

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

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, Kathryn 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.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
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, 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. These differences in elasticities contribute significantly to the unequal welfare effects of large 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.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 show that they contribute significantly to the unequal welfare effects of large import price changes.

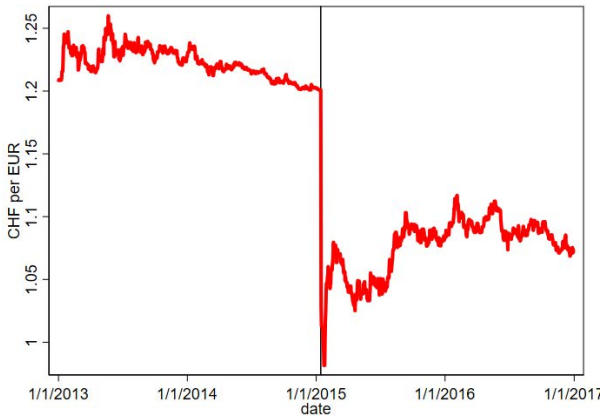
In Section 2, we begin by describing the setting in which we measure both initial expenditure shares on imported goods across households as well as unequal expenditure switching. We focus on Switzerland in a period surrounding the abrupt appreciation of the Swiss franc on January 15, 2015, which markedly 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, covering individual household purchases of food, beverages, personal care, and household supplies in supermarkets and drugstores. We merge these data with information on whether indi-

¹See, e.g., Friedman and Levinsohn (2002), Porto (2006), Fajgelbaum and Khandelwal (2016), Cravino and Levchenko (2017), Carroll and Hur (2020), and Borusyak and Jaravel (2021).

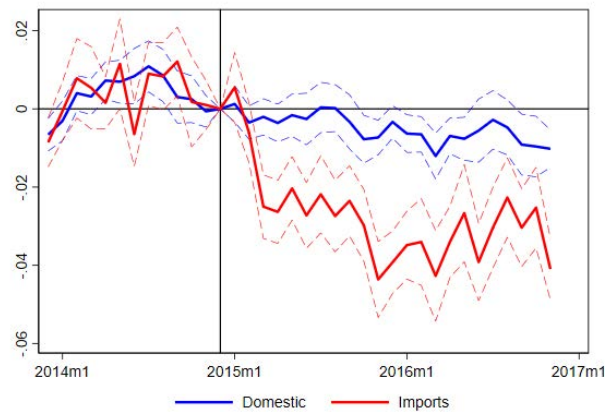
²The Swiss National Bank (SNB) had previously enforced a minimum exchange rate of 1.20 CHF per EUR. Developments abroad in late 2014 and early 2015 prompted the SNB to unexpectedly abandon this policy on January 15, 2015. The subsequent appreciation episode came after a period of remarkable exchange 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.

Figure 1: Swiss franc appreciation and resulting price changes

January 2015 Swiss franc appreciation
(relative to the Euro)



Price changes for imported and
domestically produced purchases



Notes: The left panel displays the CHF per EUR exchange rate with a vertical line at January 15, 2015; [Swiss National Bank \(2016\)](#). The right panel displays price differences relative to December 2014 separately for imports and domestically produced purchases (and associated 95% confidence intervals) from estimating equation (23); 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.

vidual barcodes 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.

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 (more specifically, 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 [Baqae and Burstein \(2021\)](#)—to provide an exact answer to this question for non-marginal price changes, taking expenditure switching into account. A household’s compensating variation in response to given income and price changes can be constructed given 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

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).

⁴For recent, alternative approaches for ex-post welfare measurement, see [Atkin et al. \(2020\)](#) and [Jaravel and Lashkari \(2021\)](#). Our approach can be used to calculate changes in welfare in response to factual or counterfactual changes in prices and income; see Section 5.

price changes on the cost of living are, therefore, shaped by differences in initial expenditure shares and differences in compensated cross-price elasticities. 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 \(Forthcoming\)](#) 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 take two approaches to estimate differences across incomes in elasticities of substitution between goods in the Homescan data, leveraging 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 (controlling for income effects). 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, in addition to income effects, 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 is denominated in EUR, using information from the goods-level survey underlying the calculation of the official Swiss import price index. Because of 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 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.

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 two. These approaches identify differences in elasticities across incomes. In our analytic results and our counterfactuals we show that the unequal welfare effects of foreign price changes are not very sensitive to the choice of the levels of these elasticities.

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 \(2007\)](#).⁵ Relative to the international trade/spatial economics literatures that estimate differences in price elasticities across income groups using either Hausman instruments or the approach developed by [Feenstra \(1994\)](#) and [Broda and Weinstein \(2006\)](#)—see, e.g., [Argente and Lee \(2021\)](#), [Handbury \(2021\)](#), [Faber and Fally \(Forthcoming\)](#)—we exploit exogenous variation in price responsiveness to an exchange rate shock and estimate larger differences in price elasticities across incomes. When we use a Hausman instrument, we find small differences as in the literature. 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).

Finally, in Section [5](#) we quantify the distributional welfare implications of factual and counterfactual changes in import prices using differences across households in initial expenditure shares and in elasticities of substitution.⁶ Given our focus on the expenditure-side effects of foreign price shocks, in our counterfactuals we assume that prices of all domestic goods move proportionately to domestic income (as in single factor trade models without imported intermediate inputs) and consider uniform changes in import prices relative to domestic prices across all sectors.⁷

Import price increases in Switzerland harm higher-income households more than lower income households for two reasons. First, higher income households have higher initial import shares (since they spend relatively more on non-grocery goods, which are more tradable than groceries).⁸ Second, they substitute away from imported goods less. For

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. In contrast, we focus on how cross-sectional variation in household income generates differential expenditure switching and welfare effects.

⁶[Bai 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.

⁷There 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. \(Forthcoming\)](#), and [Adao et al. \(2020\)](#). See, e.g., [He \(2018\)](#) and [Borusyak and Jaravel \(2021\)](#) for papers incorporating both income and expenditure-side inequality induced by trade.

⁸Variation 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).

large changes in prices, the impact of unequal expenditure-switching on welfare is substantial. For example, for a 20% increase in import prices relative to domestic prices and income, welfare of a household with income of 15,000 CHF falls by about one third less than for a household with income of 60,000 CHF. Almost half of this differential change in welfare is accounted for by differences in price elasticities. In response to a move to autarky—the focus of the quantitative trade literature—unequal expenditure switching accounts for roughly 80% of the differences in welfare across the income distribution. Following import price *reductions*, higher income households benefit more because they initially consume a higher share of imported goods than lower income households, but benefit less since they engage in less expenditure switching. For large price changes, unequal expenditure switching dominates.

2 Data and stylized facts

2.1 Data

In this section we provide an overview of the main datasets employed in the paper. Details and additional data sources are described in Appendix A.

AC Nielsen Homescan data and import status. AC Nielsen Homescan, [Nielsen Switzerland \(2016\)](#), includes information on household 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 bar code (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.

The Homescan data comes with a rich set of socioeconomic characteristics for each household, summarized in Table 12 in Appendix A for the year 2014, including the 2-digit zip code of residence, the education 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 assign each household an income level in 2014 using the relationship between household characteristics and the level of pre-tax household income estimated in the Swiss Household Panel (FORS), which includes these characteristics in addition to the level of household pre-tax annual income.

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. \(2021\)](#) collect from [codecheck.info](#). Coverage is not complete and notably excludes most fruits and vegetables, certain in-store EANs, 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, obtained from [Codecheck.info](#) between October 2015 and March 2016.⁹ We drop products for which import status is unknown.

Comparing columns 1 and 2 of Table 11 in Appendix A, we see that out of 69,088 unique goods and approximately CHF 11.1 million expenditures, there are 8,409 unique goods purchased and approximately CHF 4.2 million expenditures with known import status; the share of expenditures for which the production location is known is approximately 38%.^{10,11} 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 are purchased and the import share (at retail prices) of expenditures is 26.9%.

Expenditure and import shares by income group and sector (SFSO). 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

⁹Roughly 90% of imported goods in our data are from the European Union (EU). Our results are robust to dropping imports from other origins.

¹⁰Many of the products for which we don't observe import status appear in the Homescan data for only a short period of time. If we keep only those products that are 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.

¹¹One 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 either with the household's share of expenditure in 2014 on products for which we don't observe import status (Table 13) or with changes over time in this share (Table 14).

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

Annual income	Expenditure shares			Import shares			All
	Grocery	Other goods	Services	Grocery	Other goods	Services	
- 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 HBS using 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.

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 varies 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 income group, by aggregate sector, and by income group \times aggregate sector.¹²

Currency of invoicing of import prices at the border (SFSO). Our instrument in Section 4

¹²The import share on food and non-alcoholic beverages from the SFSO data (30.6%) is close to that in our Homescan data (26.9% in 2014) which is dominated by these product categories. The import share in the SFSO data rises to 37.9% if we include alcoholic beverages, tobacco, and personal care items. In our welfare calculations in Section 5, we use the more comprehensive SFSO import shares.

exploits variation across imported goods in 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 is denominated in EUR (out of those denominated in either EUR or CHF), using information from the good-level survey underlying the calculation of the official Swiss import price index. For additional information on these data, see [Auer et al. \(2021\)](#).

2.2 Stylized facts

SF 1 (SFSO): Initial import shares are 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. Higher income households have higher aggregate import shares in Switzerland, with the share rising monotonically from 21% to 28% between the bottom and top income groups in the SFSO data. This pattern is accounted for by two relationships. First, the import share is much higher for non-grocery goods (across all income groups) than for groceries or services, and higher income households spend a higher share of their income in this sector. Second, higher income households have a higher import share within the non-grocery goods aggregate sector and, to a lesser extent, within services.¹³

On the other hand, import shares within groceries are 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 16 in Appendix B.1.

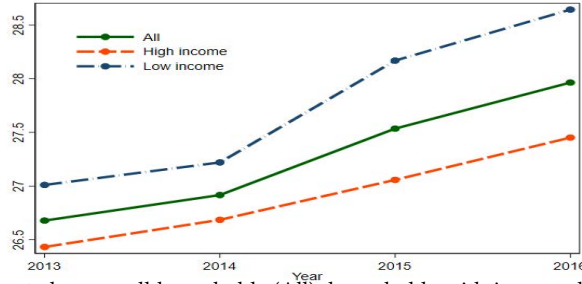
SF 2 (Homescan): The import share increased following the 2015 CHF appreciation.

The aggregate import share in the Homescan data increased from 26.9% to 27.5% between 2014 and 2015, as shown in Figure 2.¹⁴ 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

¹³This is largely because high income groups have higher budget shares on cars and cars in Switzerland tend to be imported.

¹⁴We document this and all remaining stylized facts using Homescan rather than SFSO data because the SFSO data is not available at annual frequency.

Figure 2: Aggregate import responsiveness and heterogeneity across incomes



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).

than it was in 2014 and continues increasing into 2016; see Figure 3 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.

Figure 2 displays the import share by year separately for high and low income households (those with incomes above and below our sample median of approximately 57,600 CHF). 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.

To show that this occurred within individual households—rather than from a change in the composition of expenditures across households—we complement this visual evidence by identifying within households the differential change in the import share of expenditure between each year and 2014 across household incomes. Table 2 displays our results. Import shares increased substantially less for high than low income households after the exchange rate appreciation using two measures of household income: the log of household income or an indicator for high income (defined as income above the median in our sample). For example, a household with an income three times than that of another experienced a roughly 0.7 percentage point smaller 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 fell between 2013 and 2014 before rising. These results are robust to including progressively more household-level controls interacted with year.

Table 2: Heterogeneous expenditure switching towards imports

	log(Income)			High income		
	(1)	(2)	(3)	(4)	(5)	(6)
Income 2013	-0.472* [0.266]	-0.495* [0.271]	-0.547** [0.257]	-0.213 [0.288]	-0.228 [0.299]	-0.259 [0.284]
Income 2015	-0.727** [0.272]	-0.800*** [0.290]	-0.622** [0.284]	-0.698** [0.293]	-0.782** [0.310]	-0.555* [0.293]
Income 2016	-0.953*** [0.321]	-1.069*** [0.332]	-1.006*** [0.352]	-0.383 [0.382]	-0.482 [0.417]	-0.346 [0.422]
Observations	11630	11630	11630	11630	11630	11630
Controls		X	X		X	X
Control education			X			X

Notes: Results of estimating $100 \times X_{it}^m / (X_{it}^m + X_{it}^d) = \text{FE}_t + \text{FE}_h + \sum_{y \neq 2014} \mathbb{1}_{y=t} [\beta_y \text{Inc}_h + \zeta'_y K_h] + \varepsilon_{ht}$, where FE_t and FE_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, the controls that are interacted with year are: log household size, an indicator for whether there is a child under 10, and an indicator if everyone in the HH is older than 70. In columns 3 and 6 we additionally control for an indicator for whether the HH's main earner has completed college or university 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. * $p < .1$; ** $p < .05$; *** $p < .01$

SF 4 (Homescan): Relative import prices fell following the 2015 CHF appreciation. Neither import nor domestic price changes vary 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 2.1% 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 4 in Appendix B.1. 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.

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 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 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, 2021](#))

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

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.

These results imply 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 $b_{hi}^{CV}(\mathbf{p}_h)$.¹⁵

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 given price changes we only need to know compensated budget shares as a function of prices, $b_{hi}^{CV}(\mathbf{p}_{ht})$. Along the path from t_0 to t_1 these budget shares can be constructed from initial budget shares, $b_{hi}^{CV}(\mathbf{p}_{ht_0})$, price changes, 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.¹⁶ However, in *estimating* these cross-price elasticities, income effects and taste shifters 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.¹⁷ We specify particular preferences to estimate cross-price elasticities along the initial indifference curve using the price changes induced by the 2015 CHF appreciation.

¹⁵That is, $b_{hi}^{CV}(\mathbf{p})$ corresponds to the budget shares of a fictional consumer with homothetic preferences represented by the expenditure function $e_h^{CV}(\mathbf{p}, u) = e_h(\mathbf{p}, u_{ht_0}; \zeta_{t_0})u$.

¹⁶If 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).

¹⁷In our application we only consider goods included in the Swiss CPI. Hence, our quantitative results apply to the market-good component of welfare and do not account for the welfare implications from changes in, e.g., leisure and home production.

3.2 Generalized non-homothetic CES preferences

In what follows, we restrict the general setup of Section 3.1 by imposing generalized, non-homothetic, nested CES preferences and present all derivations in Appendix C. 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

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

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

where $f_h(\cdot) > 0$.¹⁸ By Shephard's Lemma, the budget share of any good $i \in \mathcal{I}(s)$ is

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

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

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

As described in detail below, in mapping our model to the data we consider three aggregate sectors 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. In our counterfactuals, we consider uniform import price changes relative to domestic prices and income. In both approaches, $\eta_s(u)$ controls the extent of expenditure switching between domestic and imported goods.

¹⁸Here we assume that $\eta_s(u) \neq 1$ and $\sigma \neq 1$. The expenditure function must be re-written slightly when either of these elasticities equals one. These preferences reduce to nested homothetic CES if, for example, $\eta_s(u)$ is independent of u , $\zeta_i = \zeta_s = 0$ for all i and s , and $f'_h(u) > 0$.

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} . Using exact hat algebra, we can express compensated budget shares as a simple function of initial expenditure shares, initial elasticities of substitution, and changes in prices:

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

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

where hats indicate prices relative to their level in the initial period, $\hat{p}_{hi} \equiv p_{hi}/p_{hit_0}$; σ 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

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

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

We use (9) to construct changes in welfare in response to factual and counterfactual uniform import price changes (i.e. common across imported products and households) relative to domestic prices and incomes. Given these import price changes, constructing CV_h for household h requires \hat{P}_{hs} in each sector, the value of σ , and expenditure shares across sectors, b_{hst_0} . Constructing \hat{P}_{hs} in equation (8) for household h requires its import share within sector s in the initial period t_0 and elasticity of substitution in the initial period t_0 , η_{hst_0} .

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

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

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 \(2021\)](#). 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, volatility in prices benefits more (or hurts less if $\eta_{hst_0} < 1$) households that are more price sensitive.

Consider two households with common expenditure shares in the initial period. The difference in the expenditure-switching effect between these households is

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

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 4.¹⁹

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 \(Forthcoming\)](#).²⁰ 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

¹⁹Allowing for differences in initial budget shares, 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.

²⁰In Appendix C we build on [Fally \(Forthcoming\)](#) and prove that a single-sector version of these preferences can be rationalized given that the elasticities of substitution is declining in the indifference curve index.

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. 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 elasticities.

The other strong restriction we impose is that for any household there is a single elasticity, σ 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.²¹

4 Estimating the relationship between 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

$$d \log b_{hit} = d \log \zeta_{hit} + \left(\gamma_i - \frac{\partial \eta_s}{\partial u_h} u_{ht_0} \log p_{hit_0} \right) d \log u_{ht} + (1 - \eta_{hst_0}) d \log p_{hit} + \psi_{hst} \quad (11)$$

²¹In our counterfactual analysis we can introduce separate inner nests for domestic goods and imported goods. The results are invariant to the value of the elasticity of substitution within each nest since we only consider uniform price changes of imported relative to domestic goods.

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

$$d \log u_{ht} = \left(\frac{\partial \log e_h}{\partial \log u_h} \right)^{-1} \times \left(d \log \frac{I_{ht}}{P_{ht}} - \bar{\varepsilon}_{ht} \right) \quad (12)$$

where $d \log(I_{ht}/P_{ht})$ is the change in real income for household h ; $d \log P_{ht} \equiv \sum_i b_{hit_0} d \log p_{hit}$ is the household-specific weighted average of price changes across goods in all sectors; and $\bar{\varepsilon}_{ht} \equiv \sum_i (\partial \log e_h / \partial \zeta_{hi}) d \zeta_{hit}$; see Appendix C for derivations. Substituting (12) into (11) yields

$$\begin{aligned} d \log b_{hit} &= \left(\frac{\partial \log e_h}{\partial \log u_h} \right)^{-1} \times \left(\gamma_i - \frac{\partial \eta_s}{\partial u_h} u_{ht_0} \log p_{hit_0} \right) \left(d \log \frac{I_{ht}}{P_{ht}} - \bar{\varepsilon}_{ht} \right) \\ &\quad + d \log \zeta_{hit} + (1 - \eta_{hst_0}) d \log p_{hit} + \psi_{hst} \end{aligned} \quad (13)$$

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 microfounds these two restrictions). First, we assume that household h 's income elasticity for good i in sector s in the initial period (the term multiplying the change in real income in expression 13) can be expressed as the sum of a good-specific and a household-sector specific component,

$$\kappa_i + \kappa_{hs} \equiv \left(\frac{\partial \log e_h}{\partial \log u_h} \right)^{-1} \times \left(\gamma_i - \frac{\partial \eta_s}{\partial u_h} u_{ht_0} \log p_{hit_0} \right) \quad (14)$$

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

$$\eta_{hst_0} \equiv \bar{\eta}_s + \eta_s \log I_{ht_0}. \quad (15)$$

If $\eta_s < 0$, then a higher income household is less price sensitive in sector s . Assumptions (14) and (15), which are potentially testable, 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.

Under these two additional restrictions, equation (13) can be expressed as

$$d \log b_{hit} = v_{hit} + \kappa_i d \log \left(\frac{I_{ht}}{P_{ht}} \right) + [1 - \bar{\eta}_s - \eta_s \log(I_{ht_0})] d \log p_{hit} + \tilde{\psi}_{hst}. \quad (16)$$

The first term, v_{hit} , corresponds to household h 's demand shifter for good i . The second

term captures the interaction between the good-specific component of good i 's income elasticity and the change in real income for household h . 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.

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, $v_{hit} \equiv v_{it} + \widetilde{\text{FIE}}_{hst} \mathbb{I}_i^M + \tilde{v}_{hit}$, where \mathbb{I}_i^M is an indicator variable that equals one if good i is imported. This yields our baseline estimating equation

$$d \log b_{hit} = \text{FIE}_{it} + \text{FIE}_{hst}^M + \kappa_i d \log \left(\frac{I_{hit}}{P_{hit}} \right) - \eta_s \log(I_{ht_0}) d \log p_{hit} + \iota_{hit}. \quad (17)$$

In equation (17), the product-fixed effect FIE_{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 $\text{FIE}_{hst}^M \equiv \psi_{hst} + \widetilde{\text{FIE}}_{hst} \mathbb{I}_i^M$ 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 (17) 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 (17) 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 income-specific import demand shocks. In both approaches, we estimate equation (17) taking differences between 2014 and 2015. Even though these two approaches leverage entirely distinct variation to identify η_s , we find remarkably similar results.

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 $i = D$.²² In this case, whereas we can control for an aggregate import demand shock (contained in $\mathbb{F}\mathbb{E}_{it}$ in equation 17), 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

$$d \log \left(\frac{b_{hMt}}{b_{hDt}} \right) = \alpha + \kappa d \log \left(\frac{I_{ht}}{P_{ht}} \right) - \eta_s \log(I_{ht_0}) d \log \left(\frac{p_{Mt}}{p_{Dt}} \right) + \iota_{ht} \quad (18)$$

We have assumed that the relative change in import prices is common across households within groceries to link this approach to Stylized fact 3. Here, α , κ , and ι_{ht} all represent differences of the parameters in equation (17) across imported and domestic goods: $\alpha \equiv \mathbb{F}\mathbb{E}_{Mt} - \mathbb{F}\mathbb{E}_{Dt}$, $\kappa \equiv \kappa_M - \kappa_D$, and $\iota_{ht} \equiv \iota_{hMt} - \iota_{hDt}$.

We measure b_{hit} as the expenditure share 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 (the national price of a product is the average of log prices across all transactions, weighing transactions by expenditures) weighted by expenditures per product across all consumers in 2014, separately for imports and domestic goods. 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 (18) using OLS is that household-specific import demand shock deviations from the aggregate import demand shock are uncorrelated with household income between 2014 and 2015.

²²To map this 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, we use a first-order approximation of the expenditure function of any homothetic aggregator within each import status.

Table 3: Estimation of η_s in Approach 1 using equation (18)

	(1)	(2)	(3)	(4)	(5)	(6)
$\log(I_{ht_0})d \log(p_{Mt}/p_{Dt})$	2.189*** [0.554]	2.207*** [0.567]	1.838*** [0.478]	1.981*** [0.618]	2.041*** [0.672]	2.172*** [0.540]
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]
Observations	2901	2901	2901	2901	2901	2901
Baseline	X					
No winsorizing		X				
Winsorize 5%			X			
Unweighted				X		
Expenditure weight					X	
No income effects						X

Notes: The estimating equation is (18). 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 2-6 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, and in column 6 we omit income effects. *p<.1; **p<.05; ***p<.01

Results. Whereas we estimate regression (18) 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 with similar incomes. 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 3 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 ($\approx 2.189 \log 3$) lower than a household with income of 20,000 CHF. This substantial difference 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.²³ Given the granularity of this definition of a product i relative to our first approach, the household-product level data is sparse. Hence, to estimate equation (17) we aggregate product-level data across

²³Our baseline sample is the set of products that are purchased at least once per month (nationally) in the year-and-a-half before and after the CHF appreciation.

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 households level as described above.

In our baseline we use a common price change across households at the product level, $d \log p_{hit} = d \log p_{it}$. We measure the logarithm of the national product price as above: a year-specific average of transaction-specific log prices, weighing transactions by expenditures. 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. Moreover, we address the additional simultaneity concern—that household-specific demand shock deviations in the error term are correlated with observed product price changes interacted with household income—by constructing an instrument using an interaction between a product-specific cost shifter and initial household income.²⁴ Our cost shifter exploits variation across imported goods in 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 in 2014 that is denominated in EUR (out of those denominated in either EUR or CHF), which we denote by $share_{it_0}$. Because of 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 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 each border group that is denominated in EUR, $share_{it_0}$, (ii) an import indicator variable, \mathbb{I}_i^M , and (iii) the logarithm of initial household income,

²⁴The standard simultaneity concern—that taste shocks are correlated with price changes—does not apply in our baseline since the common component of price changes is absorbed by the product-time fixed effect; hence, any national demand shock is explicitly controlled for. The same applies in our robustness in which we use location-specific price changes.

Table 4: Estimation of η_s in Approach 2 using equation (17)

	(1)	(2)	(3)
	OLS	RF	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 (17). 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 intersection between import status and the share of imported goods in each border group that is denominated in EUR; 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). * $p < .1$; ** $p < .05$; *** $p < .01$

$\log(I_{ht_0})$.²⁵

Our exclusion restriction is that the pre-determined invoicing share triple interaction is not systematically correlated with the household’s product-specific demand shock, conditional on the average product-specific demand shock across households, income effects, household-time effects, and (if the product is imported) the average import demand shock for households in the same income aggregation. We provide a range of evidence consistent with this exclusion restriction in Section 4.4.

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 is 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.

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 .²⁶ In constructing changes in log expenditure shares, we winsorize changes in log expenditures at the first percentile (both in the right and left tails).

²⁵If 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.

²⁶This 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).

Finally, while our instrument varies at the level of the triple interaction between (i) the share of imported goods in each border group that is denominated in EUR, (ii) an import indicator variable, and (iii) the logarithm of initial household income, we cluster more conservatively: robust standard errors are two-way clustered at the level of household income (there are 50 such clusters) and, separately, the intersection between import status and the share (i) (there are 54 such clusters).²⁷ We revisit each of these choices in robustness.

Table 4 displays our baseline results, focusing on the parameter of interest: η_s . The first column reports results from estimating equation (17) 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 increase their expenditures by less on imported goods within border groups with a higher share of EUR-invoiced products (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.

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 and the associated F statistic is 13.1.²⁸ 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.

4.4 Robustness

Table 3 displays robustness across a range of choices in approach 1. 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. Results are robust to these choices. Finally, in column 6 we omit income effects and instead estimate variation in the uncompensated price elasticity. The uncompensated elasticity is very similar to the compensated one.

The majority of our robustness exercises focus on our second approach. The first set of exercises lend support to our causal interpretation of our estimate η_s . The second set of exercises varies specific baseline choices and shows that our baseline point estimate is

²⁷In our baseline, we cluster standard errors conservatively given that our instrument varies at the level of the triple interaction between import status, the share of products denominated in EUR in the border group, and household income. If we cluster at this level, the first-stage F statistic is well over 100.

²⁸Throughout, we report the Kleibergen-Paap Wald rk F statistic when there is only one endogenous variable.

Table 5: Robustness I supporting causal interpretation

	(1)	(2)	(3)	(4)	(5)	(6)
$\log(I_{ht_0}) \times share_{it_0} \times \Pi_i^M$	-0.140** [0.068]	-0.004 [0.066]	-0.104 [0.082]			
$\log(I_{ht_0}) \times d \log p_{it}$				1.930** [0.867]	2.191** [0.871]	2.147** [0.870]
Observations	95,325	98,652	78,800	95,325	95,325	95,325
Baseline	X			X		
Outcome period		13-14	15-16			
Additional controls I					X	
Additional controls II						X
K-P F Stat (fist stage)				13.1	15.6	17.1

Notes: Columns 1 replicates our baseline RF specification shown in column 2 of Table 4. 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). Columns 4 replicates our baseline 2SLS specification shown in column 3 of Table 4. Column 5 incorporates two additional controls interacted with year: the 2014 import share as well as the 2014 expenditure share on each border group. Column 6 additionally incorporates on more control interacted with year: the 2014 average price of each individual product. *p<.1; **p<.05; ***p<.01

robust. The third set of exercises demonstrate that the variation identifying η_s is arising within the set of imported products. The fourth and final set of exercises demonstrate that our results are robust to incorporating spatial variation in both expenditures and prices.

Robustness I: Supporting our causal interpretation of our estimate η_s . 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. This raises the possibility that differential changes in import patterns across incomes are driven in part by forces other than differences in elasticities of substitution, since relative import prices were more stable in 2013-14 and 2015-16.

Here, we begin by showing that there are no such pre- or post-trends in approach 2. Whereas our instrument predicts differential changes in import patterns across incomes between 2014-2015, it does not predict differential changes in import patterns in the absence of large changes in the exchange rate between 2013-14 or 2015-16, when the share of border prices denominated in EUR has no systematic impact on relative import price changes.

Specifically, column 1 of Table 5 replicates our baseline reduced-form specification. Column 2 of Table 5 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, as expected given the economics underlying the instrument. 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

Table 6: Robustness IIA varying baseline choices

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\log(I_{ht_0}) \times d \log p_{it}$	1.930** [0.867]	1.803** [0.729]	1.866* [0.969]	2.026* [1.121]	1.615* [0.942]	1.889** [0.889]	1.869* [0.994]
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 (fist 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 4. 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. * $p < .1$; ** $p < .05$; *** $p < .01$

three orders of magnitude smaller and statistically insignificant; see Column 2 of Table 5. Column 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.

Another concern is that the share of imported goods in each border group that is 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 each border group that is denominated in EUR with other border group or product characteristics (the 2014 import share of each border group, the 2014 expenditure share on each border group, and the 2014 average price of each individual product) does not substantially change our results. Column 4 of Table 5 replicates our baseline 2SLS estimate from Column 3 of Table 4. Columns 5 and 6 of Table 5 show that including these additional controls has little effect on results.

Robustness II: Varying baseline choices. Column 1 of Table 6 displays our baseline 2SLS estimate and the remaining columns display results from various robustness exercises. In our baseline we winsorize changes in log expenditures at the first percentile (in the top and bottom tails). In Columns 2 and 3 we instead winsorize at the 5th percentile and not at all. Our baseline sample only includes products if they were purchased at least once per month in the year-and-a-half before and after the CHF appreciation. In Column 4 we drop this sample restriction. Our baseline sample only includes products in border groups for

Table 7: Robustness IIB varying baseline choices

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\log(I_{ht_0}) \times d \log p_{it}$	1.930** [0.867]	3.551 [2.229]	2.623** [1.103]	2.269** [1.014]	1.971*** [0.622]	1.617* [0.938]	1.831** [0.846]
Observations	95,325	43,559	67,179	82,995	116,930	95,325	95,325
Baseline	X						
Horizon 3m		X					
Horizon 6m			X				
Horizon 9m				X			
Percent change					X		
No income effects						X	
All inv. currencies							X
K-P F Stat (first stage)	13.1	7.6	8.6	11.4	12.8	12.8	12.7

Notes: Column 1 replicates our baseline 2SLS estimate of η_s in column 3 of Table 4. 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 income effects. Column 7 uses an alternative instrument using the share of non-CHF invoiced border prices, including all currencies. *p<.1; **p<.05; ***p<.01

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 2015.²⁹ Each of these choices has little effect on either first-stage or second-stage results.

Column 1 of Table 7 again displays our baseline 2SLS estimate and the remaining columns display results from additional robustness exercises. In columns 2, 3, and 4 we use price changes and expenditure changes measured over the first 3, 6, and 9 months of 2014 and 2015. In all cases, changes in real income are still measured over the full year given data availability. Results remain largely stable across these specifications; the elasticity difference is larger when estimated using changes in expenditures and prices over the first 3 months, but it is not precisely estimated.

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

²⁹[Auer et al. \(2021\)](#) first calculate for each product 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 contrast, our baseline price for each product is constructed as an expenditure-weighted average price across all transactions by year.

Table 8: Robustness III using imports only

	(1)	(2)	(3)	(4)
	RF	2SLS	RF	2SLS
$\log(I_{ht_0}) \times share_{it_0} \times \mathbb{I}_i^M$	-0.140** [0.068]		-0.141* [0.068]	
$\log(I_{ht_0}) \times d \log p_{it}$		1.930** [0.867]		1.933** [0.878]
Observations	95,325	95,325	27,128	27,128
Baseline	X	X		
Imports only			X	X
K-P F Stat (fist stage)		13.1		12.8

Notes: Columns 1 and 2 replicate our baseline RF and 2SLS estimates in columns 2 and 3 of Table 4. Columns 3 and 4 display estimates of the same regressions on a sample restricted to imported goods alone (so that $\mathbb{I}_i^M = 1$ for all observations). *p<.1; **p<.05; ***p<.01

This leaves our results largely unchanged. In our baseline, we control for income effects. If we omit income effects, our estimated difference in elasticities falls; see column 6. Our baseline instrument uses the share of imported goods in each border group that is 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.

Robustness III: Imports only. In our baseline, we include both imported and domestically produced goods in our estimation sample. However, the share of imported goods in each border group that is denominated in EUR impacts the average price change for imports substantially more than for domestically produced goods. This suggests that the complier group is mostly restricted to the set of imported goods. Here, we estimate (17) using only imported goods. Results remain similar.³⁰ Columns 1 and 2 of Table 8 replicate our baseline RF and 2SLS obtained on the baseline sample of imported and domestic products. Columns 3 and 4 display results when we restrict the sample to imports. Results are largely unchanged, although we only have 26 clusters in one dimension.

Robustness IV: Incorporating spatial variation. Finally, 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 9 replicates our baseline 2SLS result from column 3 of Table 4. In the remaining columns in Table 9 we

³⁰Restricting 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. And our instrument reduces to the interaction between the log of household income and the share of imported goods in each border group that is denominated in EUR.

Table 9: Robustness IV 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 4 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 intersection between import status and the share of imported goods that is denominated in EUR and, separately, the household aggregation (income quantile \times 1-digit zip code). * $p < .1$; ** $p < .05$; *** $p < .01$

disaggregate households both across 50 income quantiles (as before) and across each of 9 one-digit zip codes in Switzerland. 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.

In Appendix B.2 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, consistent with results in the literature. We also show that this Hausman instrument may be endogenous in our particular Swiss setting (where there is little price variation across space).

5 Counterfactuals

We use equation (9) in Section 3.2 to calculate the change in household welfare—measured using CV—caused by changes in import prices. We first describe how we assign values to the statistics required to compute (9) and then present our counterfactual results.

Assigning values to sufficient statistics. In our counterfactuals, we choose the initial period t_0 to be 2014. We consider three sectors s : groceries, non-grocery goods, and services. From the SFSO data, we obtain initial expenditure shares across sectors by income group

and import shares within sector by income group in all sectors; see Table 1 in Section 2.1. Aggregating up to three sectors has no effect on our measure of overall import shares by income; we construct these shares using data on expenditures by income group and import shares within highly disaggregated consumer categories.

We impose a value of σ very close to one: $\sigma = 0.99$. This choice is consistent with the fact that there is little variation in expenditure shares between these three sectors for each income group comparing the two periods for which we have these data (2012-2014 and 2015-2018).

In our baseline, we assume that initial elasticities of substitution by income group η_{hst_0} are common across sectors. We choose differences in within-sector elasticities of substitution across incomes in the initial period to match our Homescan-based grocery estimates in Section 4: in our baseline calibration we use a value of $\eta_s = -2$ consistent with our estimates in Approach 2.

Neither of the two approaches in Section 4 identify the intercept $\bar{\eta}_s$ defined in equation (15).³¹ As we showed analytically in Section 3.2, if initial budget shares are equal across households then, up to a second-order approximation, only the difference in elasticities matters for differential welfare changes. That is, only η_s matters, not $\bar{\eta}_s$. In our numerical results we do not impose common initial budget shares and do not use a second order approximation; yet we show that our results are not substantially affected by the choice of $\bar{\eta}_s$. In our baseline we choose $\bar{\eta}_s$ so that the lowest initial elasticity of substitution across the income groups that we consider (at 120,000 CHF) is equal to 3 (and vary this in sensitivity).

Baseline results. Given our focus on the expenditure-side effects of foreign price shocks, we assume that prices of all domestic goods move proportionately to domestic income, as in single factor trade models without imported intermediate inputs.³² We consider uniform increases in import prices relative to domestic prices across all sectors ranging between 2.2% (the size of the *reduction* in 2015) and 1000% (essentially a movement to autarky). For these counterfactuals, we do not need to take a stand on the extent to which the change in relative price is driven by changes in the price of continuing goods or by changes in the number of consumed varieties within imports and within domestic consumption.

³¹See Appendix B.3 for a discussion of alternative versions of the two approaches in Section 4 that identify $\bar{\eta}_s$ (under stronger assumptions).

³²In estimating compensated price elasticities, we use actual changes in retail prices by good so we do not impose this restriction. 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.

Table 10: Baseline compensating variation

Annual income	% change in import/domestic price				
	+2.2	+10	+20	+40	+1000
1: 20,000 elasticity 6.6	-0.5	-1.8	-3.2	-4.7	-5.6
2: 60,000 elasticity 4.4	-0.5	-2.2	-4.1	-7.0	-11.1
3: 120,000 elasticity 3.0	-0.6	-2.6	-5.0	-9.1	-22.0
<hr/>					
% difference btw CV					
income groups 2 and 1	16%	22%	31%	50%	99%
income groups 3 and 1	31%	41%	57%	95%	295%
<hr/>					
Contribution heterogeneous η s					
income groups 2 and 1	8%	28%	44%	62%	79%
income groups 3 and 1	7%	26%	41%	60%	86%

Notes: Percent changes are $100 \times$ the log of the relative price change. “% difference btw CV of 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.

The top panel of Table 10 summarizes the welfare implications of import price increases for households with three different annual income levels: 20,000, 60,000, and 120,000 CHF. These incomes are within the range of incomes in our sample and additionally have substantial differences in their observed import shares.³³ All CVs are negative, since all income groups are harmed by increases in import prices.

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 share in 2014 is 21%, 24%, and 27% for households with incomes of 20,000, 60,000, and 120,000 CHF, as displayed in Table 1. Second, they have lower initial elasticities of substitution, which shape the expenditure-switching effect displayed in equation (10). The elasticities are 6.6, 4.4, and 3.0 for these households.

The bottom two panels of Table 10 highlight the increasing importance of the expenditure-switching effect as the size of the change in import prices grows. First, in the second panel of Table 10 we provide the percent difference between the CV of the lowest income group and the CV of the middle and highest income groups. For a 2.2% increase in import prices, the welfare of the middle and high income groups fall by approximately 16% and 31%, respectively, more than the lowest income group. However, as the import price increase rises, the differences in welfare changes between income groups grow substan-

³³In practice, we use an income of 60,252 instead of 60,000 CHF for the middle group. This is the cutoff separating the first and second income brackets in Table 1.

tially. For example, for a 20% import price increase, the welfare of the middle and high income group falls by approximately 31% and 57%, respectively, more than for the lower income group.

To help quantify the importance of the expenditure-switching effect, we compare our baseline results to those from an alternative counterfactual. Specifically, we impose that all income groups have a common price elasticity equal to that of households with income of 20,000 CHF (which is 6.6). The bottom panel of Table 10 displays the contribution of heterogeneous elasticities (comparing heterogeneous elasticities and initial import shares to homogeneous import shares alone) in shaping differences in welfare changes for the middle and high income groups compared to the low income group.³⁴ 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 differences in elasticities to the unequal welfare changes across income groups. For a movement to autarky, the expenditure-switching effect accounts for the vast majority of the unequal welfare effects.

In response to import price declines, the first-order and expenditure-switching effects push welfare of higher- relative to lower-income households in opposite directions. High income households benefit more because they have higher initial import shares. On the other hand, low income households benefit more because they have higher price elasticities and, hence, substitute more towards imported goods. The first channel dominates for small import price declines. However, as import price changes become larger, the second channel dominates. Table 17 in Appendix D displays the results of import price changes between -2.2% and -40%. For small import price declines, the results in Table 17 look very similar to those in our baseline Table 10, except that welfare rises rather than falls. However, as import prices decline by more, the increase in welfare for lower income households catches up with, and even surpasses that for higher income households.

Sensitivity. We consider a range of robustness exercises in Appendix D. First, in our baseline we choose $\eta_s = -2$. Table 18 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 group equal to 3, which pins down $\bar{\eta}$. As expected, the importance of heterogeneous elasticities for shaping the unequal welfare implications

³⁴For income group j , this is simply $1 - (CV_j^{\text{homog}} - CV_1^{\text{homog}})/(CV_j - CV_1)$ 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.

of foreign prices is smaller.

Second, in our baseline we choose $\bar{\eta}_s$ so that the lowest initial elasticity of substitution (that for the highest income group with income of 120,000 CHF) is equal to 3. Tables 19 and 20 report results in which we use an elasticity of substitution for the highest income group equal to 1.5 and 5, respectively. Lower levels of price elasticities imply much larger welfare losses for every income group. However, except for the movement to autarky experiment, the percentage difference in CV between income groups and the contribution of heterogeneous elasticities are not very sensitive to the level of the elasticities keeping unchanged the elasticity difference between income groups.

Third, 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 21 and 22 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 21 are very similar to those in our baseline.

6 Conclusions

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

Theoretically, we show that constructing a household's compensating variation in response to given income and price changes generically requires initial expenditure shares across products and initial compensated cross-price elasticities. Using detailed Swiss data, leveraging the 2015 Swiss franc appreciation, and imposing nested, generalized, non-homothetic CES preferences, we document that lower income households have substantially higher elasticities of substitution. 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). We show that for large changes in prices (the focus of the quantitative trade literature), unequal expenditure switching generates substantial differences in welfare across the income dis-

tribution.

References

- ADAO, R., P. CARRILLO, A. COSTINOT, D. DONALDSON, AND D. POMERANZ (2020): "International Trade and Earnings Inequality: A New Factor Content Approach," NBER Working Papers 28263, National Bureau of Economic Research, Inc.
- 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.
- 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 (2020): "Measuring Welfare and Inequality with Incomplete Price Information," NBER Working Papers 26890, National Bureau of Economic Research, Inc.
- 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.
- BAQAEE, D. AND A. BURSTEIN (2021): "Welfare and Output with Income Effects and Taste Shocks," Tech. rep., National Bureau of Economic Research.
- 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. 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 (Forthcoming): "Firm Heterogeneity in Consumption Baskets: Evidence from Home and Store Scanner Data," *Review of Economic Studies*.
- FAJGELBAUM, P. D. AND A. K. KHANDELWAL (2016): "Measuring the Unequal Gains from Trade," *The Quarterly Journal of Economics*, 131, 1113–1180.
- FALLY, T. (Forthcoming): "Generalized Separability and Integrability: Consumer Demand with a Price Aggregator," *Journal of Economic Theory*.
- 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. RODRÍGUEZ-CLARE, AND M. YI (Forthcoming): "Slicing the Pie: Quantifying the Aggregate and Distributional Effects of Trade," *Review of Economic Studies*.
- HANDBURY, J. (2021): "Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living Across U.S. Cities," *Econometrica*, 89, 2679–2715.
- 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 (2021): "Nonparametric Measurement of Long-Run Growth in Consumer Welfare," Tech. rep., Working Paper.
- 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.
- 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 NATIONAL BANK (2016): "Datenportal der Schweizerischen Nationalbank," <https://data.snb.ch/de/topics>.

Table 11: Homescan data summary statistics in 2014

	All		Known origin	
	All	All	Imported	Domestic
Number of products	69,088	8,409	4,084	4,325
Expenditures	110.7	41.9	11.3	30.6
Transactions	234.6	110.4	27.7	82.7

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

A Data appendix

Processing the Homescan data. Households record if a purchase occurs within Switzerland or in 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-January 14, 2016 (January 15, 2015 - 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 one 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 one. 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., 274 in 2014 and 613 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

Table 12: 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 location). 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.

Table 13: Relationship between household income and expenditure share on products within known import status in 2014

	(1)	(2)	(3)
log(Income)	0.253	0.282	0.389
	[0.251]	[0.241]	[0.256]
Observations	3308	3308	3308
Controls		X	X
Control education			X

Notes: This Table replicates Table 16, replacing the dependent variable with the share of household expenditure on products with known import status. *p<.1; **p<.05; ***p<.01

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.

There exist 93 different two-digit zip codes in Switzerland, which uniquely identify cities such as Basel or Zurich, or, in rural areas, smaller regions such as such "Engadin and Val Müstair." The education groups identified in the Homescan data are defined as 1=obligatory school (9 years) "obligatorische Schule", 2=Vocational Education and Training "Berufsausbildung", 3=University entrance qualification "Matura", 4=College of Higher Education "hoehere Berufsausbildung", 5=College "hoehere Fachschule", 6=University "Hochschule / Universitaet", 7= other "andere Ausbildung."

We restrict our sample to households with positive expenditures inside Switzerland (i.e., we drop transactions that occur abroad) in 2014 on products with known import

Table 14: Changes in the relationship between household income and expenditure share on products within known import status

	log(Income)			High income		
	(1)	(2)	(3)	(4)	(5)	(6)
Income 2013	-0.044 [0.168]	-0.139 [0.173]	-0.137 [0.195]	0.000 [0.261]	-0.101 [0.261]	-0.093 [0.299]
Income 2015	0.289 [0.180]	0.321 [0.196]	0.378* [0.193]	0.263 [0.225]	0.301 [0.240]	0.361 [0.242]
Income 2016	0.171 [0.332]	0.226 [0.341]	0.204 [0.358]	-0.164 [0.348]	-0.118 [0.359]	-0.171 [0.372]
Observations	11646	11646	11646	11646	11646	11646
Controls		X	X		X	X
Control education			X			X

Notes: This table replicates Table 2 replacing the dependent variable with the share of household expenditure on products with known import status. * $p < .1$; ** $p < .05$; *** $p < .01$

status; this yields a sample of 3,302 households.³⁵

Household pre-tax income. The Homescan data includes a comprehensive set of household socioeconomic characteristics, as reported in Table 12. 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 a supplementary dataset, the Swiss Household Panel compiled by the Swiss Centre of Expertise in the Social Sciences (FORS). 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 to predict the level of household income for all households in the Homescan sample. We also predict the 2014-2015 change in a household’s income following a similar procedure (using the panel structure of FORS).³⁶

We use the following socioeconomic characteristics from the Homescan and FORS databases: an indicator variable for each of the seven 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 under 10, and the number of household members over 70.³⁷

³⁵We construct the first column of Table 11 including all households with positive expenditures in 2014 without restricting to those with positive expenditures on products with known import status.

³⁶When 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 the 2013-2016 average of income. We also remove outliers of income changes.

³⁷The FORS data provides information on the Canton of residence. Cantons are more aggregated geographies than 2-digit zip codes. However, in some instances 2-digit zip codes do not map uniquely to Cantons. Of the 76 2-digit zip codes in the Homescan data, 22 map into two Cantons and 7 map into three. In these

To concord the Homescan and the FORS data, 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 and household in case 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 dataset contains information on the socioeconomic characteristics of 6658 households interviewed during January, February, September, October, November, and December 2014.

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 labour earnings, asset income, private transfers, public transfers, and social security pensions, all net of taxes.³⁸ 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 difference in response rates, which we employ; see [Kuhn \(2018\)](#) for a description.

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 datasets provided by the SFSO. In these datasets, 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.³⁹ The HBS

cases, we allocate the respective Canton fixed effects to 2-digit zip codes weighing equally the respective fixed effects.

³⁸There 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 in case the individual responses and the household responses are very inconsistent. See [Kuhn \(2018\)](#) for further explanations.

³⁹See [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.

is collected by the SFSO via a rotating and non-overlapping survey, randomly sampled throughout the year from the SFSO's register of all Swiss households. Around 250 households participate each month and record consumption expenditures during the following month for 296 HBS consumption categories. The latter include both goods and services, in categories such as "Rice", "Pasta", or "Tickets for public transport." The survey also collects data on households' socioeconomic characteristics, including income. The SFSO publishes HBS category-specific expenditure shares averaged over a three-year horizon for each of five income groups. The expenditure share data we use in our analysis covers the years 2012-2014.⁴⁰

The second data source is the disaggregate 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 disaggregate 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 2013-2016. 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.⁴¹ The resulting concordance includes 187 consumption categories.⁴²

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.⁴³ Hence,

⁴⁰Due 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.

⁴¹When 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.

⁴²This concordance is available upon request.

⁴³In our calibration, we assume that SFSO import shares are constructed omitting CB purchases (as we do in

Table 15: 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 (19). Observations weighted by expenditure in 2014. Clustered by product. The 1-digit zip code containing Zurich (which is the most populous) is omitted.

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 identical across income groups within each of the 187 consumption categories).

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

Prices across regions within Switzerland. Average prices across individual products do not systematically vary much across regions within Switzerland. To document this fact, we estimate

$$\log p_{ij} = \alpha + \text{FIE}_i + \text{FIE}_j + \varepsilon_{ij} \quad (19)$$

where $\log p_{ij}$ is the weighted average log price for domestic purchases in 2014 of product i within 1-digit-zip code j , FIE_j is a 1-digit zip-code specific fixed effect, and FIE_i is a product-specific fixed effect. We weigh observations by expenditure in 2014 and cluster by product.

Table 15 displays our estimated 1-digit-zip code fixed effects. The omitted fixed effect is for the most populous 1-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.

B Empirical appendix

our calculations in the Homescan data).

Table 16: Household income and import shares in 2014

	(1)	(2)	(3)
log(Income)	-0.062 [0.457]	-0.115 [0.453]	-0.435 [0.499]
Observations	3302	3302	3302
Controls		X	X
Control education			X

Notes: Estimation of equation (20). Controls includes the log household size, an indicator for whether there is a child under 10, and an indicator if everyone in the HH is older than 70. Education is an indicator for whether the HH's main earner has completed college or university. 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

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 (SFSO): Initial import shares are higher among higher income households.

The facts on the SFSO data are displayed in Table 1. Within the Homescan data, we estimate

$$100 \times \frac{X_h^m}{X_h^m + X_h^d} = \alpha + \beta \log(\text{Income}_h) + [\zeta' K_h] + \varepsilon_h \quad (20)$$

where X_h^m and X_h^d 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 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 16 displays the results, which are insignificantly different from zero whether or not we control for additional household characteristics.

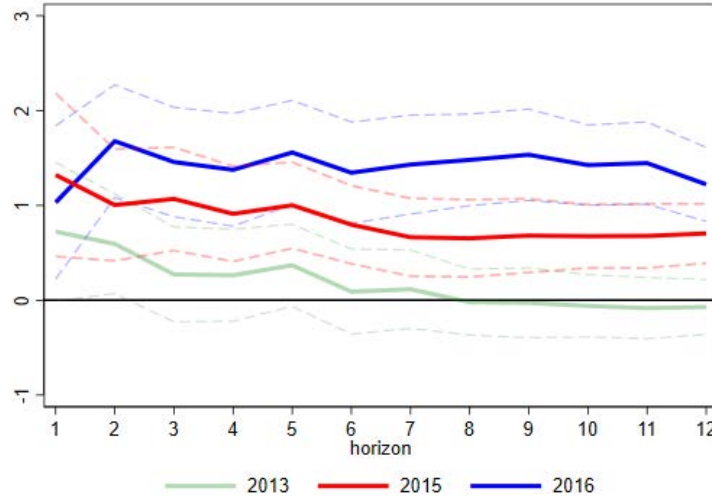
SF 2 (Homescan): The import share increased following the 2015 CHF appreciation.

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

$$100 \times \frac{X_{ht}^m}{X_{ht}^m + X_{ht}^d} = \alpha + \text{FIE}_h + \sum_{y \neq 2014} \beta_t \mathbb{I}_{y=t} + \varepsilon_{ht} \quad (21)$$

where X_{ht}^m and X_{ht}^d 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

Figure 3: Plotting β_t by time horizon from equation (21)



Notes: Estimation of equation (21) 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.

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 3 displays our estimated year fixed effects, β_t , together with their associated 95% confidence intervals when estimating regression (21) separately for each of twelve horizons, where we define horizon j as the first j months in year t and in 2014; our annual regressions are equivalent to horizon 12. Over the full year, there are no economically or statistically significant differences between 2013 and 2014. On the other hand, within households the import share was higher in 2015 than it was in 2014—the increase in the import share in 2015 is largely stable over all twelve horizons—and this persists through 2016.

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

Table 2 reports results from estimating the following household level regression

$$100 \times \frac{X_{ht}^m}{X_{ht}^m + X_{ht}^d} = \text{FE}_t + \text{FE}_h + \sum_{y \neq 2014} \mathbb{I}_{y=t} \left[\beta_t \text{Inc}_h + [\zeta'_t K_h] \right] + \varepsilon_{ht} \quad (22)$$

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 2014.⁴⁴ 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 domestic purchases.

SF 4 (Homescan): Relative import prices fell following the 2015 CHF appreciation. Neither import nor domestic price changes vary systematically with household income.

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

$$\log p_{im} = \alpha + \mathbb{F}\mathbb{E}_i + \sum_{m' \neq \text{Dec 2014}} \mathbb{I}_{m'=m} \beta_m + \varepsilon_{im} \quad (23)$$

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 good—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 move 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 household with different incomes? Separately on the sample of imported goods and domestic goods, we estimate

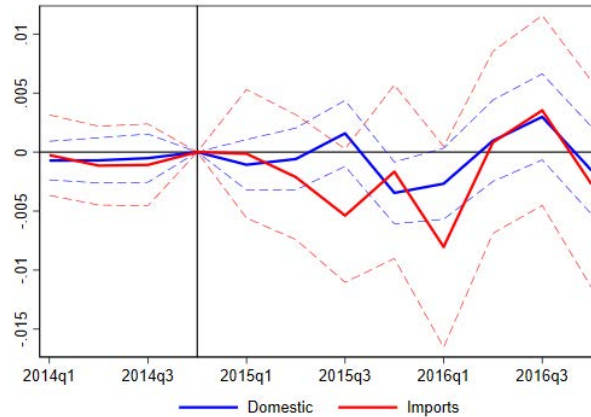
$$\log p_{ihq} = \alpha + \alpha_{ih} + \alpha_q + \sum_{y \neq 2014Q4} \mathbb{I}_{y=q} \beta_q \log(Inc_{eh}) + \varepsilon_{ihq} \quad (24)$$

where p_{ihq} is the level of the price of product i paid by household aggregation h (defined as the 50 income quantiles) in quarter q .⁴⁵ We measure this price as the geometric

⁴⁴The additional controls that are interacted with year are: log household size, an indicator for whether the HH's main earner has completed college or university, an indicator for whether there is a child under 10, and an indicator if everyone in the HH is older than 70.

⁴⁵We aggregate up from months in (23) to quarters in (24) given the finer disaggregation across incomes in (24).

Figure 4: Price changes and household income



Notes: Estimation of (24) 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 .

weighted average product price across transactions within hq , weighing by expenditures in the current quarter. We weigh observations in (24) 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 4. 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.⁴⁶

B.2 Spatial Variation

Here, we show that our cost-shock-based instrument is crucial for identification. We build on the robustness exercises in Section 4.4 in which we leverage spatial variation in changes in prices and expenditures.

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,

⁴⁶A related observation, documented in Appendix A, is that average price levels do not vary much across regions in Switzerland, either in 2014 or in 2015.

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 4. 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 (17) 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’s “Practitioner’s Guide” 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 (15). However, under stronger assumptions they can be adjusted to do so.

In our first approach in Section 4.2 using equation (18), 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 3, 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 (17) 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 combi-

nation 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).⁴⁷ For this reason, in our counterfactual analyses we impose different values of $\bar{\eta}_s$ that imply reasonable price elasticities for high income households and show that relative differences in welfare depend crucially on the value of η_s but are not substantially affected by the choice of $\bar{\eta}_s$.

C Theoretical appendix

We use a particular formulation of the nested, generalized, non-homothetic CES preferences presented in Fally (Forthcoming). In particular, utility u for household h at time t is defined as an implicit function of the consumption bundle c_{ht} and preferences ζ_{ht} by

$$f_h(u)^{\frac{\sigma-1}{\sigma}} = \sum_s (\zeta_{hs} u^{\gamma_s})^{\frac{1}{\sigma}} (c_{hst})^{\frac{\sigma-1}{\sigma}} \quad (25)$$

where

$$c_{hst} = \left(\sum_i (\zeta_{hit} u^{\gamma_i})^{\frac{1}{\eta_s(u)}} (c_{hit})^{\frac{\eta_s(u)-1}{\eta_s(u)}} \right)^{\frac{\eta_s(u)}{\eta_s(u)-1}} \quad (26)$$

$f_h(\cdot) > 0$, $\eta_s(u) \neq 1$ and $\sigma \neq 1$. These preferences reduce to nested homothetic CES if, for example, $\eta_s(u)$ is independent of u , $\zeta_i = \zeta_s = 0$, 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), where $u_{ht} \equiv v_h(p_{ht}, I_{ht}; \zeta_{ht})$ is the maximum utility achieved by household h at time t .

Deriving equation (12). The expenditure function satisfies $I_{ht} = e_h(\mathbf{p}_{ht}, u_{ht}; \zeta_{ht})$. Log-

⁴⁷In 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.

linearizing this equation at t_0 yields

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

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

$$d \log u_{ht} = \left(\frac{\partial \log e_h}{\partial \log u_h} \right)^{-1} \times \left(d \log I_{ht} - \sum_i b_{hit_0} d \log p_{hit} - \bar{\varepsilon}_{ht} \right) \quad (27)$$

This is equation (12) in the text.

Deriving equation (16). Substituting equation (27) into equation (11) yields

$$\begin{aligned} d \log b_{hit} &= \left(\frac{\partial \log e_h}{\partial \log u_h} \right)^{-1} \times \left(\gamma_i - \frac{\partial \eta_s}{\partial u_h} u_{ht_0} \log p_{hit_0} \right) \left(d \log I_{ht} - \sum_i b_{hit_0} d \log p_{hit} - \bar{\varepsilon}_{ht} \right) \\ &\quad + d \log \zeta_{hit} + (1 - \eta_{hst_0}) d \log p_{hit} + \psi_{hst} \end{aligned}$$

The previous expression and assumption (14) yield

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

The previous expression and assumption (15) yield

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

The previous expression is equation (16) given the definitions $v_{hit} \equiv d \log \zeta_{hit} - \kappa_i \bar{\varepsilon}_{ht}$ and $\tilde{\psi}_{hst} \equiv \psi_{hst} + \kappa_{hs} (d \log(I_{ht}/P_{ht}) - \bar{\varepsilon}_{ht})$.

Assumptions (14) and (15). 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} ,

$$\eta_{hst} \equiv \tilde{\eta}_s + \tilde{\eta}_s \log(u_{ht}). \quad (28)$$

If $\tilde{\eta}_s < 0$, then a household that achieves a utility 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 obtain⁴⁸

$$\frac{\partial \eta_s}{\partial u_h} u_{ht_0} \log p_{hit_0} = \tilde{\eta}_s (\log p_{it_0} + \log p_{hst_0})$$

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

$$f_h(x) = a_0 x^{a_1} \left[\sum_s \zeta_{hs} x^{\gamma_s} (P_{hs}(x))^{1-\sigma} \right]^{\frac{1}{\sigma-1}} \quad (29)$$

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

$$P_{hs}(x) = \left(\sum_{i \in \mathcal{I}(s)} \zeta_{hit_0} x^{\gamma_i} (p_{hit_0})^{1-\eta_s(x)} \right)^{\frac{1}{1-\eta_s(x)}} \quad (30)$$

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 (15), 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 (14), where $\kappa_i \equiv a_1^{-1} \gamma_i - \eta_s \log p_{it_0}$ and $\kappa_{hs} \equiv -\eta_s \log p_{hst_0}$.

Preferences can be rationalized. Here, we prove analytically that generalized non-homothetic CES preferences can be rationalized given that elasticities of substitution are defined by $\eta_{hst} \equiv \tilde{\eta}_s + \eta_s \log(u_{ht})$ with $\eta_s < 0$ as in our estimates. We build on results in Fally (Forthcoming). We focus on the case in which there is a single sector (or, equivalently, that all sectors are symmetric). In this case, direct utility can be defined implicitly by

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

where we have dropped household and time sub-indices, $\zeta_i \geq 0$ for all i and $\sum_i \zeta_i = 1$, and $\eta(u) = \bar{\eta} + \eta \log(u)$. We can re-express this as

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

⁴⁸More generally, our approach would allow for this equation to hold within each region but not across regions.

where $G_i(u) \equiv f(u) (\zeta_i u^{\gamma_i})^{\frac{1}{1-\eta(u)}}$. Given this implicit definition of the utility function, **Fally (Forthcoming)** shows that if $\eta'(u) < 0$ and $\eta(u) \neq 1$, then $du/dc_i > 0$ for all i if

$$K(u) \equiv \sum_i \exp\left(\frac{(\eta(u) - 1)^2 G'_i(u)}{\eta'(u) G_i(u)}\right) < 1$$

We now prove that this condition is satisfied under our functional form assumption for $\eta(u)$ with $\bar{\eta} > 1$ and $\eta < 0$,⁴⁹ and $f(u) = (u^{k_1})^{\frac{1}{1-\eta(u)}}$. In this case, $G_i(u) = (\zeta_i u^{\tilde{\gamma}_i})^{\frac{1}{1-\eta(u)}}$, where $\tilde{\gamma}_i \equiv \gamma_i + k_1$. Hence,

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

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

$$K(u) \equiv \sum_i \zeta_i C^{\tilde{\gamma}_i} \exp\left(\frac{\tilde{\gamma}_i (1 - \eta(u))}{C \eta'(u)}\right)$$

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 $x_i \equiv \tilde{\gamma}_i (1 - \bar{\eta}) / \eta$. Since $(1 - \bar{\eta}) / \eta > 0$, we have $\exp(x_i) < 1$ for all i if $\tilde{\gamma}_i < 0$ for all i . Hence, if $k_1 < -\min_i\{\gamma_i\}$ then $du/dc_i > 0$ for all i .⁵⁰ This condition can always be ensured to hold since the level of k_1 and γ_i are irrelevant for observable choices (which only depend on differences in γ_i) and for welfare changes given changes in prices and income.

D Quantitative appendix

⁴⁹To 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 preferences can be rationalized in a neighborhood of any u for which $\eta(u) > 1$.

⁵⁰In order to check whether this condition holds for other functional forms for $\eta(u)$ and f , we would have to do so numerically given estimates of ζ_i which we do not need for our exercises.

Table 17: Import price declines

Annual income	% change in import/domestic price			
	-2.2	-10	-20	-40
1: 20,000 elasticity 6.6	0.5	2.4	5.4	13.2
2: 60,000 elasticity 4.4	0.5	2.6	5.6	12.7
3: 120,000 elasticity 3.0	0.6	2.8	5.8	12.6
<hr/>				
% difference btw CV				
income groups 2 and 1	13%	9%	4%	-4%
income groups 3 and 1	25%	17%	9%	-4%

Notes: This table replicates the exercise in Table 10 but studying import price declines.

Table 18: Smaller differences in elasticities of substitution

Annual income	% change in import/domestic price				
	+2.2	+10	+20	+40	+1000
1: 20,000 elasticity 5.7	-0.5	-1.9	-3.3	-5.1	-6.6
2: 60,000 elasticity 4	-0.5	-2.2	-4.2	-7.2	-12.3
3: 120,000 elasticity 3	-0.6	-2.6	-5.0	-9.1	-22.0
<hr/>					
% difference btw CV					
income groups 2 and 1	16%	20%	27%	41%	86%
income groups 3 and 1	30%	38%	49%	77%	231%
<hr/>					
Contribution heterogeneous η_s					
income groups 2 and 1	6%	22%	37%	54%	76%
income groups 3 and 1	5%	20%	34%	52%	82%

Notes: This table replicates the exercise in Table 10 but imposing that $\eta_s = 1.5$ rather than $\eta_s = 2$, while maintaining that the lowest elasticity of substitution (that for the highest income group with income of 120,000 CHF), η_{hst0} , is equal to 3.

Table 19: Elasticity of substitution of high income group = 1.5

Annual income	% change in import/domestic price				
	+2.2	+10	+20	+40	+1000
1: 20,000 elasticity 5.1	-0.5	-1.9	-3.4	-5.5	-7.6
2: 60,000 elasticity 2.9	-0.5	-2.3	-4.4	-8.1	-19.9
3: 120,000 elasticity 1.5	-0.6	-2.7	-5.3	-10.4	-86.7
<hr/>					
% difference btw CV					
income groups 2 and 1	16%	22%	30%	48%	163%
income groups 3 and 1	30%	40%	54%	89%	1041%
<hr/>					
Contribution heterogeneous η s					
income groups 2 and 1	8%	28%	44%	62%	87%
income groups 3 and 1	7%	25%	41%	60%	96%

Notes: This table replicates the exercise in Table 10 but imposing that the lowest elasticity of substitution (that for the highest income group with income of 120,000 CHF), η_{hst0} , is equal to 1.5 rather than 3.

Table 20: Elasticity of substitution of high income group = 5

Annual income	% change in import/domestic price				
	+2.2	+10	+20	+40	+1000
1: 20,000 elasticity 8.6	-0.5	-1.7	-2.8	-3.8	-4.1
2: 60,000 elasticity 6.4	-0.5	-2.1	-3.8	-5.7	-7.0
3: 120,000 elasticity 5.0	-0.6	-2.5	-4.5	-7.5	-11.0
<hr/>					
% difference btw CV					
income groups 2 and 1	16%	23%	32%	50%	70%
income groups 3 and 1	31%	42%	60%	99%	168%
<hr/>					
Contribution heterogeneous η s					
income groups 2 and 1	8%	28%	44%	60%	71%
income groups 3 and 1	7%	26%	42%	60%	75%

Notes: This table replicates the exercise in Table 10 but imposing that the lowest elasticity of substitution (that for the highest income group with income of 120,000 CHF), η_{hst0} , is equal to 5 rather than 3.

Table 21: Homogeneous elasticities within the Service sector

Annual income	% change in import/domestic price				
	+2.2	+10	+20	+40	+1000
1: 20,000 elasticity 6.6	-0.5	-1.8	-3.2	-4.7	-5.7
2: 60,000 elasticity 4.4	-0.5	-2.2	-4.1	-7.0	-11.1
3: 120,000 elasticity 3.0	-0.6	-2.6	-4.9	-9.0	-21.4
<hr/>					
% difference btw CV					
income groups 2 and 1	16%	21%	29%	47%	93%
income groups 3 and 1	30%	39%	54%	89%	275%
<hr/>					
Contribution heterogeneous η_s					
income groups 2 and 1	7%	26%	42%	60%	79%
income groups 3 and 1	6%	23%	39%	58%	85%

Notes: This table replicates the exercise in Table 10 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 22: Homogeneous elasticities within the Service and Other goods sectors

Annual income	% change in import/domestic price				
	+2.2	+10	+20	+40	+1000
1: 20,000	-0.5	-1.9	-3.4	-5.5	-8.1
2: 60,000	-0.5	-2.2	-4.1	-7.0	-11.1
3: 120,000	-0.6	-2.5	-4.8	-8.3	-14.6
<hr/>					
% difference btw CV					
income groups 2 and 1	15%	18%	22%	27%	37%
income groups 3 and 1	29%	34%	40%	52%	80%
<hr/>					
Contribution heterogeneous η_s					
income groups 2 and 1	4%	15%	26%	41%	61%
income groups 3 and 1	3%	13%	23%	38%	65%

Notes: This table replicates the exercise in Table 10 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 $s = \text{Services and Other goods}$ for all h).