across sectors are essentially fixed. Table A21 reports results in which we use a much lower value of $\rho = 0.20$.

Finally, in our baseline we choose elasticities of substitution in the service sector and the other non-grocery goods sector to match those we estimated within the grocery sector; we do so because estimates of income-group-specific price elasticities are not available outside of our Homescan data on groceries. Tables A22 and A23 report results in which we impose a common price elasticity across all income groups within the service sector and within both the service and other non-grocery goods sectors, respectively. In both cases, the contribution of heterogeneous elasticities falls relative to that in our baseline. Nevertheless, since import shares within the service sector are relatively low, results in Table A22 are very similar to those in our baseline.

Table A20: Elasticity of substitution of high-income group = 5

	Import price shock					
	+2.2	+10	+20	+40	+1000	+2.2
Annual income			$\sigma = 0$			$\sigma > 0$
1: 20,000 elasticity 8.6	-0.4	-1.7	-2.8	-3.8	-4.1	-0.1
2: 60,000 elasticity 6.4	-0.5	-2.1	-3.8	-5.7	-7.0	-0.3
3: 120,000 elasticity 5.0	-0.6	-2.5	-4.5	-7.5	-11	-0.4
% difference in CV btw						
income groups 2 and 1	16	23	32	50	70	153
income groups 3 and 1	31	42	60	99	168	272
Contribution of heterogeneous ηs						
income groups 2 and 1	8	28	44	60	71	72
income groups 3 and 1	7	26	42	60	75	70

Notes: This table replicates the exercise in Table 8 but imposing the lowest elasticity of substitution (that for the highest-income household with income of 120,000 CHF), η_{hst_0}, is equal to 5 rather than 3.

Table A21: Elasticity of substitution across sectors = 0.2

	Import price shock					
	+2.2	+10	+20	+40	+1000	+2.2
Annual income			$\sigma = 0$			$\sigma > 0$
1: 20,000 elasticity 6.6	-0.4	-1.9	-3.2	-4.8	-5.8	-0.2
2: 60,000 elasticity 4.4	-0.5	-2.3	-4.2	-7.3	-11.9	-0.4
3: 120,000 elasticity 3.0	-0.6	-2.6	-5.1	-9.5	-25.2	-0.5
% difference in CV btw						
income groups 2 and 1	16	22	31	51	106	83
income groups 3 and 1	30	41	58	98	335	148
Contribution of heterogeneous ηs						
income groups 2 and 1	8	28	44	62	80	69
income groups 3 and 1	7	26	41	61	87	67

Notes: This table replicates the exercise in Table 8 but imposing $\rho = 0.2$ rather than $\rho = 0.99$.

Table A22: Homogeneous elasticities within the service sector

	Import price shock					
	+2.2	+10	+20	+40	+1000	+2.2
Annual income			$\sigma = 0$			$\sigma > 0$
1: 20,000	-0.4	-1.8	-3.2	-4.8	-5.7	-0.2
2: 60,000	-0.5	-2.2	-4.1	-7.0	-11.1	-0.4
3: 120,000	-0.6	-2.6	-4.9	-9	-21.5	-0.5
% difference in CV btw						
income groups 2 and 1	16	21	29	47	94	54
income groups 3 and 1	30	39	54	89	275	97
Contribution of heterogeneous ηs						
income groups 2 and 1	7	26	42	60	79	60
income groups 3 and 1	6	23	39	58	85	58

Notes: This table replicates the exercise in Table 8 but imposing that within the service sector all income groups have a common import elasticity equal to that of income group 2 in our baseline ($\eta_{hst_0} = 4.4$ for s = services for all h).

Table A23: Homogeneous elasticities within the service and other goods sectors

	Import price shock					
	+2.2	+10	+20	+40	+1000	+2.2
Annual income			$\sigma = 0$			$\sigma > 0$
1: 20,000	-0.4	-1.9	-3.4	-5.5	-8.1	-0.3
2: 60,000	-0.5	-2.2	-4.1	-7.0	-11.1	-0.4
3: 120,000	-0.6	-2.5	-4.8	-8.3	-14.6	-0.4
% difference in CV btw						
income groups 2 and 1	15	18	22	28	37	32
income groups 3 and 1	29	34	40	52	80	56
Contribution of heterogeneous ηs						
income groups 2 and 1	4	16	26	41	61	42
income groups 3 and 1	3	13	23	38	65	38

Notes: This table replicates the exercise in Table 8 but imposing that within the service sector and the other goods sector all income groups have a common import elasticity equal to that of income group 2 in our baseline ($\eta_{hst_0} = 4.4$ for s = services and other goods for all h).

References

AUER, R., A. BURSTEIN, AND S. LEIN (2021): "Exchange rates and prices: evidence from the 2015 Swiss franc appreciation," *American Economic Review*, 111, 1–35.

FALLY, T. (2022): "Generalized separability and integrability: Consumer demand with a price aggregator," *Journal of Economic Theory*, 203, 105471.

KUHN, U. (2018): *Collection, construction and checks of income data in the Swiss Household Panel*, University of Lausanne.

NEVO, A. (2000): "A practitioner's guide to estimation of random-coefficients logit models of demand," *Journal of economics & management strategy*, 9, 513–548.

SWISS FEDERAL STATISTICAL OFFICE (2013): "Haushaltsbudgeterhebung 2011: Kommentierte Ergebnisse und Tabellen," `https://www.bfs.admin.ch/bfs/en/home/statistics/economic-social-situation-population/income-consumption-wealth/household-budget.assetdetail.349156.html`.

——— (2014): "Household Budget Survey," `https://www.bfs.admin.ch/bfs/en/home/statistics/economic-social-situation-population/income-consumption-wealth/household-budget.html`.

——— (2016): "Consumer Price Index (December 2015 = 100): Methodological foundations," `https://www.bfs.admin.ch/bfs/en/home/statistics/prices/consumer-price-index.assetdetail.1867121.html`.

SWISS FEDERAL TAX ADMINISTRATION (2014): "Statistik der direkten Bundessteuer DBSt 2014," `https://www.estv.admin.ch/dam/estv/de/dokumente/allgemein/Dokumentation/Zahlen_fakten/Steuerstatistiken/direkte_bundessteuer/np_statistische_kennzahlen_ohne/NP_2014_mitnull.xlsx.download.xlsx/NP_2014_mitnull.xlsx`.