

NBER WORKING PAPER SERIES

UNEQUAL EXPENDITURE SWITCHING:
EVIDENCE FROM SWITZERLAND

Raphael Auer
Ariel Burstein
Sarah M. Lein
Jonathan Vogel

Working Paper 29757
<http://www.nber.org/papers/w29757>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2022, Revised November 2023

The views presented in this paper are those of the authors and not necessarily those of the Bank for International Settlements or the National Bureau of Economic Research. We thank David Atkin, David Baqaee, Kirill Borusyak, Javier Cravino, Pablo Fajgelbaum, Thibault Fally, Isabel Martínez, Katheryn Russ, Andres Santos, and David Weinstein for helpful comments. We thank Santiago Alvarez for excellent research assistance.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Raphael Auer, Ariel Burstein, Sarah M. Lein, and Jonathan Vogel. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Unequal Expenditure Switching: Evidence from Switzerland
Raphael Auer, Ariel Burstein, Sarah M. Lein, and Jonathan Vogel
NBER Working Paper No. 29757
February 2022, Revised November 2023
JEL No. E3,F1,F41

ABSTRACT

What are the unequal effects of changes in consumer prices on the cost of living? In the context of changes in import prices (driven by, e.g., changes in trade costs or exchange rates), most analyses focus on variation across households in initial expenditure shares on imported goods. However, the unequal welfare effects of non-marginal foreign price changes also depend on differences in how consumers substitute between imported and domestic goods, on which there is scant evidence. Using data from Switzerland surrounding the 2015 appreciation of the Swiss franc, we provide evidence that lower-income households have higher price elasticities. We quantify the contribution of heterogeneous elasticities for the unequal welfare effects of observed price changes between 2014–15 and for counterfactual shocks to the mean and dispersion of import price changes.

Raphael Auer
Bank for International Settlements
Postfach CH-4002
Basel
Switzerland
and CEPR
raphael.auer@bis.org

Ariel Burstein
Department of Economics
Bunche Hall 8365
Box 951477
UCLA
Los Angeles, CA 90095-1477
and NBER
arielb@econ.ucla.edu

Sarah M. Lein
Department of Business and Economics
University of Basel
Switzerland
sarah.lein@unibas.ch

Jonathan Vogel
Department of Economics
University of California at Los Angeles
8283 Bunche Hall
Mail Code 147703
Los Angeles, CA 90095
and NBER
jonathan.e.vogel@gmail.com

1 Introduction

What are the unequal effects of changes in consumer prices on the cost of living? In the context of changes in prices of imported goods (due to, e.g., changes in trade costs or exchange rates), most attempts to answer this question have focused on variation across households in initial expenditure shares on imported goods; see, e.g., [Fajgelbaum and Khandelwal \(2016\)](#), [Cravino and Levchenko \(2017\)](#), and [Borusyak and Jaravel \(2021\)](#).¹ However, the unequal welfare effects of non-marginal foreign price changes also depend on differences in how consumers substitute between imported and domestic goods (*unequal expenditure switching*), on which there is scant evidence. As noted by [Deaton \(1997, page 187\)](#):

Since my main interest here is in the distributional effects of price changes..., these [second-order] effects will change the conclusions only to the extent that the elasticities ... differ systematically between poor and rich. Although there is no reason to rule out such effects a priori, there is no reliable evidence on the topic.


In this paper, we document large differences in price elasticities across the income distribution and quantify their contribution to the unequal welfare effects of import price shocks.

In Section 2, we begin by describing the setting in which we measure both initial expenditure shares on imported goods across households as well as unequal expenditure switching. We focus on Switzerland in a period surrounding the abrupt appreciation of the Swiss franc on January 15, 2015, which markedly reduced import relative to domestic retail prices; see Figure 1.² We measure initial import exposure across the income distribution using data on expenditure shares by household income groups across 296 consumption categories (using the Swiss Household Budget Survey) and import shares across 217 slightly more aggregated consumption categories (using the disaggregated data underlying the Swiss CPI). To study unequal expenditure switching, we turn to higher-frequency and more detailed barcode-level Swiss Nielsen Homescan data, cov-

¹See also [Friedman and Levinsohn \(2002\)](#), [Porto \(2006\)](#), and [Carroll and Hur \(2020\)](#).

²The Swiss National Bank (SNB) adopted a minimum exchange rate of 1.20 CHF per EUR in September 2011. Developments abroad in late 2014 and early 2015 prompted the SNB to unexpectedly abandon this policy on January 15, 2015. The subsequent appreciation episode came after a period of remarkable exchange rate stability, was significant (the EUR/CHF appreciated by 14.7 percent by the end of June), and—in contrast to many episodes with large swings in international relative prices—occurred against the backdrop of stable economic aggregates and nominal income inequality in Switzerland. Import prices fell by more at the border than at the retail level ([Auer et al., 2021](#)); we focus on the latter given our emphasis on expenditure switching and welfare at the consumer level.

Figure 1: Swiss franc appreciation and resulting price changes



Notes: The left-hand panel displays the CHF per EUR exchange rate with a vertical line on January 15, 2015; source: [Bank for International Settlements \(2020\)](#). The right-hand panel displays retail price differences relative to December 2014 separately for imports and domestically produced purchases (and associated 95% confidence intervals) from estimating equation (A26); source: [Nielsen Switzerland \(2016\)](#). Observations are weighted by product expenditure in 2014 and confidence intervals are constructed using robust standard errors clustered at the product level.

ering individual household purchases of food, beverages, personal care, and household supplies in supermarkets and drugstores. We merge these data with information on whether individual barcode products are produced domestically or imported (as reported on product labels).³ In response to the 2015 CHF appreciation, the import share within the Homescan data rises and this increase is greater for lower-income households; this differential change in import shares across incomes is not driven by a larger decline in import prices for lower-income households. Finally, we document that the 2015 CHF appreciation was accompanied by an increase in the dispersion of price changes within the set of imported goods.

In Section 3, we characterize a set of sufficient statistics to answer our motivating question: What are the unequal welfare effects of changes in consumer prices (for instance, changes in import relative to domestic prices in our quantitative applications) through their impact on the cost of living? A large literature has addressed this question by applying a first-order approximation of the expenditure function and, hence, focusing on variation across households in initial expenditure shares on different goods. We apply known results in microeconomic theory—see, e.g., [Hausman \(1981\)](#) and more recently [Baqee and Burstein \(2023\)](#)—to provide an exact answer to this question for non-marginal

³In terms of measuring initial import shares across the income distribution, our paper is most closely related to [Borusyak and Jaravel \(2021\)](#), who also use detailed data on consumer expenditures and import shares across the full economy rather than for aggregate industries and directly observe household-specific import shares on consumer packaged goods (and motor vehicles).

price changes, taking expenditure switching into account. The sufficient statistics to calculate a household’s compensating variation in response to a given change in income and prices are initial expenditure shares across products and *compensated* cross-price elasticities (i.e. cross-price elasticities along the initial indifference curve).⁴ The unequal effects of price changes on the cost of living are shaped by differences in initial expenditure shares and differences in compensated cross-price elasticities, and these effects increase in the dispersion of price changes across goods. In practice, estimating cross-price elasticities (compensated or uncompensated) between all goods in the economy is infeasible without additional assumptions. We therefore impose nested, generalized non-homothetic CES preferences, building on Matsuyama (2019), Fally (2022) and Comin et al. (2021). Income elasticities can be non-unitary; and elasticities of substitution between goods within a sector can vary between indifference curves, but (as in standard trade models) are constant along any indifference curve.

In Section 4 we estimate how elasticities of substitution between goods in the Homescan data vary with income, taking two approaches that leverage distinct sources of variation. In our first approach, we use variation in changes in import relative to domestic expenditures between 2014 and 2015 (surrounding the appreciation) across higher- and lower-income households. Our identification assumption is that import demand shocks in 2015 are not systematically different across incomes. In our second approach, we use variation in changes in expenditures across individual barcode products and variation in product price changes. In this case, we control for product-specific demand shocks and household-specific import demand shocks. We instrument for the interaction between initial household income and the product-specific price change using an interaction between household income and a product “cost shifter.” Our cost shifter exploits variation across border groups—an aggregation of products—in the invoicing currency of imports. Specifically, we measure the share of imported goods in each border group that are denominated in EUR, using information from the goods-level survey underlying the calculation of the official Swiss import price index. Because of the stickiness of import prices at the border in their invoicing currency, Swiss retail prices of imported goods are more responsive to the appreciation if imports are denominated in EUR rather than in CHF; see Auer et al. (2021). Given additional controls and our instrumentation strategy, the exclusion restriction in the second approach is substantially weaker than in the first approach.

⁴There are well-known price indices (e.g. Törnqvist or Sato-Vartia) that incorporate information on observed expenditure shares over time and—under strong assumptions—measure welfare beyond first-order approximations. For recent alternative approaches, see Atkin et al. (2023), Jaravel and Lashkari (2023), and Baqaee et al. (2023). In contrast to all of these approaches, we calculate changes in welfare in response to counterfactual changes in prices and income; see Section 5.

In spite of these differences, we obtain very similar quantitative results across approaches. The elasticity of substitution between goods in the Homescan data is substantially lower for higher-income households: for example, the difference between the elasticities of substitution between two households, where one has an income three times that of the other, is 2.4 under the first approach and 2.1 under the second approach. These approaches identify *differences* in elasticities across incomes. To estimate the *level* of these elasticities, we make stronger identification assumptions, and the resulting estimates vary more across the two approaches. However, in our analytic and quantitative results we show that, conditional on differences in price elasticities across households, the unequal welfare effects of price changes are not very sensitive to elasticity levels.

Our estimates of higher price elasticities for lower-income households are qualitatively consistent with demand system estimates in industrial organization—see, e.g., [Berry et al. \(2004\)](#) among many others—and findings on shopping behavior in macroeconomics—see, e.g., [Kaplan and Menzio \(2016\)](#) and [Aguiar and Hurst \(2007\)](#)—all of which support [Harrod’s \(1936\)](#) Law of Diminishing Elasticity of Demand, which postulates that demand elasticities decrease in income.⁵ In the spatial economics literature, there is scant and conflicting evidence on differences in elasticities across household incomes. [Argente and Lee \(2021\)](#) and [Faber and Fally \(2022\)](#) find very small differences. [Handbury \(2021\)](#) finds differences in elasticities across incomes very similar to our estimates when controlling for the fact that higher-income households have a greater willingness to pay for quality but finds opposite results when not; our approach incorporates these differences in willingness to pay for quality (by incorporating household $\times$ barcode product-specific demand shifters that cancel out when estimating demand elasticities using changes rather than levels of household expenditure shares by income).⁶ Moreover, whereas these papers use either Hausman instruments or the approach developed by [Feenstra \(1994\)](#) and [Broda and Weinstein \(2006\)](#), we exploit exogenous variation in price responsiveness to an exchange rate shock.⁷

Finally, in Section 5 we quantify how differences in price elasticities estimated using

⁵[Bems and di Giovanni \(2016\)](#) document that a large aggregate decrease in income in Latvia reduced import shares (since high-quality imports are more income elastic), and [Coibion et al. \(2015\)](#) show that households switch expenditures toward low-price stores when local economic activity falls. Rather than focusing on expenditure switching due to changes in income, we focus on heterogeneous expenditure switching across the income distribution in response to changes in relative prices.

⁶While [Handbury \(2021\)](#) calculates regional differences in price indices abstracting from her estimated differences in price elasticities, we show that differences in elasticities are quantitatively important in shaping the unequal welfare effects of large price changes.

⁷When we use a Hausman instrument, we find small differences. In Appendix B.2 we show that the Hausman instrument may be endogenous in our Swiss context (where there is little spatial variation in price changes).

the Homescan data shape the unequal welfare effects of changes in prices.⁸ For these exercises, we consider households at three distinct income levels ranging from 20,000 to 120,000 CHF per year with respective elasticities of 6.6, 4.4, and 3. We first calculate changes in the welfare-relevant price index over grocery products in the Homescan data in response to observed price changes between 2014 and 2015, following the CHF appreciation. Prices fell on average by 1.2% between 2014 and 2015. The welfare-relevant price index (taking expenditure switching into account) declined by 1.6% for households with income of 120,000 CHF and by 2.2% for households with income of 20,000 CHF. This gap between income groups is accounted for by unequal expenditure switching between imports and domestic goods as well as between products with different price changes within the set of imports and domestic goods. If price elasticities were equal across income groups, then the gap in price indices would be small and of the opposite sign.

We then calculate changes in welfare in response to counterfactual changes in import relative to domestic prices, considering, first, only uniform price changes within each set of goods and, second, incorporating variation in price changes within each set of goods. To conduct these counterfactuals, we use our measures of import shares for each income across all consumer goods and we impose that our estimated differences in price elasticities across incomes within our Homescan data apply more broadly. To highlight the nonlinearities from expenditure switching, we consider import price shocks that are larger than the one induced by the 2015 CHF appreciation.⁹ Uniform import price increases harm higher-income households more than lower-income households in Switzerland for two reasons: (i) higher-income households have higher initial import shares (since they spend relatively more on non-grocery goods, which are more tradable than groceries) and (ii) they substitute away from imported goods less.¹⁰ For large changes in prices, the impact of unequal expenditure-switching on welfare is substantial. For example, consider

⁸Bai and Stumpner (2019) and Jaravel and Sager (2019) construct income-group and product-category-specific inflation rates and project these on changes in import penetration induced by China. Hottman and Monarch (2020) focus on differences in import price inflation rates across US households. Relative to these papers, we estimate differences in import elasticities across income groups, which we then use to quantify welfare changes for observed and counterfactual price shocks.

⁹Given our focus on the expenditure-side effects of foreign price shocks, in the counterfactuals we abstract from changes in the income distribution. There is a large empirical and theoretical literature on the impact of international trade on income inequality with multiple factors; see e.g. Burstein and Vogel (2017), Cravino and Sotelo (2019), Galle et al. (2022), and Adao et al. (2022). See, e.g., He (2018) and Borusyak and Jaravel (2021) for papers incorporating both income and expenditure-side inequality induced by trade.

¹⁰Variation of import shares with household income differs across countries depending, among other things, on whether the country is high or low income and has a comparative advantage in goods with high- or low-income elasticities. For example, Borusyak and Jaravel (2021) document that imports are flat throughout the income distribution in the US. Therefore, if we applied our counterfactuals to the US context, unequal expenditure switching would be the only channel inducing unequal welfare effects (via variation in cost of living).

a 20% uniform increase in import prices relative to domestic prices, which is not uncommon in the context of large exchange rate devaluations; see, e.g., [Burstein et al. \(2005\)](#) and [Cravino and Levchenko \(2017\)](#). In response to this shock, welfare of a household with annual income of 20,000 CHF falls by about one third less than for a household with annual income of 60,000 CHF. Almost half of this differential is accounted for by differences in price elasticities. In practice, the variance of price changes within imports and domestic goods rose during the 2015 CHF appreciation. If in our counterfactuals we additionally consider an increase in price dispersion within imported and domestic goods, then the importance of unequal expenditure switching can be substantially larger.

In summary, we make four main contributions. First, we document heterogeneity across incomes in expenditure switching in response to an exogenous shock to import prices. Second, we estimate differences in elasticities of substitution across incomes using plausibly exogenous variation resulting from this shock to import prices and its heterogeneity across product groups. Third, we apply results in welfare economics to characterize differential welfare changes due to heterogeneous price elasticities in response to foreign-induced price changes. Fourth, we show quantitatively that unequal expenditure switching contributes substantially to differences in welfare across the income distribution in response to observed Swiss price changes and to plausible counterfactual shocks to the mean and dispersion of import price changes.

2 Data and stylized facts

2.1 Data

In this section we provide an overview of the main data sets employed in the paper. Additional details and data sources are provided in [Appendix A](#).

AC Nielsen Homescan data and import status. AC Nielsen Homescan, [Nielsen Switzerland \(2016\)](#), includes information on household (HH) characteristics and shopping transactions of a demographically and regionally representative sample in Switzerland during the period surrounding the 2015 appreciation: January 2013 to December 2016. The data includes approximately 3,300 households in 2014.

Participating households record purchases—of food, beverages, personal care (health and beauty aids), and other selected general merchandise—in supermarkets and drug-stores (we refer to these goods as *groceries*). Individual products are identified by their barcode (European Article Number or EAN, which we often refer to as a *product*). In the appendix, we describe a number of adjustments we make to the data, such as dropping

likely recording errors by households. In the raw data, an observation is a transaction that includes 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 retailer. We drop all transactions that occur abroad. See [Burstein et al. \(2022\)](#) for an analysis of cross-border shopping in Switzerland using the Homescan data. We measure each product’s price (in logs) as an average of transaction-level log prices in the corresponding time period, weighing transactions by expenditures.

The Homescan data come with a rich set of socioeconomic characteristics for each household, summarized in Table [A11](#) in Appendix [A](#) for the year 2014, including the two-digit zip code of residence, the educational level of the household’s main earner, the number of household members (and the number of those under 10 and over 70 years old), and total household pre-tax annual income reported in seven bins. Given each of these characteristics, we infer a level of household pre-tax income in 2014 for each Homescan household using additional data from the Swiss Household Panel (FORS), augmented with data from the [Swiss Federal Tax Administration \(2014\)](#) (SFTA). We do so by projecting the level of 2014 pre-tax household income in FORS on a set of household characteristics available both in FORS and Homescan, including indicator variables for the pre-tax income bins available in Homescan (and which can be constructed in FORS). We then predict household income in Homescan using these coefficients and additional data on the pre-tax national household income distribution from SFTA.

We augment the Homescan data with information on whether individual products are imported or produced domestically. Whereas EANs provide information on the country in which a product has been registered, in many instances this is not the country in which the product has been produced. However, that information is disclosed on the label of each product. We use the label information that [Auer et al. \(2021\)](#) collected from [codecheck.info](#). Coverage is not complete and notably excludes most fruits and vegetables, EANs assigned by store managers locally, and goods that are only occasionally sold in grocery stores such as toys, clothing, or household electronics. Our measure of import status for each individual product is fixed over time, as obtained from [codecheck.info](#) between October 2015 and March 2016.¹¹ We drop products for which import status is unknown.

Comparing columns 1 and 2 of Table [A10](#) in Appendix [A](#), we see that out of 69,088 unique goods and expenditure of CHF 11.1 million, there are 8,409 unique goods and expenditure of CHF 4.2 million with known import status; the share of expenditure for

¹¹Roughly 90% of imported goods in our data are from the European Union. Our results are robust to dropping imports from other origins.

which the production location is known is approximately 38%.^{12,13} We further divide products with known import status into imports and domestically produced goods in columns 3 and 4. A similar number of unique imported and domestically produced goods were purchased and the import share (at retail prices) of expenditures was 26.9%.

Expenditure and import shares by income group and sector. To calculate expenditure shares and import shares by income group across each consumption category, we use two data sets provided by the Swiss Federal Statistical Office (SFSO).

The first data set, the Swiss Household Budget Survey (HBS), reports information about consumption expenditures by income group and consumption category for the periods 2012–14 and 2015–18 based on roughly 250 households per month randomly selected from a large and representative registry. At the lowest level of disaggregation, there are 296 consumption categories for goods and services, such as “rice”, “pasta”, or “tickets for public transport.” The SFSO collects expenditures on these consumption categories separately for each of five income groups. We use data for the pre-appreciation period 2012–14 to construct sectoral expenditure shares for each group. While we construct these sectoral expenditure shares by income group for each of the 296 disaggregated consumption categories, Table 1 displays expenditure shares in the aggregate and separately for each income group, aggregated up to three sectors: groceries (matching as close as possible our Homescan goods), non-grocery goods, and services.

The second data set contains a cross section of import shares by disaggregated product category, obtained by the SFSO via firm surveys published in 2016 based on information from previous years. These shares, used by the SFSO to calculate a CPI for imported goods, are available at a similar disaggregation to the ones in the HBS data. We combine these import shares—which vary across disaggregated category—with the HBS data—which vary across disaggregated category and income group—to construct import shares by income group within each of our three aggregate sectors. To do so, we assume that different income groups have common import shares within each disaggregated product category, an assumption we do not impose in the Homescan data since we observe the import status of individual products. Table 1 displays the resulting import shares by

¹²Many of the products for which we do not observe import status appear in the Homescan data for only a short period of time. If we keep only those products that were purchased at least once per year between 2013 and 2016, the share of expenditure on goods with known origin is 47% instead of the 38% we observe in our baseline sample.

¹³One concern might be that expenditure on products for which we do not observe import status is correlated with household income. However, in Appendix A we show 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.

Table 1: Expenditure and import shares by income group and sector

Annual income	Expenditure shares			Import shares			
	Grocery	Other goods	Services	Grocery	Other goods	Services	All
– 60,252	20.1	18.5	61.4	36.6	66.9	2.2	21.1
60,252 – 88,032	18.6	21.6	59.8	37.2	71.1	3.3	24.2
88,032 – 119,736	18.0	23.4	58.6	36.6	72.6	3.6	25.7
119,736 – 164,244	17.1	24.3	58.6	37.4	74.7	4.2	27.0
164,244 –	15.1	25.6	59.3	40.2	75.3	5.1	28.3
All	17.2	23.5	59.3	37.9	73.3	4.0	26.1

Notes: Expenditure shares by income range and sector—aggregated to groceries, other non-grocery goods, and services—are from SFSO using the years 2012–14. Import shares are constructed from import shares in disaggregated product categories and expenditure shares by income across these categories. The final row represents the value of each column across all households and the final column represents the import share of each income group across all sectors.

income group, by aggregate sector, and by income group $\times$ aggregate sector.¹⁴

Currency of invoicing of import prices at the border. Our instrument in Section 4 exploits variation across imported goods in the invoicing currency of prices at the border. We match individual barcode products in the Homescan data to groups of imported products at the border (*border groups*) and measure the share of imported products in each border group in 2014 that are denominated in EUR (out of those denominated in either EUR or CHF), using information from the goods-level survey underlying the calculation of the official Swiss import price index. For additional information, see [Auer et al. \(2021\)](#).

2.2 Stylized facts


In this section we present our five stylized facts (SFs), indicating in brackets for each fact the data set that we use to calculate it. Facts 2–5 use Homescan rather than SFSO data because the SFSO data are not available at annual frequency. In Appendix B.1 we provide additional information.

SF 1 (SFSO): Import shares before the 2015 CHF appreciation were higher among higher-income households.

The right-most column of Table 1 displays the aggregate import share—across all consumption categories—for each of the five income groups in the SFSO data (for the 2012–14 period). Higher-income households had higher aggregate import shares in Switzerland, with the share rising monotonically from 21% to 28% between the bottom and top income

¹⁴The aggregate import share for groceries in the SFSO sample (37.9%) is higher than in the Homescan data (26.9%). In Appendix A, we show that this is due to expenditures in the SFSO data being comparatively concentrated in categories with high import shares. Applying common expenditure weights across categories in the SFSO and Homescan data yields more similar aggregate import shares.

Figure 2: Aggregate import responsiveness and heterogeneity across incomes in Homescan data



Notes: Import shares by year aggregated across all households (All), households with incomes below our sample median (Low income), and households with income above our sample median (High income). Point estimates and ten percent confidence intervals for each income group and year are calculated by regressing separately across household h within the low-income group and the high-income group $100 \times \frac{X_{hM}}{X_{hM} + X_{hD}}$ on a constant, weighing h by $X_{hM} + X_{hD}$ and clustering by income quantile (of which there are fifty).

groups in the SFSO data. This pattern is accounted for by two relationships. First, the import share was much higher for non-grocery goods (across all income groups) than for groceries or services, and higher-income households spent a higher share of their income in this sector. Second, higher-income households had a higher import share within the non-grocery goods aggregate sector and, to a lesser extent, within services.¹⁵

On the other hand, import shares within groceries were not strongly correlated with income. This is evident in the SFSO data from Table 1. The same (weak) relationship holds within our product-level Homescan data, as shown in Table A14 in Appendix B.1.

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

The aggregate import share in the Homescan data increased from 26.9% to 27.5% between 2014 and 2015, as shown in Figure 2. To show that this occurred within individual households—rather than from a change in the composition of expenditures across households—we regress each household's import share of expenditure in each year on household effects and year effects, excluding the year 2014. These year effects identify the change, within households, in the import share of expenditures between each year and 2014. We find 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 and continues increasing into 2016; see Figure A3 in Appendix B.1, which reports the rise in year effects estimated separately over all monthly time horizons within the year (we define horizon j as the first j months in year t and in 2014).

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

¹⁵This is largely because high-income groups have higher budget shares on cars and cars in Switzerland tend to be imported.

Table 2: Heterogeneous expenditure switching towards imports

	log(Income)			High income		
	(1)	(2)	(3)	(4)	(5)	(6)
Income 2013	-0.472* [0.266]	-0.489* [0.265]	-0.496* [0.261]	-0.213 [0.288]	-0.179 [0.271]	-0.184 [0.278]
Income 2015	-0.727** [0.272]	-0.813*** [0.279]	-0.835*** [0.291]	-0.698** [0.293]	-0.750** [0.316]	-0.778** [0.325]
Income 2016	-0.953*** [0.321]	-0.970*** [0.339]	-1.002*** [0.336]	-0.383 [0.382]	-0.284 [0.424]	-0.327 [0.433]
Observations	11630	11630	11630	11630	11630	11630
Control size		X	X		X	X
All controls			X			X

Notes: Results of estimating $X_{hMt}/(X_{hMt} + X_{hDt}) = \mathbb{F}\mathbb{E}_t + \mathbb{F}\mathbb{E}_h + \sum_{y \neq 2014} \mathbb{I}_{y=t} [\beta_t Inc_h + [\zeta'_t K_h]] + \varepsilon_{ht}$, where X_{hMt} and X_{hDt} are household h 's expenditure on imports and domestic goods in year t ; $\mathbb{F}\mathbb{E}_t$ and $\mathbb{F}\mathbb{E}_h$ are year and household fixed effects; Inc_h is either the logarithm of household income (columns 1–3) or an indicator that equals one if HH income is greater than the median in our sample of 57,647 CHF (columns 4–6); Income y displays the coefficient on β_y . In columns 2 and 5, we control for HH size interacted with year. In columns 3 and 6 we additionally include an indicator for whether there is a child under 10 and an indicator if everyone in the HH is older than 70, each interacted with year. 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 the income quantile and the household's share of expenditure in 2014 within its income quantile. * $p < .1$; ** $p < .05$; *** $p < .01$

Figure 2 displays the import share by year separately for high- and low-income households (those with annual incomes above and below our sample median of approximately 57,600 CHF).¹⁶ Between 2014 and 2015 the import share of low-income households increased more than that of high-income households. The difference between the import shares of low- and high-income households expanded from approximately 0.5 to 1.1 percentage points between 2014 and 2015. This gap rose further to 1.2 percentage points in 2016. Figure 2 shows no evidence of pre-trends in the difference between low- and high-income import shares.

This greater increase in import shares among low-income households occurred within individual households, rather than from a change in the composition of expenditures across households. To show this, we estimate the differential change across household incomes in the import share of expenditures between each year and 2014 at the household level, controlling for household fixed effects. To measure income we use either the log of household income or an indicator for household incomes above the median.

Table 2 shows that import shares increased substantially less for high- than low-income households after the exchange rate appreciation. For example, a household with an income three times than that of another experienced a roughly 0.7 percentage point smaller

¹⁶The higher import share for low-income households in 2014 displayed in Figure 2 appears inconsistent with Stylized Fact 1. However, the difference is not statistically significant, as revealed by the wide confidence in the figure and the results in Table A14.

increase in its import share between 2014 and 2015. This is not a continuation of pre-existing trends: the coefficients in Table 2 indicate that the gap between the import shares of low- and high-income households did not rise between 2013 and 2014 (see the first row of columns 1 and 4) before rising between 2014 and 2015. Changes between 2014 and 2015 also do not appear to be reversion to the mean: the coefficient in row 1 of column 4 is insignificantly different from zero and the coefficient in row 1 of column 1, while marginally significant, is much smaller than in row 2. These results are robust to including progressively more household-level controls interacted with year; see columns 2, 3, 5, and 6.

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.

Regressing individual product prices on product fixed effects and month fixed effects (omitting a fixed effect for the month preceding the CHF appreciation) separately for imports and domestic goods, we find that import prices fell by approximately 2% relative to domestic prices following the appreciation (averaging the change between December 2014 and each month in 2015); see Figure 1 in the Introduction.

In Appendix B.1, we run a related regression that separates prices paid for each product by household income. We find no differential changes in import or domestic prices across income, either economically or statistically; see Figure A4 in Appendix B.1.¹⁷ This pattern is robust to further disaggregating prices by region in Switzerland and including region fixed effects in the regression. This implies that the pattern of heterogeneous changes in import shares across households with different incomes described in Stylized Fact 3 is not driven by lower-income households facing greater declines in relative import prices.

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.

Table 3 displays the standard deviation across barcode products of annual log price changes within the set of imported goods, domestic goods, and all goods between years t and $t + 1$ for $t = 2013, 2014$, and 2015. The standard deviation of log price changes rose from 5.8% to 7.6% for imported and from 4.2% to 4.8% for domestic goods between 2013–

¹⁷We also find no relationship between income and the level of the price paid within individual barcode products. Relatedly, defining products as aggregations of barcodes, Handbury (2021) finds that most of the variation in prices paid across incomes is accounted for by differences in products purchased.

Table 3: Standard deviation of log price annual changes across barcode products

	Imports	Domestic	All
2013–14	0.058	0.042	0.047
2014–15	0.076	0.048	0.058
2015–16	0.061	0.041	0.047

Notes: Expenditure-weighted standard deviation of annual log price change across barcode products between t and $t + 1$ for $t = 2013$, 2014, and 2015, for imported goods (column 1), domestic goods (column 2), and all goods (column 3) .

14 and 2014–15; each approximately reverts to its initial level between 2015–16. We show that these patterns are robust to various alternative choices in Table A16 in Appendix B.1.

We will show that this increase in the dispersion of price changes—together with greater expenditure switching of lower-income households—contributes to the unequal welfare effects of the 2015 CHF appreciation.

3 Welfare impacts of price changes

Our objective is to construct a measure of the change in welfare for different households, starting from an initial observed period, in response to factual or counterfactual changes in income and in the prices of consumption goods. In Section 3.1 we provide sufficient statistics for this calculation under general preferences, building on results in micro theory. In Section 3.2 we restrict these preferences to a particular parametric form, which we use in the following sections to estimate differences in elasticities of substitution across households and to perform our quantitative applications.

3.1 General formulation

Household h 's preferences over N consumption goods, indexed by ζ (taste shifters), can be represented by the expenditure function $e_h(\mathbf{p}, u; \zeta)$, which indicates the minimum cost of achieving utility u given a vector of prices $\mathbf{p}$. The associated budget share on good i is denoted by $b_{hi}(\mathbf{p}, u; \zeta)$, which by Shephard's lemma equals $\partial \log e_h(\mathbf{p}, u; \zeta) / \partial \log p_i$. Given income I (which we assume is equal to expenditures), the indirect utility function is $v_h(\mathbf{p}, I; \zeta)$.

We consider a change in household h 's income from I_{ht_0} to I_{ht_1} , prices, from $\mathbf{p}_{ht_0}$ to $\mathbf{p}_{ht_1}$, and taste shifters from ζ_{ht_0} to ζ_{ht_1} . Our welfare measure is the compensated variation

evaluated at initial preferences, CV_h , which is implicitly defined by

$$v_h(\mathbf{p}_{ht_0}, I_{ht_0}; \zeta_{ht_0}) = v_h(\mathbf{p}_{ht_1}, e^{-CV_h} I_{ht_1}; \zeta_{ht_0}).$$

In words, CV_h is the reduction in income (in logs) under the final budget set that makes the household with initial preferences equally well-off as under the initial budget set. Household h is better off under the final than the initial budget set if and only if $CV_h > 0$.

We can express CV_h using the expenditure function as

$$CV_h = \log \left(\frac{e_h[\mathbf{p}_{ht_1}, v_h(\mathbf{p}_{ht_1}, I_{ht_1}; \zeta_{ht_0}); \zeta_{ht_0}]}{e_h[\mathbf{p}_{ht_1}, v_h(\mathbf{p}_{ht_0}, I_{ht_0}; \zeta_{ht_0}); \zeta_{ht_0}]} \right) = \log \left(\frac{I_{ht_1}}{I_{ht_0}} \right) - \log \left(\frac{e_h[\mathbf{p}_{ht_1}, u_{ht_0}; \zeta_{ht_0}]}{e_h[\mathbf{p}_{ht_0}, u_{ht_0}; \zeta_{ht_0}]} \right) \quad (1)$$

where the second equality uses the fact that $e_h[\mathbf{p}_{ht}, v_h(\mathbf{p}_{ht}, I_{ht}; \zeta_{ht_0}); \zeta_{ht_0}] = I_{ht}$ and where $u_{ht} \equiv v_h(\mathbf{p}_{ht}, I_{ht}; \zeta_{ht})$ represents utility achieved under the time t budget set and preferences. Welfare changes equal the change in household nominal income deflated by the change in the expenditure function in response to changes in prices, evaluated along the initial indifference curve. The deflator is evaluated at the initial indifference curve because, by the definition of CV_h , the income compensation it receives at t_1 leaves the household on that indifference curve.

To understand what one needs to know in order to construct the price deflator in (1), consider any smooth path of prices from $\mathbf{p}_{ht_0}$ to $\mathbf{p}_{ht_1}$, where t indexes “time” (or, more generally, increments along which prices change between two points t_0 and t_1). Using Shephard’s lemma, (1) can be expressed as (see Lemma 1 in [Baqaee and Burstein, 2023](#))

$$CV_h = \log \left(\frac{I_{ht_1}}{I_{ht_0}} \right) - \int_{t_0}^{t_1} \sum_i b_{hi}^{CV}(\mathbf{p}_{ht}) \frac{d \log p_{iht}}{dt} dt, \quad (2)$$

where $b_{hi}^{CV}(\mathbf{p}_h) \equiv b_{hi}(\mathbf{p}_h, u_{ht_0}; \zeta_{ht_0})$ represents household h ’s budget share on good i at prices $\mathbf{p}_h$ along its initial indifference curve. Equation (2) implies that welfare changes for a consumer with non-homothetic preferences that are subject to taste shocks are identical to welfare changes for a fictional consumer with homothetic and stable preferences with budget shares as a function of prices given by $b_{hi}^{CV}(\mathbf{p}_h)$.¹⁸

Discussion. Equations (1) and (2) hold globally—for changes in prices and incomes of any size. According to equation (2), in order to measure CV (for given price changes) we only need to know compensated budget shares as a function of prices, $b_{hi}^{CV}(\mathbf{p}_{ht})$. Given

¹⁸That is, $b_{hi}^{CV}(\mathbf{p})$ corresponds to the budget shares of a fictional consumer with homothetic preferences represented by the expenditure function $e_h^{CV}(\mathbf{p}, u) = e_h(\mathbf{p}, u_{ht_0}; \zeta_{t_0})u$.

the path of prices from t_0 to t_1 , these budget shares can be constructed from initial budget shares, $b_{hi}^{CV}(\mathbf{p}_{ht_0})$, and cross-price elasticities between all goods along the initial indifference curve. *Given* these cross-price elasticities, measuring CV_h does not require income elasticities or taste shifters.¹⁹ However, in *estimating* these cross-price elasticities, shifts in demand induced by income effects and taste shocks cannot be ignored, as we discuss in Section 4.

In our quantification of the welfare impacts of factual or counterfactual changes in prices, we directly measure initial budget shares over consumption goods in our Swiss data. We specify particular preferences to estimate cross-price elasticities along the initial indifference curve using the price changes induced by the 2015 CHF appreciation.

3.2 Non-homothetic CES preferences

In what follows, we restrict the general setup of Section 3.1 by imposing non-homothetic CES preferences. There are multiple sectors, indexed by s , and within each sector there is a fixed set of goods, indexed by $i \in \mathcal{I}(s)$, some imported and some produced domestically.

The expenditure function is given by

$$e_h(\mathbf{p}_{ht}, u; \zeta_{ht}) = f_h(u) \left[\sum_s \zeta_{hst} u^{\gamma_s} (P_{hst})^{1-\rho} \right]^{\frac{1}{1-\rho}} \quad (3)$$

$$P_{hst} = \left(\sum_{i \in \mathcal{I}(s)} \zeta_{hit} u^{\gamma_i} (p_{hit})^{1-\eta_s(u)} \right)^{\frac{1}{1-\eta_s(u)}} \quad (4)$$

where $f_h(\cdot) > 0$ and $\rho, \eta_s(\cdot) \in [0, 1) \cup (1, \infty)$.²⁰ By Shephard's lemma, the budget share of any good $i \in \mathcal{I}(s)$ is

$$b_{hit} = \frac{\zeta_{hit} u^{\gamma_i} p_{hit}^{1-\eta_s(u)}}{\sum_{i' \in \mathcal{I}(s)} \zeta_{hi't} u^{\gamma_{i'}} p_{hi't}^{1-\eta_s(u)}} \times b_{hst} \quad (5)$$

¹⁹If we used the equivalent (rather than compensating) variation under final (rather than initial) preferences, then computing welfare changes requires budget shares as a function of prices along the final (rather than initial) indifference curve. Since in our applications we consider the welfare implications for Swiss consumers of counterfactual price changes starting in 2014, it is more convenient to focus on CV (which requires estimates of price elasticities in 2014) rather than EV (which requires budget shares and estimates of price elasticities in an unobserved initial period such as autarky).

²⁰These preferences reduce to nested homothetic CES if, for example, $\eta_s(u)$ is independent of u , $\gamma_i = \gamma_s = 0$ for all i and s , and $f'_h(u) > 0$. See Appendix C for additional information on these preferences.

where $b_{hst} \equiv \sum_{i \in \mathcal{I}(s)} b_{hit}$ is the share of sector s in h 's budget at time t , given by

$$b_{hst} = \frac{\zeta_{hst} u_{ht}^{\gamma_s} P_{hst}^{1-\rho}}{\sum_{s'} \zeta_{hs't} u_{ht}^{\gamma_{s'}} P_{hs't}^{1-\rho}}. \quad (6)$$

As described in detail below, in mapping our model to the data we consider three aggregate sectors s listed in Table 1. Within each sector we either map each i to a homothetic aggregator across domestic products and a homothetic aggregator across imported products (we do not introduce notation for these aggregators to simplify presentation) or we map each i to individual barcode products. In the first approach, $\eta_s(u)$ is the elasticity of substitution between the aggregate domestic good and the aggregate imported good in sector s . In the second approach, $\eta_s(u)$ is the elasticity of substitution between any pair of barcode products in sector s , irrespective of import status.

Welfare changes. Compensated budget shares $b_{hi}^{CV}(\mathbf{p}_h)$ are obtained by fixing utility at u_{ht_0} and tastes at ζ_{hit_0} and ζ_{hst_0} . We express compensated budget shares as a simple function of initial expenditure shares, initial elasticities of substitution, and changes in prices:

$$b_{hi}^{CV}(\mathbf{p}_h) = b_{hit_0} \times \left(\frac{\hat{p}_{hi}}{\hat{P}_{hs}} \right)^{1-\eta_{hst_0}} \times \frac{\hat{P}_{hs}^{1-\rho}}{\sum_{s'} b_{hs't_0} \hat{P}_{hs'}^{1-\rho}} \quad (7)$$

$$\hat{P}_{hs} \equiv \left(\sum_{i \in \mathcal{I}(s)} \frac{b_{hit_0}}{b_{hst_0}} (\hat{p}_{hi})^{1-\eta_{hst_0}} \right)^{\frac{1}{1-\eta_{hst_0}}} \quad (8)$$

where $\hat{x} \equiv x/x_{t_0}$ for any x ; ρ is the elasticity of substitution along the initial indifference curve between sectors, which we assume is common across all households and constant; and η_{hst_0} is the elasticity of substitution along the initial indifference curve for household h between goods within sector s .

Given compensated budget shares, the expression for welfare changes in the general setup, (2), simplifies to

$$CV_h = \log(\hat{I}_h) - \frac{1}{1-\rho} \log \left[\sum_s b_{hst_0} (\hat{P}_{hs})^{1-\rho} \right] \quad (9)$$

where $\hat{P}_{hs}$ is defined by equation (8).

We use (9) to construct changes in welfare in response to factual and counterfactual income and price changes. Constructing CV_h for household h requires the value of the elasticity of substitution between sectors, ρ , expenditure shares across sectors in the ini-

tial period, b_{hst_0} , income changes, $\hat{I}_h$, and sectoral price changes $\hat{P}_{hs}$. Constructing $\hat{P}_{hs}$ in equation (8) for household h requires expenditure shares within sector s in the initial period, b_{hit_0} , and the elasticity of substitution in the initial period t_0 , η_{hst_0} .

To a second-order approximation, and setting $\rho \rightarrow 1$, equation (9) can be expressed as

$$CV_h \approx \underbrace{\log(\hat{I}_h) - \mathbb{E}_{b_{ht_0}} [\log \hat{P}]}_{\text{First-order effect}} + \underbrace{\frac{1}{2} \sum_s b_{hst_0} (\eta_{hst_0} - 1) \text{Var}_{b_{ht_0}|s} [\log \hat{P}]}_{\text{Expenditure-switching effect}} \quad (10)$$

where $\mathbb{E}_{b_{ht_0}} [\log \hat{P}]$ is the budget-share weighted average of log price changes and where $\text{Var}_{b_{ht_0}|s} [\log \hat{P}]$ is the budget-share weighted variance of log price changes within sector s ; see [Baqaee and Burstein \(2023\)](#). The approximation error in expression (10) vanishes as price changes become smaller and as $\eta_{hst_0} \rightarrow 1$. The literature on the unequal effects of price changes has largely focused on the first-order effects in equation (10). The expenditure-switching effect, which is the focus of our paper, raises welfare if the elasticity of substitution η_{hst_0} is greater than one, and is increasing in η_{hst_0} . That is, households that are more price sensitive benefit more from volatility in prices (or lose less if $\eta_{hst_0} < 1$).

Unequal expenditure switching. Here, we provide two special cases that highlight the importance of differences in elasticities, $\eta_{h'st_0} - \eta_{hst_0}$, for differences in $CV_{h'}$ and CV_h . In both cases, we set $\rho \rightarrow 1$.

First, consider two households with common expenditure shares in the initial period and arbitrary changes in product prices. To a second-order approximation, the difference in the expenditure switching effect between these households is

$$\frac{1}{2} \sum_s b_{hst_0} (\eta_{h'st_0} - \eta_{hst_0}) \text{Var}_{b_{ht_0}|s} [\log \hat{P}] \quad (11)$$

which depends on the difference in their elasticities of substitution in the initial period, $\eta_{h'st_0} - \eta_{hst_0}$. Conditional on this difference, the levels of these elasticities do not matter to a second-order approximation. It is precisely this difference in elasticities that we estimate in Section 4.²¹

Second, consider two households with (potentially) different expenditure shares within

²¹ Allowing for differences in initial budget shares, the differences in the expenditure-switching effect includes the additional term

$$\frac{1}{2} \sum_s (\eta_{h'st_0} - 1) \left\{ b_{h'st_0} \text{Var}_{b_{h't_0}|s} [\log \hat{P}] - b_{hst_0} \text{Var}_{b_{ht_0}|s} [\log \hat{P}] \right\}$$

This additional term depends on the level of the elasticity. However, since the differences in price variances across households is small in our quantitative application, this additional term is small.

sectors and assume that the distribution of log price changes, $\log \hat{p}_{hit}$, in each sector is normal with mean μ_s and standard deviation σ_s . In this case, equation (10) is exact. Moreover, if expenditure shares across sectors are common across households, then the difference in the expenditure switching effect between them is given by equation (11), where $\text{Var}_{b_{hit_0}|s} [\log \hat{P}]$ is simply σ_s^2 . Conditional on differences in $\eta_{h'st_0} - \eta_{hst_0}$, the levels of these elasticities do not matter, now globally.

Discussion of preferences. The non-homothetic CES preferences we consider are general in a number of ways. First, they allow for non-unitary income elasticities that vary across goods within sectors as can be seen in equation (5), driven by differences in γ_i across goods and the dependence of price elasticities on u , and also across sectors as can be seen in equation (6). As shown in equation (2) and discussed in Section 3.1 in the general formulation, income elasticities play no role in the construction of the CV conditional on knowing initial expenditure shares and compensated cross-price elasticities. Second, these preferences allow for elasticities of substitution that vary across households as a function of utility u_h , as in Fally (2022).²² As shown in equation (9), calculating CV_h requires values for these elasticities of substitution in the initial period.

Contrary to this generality, these preferences impose strong restrictions. Elasticities of substitution are constant along any indifference curve as in standard CES models. We make this assumption for three reasons. First, we estimate these elasticities of substitution leveraging the 2015 Swiss franc appreciation, which does not contain sufficient price variation to estimate them globally. Second, this restriction has an appealing theoretical property: it implies that the integral over prices in equation (2) simplifies substantially, as shown in equation (9). It additionally implies that compensated budget shares in equations (7) and (8) and the CV in equation (9) for a particular household are identical to those in a model in which the household has homothetic and stable CES preferences with household-specific, exogenously given, and constant demand shifters and elasticities.²³ Third, this restriction implies that only a small subset of preference parameters are required for measuring CV, as opposed to other demand systems, e.g. the Almost Ideal Demand System, in which the same parameters control both income and cross-price elas-

²²Fally (2022) establishes sufficient conditions for the rationalization of non-homothetic CES demand when the elasticity of substitution is a decreasing function of u , which is the empirically relevant case in our data. In Appendix C we show—under certain assumptions—that these conditions are satisfied under the parameterization of $\eta(u)$ that we assume to derive our main estimating equation. We also describe numerical simulations for which the expenditure function is monotonically increasing in u for a wide range of parameters.

²³It is standard to calculate changes in price indices across households imposing homothetic CES preferences with demand shifters and elasticities that vary across households but are fixed in the counterfactuals (see, e.g., Handbury 2021). Our results show that this is equivalent to calculating changes in the welfare-relevant deflator when preferences are generalized non-homothetic CES.

ticities.

The other strong restriction we impose is that for any household there is a single elasticity, ρ , that shapes substitution between sectors and a single elasticity, η_{hst} , that shapes patterns of substitution between goods within sector s . This dramatically reduces the dimensionality of the problem. This formulation is equivalent to one with additional nests (e.g. product categories) under the assumption that the elasticity of substitution within nests is equal to the one between nests.

4 Elasticities of substitution and income

In this section we estimate differences in compensated price elasticities across incomes using the Homescan data, where we observe household-product-specific expenditure shares and prices.

4.1 Estimating equation

To estimate how elasticities of substitution vary with income, we must take into account that changes over time in budget shares reflect not only price changes but also income effects and demand shifters. For any continuing good, differentiating equation (5) at t_0 yields

$$d \log b_{hit} = d \log \zeta_{hit} + \left(\gamma_i - \frac{\partial \eta_s}{\partial u_h} u_{ht_0} \log p_{hit_0} \right) d \log u_{ht} + (1 - \eta_{hst_0}) d \log p_{hit} + \psi_{hst} \quad (12)$$

where $\psi_{hst} \equiv d \log \left(b_{hst} / \sum_{i' \in \mathcal{I}(s)} \zeta_{hi't} u_{ht}^{\gamma_{i'}} p_{hi't}^{1-\eta_s(u_{ht})} \right)$ and all derivatives (in the previous and subsequent equations) are evaluated at t_0 . Differentiating $I_{ht} = e_h(\mathbf{p}_{ht}, u; \zeta_{ht})$,

$$d \log u_{ht} = \left(\frac{\partial \log e_h}{\partial \log u_h} \right)^{-1} \times \left(d \log \frac{I_{ht}}{P_{ht}} - \bar{\varepsilon}_{ht} \right) \quad (13)$$

where $d \log(I_{ht}/P_{ht})$ is the change in income deflated by a household-specific weighted average of price changes across goods in all sectors $d \log P_{ht} \equiv \sum_i b_{hit_0} d \log p_{hit}$, and $\bar{\varepsilon}_{ht} \equiv \sum_i (\partial \log e_h / \partial \zeta_{hi}) d \zeta_{hit}$ is the shift in the expenditure function due to taste shifters; see Appendix C for derivations. We refer to $d \log(I_{ht}/P_{ht})$ as the change in real income for

household h . Substituting (13) into (12) yields

$$d \log b_{hit} = \left(\frac{\partial \log e_h}{\partial \log u_h} \right)^{-1} \times \left(\gamma_i - \frac{\partial \eta_s}{\partial \log u_h} \log p_{hit_0} \right) \times \left(d \log \frac{I_{ht}}{P_{ht}} - \bar{\epsilon}_{ht} \right) + d \log \zeta_{hit} + (1 - \eta_{hst_0}) d \log p_{hit} + \psi_{hst} \quad (14)$$

To estimate how elasticities of substitution vary with initial income, we impose two restrictions (in Appendix C we provide a cardinalization of the utility function that microfound these two restrictions). These restrictions play no role for our counterfactual welfare calculations conditional on estimates of elasticities of substitution; we impose these restrictions to facilitate the estimation of these elasticities. First, we assume that household h 's income elasticity for good i at t_0 , $\partial \log b_{hit} / \partial \log I_h$, can be expressed as the sum of a good i -specific term that is common for all households, which we denote by κ_i , and a household-sector-specific term. This assumption holds if the term multiplying the change in real income in expression (14) evaluated at t_0 can be written as

$$\left(\frac{\partial \log e_h}{\partial \log u_h} \right)^{-1} \times \left(\gamma_i - \frac{\partial \eta_s}{\partial \log u_h} \log p_{hit_0} \right) = \kappa_i + \kappa_{hs} \quad (15)$$

Second, we assume a log-linear relation between the elasticity of substitution in sector s and household income in the initial period,

$$\eta_{hst_0} \equiv \bar{\eta}_s + \eta_s \log I_{ht_0}. \quad (16)$$

If $\eta_s < 0$, then a higher-income household is less price sensitive in sector s at t_0 .

Under these two additional restrictions, equation (14) can be expressed as

$$d \log b_{hit} = v_{hit} + \kappa_i d \log \left(\frac{I_{ht}}{P_{ht}} \right) + [1 - \bar{\eta}_s - \eta_s \log(I_{ht_0})] d \log p_{hit} + \tilde{\psi}_{hst}. \quad (17)$$

The first term, $v_{hit} \equiv d \log \zeta_{hit} - \kappa_i \bar{\epsilon}_{ht}$, corresponds to household h 's demand shifter for good i due to taste shocks. The second term captures the interaction between the good i -specific component of the income elasticity and the change in real income for household h , giving rise to a demand shifter for good i due to income effects. The third term corresponds to the compensated price elasticity for good i in the initial period interacted with the change in the price of good i . The last term, $\tilde{\psi}_{hst}$, groups all factors that vary at the sector $\times$ household level.²⁴

²⁴Setting changes in income, tastes, and prices of goods $j \neq i$ equal to zero, equation (17) resembles the familiar Slutsky equation relating Marshallian, Hicksian, and income elasticities. Our baseline approach to

We can decompose the demand shifter v_{hit} —without loss of generality—into the component of the demand shock for good i that is common across all households, a demand shock for imports that varies freely across households, and a household-good-specific deviation from these. Specifically, $v_{hit} \equiv v_{it} + \widetilde{\text{FE}}_{hst} \mathbb{I}_i^M + \tilde{v}_{hit}$, where $\mathbb{I}_i^M$ is an indicator variable that equals one if good i is imported. This yields our baseline estimating equation

$$d \log b_{hit} = \text{FE}_{it} + \text{FE}_{hst}^M + \kappa_i d \log \left(\frac{I_{ht}}{P_{ht}} \right) - \eta_s \log(I_{ht_0}) d \log p_{hit} + \iota_{hit}. \quad (18)$$

In equation (18), the product-fixed effect FE_{it} is the sum of the average product-specific demand shock across households, v_{it} , and the common impact of the average price change for product i across households, $(1 - \bar{\eta}_s) d \log p_{it}$; the term $\text{FE}_{hst}^M \equiv \psi_{hst} + \widetilde{\text{FE}}_{hst} \mathbb{I}_i^M$ is a household $\times$ import status fixed effect; and finally, the term $\iota_{hit} \equiv \tilde{v}_{hit} + (1 - \bar{\eta}_s)(d \log p_{hit} - d \log p_{it})$ is a residual that includes both the household's demand-shock deviation for product i (relative to the average across households and, if i is imported, the household's average demand shock for imported goods) as well as the common effect of the household-specific deviation in the change in product i 's price relative to its average change across households.

We estimate variation in elasticities of substitution across households, η_s , using equation (18) in two different ways leveraging distinct sources of variation. In our first approach, we use variation in changes in import relative to domestic expenditures across higher- and lower-income households, similar to the variation in Stylized Fact 3. In our second approach, we use variation in changes in expenditures across individual barcode products and variation in product price changes across aggregations of higher- and lower-income households. The advantages of the first approach are simplicity, the ability to estimate equation (18) at the household level, and the straightforward connection to Stylized Fact 3. We view this as a first pass. The benefit of the second approach is that it substantially relaxes our identification assumption: it is valid in the presence of entry and exit of products and in the presence of import demand shocks that vary systematically with income (which we cannot a priori rule out, even though the 2015 CHF appreciation was triggered by a policy response to foreign events and took place in the context of a stable Swiss economy both in terms of aggregates and nominal income inequality). In both approaches, we estimate equation (18) taking differences between 2014 and 2015. Even though these two approaches leverage entirely distinct sources of variation to identify η_s , they yield remarkably similar results.

estimating differences in Hicksian price elasticities does not require estimating income elasticities for each good using cross-sectional data. In Appendix B.2 we consider an alternative procedure that relaxes restriction (15) but requires first estimating income elasticities in the cross section under additional assumptions.

4.2 Approach 1: Import and domestic expenditures by household

In our first approach, we assume that there are only two goods within groceries: an imported good $i = M$ and a domestic good $i = D$.²⁵ In this case, whereas we can control for an aggregate import demand shock (contained in $\mathbb{F}\mathbb{E}_{it}$ in equation 18), we cannot control for a household-specific import demand shock, since there is only one imported good. Hence, $\mathbb{F}\mathbb{E}_{hst}^M$ reduces to a household effect. Since there are only two goods, we take differences across the imported and domestic goods and, since there are only two time periods, we also replace the time effect with a constant and estimate

$$d \log \left(\frac{b_{hMt}}{b_{hDt}} \right) = \alpha + \kappa d \log \left(\frac{I_{ht}}{P_{ht}} \right) - \eta_s \log(I_{ht_0}) d \log \left(\frac{p_{Mt}}{p_{Dt}} \right) + l_{ht} \quad (19)$$

Here, α , κ , and l_{ht} all represent differences of the parameters in equation (18) across imported and domestic goods: $\alpha \equiv \mathbb{F}\mathbb{E}_{Mt} - \mathbb{F}\mathbb{E}_{Dt}$, $\kappa \equiv \kappa_M - \kappa_D$, and $l_{ht} \equiv l_{hMt} - l_{hDt}$.

We measure b_{hDt} and b_{hMt} as the expenditure shares on domestic and imported goods within each individual household. The price changes for imported and domestic goods are measured as weighted averages of annual changes in national prices of products weighted by expenditures per product across all consumers in 2014, separately for imports and domestic goods. By using a single national relative import price change, our estimating equation in approach 1 is very similar to that used in columns 1–3 of Table 2, where we documented Stylized Fact 3. We measure household h 's inflation rate (across all sectors), $d \log P_{ht}$, using disaggregated price data in the CPI as measured by the SFSO (these price changes are common across households) and income group-specific expenditure shares across these disaggregated categories. We measure annual changes in nominal income by household using a Swiss household panel on income (FORS); details are available in Appendix A.

Identification. We identify differences in elasticities across household incomes from changes in import expenditure shares that are correlated with household income. We do not instrument for price changes since we have only one value of $d \log (p_{Mt}/p_{Dt})$. The identification assumption estimating regression (19) using OLS is that household-specific import demand shock deviations from the aggregate import demand shock between 2014 and 2015 are uncorrelated across household incomes in 2014. We relax this restriction in Ap-

²⁵To map this first approach to our data (in which there are multiple imported and domestic goods), we assume that each of the two goods is itself a stable homothetic aggregator across a fixed set of imported and domestic varieties. In constructing import and domestic prices within groceries to use in the estimation, we use a first-order approximation of the expenditure function of any homothetic aggregator within each import status.

Table 4: Estimation of η_s in Approach 1 using equation (19)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\log(I_{ht0})d \log(p_{Mt}/p_{Dt})$	2.189*** [0.554]	2.207*** [0.567]	1.838*** [0.478]	1.981*** [0.618]	2.041*** [0.672]	2.172*** [0.540]	2.361*** [0.591]
Constant	0.553*** [0.129]	0.557*** [0.132]	0.469*** [0.110]	0.492*** [0.144]	0.516*** [0.157]	0.550*** [0.126]	0.581*** [0.131]
Observations	2901	2901	2901	2901	2901	2901	2901
Baseline	X						
No winsorizing		X					
Winsorize 5%			X				
Unweighted				X			
Expenditure weight					X		
No income effects						X	
HH size							X

Notes: The estimating equation is (19). We report $-\eta_s$ and α . Observations are households and the dependent variable is the log change in import-relative-to-domestic expenditures across all Homescan products between 2014 and 2015. In our baseline in column 1, robust standard errors are clustered by 50 household bins according to household income in 2014, observations are weighted by the product of the number of households in each of the 50 bins and the share of each household's expenditures among households within that bin, and we winsorize the dependent variable at the first percentile in both tails. Columns 2–7 each make one change relative to the baseline in column 1. In column 2 we do not winsorize, in column 3 we instead winsorize at the 5th percentile, in column 4 we do not weigh observations, in column 5 instead weigh observations by expenditures in 2014, in column 6 we omit changes in real income from the regression specification, and in column 7 we control for household size. *p<.1; **p<.05; ***p<.01

proach 2.

Results. Whereas we estimate regression (19) at the household level, we cluster standard errors by 50 income bins defined by quantiles of the household income distribution in our sample in 2014. We do so to allow for the possibility of correlated imported demand shocks across households in the same income bin; however, as stated above, we continue to require that import demand shocks between 2014 and 2015 across income bins are not systematically related to income. We weigh observations (households) by the product of the number of households in each of the 50 bins and the share of household h 's expenditures among households within that bin and winsorize the dependent variable at the first percentile in both tails. We revisit each of these choices in robustness.

The first column of Table 4 displays our baseline results. In all tables we report the estimated coefficient, which is $-\eta_s$. We find $\eta_s = -2.19$, which implies a substantially lower elasticity of substitution for higher-income households. For example, the elasticity of substitution of a household with 2014 income of 60,000 CHF is approximately 2.4 ($\approx 2.19 \times \log 3$) lower than a household with income of 20,000 CHF. This gap shapes the non-linear effects of import price changes in our quantification.

4.3 Approach 2: Product-level expenditures by income group

In our second approach, each i is an individual barcode product.²⁶ Given the granularity of this definition of a product i relative to our first approach, the household-product-level data are sparse. Hence, to estimate equation (18) we aggregate product-level data across groups of households, as is standard in demand estimation. Specifically, we group households into 50 income bins defined by quantiles of the household income distribution in our sample in 2014; $h \in \{1, \dots, 50\}$ now denotes the income bin. Within each bin, we take the median value of 2014 income and the median annual change in nominal income between 2014 and 2015 across individual households. We measure inflation and changes in income at the household level as described above.

In our baseline we use a common price change across households at the product level, $d \log p_{hit} = d \log p_{it}$. In robustness we consider a more disaggregated household aggregation that incorporates spatial variation; in this case, we measure a common $d \log p_{hit}$ for all households within a one-digit zip code and allow these price changes to vary across space.

Identification. In our second approach, we identify differences in elasticities across households from differences across the income distribution in the relationship between changes in expenditure shares and prices at the barcode product level. In this case, we can explicitly incorporate import demand shocks that are specific to each of the fifty household aggregates.

There are two remaining endogeneity concerns. The first is measurement error, which generates attenuation bias. The second is an economic argument for endogeneity. Suppose that high-income households are less price sensitive, consistent with our findings in Approach 1. Consider a product that faces demand shocks that are higher for higher-income households (we control for the average demand shock across households, so only deviations from the average remain in the residual). In response, the firm will face a more inelastic demand and will, therefore, charge a higher markup. Hence, there is a positive correlation between household-specific demand shock deviations (i.e., the residual) and the interaction between product-specific price changes and household initial income (the independent variable of interest). This implies that under the hypothesis that high-income households are less price sensitive, OLS is upward biased.

We address these concerns by constructing an instrument using an interaction between a product-specific cost shifter and initial household income. Our cost shifter ex-

²⁶Our baseline sample is the set of products that were purchased at least once per month (nationally) in the year-and-a-half before and after the CHF appreciation.

exploits variation across imported goods in the invoicing currency of prices at the border. As described in Section 2.1, we match products to border groups and measure the share of imported products in each border group that are denominated in EUR (out of those denominated in either EUR or CHF) in 2014, which we denote by $share_{it_0}$. Because of the stickiness of import prices at the border in their invoicing currency, Swiss retail prices of imported goods are more responsive to the CHF appreciation if imports are denominated in EUR rather than in CHF; see Auer et al. (2021).

Since the expected reduction in Swiss retail prices in response to the CHF appreciation is greater for imported products that belong to border groups with a higher fraction of border prices invoiced in EUR, we construct our instrument as the interaction between (i) the share of imported goods in the corresponding border group that are denominated in EUR, $share_{it_0}$, (ii) an import indicator variable, $\mathbb{I}_i^M$, and (iii) the logarithm of initial household income, $\log(I_{ht_0})$. If we restrict our sample to imported goods, as we do in robustness, then the instrument is the interaction between (i) and (iii) alone: $\log(I_{ht_0})share_{it_0}$. In this case, we leverage the fact that the expected reduction in Swiss retail prices among imported goods in response to the CHF appreciation is greater for those goods belonging to border groups with a higher fraction of border prices invoiced in EUR.²⁷

In some border groups, the number of border price observations denominated in either EUR or CHF with which to construct (i)—the share of imported goods that are denominated in EUR (out of those denominated in EUR or CHF)—is small and, therefore, the share is unreliable. Hence, in our baseline we restrict the sample of products to those in border groups with more than 28 border price observations in 2014 and vary this cutoff in robustness. This leaves 35 border groups in our baseline estimation sample, 7 of which have no imported goods denominated in EUR (i.e., $share_{it_0} = 0$ for all products i in these 7 border groups).

Results. In our baseline, we weigh observations by the product of the number of households in each aggregation h and the share of expenditures among households within that aggregation on product i .²⁸ In constructing changes in log expenditure shares, we winsorize changes in log expenditures at the first percentile (both in the right and left tails). Finally, we two-way cluster standard errors at the level of household income (there are 50

²⁷There is a large literature (see, e.g., Gopinath and Itskhoki 2022) arguing that the data on invoicing currency of export and import prices are consistent with models in which firms' invoicing currency choices are based on desired pass-through to exchange rate movements. A sufficient condition for our exclusion restriction is that heterogeneity in *anticipated* relative demand shocks across the income distribution between 2014 and 2015 does not shape pre-shock invoicing currency choices.

²⁸This approach puts equal weight on each underlying household rather than giving a higher weight to those household aggregations with higher expenditures (since our objective is to estimate how price sensitivities vary with income).

Table 5: Estimation of η_s in Approach 2 using equation (18)

	(1) OLS	(2) RF	(3) 2SLS
$\log(I_{ht_0}) \times d \log p_{it}$	0.018 [0.134]		1.930** [0.867]
$\log(I_{ht_0}) \times share_{it_0} \times \mathbb{I}_i^M$		-0.140** [0.068]	
Observations	95,325	95,325	95,325
K-P F Stat (first stage)			13.1

Notes: The estimating equation is (18). Observations are barcode product $\times$ household aggregates, where households are aggregated into 50 bins according to initial income. The dependent variable is the log change in expenditures between 2014 and 2015. Column 1 reports OLS results, column 2 reports reduced-form results in which we replace $\log(I_{ht_0})d \log p_{it}$ with $\log(I_{ht_0})share_{it_0}\mathbb{I}_i^M$, and column 3 reports 2SLS results in which we instrument for $\log(I_{ht_0})d \log p_{it}$ with $\log(I_{ht_0})share_{it_0}\mathbb{I}_i^M$. Robust standard errors are two-way clustered at the level of household income bin and, separately, the interaction between import status and the share of imported goods denominated in EUR in the border group; observations are weighted by the product of the number of households in each aggregation and the share of expenditures among households within that aggregation on product i ; and we winsorize changes in log expenditures at the first percentile (both in the right and left tails). * $p < .1$; ** $p < .05$; *** $p < .01$

such clusters) and, separately, the interaction between import status and the value of the share of imported goods denominated in EUR in the corresponding border group (there are 54 such clusters). We revisit each of these choices in robustness.

Table 5 displays our baseline results, focusing on the parameter of interest: η_s . The first column reports results from estimating equation (18) using OLS, where we find an economically small and statistically insignificant estimate. Column 2 reports results from estimating the reduced-form specification, in which we replace the change in product price interacted with the logarithm of initial household income with the instrument. We find that, between 2014 and 2015, higher-income households increased their expenditures by less on imported goods within border groups with a higher share of products invoiced in EUR (those with a larger decline in border and retail prices in response to the 2015 CHF appreciation) conditional on real income changes, import demand shocks that vary freely across household income groups, and other covariates. This is the expected sign of the reduced-form relationship.

Column 3 reports the baseline version of our main empirical result, the two-stage least squares estimate of η_s . The first-stage coefficient is -0.073 (implying that, on average, the price of an imported product in a border group that is entirely invoiced in EUR fell by 7.3% more than a product in a border group entirely invoiced in CHF in response to the roughly 14% appreciation of the CHF) and the associated F statistic is 13.1.²⁹ The second-stage coefficient of interest, $\eta_s = -1.93$, is very similar to the estimate in our first approach, which leverages an entirely distinct source of identification. The bias of the

²⁹Throughout, we report the Kleibergen-Paap Wald rk F statistic when there is only one endogenous variable.

OLS estimate is as expected both due to measurement error and the economic argument described above (recall that we display the negative value of the OLS coefficient, since this is the structural parameter of interest).

4.4 Robustness and sensitivity

Table 4 displays robustness across a range of choices in our first approach. In column 2 we do not winsorize the dependent variable; in column 3 we instead winsorize at the fifth percentile. In column 4 we do not weigh observations and in column 5 we instead weigh by household expenditure in 2014. In column 6 we omit changes in real income from the estimating equation.³⁰ Finally, in column 7 we control for household size in case it is correlated with income and elasticities vary with household size. Results are largely robust to these choices.

The majority of our robustness exercises focus on our second approach. Most of these—such as incorporating spatial variation into the estimation, varying choices in the construction of household income, and determining whether particular income groups drive our results—are contained in Appendix B.2. Here, we describe a small number of sensitivity and robustness exercises, with results displayed in Table 6.

Recall from Table 2 in Section 2.2 that the gap between the import shares of low- and high-income households fell between 2013–14 and rose both between 2014–2015 and 2015–2016. Changes between 2013–14 might suggest pre-existing trends that would call into question our baseline results. Changes between 2015–16 might suggest mechanisms generating lags in expenditure-switching responses.

Here, we begin by showing that there are no such pre-trends in Approach 2. Column 1 of Table 6 replicates our baseline reduced-form specification. Column 2 of Table 6 documents an absence of pre-existing trends in the reduced-form specification; we cannot study pre-trends in the structural specification since our instrument has no power before the CHF appreciation, given that the CHF-EUR exchange rate is stable between 2013–14. Whereas in our baseline we obtain a coefficient of -0.14 that is significant at the 5% level, running the same regression but replacing changes in expenditure shares between 2014–15 with changes between 2013–14 yields a coefficient that is three orders of magnitude smaller and statistically insignificant; see column 2 of Table 6. We interpret this lack of differential pre-existing trends as strengthening the structural interpretation

³⁰In additional sensitivity on income effects in Appendix B.2, mentioned in footnote 24, we use an alternative approach that relaxes restriction (15), but requires estimating cross-sectional income elasticities under strong assumptions. We apply this in Approach 1 and show that results are robust in Table A17.

Table 6: Robustness of Approach 2

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\log(I_{ht0}) \times share_{it0} \times \mathbb{I}_i^M$	-0.14** [0.07]	-0.00 [0.07]	-0.10 [0.08]	-0.14* [0.07]				
$\log(I_{ht0}) \times d \log p_{it}$					1.93** [0.87]	1.93** [0.88]	2.19** [0.87]	2.15** [0.87]
Observations	95,325	98,652	78,800	27,128	95,325	27,128	95,325	95,325
Baseline	X				X			
Outcome period		13-14	15-16					
Imports only				X		X		
Additional controls I							X	X
Additional controls II								X
F Stat (first stage)					13.1	12.8	15.6	17.1

Notes: Column 1 replicates our baseline RF specification shown in column 2 of Table 5. Columns 2 and 3 report the same specification, but in which the outcome variable is defined over the period 2013–14 (column 2) and 2015–16 (column 3). Column 4 displays estimates of the RF specification on a sample restricted to imported goods alone (so that $\mathbb{I}_i^M = 1$ for all observations). Column 5 replicates our baseline 2SLS specification shown in column 3 of Table 5. Column 6 displays estimates of the 2SLS specification on a sample restricted to imported goods alone. Column 7 incorporates two additional controls interacted with year: the 2014 import share as well as the 2014 expenditure share on each border group. Column 8 additionally incorporates one more control interacted with year: the 2014 average price of each individual product. In columns 5–8 we report the KP F statistic. *p<.1; **p<.05; ***p<.01

of our baseline results.³¹

Another concern is that the share of imported goods in each border group that are denominated in EUR is correlated with some other product characteristic and that this other product characteristic is driving the differential patterns of substitution for higher- and lower-income households. Here we show that controlling for additional triple interactions in which we replace the share of imported goods in the border group that are denominated in EUR with other border group or product characteristics (the 2014 import share of each border group, the 2014 expenditure share of each border group, and the 2014 average price of each individual product) does not substantially change our results. Column 5 of Table 6 replicates our baseline 2SLS estimate from column 3 of Table 5. Columns 7 and 8 of Table 6 show that including these additional controls has little effect on results.

Finally, in our baseline, we include both imported and domestically produced goods in our estimation sample. However, the instrumented change in price is common for all domestic goods. Columns 4 and 6 display reduced-form and 2SLS results when we restrict the sample to imports.³² Results are largely unchanged, although we only have

³¹Column 3 similarly replicates our baseline reduced-form specification using changes in expenditure shares between 2015–16. As in the 2013–14 period, we find statistically insignificant results. However, the estimated coefficient is closer to the baseline value. This could represent evidence of dynamics in expenditure-switching responses, which we do not explore further.

³²Restricting the sample to imported goods requires replacing our household-import indicator fixed effect with a household fixed effect, since there is no need (or possibility) to control for differential import demand shocks across income groups. Moreover, our instrument reduces to the interaction between the log of household income and the share of imported goods denominated in EUR in the border group.

26 clusters in one dimension.

4.5 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; see Appendix B.3 for details.

In Approach 1, we can identify $\bar{\eta}_s$ if we assume that the average import demand shifter v_{it} is zero between 2014 and 2015. This yields $\bar{\eta}_s \approx 26.6$. In Approach 2, we can identify $\bar{\eta}_s$ if we do not control for the average product-specific demand shock v_{it} and, instead, move it to the residual. This yields $\bar{\eta}_s = 20.87$. The first (second) approach implies that the elasticity of substitution is 4.92 (1.76) for a household with income of 20,000 CHF and that this elasticity remains positive for all household incomes below approximately 190,000 (50,000) CHF.

To sum up, we impose weaker assumptions estimating η_s than $\bar{\eta}_s$ and our estimates of η_s are much more similar across approaches 1 and 2 than are our estimates of $\bar{\eta}_s$. For these reasons, in our quantitative analyses in Section 5, we present results for a range of $\bar{\eta}_s$ and show that differences in welfare across incomes do not depend crucially on these values within a broad range, consistent with our analytic results in Section 3.2.

5 Quantification

In this section, we use our estimates from Section 4 to assess the role of heterogeneous expenditure switching in shaping the welfare implications—measured using equation (9)—of factual and counterfactual changes in prices. In Section 5.1 we quantify changes in the welfare-relevant price index for groceries using observed price changes in the Homescan data. In Section 5.2 we consider changes in welfare in response to counterfactual changes in import relative to domestic prices across all consumer goods.

We make the following choices in both sections. The initial period, t_0 , is 2014. Differences in within-grocery elasticities of substitution across incomes, $\eta_s = -2$, match our Homescan-based grocery estimates in Section 4. We consider households at three income levels within the range of incomes in our sample—20,000, 60,000, and 120,000 CHF—and choose $\bar{\eta}_s$ so that the lowest elasticity of substitution across the income groups that we consider (at 120,000 CHF) is equal to 3.

Table 7: Welfare-relevant grocery price changes 2014–15

Income	Heterogeneous elasticities				Homogeneous elasticities			
	Elasticity	(1) 1st-order	(2) Switching	(3) Exact	Elasticity	(4) 1st-order	(5) Switching	(6) Exact
1: 20,000	6.6	-1.1	-1.0	-2.2	6.6	-1.1	-1.0	-2.2
2: 60,000	4.4	-1.2	-0.6	-1.7	6.6	-1.2	-0.9	-2.2
3: 120,000	3.0	-1.3	-0.3	-1.6	6.6	-1.3	-0.9	-2.3

Notes: This table displays changes in the welfare-relevant price index, $\hat{P}_{hs}$ in equation (8), for the grocery sector in response to observed price changes of individual products in the Homescan data between 2014 and 2015. Rows 1–3 display results for households with incomes of 20,000 CHF, 60,000 CHF, and 120,000 CHF. Columns 1–3 use heterogeneous elasticities (6.6, 4.4, and 3) whereas columns 4–6 impose common elasticities (all set to 6.6). Columns 1 and 4 display the first-order effects, columns 2 and 5 display the second-order effect, and columns 3 and 6 display the exact change.

5.1 Heterogeneous effects of observed Homescan price changes

In this section we quantify changes in the welfare-relevant price index, $\hat{P}_{hs}$ in equation (8), for the grocery sector in response to observed price changes of individual products in the Homescan data between 2014 and 2015. We do not include non-groceries because we only observe price changes at the CPI level, which is much more aggregated than in the Homescan data, where we observe price changes at the barcode level. We measure these barcode-specific national price changes as described in Section 4.3.³³

To calculate $\hat{P}_{hs}$ in equation (8), we need expenditure shares and elasticities of substitution by income group in 2014. To measure expenditure shares, we divide households into three equal-sized bins based on 2014 income: the lowest, middle, and top bins contain households with 2014 annual income of 20,000 CHF, 60,000 CHF, and 120,000 CHF. For each bin, we calculate the expenditure share on each product and assign this expenditure share to the corresponding household; in sensitivity analysis we show that our main results are robust to using common expenditure shares across all households. The elasticities required to compute the price index are calculated as described above, yielding elasticities of 6.6, 4.4, and 3 for households with incomes of 20,000 CHF, 60,000 CHF, and 120,000 CHF, respectively; in sensitivity analysis we show that our main results are robust to varying $\bar{\eta}_s$ so that the elasticity for households with incomes of 120,000 CHF ranges between 1.5 and 5.

The left-hand panel of Table 7 contains our baseline results with heterogeneous elasticities. Column 1 displays the first-order welfare effect of price changes in groceries (using the second-order approximation in equation 10), which is simply the expenditure-share-weighted average of product price changes within groceries. These range from -1.1

³³To reduce the role of abnormally large price changes on the price index, we drop products with year-to-year price ratios above 3 or below 1/3. This has almost no impact on the 2014–15 results.

percent for households with incomes of 20,000 CHF to -1.3 percent for households with incomes of 120,000 CHF. Column 2 displays the expenditure-switching welfare effect of price changes, which is $(1 - \eta_{hst_0})$ times half the expenditure-share-weighted variance of price changes within groceries. Whereas these effects are smaller than the first-order terms (they range from -1.0 percent for households with incomes of 20,000 CHF to -0.3 percent for households with incomes of 120,000 CHF), their variation across incomes is larger. These differences in expenditure switching are driven almost entirely by differences in elasticities; in particular, the weighted variance is very similar across income groups. Column 3 shows the full non-linear effect, which is very similar to the sum of the first-order and expenditure-switching effects. The price change for low-income households is -2.2 percent, which is about 50% larger than the price change for high-income households. This gap is almost identical to the gap in the expenditure-switching effect.

Another way to see the importance of the expenditure-switching effect is to set all elasticities to be equal. In the right-hand panel of Table 7, we display results in which we impose the elasticity of the lowest-income households, 6.6, for all three household groups. The first-order effect, displayed in column 4, is obviously unchanged. However, now the expenditure-switching effect is very similar across households, unlike in column 2; it is not identical across households because of small differences in the expenditure-share-weighted variance of price changes. Price indices across income groups, displayed in column 6, differ by much less than those under heterogeneous elasticities reported in column 3; moreover, these differences are of the opposite sign.

Sensitivity. We provide a range of additional results in Appendix D. First, we display welfare-relevant grocery price changes the year before the CHF appreciation (2013–14). The variance of price changes between 2014–15 is one-and-a-half times the variance of price changes between 2013–14, as implied by Stylized Fact 5 above. Hence, the gap between income groups in the expenditure-switching effect is similarly one-and-a-half times larger in 2014–15 than in 2013–14.

We also display results imposing common expenditure shares across households. Whereas the first-order effects are, obviously, now identical across households, the second-order effects are little changed from our baseline. Finally, we also display results 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.

5.2 Heterogeneous effects of counterfactual import price changes

In this section we quantify the effect of counterfactual changes in import prices across all consumer goods. Our focus here on import price changes contrasts with our focus in Section 5.1 on observed price changes between 2014 and 2015, which reflect not only the CHF appreciation but also price changes that would have occurred in its absence. Our focus here on all consumer goods contrasts with our focus in Section 5.1 on groceries alone, where import shares do not vary systematically with income and the first-order effects of import price changes are mechanically very similar across households. At the aggregate level, import shares in 2014 are higher among higher-income households, as shown in Stylized Fact 1, yielding heterogeneous first-order effects of import price changes. Finally, to highlight the non-linearities induced by expenditure switching, we consider larger import price shocks like those induced by much larger exchange rate changes (see, e.g., Cravino and Levchenko, 2017) or a movement to autarky.

To model counterfactual price changes, we assume that the price change of any imported $j = M$ or domestic $j = D$ product i in any sector s is given by

$$\log \hat{p}_i = \log \hat{p}_j + \sigma_j \epsilon_i \quad (20)$$

where $\log \hat{p}_j$ is the *uniform component* of price changes across all imported or domestic products and $\sigma_j \epsilon_i$ is the product-level *idiosyncratic component* of price changes, with $\epsilon_i \sim \mathcal{N}(0, 1)$.

Given our focus on the expenditure-side effects of foreign price shocks, we assume that the log change in income for all households, $\hat{I}_h$, is equal to the average change in the log price of domestic goods, $\hat{p}_D$, as in single-factor trade models without imported intermediate inputs.³⁴

Under these assumptions, changes in welfare in (9) are given by

$$CV_h = -\frac{1}{1-\rho} \log \left[\sum_s b_{hst_0} \left(\sum_{j \in M, D} \frac{b_{h j st_0}}{b_{hst_0}} \left[\frac{\hat{p}_j}{\hat{p}_D} e^{\frac{1}{2} \sigma_j^2 (1-\eta_{hst_0})} \right]^{1-\eta_{hst_0}} \right)^{\frac{1-\rho}{1-\eta_{hst_0}}} \right] \quad (21)$$

where $b_{h j st_0} / b_{hst_0}$ is the share of expenditure in sector s on either imports $j = M$ or domestic goods $j = D$. According to (21), a higher variance σ_j^2 increases CV_h (for $\eta_{hst_0} > 1$), in a similar way that a lower $\log \hat{p}_j$ does, and this effect is stronger the larger η_{hst_0} is. Whereas

³⁴In estimating compensated price elasticities, we do not impose this restriction, but instead use actual changes in retail prices by good. In our counterfactuals, if all domestic goods have a common imported intermediate share, then differences between households in welfare changes do not depend on the value of this share, for any given change in import to domestic prices.

we use equation (21) in our analysis below, we gain further intuition in the special case in which $\sigma_j = \sigma$ and $\eta_{hst_0} = \eta_{ht_0}$, where changes in welfare are given by

$$CV_h = - \underbrace{\frac{1}{1-\rho} \log \left[\sum_s b_{hst_0} \left(\frac{b_{hMst_0}}{b_{hst_0}} \left(\frac{\hat{p}_M}{\hat{p}_D} \right)^{1-\eta_{ht_0}} + \frac{b_{hDst_0}}{b_{hst_0}} \right)^{\frac{1-\rho}{1-\eta_{ht_0}}} \right]}_{\text{uniform component}} + \underbrace{\frac{1}{2} \sigma^2 (\eta_{ht_0} - 1)}_{\text{idiosyncratic component}} \quad (22)$$

In this case, changes in welfare are additively separable in the uniform and idiosyncratic components of price changes.

We consider three sectors s : groceries, non-grocery goods, and services. For each of the three income groups (20,000, 60,000, and 120,000 CHF), we construct import shares in each of these aggregate sectors and expenditure shares across them using data on expenditures and import shares within highly disaggregated consumer categories in the SFSO data (see Table 1 in Section 2.1).³⁵ Aggregating up to three sectors has no effect on our measure of overall import shares by income. We impose a value of ρ very close to one, $\rho = 0.99$, and vary this parameter in sensitivity analysis. We choose values for η_{hst_0} as described above.

We quantify the impact of import price shocks, $\Delta \equiv \log \hat{p}_M - \log \hat{p}_D$, for different values of $\Delta > 0$ ranging from $\Delta = 2.2\%$ (the size of the *reduction* of import prices relative to domestic prices in 2015) to 1,000% (a movement to autarky, which is a focus of the quantitative trade literature). We begin by imposing $\sigma_j^2 = 0$, so that only the uniform component in equation (22) is active.

Uniform price changes within M and D. The first panel of Table 8 reports the welfare implications for each household of import price increases of various sizes.

Higher-income groups are harmed more by import price increases for two reasons. First, they have higher import shares, which shape the first-order effect displayed in equation (10). The import shares in 2014 are 21%, 24%, and 27% for households with incomes of 20,000, 60,000, and 120,000 CHF, respectively, as displayed in Table 1. Second, they have lower initial elasticities of substitution, which shape the expenditure-switching effect.

The bottom two panels of Table 8 highlight the increasing importance of the expenditure-switching effect as the size of the change in import prices Δ grows. The middle panel displays the percent difference between the CV of the lowest-income household and the CV of the middle- and highest-income households. For a 10% import price increase, the welfare of the middle- and high-income households fall by approximately 22% and 41%,

³⁵In practice, in assigning import shares in each of these three sectors and expenditure shares across them for our household with income of 60,000 CHF, we use an income of 60,252 CHF instead of 60,000 CHF. This is the cutoff separating the first and second income brackets in Table 1.

Table 8: Compensating variation of counterfactual import price shocks

Annual income	Import price shock					+2.2 $\sigma > 0$
	+2.2	+10	+20 $\sigma = 0$	+40	+1000	
1: 20,000 elasticity 6.6	-0.4	-1.8	-3.2	-4.7	-5.6	-0.2
2: 60,000 elasticity 4.4	-0.5	-2.2	-4.1	-7.0	-11.1	-0.4
3: 120,000 elasticity 3.0	-0.6	-2.6	-5.0	-9.1	-22	-0.5
% difference in CV btw						
income groups 2 and 1	16	22	31	50	99	83
income groups 3 and 1	30	41	57	95	295	148
Contribution of heterogeneous η s						
income groups 2 and 1	8	28	44	62	79	69
income groups 3 and 1	7	25	41	60	86	67

Notes: Percent changes are $100 \times$ the log of the relative price change. “% difference in CV btw income group j and 1” is $(CV_j - CV_1)/CV_1$ for income group j . “Contribution of heterogeneous elasticities” is $1 - (CV_j^{\text{homog}} - CV_1^{\text{homog}})/(CV_j - CV_1)$ where CV_j^{homog} is the compensating variation of income group j in our alternative counterfactual in which elasticities are common across income groups and set to the value for a household with income of 20,000 CHF. All columns but the last impose $\sigma_j^2 = 0$. In the final column, we set σ_j^2 for $j = D$ and $j = M$ to match the observed increase in idiosyncratic volatilities between 2013–14 and 2014–15.

respectively, more than for the low-income household. When import prices rise by more, the differences in welfare changes between incomes and the contribution of heterogeneous elasticities to these differences grow substantially.

To quantify the importance of the expenditure-switching effect, we consider an alternative parameterization in which we impose a common price elasticity across incomes equal to that of households with income of 20,000 CHF (which is 6.6). The bottom panel of Table 8 displays the contribution of heterogeneous elasticities (comparing heterogeneous elasticities and import shares to heterogeneous import shares alone) in shaping differences in welfare changes for the middle- and high-income groups compared to the low-income group.³⁶ Differences in elasticities between the low- and high-income groups explain only 8% (middle vs low income) and 7% (high vs low income) of the difference in welfare changes when the import price rises by 2.2%. However, when the import price rises by 20%, differences in elasticities explain 44% and 41% of the differences in welfare changes. The larger is the increase in import prices, the higher is the contribution of dif-

³⁶For income group j , this is simply $1 - (CV_j^{\text{homog}} - CV_1^{\text{homog}})/(CV_j - CV_1)$ where CV_j^{homog} is CV of income group j in our alternative parameterization with homogeneous elasticities. Another way to quantify the contribution of heterogeneous elasticities is to compare results with heterogeneity in both elasticities and import shares to results with heterogeneity in elasticities alone. These results are very similar to those reported in the bottom panel of Table 8.

ferences in elasticities to the unequal welfare changes across incomes. For a movement to autarky, the expenditure-switching effect accounts for the vast majority (79% and 86%) of the unequal welfare effects.

Dispersed price changes within M and D. To this point in Section 5.2, we have considered the uniform component of import price shocks $\Delta \equiv \log \hat{p}_M - \log \hat{p}_D$ imposing a zero variance for idiosyncratic price changes within imported $j = M$ and domestic goods $j = D$. Recall from equation (22) that if $\sigma_M^2 = \sigma_D^2$, the welfare impact due to idiosyncratic price changes is additively separable from the uniform component. To evaluate the overall effect, we must calibrate σ_j^2 .

We first consider an import price shock of size $\Delta = 2.2\%$, which is the size of the average decline in import relative to domestic prices observed between 2014–15. Rather than setting σ_j^2 to the observed variance of price changes between 2014–15 (which includes price changes that would have occurred in the absence of the import price shock), we set it to the observed *increase* in the variance of price changes between 2013–14 and 2014–15.³⁷

The final column of Table 8 shows that the welfare loss of the lowest-income household is smaller than for the middle- and highest-income households, and more than two-thirds of the differences across incomes is driven by heterogeneity in η s across households. These results contrast with the first column of Table 8 (no idiosyncratic price changes), where relative differences in welfare across households are smaller and mostly driven by heterogeneous import shares. Intuitively, the price volatility associated with the import price shock reduces its welfare costs for all households, but does so disproportionately for more elastic households.

Finally, we consider a larger import price shock, of size $\Delta = 10\%$. Since the idiosyncratic price volatility generated by such a shock is unobserved, we consider a wide range of volatilities. Each column in Table 9 considers a value of σ_j^2 that is a factor x of our calibrated variance under the $\Delta = 2.2\%$ shock for various values of x . The first column ($x = 0$) corresponds to the second column of Table 8, where we set the idiosyncratic price variance to zero. The second column ($x = 1$) imposes the same idiosyncratic variance as under the smaller, $\Delta = 2.2\%$ shock. As x increases from 0 to 3, differences in wel-

³⁷ To motivate why we set the counterfactual variance equal to the difference in variance between 2013–14 and 2014–15, suppose that the idiosyncratic component of price changes is the sum of a component induced by the import price shock, $\sigma_{1j}\epsilon_{1i}$, and a component that is orthogonal to the import price shock, $\sigma_{2j}\epsilon_{2i}$, where ϵ_{1i} and ϵ_{2i} are i.i.d. and normally distributed. The sum of the two components can be written as in equation (20), where $\sigma_j^2 = \sigma_{1j}^2 + \sigma_{2j}^2$. For our counterfactual import price shocks, we set $\sigma_{2j} = 0$. To assign σ_{1j} for a 2.2% import price shock, we assume that the variance of price changes between 2013–14 equals σ_{2j}^2 and the variance of price changes between 2014–15 equals $\sigma_{1j}^2 + \sigma_{2j}^2$. Thus, σ_{1j}^2 equals the variance of price changes between 2014–15 minus the variance of price changes between 2013–14. Specifically, from Table 3, we obtain $\sigma_M^2 = 0.0058 - 0.0033 = 0.0025$ and $\sigma_D^2 = 0.0023 - 0.0018 = 0.0005$.

Table 9: Compensating variation of a 10% import price shock

Annual income	Ratio of variance of idiosyncratic price changes to the calibrated variance with a 2.2% shock					
	0	1	1.5	2	2.5	3
1: 20,000 elasticity 6.6	-1.8	-1.6	-1.5	-1.4	-1.3	-1.1
2: 60,000 elasticity 4.4	-2.2	-2.1	-2	-1.9	-1.9	-1.8
3: 120,000 elasticity 3.0	-2.6	-2.5	-2.4	-2.4	-2.3	-2.3
% difference in CV btw						
income groups 2 and 1	22	30	35	40	47	55
income groups 3 and 1	41	55	63	73	85	99

Notes: All columns display results that correspond to the top two panels of the second column of Table 8 (using a 10% import price shock). But instead of setting $\sigma_D^2 = \sigma_M^2 = 0$, we set $\sigma_D^2 = x \times 0.0005$ and $\sigma_M^2 = x \times 0.0025$ (as described in footnote 37), for values of x displayed at the top of each column. Column 1 corresponds to column 2 of Table 8.

fare between incomes grow. This growth is entirely driven by the expenditure-switching effect.

Sensitivity. We provide a range of additional results in Appendix D. First, we consider import price declines rather than increases. In this case, high-income households benefit more from the first-order effect (they have higher initial import shares) whereas low-income households benefit more from the expenditure-switching effect (they have higher price elasticities). If $\sigma_j = 0$, the expenditure-switching effect dominates for large import price declines; it also dominates for our smallest import price decline of 2.2% if σ_j is calibrated to match its observed increase in 2014–15. Second, we set $\eta_s = -1.5$, which is at the lower end of our estimates of differences in elasticities of substitution across incomes, rather than $\eta_s = -2$. This slightly reduces differences in welfare across incomes induced by expenditure switching. Third, we vary $\bar{\eta}_s$ so that the lowest elasticity of substitution (that for the highest-income household, with income of 120,000 CHF) is equal to 1.5 or 5, instead of equal to 3. This leaves the differences in welfare across households largely unchanged (except for the extreme shock to import prices that essentially results in autarky). Fourth, we set the elasticity of substitution between sectors to $\rho = 0.2$ instead of $\rho = 0.99$, which does not have a strong impact on our results. Finally, we consider alternative assumptions on the value of η_s in non-grocery sectors s . Even when households share common elasticities within the service sector we obtain very similar results.

6 Conclusions

In this paper we revisit a classic question: what are the distributional implications of changes in foreign prices? We focus on differential changes in costs of living across households.

Theoretically, we show that differences across households in compensating variation in response to given income and price changes are shaped by initial expenditure shares across products and initial compensated cross-price elasticities. Empirically, we use detailed Swiss data to document that lower-income households engaged in significantly more expenditure switching towards imported goods in response to the 2015 Swiss franc appreciation. Leveraging these data and imposing generalized, non-homothetic CES preferences, we estimate substantially higher elasticities of substitution for lower-income households.

Import price increases in Switzerland harm higher-income households more than lower-income households both because higher-income households have higher initial import shares (the standard channel considered by the literature) and because they engage in less expenditure switching between imported and domestic goods (a channel from which the literature has abstracted). Quantitatively, we show that for large and dispersed price changes, unequal expenditure switching generates substantial differences in welfare across the income distribution.

Unequal expenditure switching can be relevant for the distributional consequences of high-inflation episodes, if these coincide with a rise in the dispersion of price changes (for evidence on the relation between inflation and price dispersion, see, e.g., [Alvarez et al., 2019](#)). We leave this for future research.

Data Availability Statement

The code and the publicly available data underlying this paper are available on Zenodo at <https://doi.org/10.5281/zenodo.8291248>. The Nielsen Homescan data is proprietary and available for purchase. The FORS data is not publicly available but can be requested free of charge.

References

ADAO, R., P. CARRILLO, A. COSTINOT, D. DONALDSON, AND D. POMERANZ (2022): “Imports, exports, and earnings inequality: Measures of exposure and estimates of incidence,” *The Quarterly Journal of Economics*, 137, 1553–1614.

- AGUIAR, M. AND E. HURST (2007): "Measuring Trends in Leisure: The Allocation of Time Over Five Decades," *The Quarterly Journal of Economics*, 122, 969–1006.
- ALVAREZ, F., M. BERAJA, M. GONZALEZ-ROZADA, AND P. A. NEUMEYER (2019): "From hyperinflation to stable prices: Argentina's evidence on menu cost models," *The Quarterly Journal of Economics*, 134, 451–505.
- ARGENTE, D. AND M. LEE (2021): "Cost of Living Inequality During the Great Recession," *Journal of the European Economic Association*, 19, 913–952.
- ATKIN, D., B. FABER, T. FALLY, AND M. GONZALEZ-NAVARRO (2023): "Measuring Welfare and Inequality with Incomplete Price Information," *The Quarterly Journal of Economics*, forthcoming.
- 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.
- BAI, L. AND S. STUMPNER (2019): "Estimating US Consumer Gains from Chinese Imports," *American Economic Review: Insights*, 1, 209–24.
- BANK FOR INTERNATIONAL SETTLEMENTS (2020): "US dollar exchange rates," <https://www.bis.org/statistics/xrusd.htm?m=2675>.
- BAQAEE, D. R. AND A. BURSTEIN (2023): "Welfare and output with income effects and taste shocks," *The Quarterly Journal of Economics*, 138, 769–834.
- BAQAEE, D. R., A. BURSTEIN, AND Y. KOIKE-MORI (2023): "Measuring Welfare by Matching Households across Time," *The Quarterly Journal of Economics*, forthcoming.
- BEMS, R. AND J. DI GIOVANNI (2016): "Income-Induced Expenditure Switching," *American Economic Review*, 106, 3898–3931.
- BERRY, S., J. LEVINSOHN, AND A. PAKES (2004): "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market," *Journal of Political Economy*, 112, 68–105.
- BORUSYAK, K. AND X. JARAVEL (2021): "The Distributional Effects of Trade: Theory and Evidence from the United States," Tech. rep., Mimeo: London School of Economics.
- BRODA, C. AND D. E. WEINSTEIN (2006): "Globalization and the Gains From Variety," *The Quarterly Journal of Economics*, 121, 541–585.
- BURSTEIN, A., M. EICHENBAUM, AND S. REBELO (2005): "Large Devaluations and the Real Exchange Rate," *Journal of Political Economy*, 113, 742–784.
- BURSTEIN, A., S. LEIN, AND J. VOGEL (2022): "Cross-border shopping: evidence and welfare implications for Switzerland," Tech. rep., Mimeo: UCLA.
- BURSTEIN, A. AND J. VOGEL (2017): "International trade, technology, and the skill premium," *Journal of Political Economy*, 125, 1356–1412.
- CARROLL, D. R. AND S. HUR (2020): "On the heterogeneous welfare gains and losses from trade," *Journal of Monetary Economics*, 109, 1–16.

- COIBION, O., Y. GORODNICHENKO, AND G. H. HONG (2015): "The cyclicalities of sales, regular and effective prices: Business cycle and policy implications," *American Economic Review*, 105, 993–1029.
- COMIN, D., D. LASHKARI, AND M. MESTIERI (2021): "Structural Change With Long-Run Income and Price Effects," *Econometrica*, 89, 311–374.
- CRAVINO, J. AND A. A. LEVCHENKO (2017): "The Distributional Consequences of Large Devaluations," *American Economic Review*, 107, 3477–3509.
- CRAVINO, J. AND S. SOTELO (2019): "Trade-Induced Structural Change and the Skill Premium," *American Economic Journal: Macroeconomics*, 11, 289–326.
- DEATON, A. (1997): *The Analysis of Household Surveys. A Microeconometric Approach to Development Policy*, Johns Hopkins University Press for the World Bank.
- FABER, B. AND T. FALLY (2022): "Firm heterogeneity in consumption baskets: Evidence from home and store scanner data," *The Review of Economic Studies*, 89, 1420–1459.
- FAJGELBAUM, P. D. AND A. K. KHANDELWAL (2016): "Measuring the Unequal Gains from Trade," *The Quarterly Journal of Economics*, 131, 1113–1180.
- FALLY, T. (2022): "Generalized separability and integrability: Consumer demand with a price aggregator," *Journal of Economic Theory*, 203, 105471.
- FEENSTRA, R. C. (1994): "New Product Varieties and the Measurement of International Prices," *American Economic Review*, 84, 157–177.
- FRIEDMAN, J. AND J. LEVINSOHN (2002): "The Distributional Impacts of Indonesia's Financial Crisis on Household Welfare: A "Rapid Response" Methodology," *The World Bank Economic Review*, 16, 397–423.
- GALLE, S., A. RODRÁGUEZ-CLARE, AND M. YI (2022): "Slicing the Pie: Quantifying the Aggregate and Distributional Effects of Trade," *The Review of Economic Studies*, 90, 331–375.
- GOPINATH, G. AND O. ITSKHOKI (2022): "Chapter 2 - Dominant Currency Paradigm: a review," in *Handbook of International Economics: International Macroeconomics, Volume 6*, ed. by G. Gopinath, E. Helpman, and K. Rogoff, Elsevier, vol. 6 of *Handbook of International Economics*, 45–90.
- HANDBURY, J. (2021): "Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living Across U.S. Cities," *Econometrica*, 89, 2679–2715.
- HARROD, R. F. (1936): *Trade cycle. An essay*, Oxford University Press, London.
- HAUSMAN, J. A. (1981): "Exact Consumer's Surplus and Deadweight Loss," *American Economic Review*, 71, 662–676.
- HE, Z. (2018): "Trade and Real Wages of the Rich and Poor: Cross-Region Evidence," Tech. rep., Mimeo.
- HOTTMAN, C. J. AND R. MONARCH (2020): "A matter of taste: Estimating import price inflation across U.S. income groups," *Journal of International Economics*, 127, 103382.

- JARAVEL, X. AND D. LASHKARI (2023): "Measuring Growth in Consumer Welfare with Income-Dependent Preferences: Nonparametric Methods and Estimates for the United States," *The Quarterly Journal of Economics*, forthcoming.
- JARAVEL, X. AND E. SAGER (2019): "What are the Price Effects of Trade? Evidence from the U.S. and Implications for Quantitative Trade Models," CEPR Discussion Papers 13902, C.E.P.R. Discussion Papers.
- KAPLAN, G. AND G. MENZIO (2016): "Shopping Externalities and Self-Fulfilling Unemployment Fluctuations," *Journal of Political Economy*, 124, 771–825.
- KUHN, U. (2018): *Collection, construction and checks of income data in the Swiss Household Panel*, University of Lausanne.
- MATSUYAMA, K. (2019): "Engel's law in the global economy: Demand-induced patterns of structural change, innovation, and trade," *Econometrica*, 87, 497–528.
- NEVO, A. (2000): "A practitioner's guide to estimation of random-coefficients logit models of demand," *Journal of economics & management strategy*, 9, 513–548.
- NIELSEN SWITZERLAND (2016): "Homescan Data Switzerland, Jan 2012-May2016," <https://www.nielsen.com/ch/de/contact-us/>.
- PORTO, G. G. (2006): "Using survey data to assess the distributional effects of trade policy," *Journal of International Economics*, 70, 140–160.
- 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.

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

Table A10: 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 are denominated in hundreds of thousands of CHF and Transactions are denominated in ten thousands.

Table A11: 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.

the 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.^{S38}

A.2 Household pre-tax income

Overview. The Homescan data includes a comprehensive set of household socioeconomic characteristics, as reported in Table A11. 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

^{S38}We construct the first column of Table A10 including all households with positive expenditures in 2014 without restricting to those with positive expenditures on products with known import status.

Swiss Household Panel compiled by the Swiss Centre of Expertise in the Social Sciences (henceforth 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.^{S39} We also predict the change in each household’s income between 2014–15 following a similar procedure (using the panel structure of FORS).^{S40}

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.^{S41} 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.

^{S39}In practice, the predicted level of household income falls within the relevant Homescan income bin for each household.

^{S40}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.

^{S41}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.

SFTA data. We additionally use data on the distribution of annual taxable household income 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 15,000, 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.^{S42} 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.^{S43}

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

^{S42}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.

^{S43}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 A12: Relationship between household income and expenditure share on products with known import status in 2014

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

Notes: Estimation of equation (A23), 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. * $p < .1$; ** $p < .05$; *** $p < .01$

estimated income fixed effects using SFTA data).

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 (A23), but replace the dependent variable with the share of household expenditure on products with known import status. Table A12 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.^{S44} 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

^{S44}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.

for each of five income groups. The expenditure share data we use in our analysis covers the years 2012–14.^{S45}

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.^{S46} The resulting concordance includes 187 consumption categories.^{S47}

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 in-

^{S45}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.

^{S46}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.

^{S47}This concordance is available in the replication material.

clude 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.

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 A13. 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 A13) and 26.2% in the Homescan data (the Homescan import share reported in column 4 of Table A13). 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 A13, 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

Table A13: Import shares in groceries using SFSO and Homescan data

	All categories	Unweighted	Matched categories	
	Own weights		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.

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

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

$$100 \times \frac{X_{hM}}{X_{hM} + X_{hD}} = \alpha + \beta \log(\text{Income}_h) + [\zeta' K_h] + \varepsilon_h \quad (\text{A23})$$

where X_{hM} and X_{hD} are household h 's expenditure on imports and domestic goods in 2014, $\log(\text{Income}_h)$ 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

Table A14: Household income and import shares in Homescan in 2014

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

Notes: Estimation of equation (A23). 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

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 A14 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.


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

$$100 \times \frac{X_{hMt}}{X_{hMt} + X_{hDt}} = \alpha + \text{FIE}_h + \sum_{y \neq 2014} \beta_t \mathbb{I}_{y=t} + \varepsilon_{ht} \quad (\text{A24})$$

where X_{hMt} and X_{hDt} are expenditures on imports and domestic goods for household h in year t , FIE_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 A3 displays our estimated year fixed effects, β_t , together with their associated 95% confidence intervals when estimating regression (A24) 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

Figure A3: Plotting β_t by time horizon from equation (A24)



Notes: Estimation of equation (A24) 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.

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.

Table 2 reports results from estimating the following household level regression

$$100 \times \frac{X_{hMt}}{X_{hMt} + X_{hDt}} = \mathbb{F}\mathbb{E}_t + \mathbb{F}\mathbb{E}_h + \sum_{y \neq 2014} \mathbb{I}_{y=t} \left[\beta_t \text{Inc}_h + [\zeta'_t K_h] \right] + \varepsilon_{ht} \quad (\text{A25})$$


where $\mathbb{F}\mathbb{E}_h$ and $\mathbb{F}\mathbb{E}_t$ 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.^{S48} 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

^{S48}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 A4: Price changes and household income



Notes: Estimation of (A27) 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 .

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 separately for domestic and imported goods,

$$\log p_{im} = \alpha + \mathbb{F}\mathbb{E}_i + \sum_{m' \neq \text{Dec 2014}} \mathbb{I}_{m'=m} \beta_m + \varepsilon_{im} \quad (\text{A26})$$

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

$$\log p_{ihq} = \alpha + \alpha_{ih} + \alpha_q + \sum_{y \neq 2014Q4} \mathbb{I}_{y=q} \beta_q \log(\text{Income}_h) + \varepsilon_{ihq} \quad (\text{A27})$$

where p_{ihq} is the level of the price of product i paid by household aggregation h (de-

Table A15: A lack of systematic price variation across space

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

Notes: Estimation of equation (A28). Observations weighted by expenditure in 2014. Clustered by product. The one-digit zip code containing Zurich (which is the most populous) is omitted.

financed as the 50 income quantiles) in quarter q .^{S49} We measure this price as the geometric weighted average product price across transactions within hq , weighing by expenditures in the current quarter. We weigh observations in (A27) 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 A4. 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

$$\log p_{ij} = \alpha + \text{FE}_i + \text{FE}_j + \varepsilon_{ij} \quad (\text{A28})$$

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

Table A15 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

^{S49}We aggregate up from months in (A26) to quarters in (A27) given the finer disaggregation across incomes in (A27).

Table A16: 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.

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 ratios 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 A16.

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 pe-

riod. 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

$$d \log b_{hit} = \text{FE}_{it} + \text{FE}_{hst}^M + \kappa_{hi} d \log \left(\frac{I_{ht}}{P_{ht}} \right) - \eta_s \log(I_{ht_0}) d \log p_{hit} + \iota_{hit}. \quad (\text{A29})$$

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: $\kappa_h^{MD} \equiv \kappa_{hM} - \kappa_{hD}$.^{S50} Given an estimate of κ_h^{MD} , we obtain η_s from a modified version of equation (19):

$$d \log \left(\frac{b_{hMt}}{b_{hDt}} \right) - \kappa_h^{MD} d \log \left(\frac{I_{ht}}{P_{ht}} \right) = \alpha - \eta_s \log(I_{ht_0}) d \log \left(\frac{p_{Mt}}{p_{Dt}} \right) + \iota_{ht} \quad (\text{A30})$$

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

$$\log \left(\frac{b_{hMt_0}}{b_{hDt_0}} \right) = \beta_0 + \beta_1 \log(I_{ht_0}) + \beta_2 \log(I_{ht_0})^2 + X'_h \gamma + \iota_{ht_0} \quad (\text{A31})$$

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 (A31) is

^{S50}In Approach 2, we would have to estimate thousands of income elasticities at the barcode product level.

Table A17: Robustness of Approach 1: Using cross-sectional Engel curves

	Linear			Quadratic		
	(1)	(2)	(3)	(4)	(5)	(6)
$\log(I_{ht0})d \log(p_{Mt}/p_{Dt})$	2.17*** [0.54]	2.16*** [0.54]	2.17*** [0.54]	2.09*** [0.54]	2.06*** [0.54]	2.08*** [0.54]
Observations	2901	2901	2901	2901	2901	2901
Control size		X	X		X	X
All controls			X			X

Notes: Results of estimating equation (A30) using estimates of κ_h^{MD} estimated using equation (A31) (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 (A30) is clustered and weighted as in the baseline of Approach 1. The regression (A31) 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 A18: Robustness of Approach 2: Varying baseline choices

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\log(I_{ht0}) \times d \log p_{it}$	1.93** [0.87]	1.80** [0.73]	1.87* [0.97]	2.03* [1.12]	1.61* [0.94]	1.89** [0.89]	1.87* [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 $d \log p_{it}$ as the log price change between 2015 and the fourth quarter of 2014. *p<.1; **p<.05; ***p<.01

very similar to regression (A23), where the dependent variable is $\frac{b_{hMt0}}{b_{hMt0} + b_{hDt0}}$ rather than $\log\left(\frac{b_{hMt0}}{b_{hDt0}}\right)$.

Table A17 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 the left-hand side of equation (A30) depends on a coefficient estimated in the first step of the procedure.

Varying baseline choices. Column 1 of Table A18 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

Table A19: Robustness of Approach 2: Varying baseline choices

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\log(I_{ht0}) \times d \log p_{it}$	1.93** [0.87]	3.55 [2.23]	2.62** [1.10]	2.27** [1.01]	1.97*** [0.62]	1.62* [0.94]	1.83** [0.85]	2.35* [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 $\log(I_{ht0})d \log p_{it}$. The unreported SW F stat on the household-size interaction is over 14. * $p < .1$; ** $p < .05$; *** $p < .01$

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 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.^{S51} Each of these choices has little effect on either first-stage or second-stage results.

Column 1 of Table A19 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.

^{S51}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.

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 income. 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 A20.

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 A20.

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 A21 and A22 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.

Table A20: Robustness of Approaches 1 and 2 to inferring household income

	Without using FORS data		Without using SFTA data	
	(1)	(2)	(3)	(4)
	Approach 1	Approach 2	Approach 1	Approach 2
$\log(I_{ht0})d \log(p_{Mt}/p_{Dt})$	2.14*** [0.54]		2.60*** [0.72]	
$\log(I_{ht0}) \times d \log p_{it}$		2.29*** [0.72]		2.27** [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 A21: Robustness of Approach 1 dropping household income ranges

	(1)	(2)	(3)	(4)
	Drop lowest	Drop 2 lowest	Drop 2 highest	Drop highest
$\log(I_{ht0})d \log(p_{Mt}/p_{Dt})$	1.58 [1.29]	0.72 [1.55]	2.52*** [0.42]	2.52*** [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

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 A23 replicates our baseline 2SLS result from column 3 of Table 5. In the remaining columns in Table A23 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 correspondingly 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.

Table A22: Robustness of Approach 2 dropping household income ranges

	(1)	(2)	(3)	(4)
	Drop lowest	Drop 2 lowest	Drop 2 highest	Drop highest
$\log(I_{ht_0}) \times d \log p_{it}$	2.60 [1.84]	2.29 [1.73]	1.55 [1.21]	1.69** [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 A23: Robustness of Approach 2: Incorporating spatial variation

	(1)	(2)	(3)
$\log(I_{ht_0}) \times d \log p_{it}$	1.930** [0.867]	2.170*** [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 (first 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

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 $h \in \{1, \dots, 50\}$ living in a particular one-digit zip code $j \in \{1, \dots, 9\}$, 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 v_{it} is zero between 2014 and 2015, then $\bar{\eta}_s$ is identified from the constant α as $\bar{\eta}_s = 1 - \alpha / (d \log (p_{Mt} / p_{Dt}))$. Given $d \log (p_{Mt} / p_{Dt}) = -0.0216$ 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 v_{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).^{S52}

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

$$f_h(u)^{\frac{\rho-1}{\rho}} = \sum_s (\zeta_{hst} u^{\gamma_s})^{\frac{1}{\rho}} (c_{hst})^{\frac{\rho-1}{\rho}}, \quad (\text{A32})$$

where

$$c_{hst} = \left(\sum_{i \in \mathcal{I}(s)} (\zeta_{hit} u^{\gamma_i})^{\frac{1}{\eta_s(u)}} (c_{hit})^{\frac{\eta_s(u)-1}{\eta_s(u)}} \right)^{\frac{\eta_s(u)}{\eta_s(u)-1}}, \quad (\text{A33})$$

^{S52}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.

$f_h(\cdot) > 0$ 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 $e(p_{ht}, u_{ht}; \zeta_{ht}) = I_{ht}$. We discuss below conditions that ensure that the expenditure function is monotonic in u .

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

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

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

$$d \log u_{ht} = \left(\frac{\partial \log e_h}{\partial \log u_h} \right)^{-1} \times \left(d \log I_{ht} - \sum_i b_{hit_0} d \log p_{hit} - \bar{\epsilon}_{ht} \right) \quad (\text{A34})$$

This is equation (13) in the text.

Deriving equation (17). Substituting equation (A34) into equation (12) yields

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

The previous expression and assumption (15) yield

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

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

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

The previous expression is equation (17) given the definitions $v_{hit} \equiv d \log \zeta_{hit} - \kappa_i \bar{\epsilon}_{ht}$ and $\tilde{\psi}_{hst} \equiv \psi_{hst} + \kappa_{hs} (d \log(I_{ht}/P_{ht}) - \bar{\epsilon}_{ht})$. The demand shifter v_{hit} combines the taste shifter for good i , $d \log \zeta_{hit}$, and the change in utility due to taste shifters, $\bar{\epsilon}_{ht}$ 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} ,

$$\eta_{hst} \equiv \tilde{\eta}_s + \tilde{\eta}_s \log(u_{ht}). \quad (\text{A35})$$

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

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

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

$$f_h(x) = a_0 x^{a_1} \left[\sum_s \zeta_{hst} x^{\gamma_s} (P_{hs}(x))^{1-\rho} \right]^{\frac{1}{\rho-1}} \quad (\text{A36})$$

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

$$P_{hs}(x) = \left(\sum_{i \in \mathcal{I}(s)} \zeta_{hit_0} x^{\gamma_i} (p_{hit_0})^{1-\eta_s(x)} \right)^{\frac{1}{1-\eta_s(x)}} \quad (\text{A37})$$

In this case, $e_h(\mathbf{p}_{ht_0}, u_{ht_0}; \zeta_{ht_0}) = I_{ht_0} = a_0 \times u_{ht_0}^{a_1}$ and $\partial \log e_h / \partial \log u_h = a_1$ when evaluated at t_0 . These cardinalization assumptions imply equation (16), where $\tilde{\eta}_s \equiv \tilde{\eta}_s - a_1^{-1} \tilde{\eta}_s \log(a_0)$ 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 (A33) 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 (A32)–(A33). 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 (A32) is

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

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 $G_i(u) \equiv f(u) (\zeta_i u^{\gamma_i})^{\frac{1}{1-\eta(u)}}$, and re-express the utility function as

$$1 = \sum_i (c_i / G_i(u))^{\frac{\eta(u)-1}{\eta(u)}} \quad (\text{A38})$$

the expenditure function as

$$e(p, u; \zeta) = \left(\sum_i (G_i(u) p_i)^{1-\eta(u)} \right)^{\frac{1}{1-\eta(u)}}, \quad (\text{A39})$$

and demand for good i as

$$\frac{p_i c_i}{I} = \left(\frac{G_i(u) p_i}{I} \right)^{1-\eta(u)} \quad (\text{A40})$$

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 (A40) with $\eta'(u) < 0$ to be integrable is

$$K(u) \equiv \sum_i \exp \left(\frac{(\eta(u) - 1)^2 G_i'(u)}{\eta'(u) G_i(u)} \right) < 1. \quad (\text{A41})$$

The proof of Proposition 4 in Fally (2022) shows that if condition (A41) is satisfied, then there is a unique solution u in (A38) and u in $e(p, u; \zeta) = I$, and that around each of those values of u the expenditure function is increasing in u .

We prove that (A41) is satisfied under our functional form assumption $\eta(u) = \bar{\eta} + \eta \log(u)$ with $\bar{\eta} \neq 1$, $\eta < 0$, and $f(u) = (u^{k_1})^{\frac{1}{1-\eta(u)}}$. In this case, $G_i(u) = (\zeta_i u^{\gamma_i})^{\frac{1}{1-\eta(u)}}$,

where $\tilde{\gamma}_i \equiv \gamma_i + k_1$.^{S53} Hence,

$$\frac{G'_i}{G_i} = \log(\zeta_i u^{\tilde{\gamma}_i}) + \frac{(1 - \eta(u))}{\eta'(u)} \frac{\tilde{\gamma}_i}{u}$$

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 $k_1 < -\max_i\{\gamma_i\}$ then condition (A41) 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 (A41) is satisfied. For any $\bar{\eta} \neq 1$, condition (A41) 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 (A36) 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 (A36) 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

^{S53} 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 (A41) holds in a neighborhood of any u for which $\eta(u) > 1$.

Table A24: Welfare-relevant grocery price changes: Additional results I

Annual income	2013–14 Heterogeneous elasticities			2014–15 Common exp. shares		
	1st-order	Switching	Exact	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 A25: 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.

in the monotonic region of the expenditure function.

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 A24 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 A25 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 A26: Import price declines

		Import price shock				
		-2.2	-10	-20	-40	-2.2
Annual income		$\sigma = 0$				$\sigma > 0$
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 η_s because the first-order and higher-order effects move in opposite directions.

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 A26 displays the results.

Second, in our baseline we choose $\eta_s = -2$. Table A27 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 A28 and A29 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

Table A27: Smaller differences in elasticities of substitution

		Import price shock					+2.2 $\sigma > 0$
		+2.2	+10	+20	+40	+1000	
Annual income		$\sigma = 0$					
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 η s							
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_{hst0} , is equal to 3.

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 A30 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 A31 and A32 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 A31 are very similar to those in our baseline.

Table A28: Elasticity of substitution of high-income group = 1.5

Annual income	Import price shock					+2.2 $\sigma > 0$
	+2.2	+10	+20 $\sigma = 0$	+40	+1000	
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 η s						
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_{hst0} , is equal to 1.5 rather than 3.

Table A29: Elasticity of substitution of high-income group = 5

Annual income	Import price shock					+2.2 $\sigma > 0$
	+2.2	+10	+20 $\sigma = 0$	+40	+1000	
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 η s						
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_{hst0} , is equal to 5 rather than 3.

Table A30: Elasticity of substitution across sectors = 0.2

Annual income	Import price shock					+2.2 $\sigma > 0$
	+2.2	+10	+20	+40	+1000	
	$\sigma = 0$					
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 η s						
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 A31: Homogeneous elasticities within the service sector

Annual income	Import price shock					+2.2 $\sigma > 0$
	+2.2	+10	+20	+40	+1000	
	$\sigma = 0$					
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 η s						
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 $s = \text{services}$ for all h).

Table A32: Homogeneous elasticities within the service and other goods sectors

Annual income	Import price shock					+2.2 $\sigma > 0$
	+2.2	+10	+20 $\sigma = 0$	+40	+1000	
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 η s						
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).