NBER WORKING PAPER SERIES

THE DISTRIBUTIONAL CONSEQUENCES OF LARGE DEVALUATIONS

Javier Cravino
Andrei A. Levchenko

Working Paper 23409
http://www.nber.org/papers/w23409

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2017

We are grateful to Mark Aguiar, David Atkin, Ariel Burstein, Yu Chin Chen, Michael Devereux, Natalia Ramondo, Daniele Siena, and seminar and workshop participants at several institutions for helpful suggestions, and to Laurien Gilbert and Nitya Pandalai-Nayar for excellent research assistance. We would especially like to thank Christian Ahlin and Mototsugu Shintani for sharing the digitized pre-April 1995 Mexican consumer price data. Financial support from Michigan Institute for Teaching and Research in Economics (MITRE) and from the National Science Foundation under grant SES-1628879 is gratefully acknowledged. Levchenko would also like to thank the University of Zurich, the UBS Center for Economics in Society, and the Excellence Foundation Zurich for their hospitality during the 2014-15 academic year and financial support for this project. Email: jcravino@umich.edu, alev@umich.edu. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

The Distributional Consequences of Large Devaluations
Javier Cravino and Andrei A. Levchenko
NBER Working Paper No. 23409
May 2017
JEL No. F31,F61

ABSTRACT

We study the impact of large exchange rate devaluations on the cost of living at different points on the income distribution. Poor households spend relatively more on tradeable product categories, and consume lower-priced varieties within categories. Changes in the relative price of tradeables and of lower-priced varieties affect the cost of living of low-income relative to high-income households. We quantify these effects following the 1994 Mexican devaluation and show that they can have large distributional consequences. Two years post-devaluation, the cost of living for the bottom income decile rose 1.48 to 1.62 times more than for the top income decile.

Javier Cravino
Department of Economics
University of Michigan
611 Tappan Street, Lorch Hall 365B
Ann Arbor, MI 48109
and NBER
jcravino@umich.edu

Andrei A. Levchenko
Department of Economics
University of Michigan
611 Tappan Street
Ann Arbor, MI 48109
and NBER
alev@umich.edu

A online appendix is available at http://www.nber.org/data-appendix/w23409

Large exchange rate devaluations are associated with dramatic changes in relative prices. In the aftermath of a devaluation, the price of tradeable goods “at the dock” moves one-for-one with the exchange rate, the retail price of tradeable goods increases, though less than the exchange rate, while non-tradeable goods’ prices are relatively stable.[1] A clear illustration of such relative price movements is presented in Figure 1, which plots the evolution of these prices following the 1994 Mexican devaluation. The retail price of tradeables is much closer to the price of non-tradeables than to prices of tradeables at the dock, consistent with the importance of local distribution costs in retail prices.

Figure 1: Price changes during the 1994 Mexican devaluation

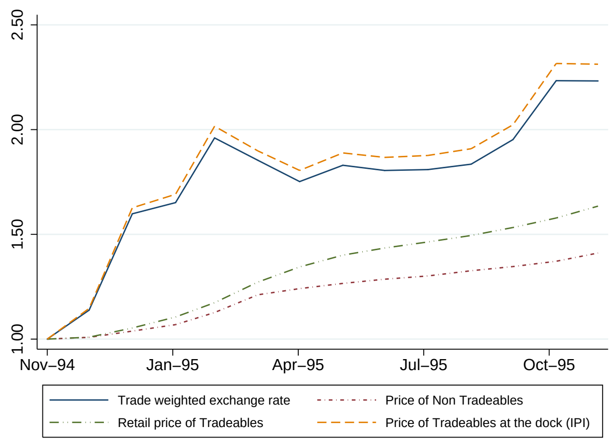

Notes: This figure plots the trade-weighted nominal exchange rate, the import price index, and the consumption price indices of tradeables and non-tradeables following the November 1994 peso devaluation, each rebased to November 1994.

This paper studies the distributional consequences of such relative price movements. It is well known that households at different income levels consume very different baskets of goods.[2] We distinguish two types of differences, which we label *Across* and *Within*. Across product categories, low-income households spend relatively more on tradeables (such as food), while high-income households spend relatively more on non-tradeables (such as personal services). Within product categories, low-income households spend relatively more on lower-end goods purchased from lower-end retail outlets. Changes in the relative price of tradeables and of low-priced varieties following a large devaluation will thus affect households differentially, generating a distributional welfare impact.

[1] These patterns were first documented by Burstein et al. (2005) for 5 large devaluations. In summarizing the literature, Burstein and Gopinath (2015) extend these findings to include more devaluation episodes.

[2] This was documented as early as the 19th century by Engel (1857, 1895, "Engel’s Law"). For recent evidence using household surveys from multiple countries, see Almås (2012).

We measure the magnitude of these two effects during the 1994 Mexican devaluation. For this episode, we combine two sources of detailed microdata that are key for studying these mechanisms. The first is household-level expenditures on detailed product categories from the Mexican household surveys both immediately before and after the crisis. The second is monthly data on unique product-outlet level prices that the Bank of Mexico uses to construct the consumer price index. In what follows, we refer to a unique product-outlet combination as a variety. Crucially, the consumption categories in the household survey can be matched to the product categories for which the Bank of Mexico collects price data. Indeed, these datasets are the two principal inputs underlying the official Mexican CPI.

We first calculate an income-specific price index that captures the *Across* effect by weighting price indices for disaggregated consumption categories with income-specific expenditure shares from the 1994 household expenditure survey. According to this index, in the 2 years following the devaluation the consumers in the bottom decile of the Mexican income distribution experienced cost of living increases about 1.25 times larger than the consumers in the top income decile. The increase in the price index was 95% for households in the poorest decile, compared to 76% for households in the richest decile. The effect is monotonic across all income deciles.

We then compute an income-specific price index that captures the *Within* effect using the unique product-outlet level price data and household expenditure data. First, we use the household survey data to show that high-income households tend to pay higher unit values within detailed product categories (i.e. both the high- and low-income households buy bread, but the high-income households pay more per kilo). This evidence supports the notion that households at the top of the income distribution purchase higher-priced varieties. We then compute a Within price index by assuming that all consumers have the same expenditure shares across product categories, but that within each category, the high-income households consume the more expensive varieties, and the low-income the less expensive ones. In our benchmark index, the Within effect implies that inflation for the lower-income consumers was between 14 and 22 percentage points higher than for the higher-income consumers. We supplement the Within effect results for Mexico using the Economist Intelligence Unit CityData on store prices in a sample of several emerging market devaluations.

The Across and Within effects are roughly additive, reinforcing each other. Our preferred estimate of the price index that combines these two effects implies that the households in the bottom decile of the Mexican income distribution experienced increases in the cost of living between 1.48 and 1.62 times higher than the households in the top decile in the two years that follow the devaluation. Absent any changes in nominal income, our combined price index implies a decline in real income of about 50% for households in the bottom decile compared to about 40% for households in the top decile. The main finding is thus that both the Across and the Within distributional effects were large and economically significant in the 1994 Mexican devaluation.

Understanding why the observed price changes are anti-poor requires an account of the mechanisms behind the relative price changes that follow a large devaluation. We show that the poor spend a higher fraction of their income on tradeable product categories, and among tradeables, on categories with a systematically lower non-tradeable component. This is primarily driven by differences in distribution margins rather than by differences in the prevalence of local goods across categories. As the relative price of tradeables to non-tradeables increases following the devaluation, the prices paid by the low-income households rise by proportionally more than those paid by the high-income households. This mechanism provides an account of the Across effect.

We then evaluate whether the leading explanations for incomplete exchange rate pass-through into retail prices are consistent with the observed relative price changes within product categories.[3] First, if cheaper varieties have lower distribution margins, their relative price will increase following a devaluation. We show in a simple flexible price framework that differences in distribution margins account well for the observed differences in price changes across varieties. Second, if some varieties are not traded internationally but only produced and sold locally, the price of these varieties may fall relative to imported ones. If this is the case and imported varieties are more expensive than local ones, then the price of the expensive varieties should actually increase by more than cheap varieties following the devaluation. This is at odds with the relative price movements we document. Third, if markups of higher-quality varieties fall by more following a devaluation, we should expect the relative prices of expensive varieties to decrease.[4] This type of effect is consistent with the relative price changes observed in our data.

Our analysis is expressly about the differences in consumption price levels for households of different incomes, and is silent on how nominal income changed across the income distribution. As such, our results can be interpreted as differences in the compensating variation of changes in the consumption price level across the income distribution. That is, we answer the question, by how much should the nominal income of different households have changed to leave everyone relatively as well off as before? Our

[3]See e.g. Burstein et al. (2005); Burstein and Gopinath (2015).

[4]This assumes that prices are increasing in product quality. See Auer et al. (2014) and Antoniades and Zaniboni (2015) for empirical evidence that exchange rate pass-through is lower for high-quality products.

results can be benchmarked to existing studies of how incomes changed during the Mexican devaluation. According to Mexico's National Statistical Institute (INEGI) there was not much differential impact in the decline in income per capita across deciles over this period, with incomes falling by 29% in inflation-adjusted terms for the highest income decile, and by 27% for the lowest decile.[5] Lopez-Acevedo and Salinas (2000) report a modest decrease in income inequality over this period using the same household survey that we use in this paper, which is a repeated cross-section. Using a panel survey of wages (ENEU), Maloney et al. (2004) report that median real wages fell by 30%, but that there was not much differential impact across education groups (which can serve as a rough proxy for income). Using the ENEU, Verhoogen (2008) shows that inequality, measured by the 90-10 income ratio or the white-blue collar wage gap actually increased over the 1994-1996 period, and more broadly did not experience any change in its (upward) trend. All in all, available evidence suggests that it is unlikely that a large pro-poor change in nominal incomes could have erased the anti-poor price changes that we document.[6]

Our paper belongs to the literature on large devaluations, surveyed by Burstein and Gopinath (2015). This literature has highlighted that pass-through into retail prices is incomplete in part because consumer prices include a large non-traded component – the distribution margin. Goldberg and Campa (2010) document the heterogeneity in distribution margins across sectors. We study a pattern that has until now been ignored in the exchange rate literature: the importance of the non-traded component in the total consumption basket varies systematically across the income distribution, both across and within detailed product categories. Some evidence on what we label the Across effect is provided by Friedman and Levinsohn (2002) and Levinsohn et al. (2003) for Indonesia's 1998 depreciation, Kraay (2008) for the Egyptian 2000-05 depreciation, and de Carvalho Filho and Chamon (2008) for Brazil and Mexico over the period 1980-2006. Our paper examines the Across effect more systematically and relates it to the interaction between distribution margin heterogeneity and differences in consumption baskets.

Our paper is also related to a large and growing literature in international trade that models demand non-homotheticities and examines the distributional impact of economic integration across consumers (see, e.g. Fajgelbaum et al., 2011; Fajgelbaum and Khandelwal, 2016; Atkin et al., 2016). The closest to ours are papers by Porto (2006) and Faber

[5]See Encuesta Nacional de Ingresos y Gastos de los Hogares (ENIGH), Síntesis histórica, 1992-2008.

[6]Changes in asset values/incomes are more difficult to ascertain, but available evidence suggests that assets of the poor suffered larger losses than those of the rich. Halac and Schmukler (2004) document that in a sample of Latin American crises that includes Mexico in 1994, larger depositors and larger borrowers suffered less than small ones. Lopez-Acevedo and Salinas (2000) document that changes in capital and financial income during the Mexican crisis favored the top income decile households.

(2014). Porto (2006) uses household consumer expenditure data in Argentina following Mercosur to trace the distributional impact of this regional trade agreement on different consumers. The analysis incorporates the Across effect but not the Within effect. Faber (2014) shows that following NAFTA, intermediate inputs used in production of higher-quality varieties became cheaper in Mexico, and richer consumers benefited more – a type of Within effect that is differential across product categories according to their intensity of imported input use. Relative to these papers, that focus on long-run changes, we examine the relatively short-run effects following large devaluations. Our paper is the first, to our knowledge, to combine the analysis of Across and Within effects.

The rest of the paper is organized as follows. Section 2 illustrates the distributional effects of relative price changes when consumption baskets differ across consumers. Section 3 describes the data and the main results. Section 4 discusses the possible mechanisms for the main findings, with an emphasis on variation in distribution margins, and Section 5 concludes.

2 Conceptual framework

Let the indirect utility of a household h be denoted by V_t^h, and let $\widehat{x_t} \equiv x_t/x_{t_0} - 1$ denote the cumulative growth rate of variable x_t between some base period t_0 and time t. The proportional change in welfare following a change in income and the vector of prices is to a first approximation given by

$$\widehat{V_t}^h = \widehat{W}_t^h - \sum_{g \in G} \omega_g^h \widehat{P}_{g,t}, \tag{1}$$

where W_t^h is nominal income, g indexes goods, ω_g^h are household-specific expenditure shares, and $\widehat{P}_{g,t}$ are good-specific price changes. To illustrate the distributional effects of a change in prices across households, it helps to write (1) as:

$$\widehat{V}_t^h = \underbrace{\widehat{W}_t^h - \sum_{g \in G} \omega_g \widehat{P}_{g,t}}_{\text{homothetic-utility } \widehat{V}} - \underbrace{\sum_{g \in G} \widehat{P}_{g,t}(\omega_g^h - \omega_g)}_{Cov(\widehat{P}_{g,t}, \omega_g^h - \omega_g)}, \tag{2}$$

where ω_g is the economy-wide share of spending on good g. The first term of this expression is the change in welfare that we would obtain if utility were homothetic and every h had the same consumption basket. The second term captures the distributional impact across households. The term is reminiscent of a (negative) covariance between price changes and household-level relative spending shares. If the pattern of price changes across g is positively correlated with h's relative spending shares, then h suffers more from this vector of price changes than the average household, because prices go up on average more in goods that the household consumes more of.

Consider an example in which there are two households, rich and poor, $h = r, p$, and two goods, tradeables and non-tradeables: $g = T, NT$. Suppose further that the poor have higher expenditure shares in tradeables: $\omega_T^p > \omega_T > \omega_T^r$. If an exchange rate depreciation leads to a higher increase in the price of tradeables than in the price of non-tradeables – $\widehat{P}_{T,t} > \widehat{P}_{NT,t}$ – then the last term in (2) will be negative for the poor and positive for the rich. This is the simplest version of what in the empirical analysis below we refer to as the Across effect.

To illustrate the Within effect, suppose instead that the two goods were an expensive variety and a cheap variety: $g = E, C$, and the poor consumed a higher share of the cheap variety than the rich, $\omega_C^p > \omega_C > \omega_C^r$. If the price of the cheap variety increased by more after a devaluation, $\widehat{P}_{C,t} > \widehat{P}_{E,t}$, we would once again have an anti-poor distributional effect.

The discussion above underscores the point that there is no fundamental difference in how the Across and Within effects work. Both are driven by the covariance of price changes and relative spending shares across the income distribution. Because they have different data requirements, it is still convenient to separate them in the empirical analysis. Note also that the expression (1) has a natural compensating variation interpretation: in response to a given vector of price changes $\widehat{P}_{g,t}$, a compensating variation for household h is a change in income $\widehat{W}_t^h$ that leaves welfare unchanged ($\widehat{V}_t^h = 0$). Thus, while we state the empirical results in terms of changes in household-level costs of living indices $\widehat{P}_t^h$, they can equivalently be stated in terms of the heterogeneity in the compensating variation across households.

2.1 Within and Across effects: definitions and measurement

This section defines the Across, Within, and Combined price indices. Let there be G goods categories indexed by g, and let each g contain varieties indexed by v_g. Households spend different shares of their income both across goods categories g, and across varieties v_g within each g. The change in the aggregate price index is defined by:

$$\widehat{P}_t \equiv \sum_{g \in G} \omega_g \widehat{P}_{g,t}, \tag{3}$$

where $\omega_g \equiv \frac{\sum_h P^h_{g,t_0} q^h_{g,t_0}}{\sum_h \sum_g P^h_{g,t_0} q^h_{g,t_0}}$ is the economy-wide expenditure share on good g at some base period t_0, and

$$\widehat{P}_{g,t} \equiv \frac{1}{V_g} \sum_{v_g \in g} \widehat{P}_{v_g,t} \tag{4}$$

is the change in the price index for good category g that has V_g varieties. $\widehat{P}_t$ is the change in the CPI as it would be constructed by national statistical agencies.

The change in the household-specific price index is given by:

$$\widehat{P}^h_t \equiv \sum_{g \in G} \omega^h_g \widehat{P}^h_{g,t}, \tag{5}$$

where $\omega^h_g \equiv \frac{P^h_{g,t_0} q^h_{g,t_0}}{\sum_g P^h_{g,t_0} q^h_{g,t_0}}$ is now the share of household h's expenditures that go towards good category g, and $\widehat{P}^h_{g,t}$ is the change in the price sub-index of good g. It varies across households because they consume different varieties:

$$\widehat{P}^h_{g,t} \equiv \sum_{v_g} s^h_{v_g} \widehat{P}_{v_g,t}, \tag{6}$$

where $s^h_{v_g}$ is household h's share of expenditures in variety v_g within the good category g, and $\widehat{P}_{v_g,t}$ is the (non-household-specific) change in the price of variety v_g of good g. $\widehat{P}^h_{g,t}$ can vary across households if households of different incomes consume different goods within each good category g. This would happen, for instance, if the richer households consume systematically higher-priced varieties within each g.

We define the *Across* change in the price index for household h as:

$$\widehat{P}^h_{Across,t} \equiv \sum_{g \in G} \omega^h_g \widehat{P}_{g,t}, \tag{7}$$

and the *Within* change in the price index for household h as:

$$\widehat{P}^h_{Within,t} \equiv \sum_{g \in G} \omega_g \widehat{P}^h_{g,t}. \tag{8}$$

In words, $\widehat{P}^h_{Across,t}$ is the change in the cost of living for a hypothetical household that has h's expenditure shares across g, and faces the unweighted average price change across all varieties within each g. By contrast, $\widehat{P}^h_{Within,t}$ is the change in the cost of living for a hypothetical household that has aggregate consumption shares across goods g, but consumes household h's varieties within each good g.

Using these expressions, the change in the price index of household h is:

$$\widehat{P}_t^h = \underbrace{\sum_{g \in G} \omega_g^h \widehat{P}_{g,t}}_{\widehat{P}_{Across,t}^h} + \underbrace{\sum_{g \in G} \omega_g \widehat{P}_{g,t}^h}_{\widehat{P}_{Within,t}^h} + \underbrace{\sum_{g \in G} \left(\omega_g^h - \omega_g\right)\left(\widehat{P}_{g,t}^h - \widehat{P}_{g,t}\right)}_{\widehat{P}_{Cov,t}^h} - \underbrace{\sum_{g \in G} \omega_g \widehat{P}_{g,t}}_{\widehat{P}_t}.$$

The third term, labeled $\widehat{P}_{Cov,t}^h$, is a "covariance" across goods between how different price changes are for h relative to the average and how different h's expenditure share relative to the average. It is not formally a covariance because $\widehat{P}_{g,t}$ is not the mean across goods, but rather the mean across varieties within g, and ω_g is not the mean across goods but an expenditure-weighted average across households. The "covariance" will be positive when h experiences large deviations from the mean in its household-specific price in its relatively large expenditure categories.

The difference in the change of the price indices of two households h and h' at different points in the income distribution is given by

$$\Delta\widehat{P}_t = \Delta\widehat{P}_{Across,t} + \Delta\widehat{P}_{Within,t} + \Delta\widehat{P}_{Cov,t},$$

where $\Delta\hat{x}_t \equiv \hat{x}_t^h - \hat{x}_t^{h'}$ denotes a cross-sectional rather than a time difference. The difference in $\widehat{P}_t{}^h$ is the sum of the differences in the Across and Within indices and the covariance term. Section 3 calculates $\Delta\widehat{P}_t$, $\Delta\widehat{P}_{Across,t}$ and $\Delta\widehat{P}_{Within,t}$ following the 1994 Mexican devaluation and shows that the covariance term is quantitatively small.

3 Price changes during the 1994 Mexican devaluation

This section quantifies the distributional consequences of the 1994 Mexican devaluation. After describing the data sources, we report the Across, Within, and Combined effects. We conclude the section by recalculating price indices under alternative assumptions to show the robustness of the results.

3.1 Data description

The analysis uses two main data sources. The first is monthly data on unique product-outlet level prices that the Bank of Mexico uses to construct the consumer price index. The second is household-level expenditure data on detailed product categories from the

Mexican household surveys both immediately before and after the crisis. Our baseline indices incorporate price and expenditure data from all regions in Mexico.[7]

3.1.1 Mexican data on consumer prices

The Mexican micro data on consumer prices are collected by the Bank of Mexico with the purpose of computing the Consumer Price Index. Since January 1994, the prices that underlie the construction of the CPI are published monthly in the Diario Oficial de la Federacion (DOF), the official bulletin of the Mexican government. Each price quote in the DOF corresponds to a 'specific' variety, which is a unique product-city-outlet combination that can be traced through time. An exact product description – e.g. Kellogg's, Corn Flakes, 500gr box – for each variety was published in the April 1995 DOF. Unfortunately, outlet identifiers are not available in the data for this time period. The varieties are grouped into 313 'generic' categories – e.g. Cereal in Flakes – representing the goods and services consumed in Mexico. For most generic product categories, the price quotes for the specific varieties are expressed in common units. For example, the prices of varieties within the category Cereal in Flakes are quoted per kilo of cereal. These micro price data from the DOF have been used previously by Ahlin and Shintani (2007) and Gagnon (2009).

We focus on a sample of 28,675 specific varieties grouped into 284 generic categories that can be observed continuously in 35 municipalities throughout Mexico from January 1994 to December 1996.[8] For each specific variety, we observe its monthly price, its generic category, the city in which it is sold and the units in which prices are quoted. The DOF also publishes the specific varieties that are added because of product substitutions, or changes in the outlets that are being sampled by the price inspectors. We focus on the specific varieties that can be observed continuously through our sample. Appendix Table A4 reports the 284 generic categories.

3.1.2 Mexican household surveys

We use the Mexican household surveys, Encuesta Nacional de Ingresos y Gastos de los Hogares (ENIGH) for 1994 and 1996 to obtain consumption expenditures across consumption categories by household. The key variables that come from this dataset are the household's city, income, and total expenditures in 597 detailed product categories. Crucially, the product categories in the ENIGH can be mapped to the 331 generic good

[7] Appendix B reports results restricting attention to relative price changes within Mexico City only.

[8] There was a revision in April 1995, in which some of the generic categories were changed.

categories used to calculate the CPI – in fact, the weights used to compute the official CPI are derived from the ENIGH. In addition, for some product categories the ENIGH reports the total quantity of the good consumed by each household. We combine the total quantities with the expenditure data to compute the unit value paid by each household in each product category.

The top panel of Appendix Table A5 reports the average quarterly income in Mexico in each income decile, in pesos. The income of the average household in the top income decile was more than six times higher than the average household in the median decile, and 23 times higher than the average household in the bottom decile. The bottom panel of Appendix Table A5 reports the consumption expenditure shares in the 8 1-digit CPI categories by income decile.

3.2 The Across effect

We calculate the Across price index in equation (7), reproduced here to facilitate exposition:

$$\widehat{P}^h_{Across,t} \quad = \quad \sum_{g \in G} \omega^h_g \widehat{P}_{g,t}.$$

The category-level price indices $\widehat{P}_{g,t}$ aggregate the micro prices from the DOF according to equation (4). We define the product categories G for two alternative levels of disaggregation for which the Bank of Mexico computes consumer price indices: at the 1-digit level (8 good categories listed in Appendix Table A5), and at the 9-digit level (284 categories listed in Appendix Table A4). The expenditure shares ω^h_g for the product categories come from the 1994 household expenditure survey. In particular, we sort households into income deciles and compute the expenditure shares of each decile in each of the G product categories. The price indices are normalized to 1 in October 1994, the month before the devaluation.

Tables 1a and 1b report the resulting price indices for different deciles of the income distribution when the product categories are defined at the 1- and 9-digit levels of disaggregation. Our aggregate price index closely follows the official inflation rate computed by the Bank of Mexico.[9] Changes in $\widehat{P}^h_{Across,t}$ differ dramatically across the income distribution in the two years following the devaluation. The Across price index computed at the 1-digit level of disaggregation increased by 87 percent for the households in the

[9]Differences in the two indices arise in part because the official Mexican CPI used expenditure weights from the 1977 survey prior to the 1995 revision.

bottom decile, compared to only 79 percent for households in the top decile. The relation between the change in the indices and household income decile is monotonic, with households of lower income experiencing higher inflation in this period.

The difference in the price indices is more dramatic when $\widehat{P}^h_{Across,t}$ is computed at the 9-digit level of disaggregation. The change in the 9-digit Across price index was 95 percent for households in the bottom decile, compared to 76 percent for the top decile. Two years after the devaluation, inflation for the bottom decile was 1.25 times higher than inflation for the top decile due to differences in household expenditure shares across product categories.

Table 1: The Across price index by income decile, 1994 weights

(a) 1-Digit

	Income Decile											
	1	2	3	4	5	6	7	8	9	10	Aggregate	Official
Oct. 94	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Oct. 95	1.48	1.47	1.47	1.47	1.47	1.46	1.46	1.46	1.45	1.44	1.45	1.49
Oct. 96	1.87	1.86	1.85	1.85	1.84	1.83	1.83	1.82	1.81	1.79	1.82	1.88

(b) 9-Digit

	Income Decile										
	1	2	3	4	5	6	7	8	9	10	Aggregate
Oct. 94	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Oct. 95	1.51	1.50	1.49	1.49	1.48	1.47	1.47	1.46	1.45	1.42	1.45
Oct. 96	1.95	1.91	1.89	1.88	1.86	1.84	1.83	1.82	1.81	1.76	1.82

Note: These tables report the Across price indices defined in equation (7) for different income deciles. Table 1a computes the price index using 8 1-Digit product categories for G, while Table 1b computes the price index using 284 9-Digit product categories for G. The expenditure weights come from the 1994 household survey.

We next compute the Across price indices at the household level. Appendix Figure A1 plots the quadratic and the local polynomial fit of $\widehat{P}^h_{Across,t}$ for October 1996 computed at the 9-digit level of disaggregation, for households of different income levels. The figure confirms that the relation shown in Tables 1a and 1b between inflation and income is monotonic. The price difference between the richest and poorest household exceeds 25 percentage points. The confidence intervals show that the difference in price indices between the top and the bottom of the income distribution is strongly statistically significant.

3.3 The Within effect

The Within price index is defined by equation (8), reproduced here for convenience:

$$\widehat{P}^h_{Within,t} \quad = \quad \sum_{g \in G} \omega_g \widehat{P}^h_{g,t}.$$

We weight the generic product categories g with aggregate expenditure weights ω_g computed from the household expenditure survey, and allow for differences in the price indices that households face for each generic category: $\widehat{P}^h_{g,t} \equiv \sum_{v_g \in g} s^h_{v_g} \widehat{P}_{v_g,t}$. Differences in the price indices $\widehat{P}^h_{g,t}$ stem from differences in the expenditure shares $s^h_{v_g}$ across the different varieties v_g within each product category g.

While we can observe the price change $\widehat{P}_{v_g,t}$ of every specific variety in the DOF, it is important to emphasize that the expenditure shares of each household $s^h_{v_g}$ are not observable. Appendix A uses data from the 1994 and 1996 household expenditure surveys to document that within narrow product categories, richer households tend to purchase more expensive varieties. We link expenditure shares $s^h_{v_g}$ to household income following this evidence, and assume that high-income households consume high-priced varieties while low-income households consume low-priced varieties. Section 3.5.1 below performs two additional exercises that employ information on spending patterns to construct alternative versions of the Within price index.[10]

We classify varieties as high- or low-priced using two alternative criteria. First, we split varieties according to whether their average price between January 1994 and October 1994 – the 10 months prior to the devaluation for which we have data – was above or below the average price of the median good in the generic category. Second, we split the January 1994-October 1994 average prices into quartiles in each generic category, and focus on products that are in the highest vs. the lowest quartiles. Focusing on the 10-month average (January 1994-October 1994) as the base period in which we classify varieties into high- or low- price bins, as opposed to the price in one particular month, has the advantage that temporary sales are less likely to be identified as low prices. Appendix 3.5 shows that using January 1994 as our base period does not significantly affect our results.

One potential concern with this procedure is that high and low pre-devaluation prices

[10]Note that the distinction between the Across and Within effects is driven purely by data availability considerations. An alternative approach would be to carry out the entire analysis at a higher level of aggregation, such that we can always observe expenditure shares. In a sense, Tables 1a and 1b already do that by comparing the price indices obtained under the coarsest product classification (8 categories) and the finest product classification (about 300 product categories) for which expenditure shares are observable. Moving to a more disaggregated level increases the disparity in the cost of living changes between the high- and low-income households, suggesting that the anti-poor pattern in price changes manifests itself at multiple levels of product disaggregation.

may not reflect differences in product attributes (such as the type of retail outlet), but may come simply from price dispersion due to staggered price adjustment. If some prices are low at the beginning of the sample because they have not been adjusted in a long time, a large increase in these prices may simply reflect that the price is finally being adjusted. To avoid this concern, we limit our analysis to specific varieties for which we see a price change between January 1994, our base month, and October 1994, the month prior to the devaluation. For this sample of products, we can be more confident that changes in prices that occur after October 1994 are not due to the firms resetting old prices.

Finally, the Within price index from equation (8) can only be computed for those product categories in which identical goods can be observed continuously through time. Unfortunately, this is not feasible for every category, since some categories were discontinued in the April 1995 revision of the consumer price index. As a consequence, only 284 of the 331 generic categories can be traced before March 1995. The continuing categories account for 82 percent of the expenditures. In addition, there are some generic categories, most prominently apparel, for which the micro price quotes are based on 'samples' of products, as opposed to unique individual products. After excluding these product categories, there are 231 categories in which identical products can be observed continuously through time, accounting for 55 percent of total consumption expenditures.[11] To compute a price index that reflects the importance of the Within effect for the entire economy we need to take a stand on how the relative price of cheap vs. expensive varieties changed for the missing categories.

With this in mind, we compute the Within price index under two limiting assumptions. First, we take a conservative approach and assume that the relative price of cheap vs. expensive varieties remained constant for the missing generic categories. In this case, the Within price index is given by:

$$\widehat{P}^h_{Within,t} = \sum_{g \in G_M} \omega_g \widehat{P}^h_{g,t} + \sum_{g \in G_U} \omega_g \widehat{P}_{g,t}, \tag{9}$$

where G_M is the set of categories for which identical varieties are measured continuously through time, G_U is the set of categories for which identical goods cannot be measured continuously through time, and $\widehat{P}_{g,t}$ is the change in the aggregate price index for the goods in category g. Second, we make the opposite assumption that the change in the relative price of cheap vs. the expensive varieties for the unmeasured categories was equal to the (weighted) average change of the price of cheap and expensive varieties that

[11]For the median category, we can trace 69 different price quotes through time, and the initial ratio of the maximum to the minimum price within the median category is 4.7.

we do observe. In particular, we assume that for each category $g \in G_U$, the price index is $\widehat{P}^h_{g,t} = \widehat{P}_{g,t} \times \frac{\sum_{g \in G_M} \omega_g \widehat{P}^h_{g,t}}{\sum_{g \in G_M} \omega_g \widehat{P}_{g,t}}$. In this case, the Within price index is given by:

$$\widehat{P}^h_{Within,t} = \sum_{g \in G_M} \omega_g \widehat{P}^h_{g,t} + \sum_{g \in G_U} \omega_g \widehat{P}_{g,t} \frac{\sum_{g \in G_M} \omega_g \widehat{P}^h_{g,t}}{\sum_{g \in G_M} \omega_g \widehat{P}_{g,t}}. \quad (10)$$

Figure 2 plots the evolution of the Within price indices computed when we sort goods relative to the median price within each product category. The price indices for high vs. low prices are very close to each other before the October 1994 devaluation. Following the devaluation, the price indices start to diverge.

Figure 2: The Within price indices

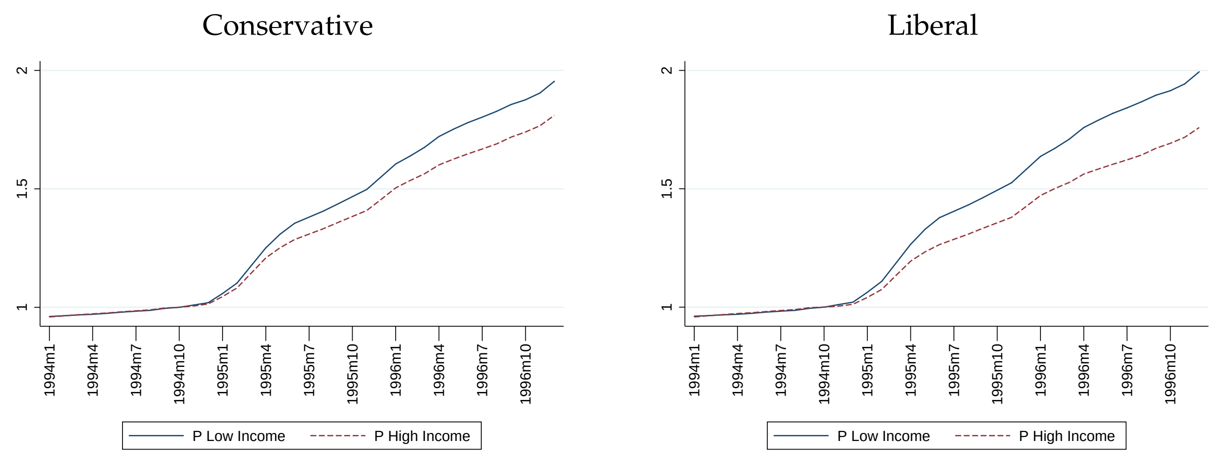

Notes: This figure plots the Within price indices for consumers that buy the varieties priced above ("P High Income") and below ("P Low Income") the median price within each product category. The Conservative price indices are defined in (9), and the Liberal indices in (10).

The values for the resulting price indices are reported in Table 2. Columns 1-2 and 4-5 report the price indices when we sort varieties based on whether their average price prior to the devaluation was below and above the median. Even according to our most conservative price index, inflation was substantially higher for the varieties that were initially below the median: by October 1996, the price index composed of these varieties increased by 14 percentage points more than the price index of varieties initially above the median. According to the 'Liberal' index, the difference in inflation between these price indices was 22 percent. Columns 3-4 and 7-8 show the price indices based on varieties that were in the top and bottom quartiles of the price distribution as of the January-October 1994 period. By October 1996, inflation was between 21 and 35 points higher, depending on the choice of the price index, for varieties in the cheapest quartile relative to the most expensive quartile. This shows that the welfare losses from exchange rate depreciations for poor households can be significantly higher due to the Within effect.

Table 2: The Within price index

	Conservative				Liberal			
	Below Median	Above Median	Quart. 1	Quart. 4	Below Median	Above Median	Quart. 1	Quart. 4
Oct. 94	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Oct. 95	1.50	1.41	1.52	1.39	1.53	1.38	1.57	1.34
Oct. 96	1.88	1.74	1.92	1.71	1.91	1.69	1.99	1.64

Note: These tables report the Within price indices defined in equation (8). The left panel reports the Conservative price indices (equation 9), while the right panel reports the Liberal price indices (equation 10). Columns labeled Below/Above Median report the price indices for consumers that buy the varieties priced above/below the median price in each product category. Columns labeled Quart. 1/4 report the price indices for consumers that buy varieties with prices in the 1/4th quartiles of the price distribution within each product category.

3.4 The Combined effect

This section computes the Combined price index, defined in equation (5) and reproduced here for convenience:

$$\widehat{P}_t^h = \sum_{g \in G} \omega_g^h \widehat{P}_{g,t}^h.$$

This index combines the two mechanisms captured by the Across and Within price indices computed above. Since we do not observe the varieties consumed by each household, we report the comparison of a hypothetical low-income and a hypothetical high-income household. The low-income household is defined as one that has across-goods expenditure shares ω_g^h of a household in the bottom income decile, and on top of that consumes the cheaper varieties within each g. The high-income household has ω_g^h's of the top income decile, and within each g consumes the more expensive varieties.

As discussed in Section 3.3, the indices $\widehat{P}_{g,t}^h$ cannot be computed for all product categories. We proceed as above, and compute the Combined price index under the two limiting assumptions from the previous section. In particular, in the conservative version there is no Within effect in categories where it cannot be directly measured:

$$\widehat{P}_t^h = \sum_{g \in G_M} \omega_g^h \widehat{P}_{g,t}^h + \sum_{g \in G_U} \omega_g^h \widehat{P}_{g,t}, \tag{11}$$

while in the liberal version the Within effect is equally strong in the unmeasured categories as it is in measured ones:

$$\widehat{P}_t^h = \sum_{g \in G_M} \omega_g^h \widehat{P}_{g,t}^h + \sum_{g \in G_U} \omega_g^h \widehat{P}_{g,t} \frac{\sum_{g \in G_M} \omega_g^h \widehat{P}_{g,t}^h}{\sum_{g \in G_M} \omega_g^h \widehat{P}_{g,t}}. \tag{12}$$

Figure 3 plots the month-to-month evolution of the Combined price index under the two alternative assumptions, computed when the high-income household consumes varieties priced above the median, and the poor household below the median within each product category. Note that the price indices for the two households are very close to each other before the October 1994 devaluation, after which they start to diverge.

Figure 3: The Combined price indices

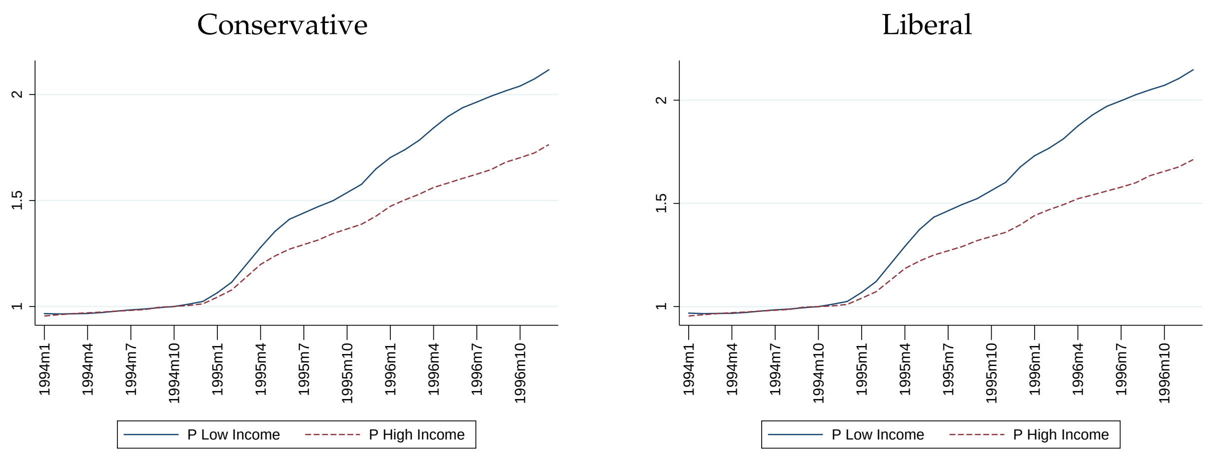

Notes: This figure plots the Combined price indices. The Conservative price indices are defined in (11), and the Liberal indices in (12). The Combined indices are depicted for consumers that buy the varieties priced above and below the median price within each product category.

The corresponding price indices are reported in Table 3. The difference in inflation faced by high- and low-income households is startling. According to the most conservative index, if we split varieties according to median prices, the change in price two years after the devaluation was 34 percentage points higher for the poorest households compared to the richest ones. Under the liberal index, inflation for the poorest households was 41 percentage points higher than for the richest households. The following subsection shows that the magnitude of these results is robust to a number of alternative assumptions used to build the price indices.

Table 3: The Combined price index

	Conservative				Liberal			
	Below Median	Above Median	Quart. 1	Quart. 4	Below Median	Above Median	Quart. 1	Quart. 4
Oct. 94	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Oct. 95	1.58	1.39	1.60	1.38	1.60	1.36	1.64	1.34
Oct. 96	2.04	1.70	2.08	1.68	2.07	1.66	2.13	1.62

Note: These tables report the Combined price indices defined in equation (5). The left panel reports the price indices under the Conservative assumptions (equation 11), while right panel reports the Liberal price indices (equation 12). Columns labeled Below/Above Median report the price indices for consumers that buy the varieties priced above/below the median price in each product category. Columns labeled Quart. 1/4 report the price indices for consumers that buy varieties with prices in the 1/4th quartiles of the price distribution within each product category.

3.5 Robustness

This section presents two sets of robustness checks. First, we provide two alternative measurements of the Within effect, in which differences in expenditure patterns across households are benchmarked to different data sources. Second, we evaluate whether differences in substitution possibilities across high- and low-income households exacerbate or dampen the welfare implications of our findings. Appendix B collects additional robustness checks, including: (i) alternative assumptions for calculating the baseline Within effect; (ii) restricting attention to consumers and prices in Mexico City; and (iii) 'placebo' experiments to show that the Within effect is not present in non-devaluation periods. Additionally, Appendix C discusses evidence based on an entirely different data source, the Economist Intelligence Unit CityData.

3.5.1 Additional measurement of expenditure shares

The main limitation of the Within price index is that variety-level expenditures by Mexican households are not directly observed. As a result, the baseline Within effect is based on hypothetical households consuming varieties above and below median in each product category. Unfortunately, there are no available data sources for variety-level expenditure over this period in Mexico. This subsection contains two exercises that adopt alternative approaches to model the within-category expenditure shares to construct the Within price indices.

Matching estimated differences in prices paid by high- and low-income households This exercise uses data from the Mexican household expenditure surveys to match varieties to households in the top vs. the bottom income decile. We proceed in three steps. First, for each household in the survey, we compute the unit values in each product category as the ratio of expenditures in the category divided by quantity consumed. Second, for each product category with available unit value data, we obtain the log difference in unit values paid by households in the highest and the lowest income deciles. Third, we combine these estimates with the DOF data and, starting from the variety that has the median price in each category find the two prices that are closest to being at a log difference corresponding to the unit value observed in the survey. Further details of unit value differences estimation are described in Appendix A.

This procedure has the advantage of being based on the actual differences in unit values paid by high- vs. low-income households in each g. As such, it captures the heterogeneity in the consumption patterns across the income distribution for different goods: there may be some g in which the high- and the low-income households consume similar unit values on average, while in other g the unit values of different households are vastly different. There are two caveats, however. First, while there are infinitely many bundles of goods that would give the same unit values, this procedure uses only two varieties per product category. Second, the expenditure survey only contains unit value data for a limited set of products, an thus we can only compute the indices for a bundle of goods that accounts for 20 percent of consumption expenditures (as opposed to 55 percent in our baseline procedure).

Appendix Table A6 reports the resulting Within price indices. The magnitude of the liberal Within effect is slightly larger than our baseline when using the above/below the median prices of the varieties. Note that the conservative Within effect is mechanically lower than in the baseline (0.05 two years after the devaluation vs. 0.13 in Table 2), since the categories for which we can compute the Within effect with this alternative methodology comprise a lower share of consumption expenditures (0.20 vs. 0.55), and the conservative calculation attributes zero Within effect to unmeasured categories.

Matching expenditure shares from US scanner data This exercise uses scanner data for the United States to compute expenditure weights in high- and low-priced varieties for households across the income distribution. In particular, we use the Nielsen HomeScan database described by Broda and Weinstein (2010) and the large literature that followed. This database contains barcode-level purchases by about 50,000 US households in 23 cities in grocery stores, drug stores, and general merchandise stores. The barcode items are divided into about 1,200 product modules, which are fairly specific.[12] We use data for 2006 (earliest year of HomeScan available to us). Within each product module, we express all prices in common units (per ounce or per item), and rank barcode-store combinations according to price.[13] We then compute the expenditure shares of high- and low-income households in the survey on expensive and cheap varieties, and use those expenditure shares to construct alternative Within price indices.

This exercise comes with a number of caveats. First, the scanner data are for a different country and time period. Second, these data only cover grocery and household merchandise, expenditure on which accounts for on average less than 7% of pre-tax household income in these data. Third, the range and reliability of household income data in HomeScan is limited. The income variable is household income 2 years prior to the year the scanner data were collected. The income information comes in ranges, with the highest income category being $200,000 or above. We compare the expenditures of households in this high-income category to the households with reported household income below $20,000. Note that the income disparity between the high- and low-income households in HomeScan is smaller than the one between the top and bottom deciles in Mexico, which was about 23-fold in 1994 (Appendix Table A5). We found that while the high-income category is reasonably homogeneous, the low-income category is highly heterogeneous and includes households that are not low-income in permanent-income terms, such as younger households and students. For these households, the fact that income is reported with a 2-year lag potentially injects substantial noise. To partly address this issue, we focus on married households with heads between 30 and 65 in our analysis.

Appendix Figure A3 plots the shares of expenditure by high- and low-income households on items that belong in each price decile within their product module. It is indeed the case that lower-income households spend disproportionately on lower-priced items, and high-income households on higher-priced items within modules. The shares are monotonic: the highest expenditure share for the high-income consumers is in the 10th price decile, and shares decline moving down deciles. On the flip side, the highest expenditure share for the low-income consumers is in the bottom decile, and shares decline moving up deciles. All in all, 78% of expenditure by high-income households is on items above the median price, and 61% of expenditure of the low-income households is on items below the median price. Given the considerations mentioned above, it

[12]For example, there are 18 different product modules of cheese, such as "Cheese, Grated," or "Cheese, Processed, Snack."

[13]This requires restricting attention to product modules in which we are confident that the items are comparable. For example, we can rank prices per ounce in product module "Tomato Puree," but not in "Frozen Novelties," and thus we use the former but not the latter product module.

is not surprising that the expenditure shares in the scanner data are less stark than our assumption that the high-income households consume only items above the median and low-income households below the median. Product categories that we normally think of as more differentiated by quality exhibit expenditure patterns very much in line with our assumption. For instance, in Men's Toiletries, Photographic Supplies (that includes cameras), and Wine, nearly all the expenditures by the low-income households is on items below the median price, and 100% of expenditure by the high-income households is on items above the median price.

We use the expenditure shares observed in the US scanner data to construct the Within effect for the Mexican devaluation. Instead of assuming that the high-income households have equal expenditure shares on all items above the median price as in the baseline, we assign to the high-income households the expenditure shares in each decile reported in Appendix Figure A3. Then, we compute the Within price index of that household by tracking prices in each price decile following the devaluation. Formally, the household-specific price change in product category g for household h is:

$$\widehat{P}^{h}_{g,t} \equiv \sum_{pdec=1}^{10} s^{h}_{pdec} \widehat{P}_{g,pdec,t}, \tag{13}$$

where s^{h}_{pdec} is the expenditure share by household h on items whose price is in price decile $pdec$, that comes from the HomeScan data and reported in Appendix Figure A3. $\widehat{P}_{g,pdec,t}$ is the average price change of items in product category g that belong to the decile $pdec$ of prices in that product category. Then, the Within effect aggregates these household- and product category-specific prices as in the baseline, equation (8).

The results are presented in the bottom panel of Appendix Table A6. The basic finding in confirmed. The cost of living inflation for the low-income households was 7 percentage points higher according to the Conservative Within price index, and 12 percentage points higher according to the Liberal one. While the magnitudes are smaller than in the baseline, this is not surprising: for the reasons outlined above the differences in expenditure patterns between high-income and low-income households in the HomeScan data are likely to be attenuated relative to what is likely the case in Mexico. Nonetheless, even when we apply these relatively modest expenditure differences to the Mexican devaluation experience, the Within effect continues to be noticeably anti-poor.

3.5.2 Differences in substitution possibilities across households

Substitution bias and the Across effect One well-known limitation of Laspeyres price indices is that they overstate how price changes affect welfare due to the substitution bias (see, e.g. Hausman, 2003). In particular, differences in the measured price index changes for high- and low-income households may not necessarily translate into differences in welfare if poor households are better able to substitute consumption across categories in response to price changes. With this in mind, we recalculate the Across price indices using expenditure weights from the 1996 household survey. The price index based on end-of-period weights is likely to understate the true welfare effects of the price changes.

The price indices under 1996 weights are reported in Appendix Tables A7a and A7b. The magnitude of the observed inflation differences between income deciles is similar to that obtained under the 1994 weights: inflation for the poorest decile is 18 percentage points higher than inflation for the richest decile. We conclude that the ability to substitute towards cheaper categories did not substantially mitigate the disparity in the welfare losses between rich and poor households arising from differences in expenditure shares across product categories.

Substitution bias and the Within effect The Within effect measured in the previous section was also computed using Laspeyres price indices, and hence subject to the substitution bias. If low-income households are better able to substitute away from high-inflation varieties than high-income households, our Within indices will overstate the distributional impact of the devaluation. Unfortunately, we cannot conduct a robustness exercise analogous to the one above for the Within effect, as we do not observe expenditure shares for the different varieties within product categories either before or after the devaluation.

To evaluate whether differences in substitution possibilities for high- vs low- income households can overturn the Within effect, we simulate changes in expenditures assuming a CES demand structure across varieties within each good and using our price data. In particular, let the share of expenditures by household h on variety v_g of good g be given by the CES functional form:

$$s^h_{v_g,t} = \frac{a^h_{v_g} p^{1-\sigma_g}_{v_g,t}}{\sum_{v'_g \in g} a^h_{v'_g} p^{1-\sigma_g}_{v'_g,t}}, \tag{14}$$

where $a^h_{v_g}$ is a taste shifter for variety v_g in household h's preferences, and σ_g is the elasticity of substitution between varieties of product category g. The preference shifters $a^h_{v_g}$ capture, in reduced form, the notion that different households prefer different varieties, perhaps in a systematic way – such as the high-income households preferring higher-quality varieties. They are treated as free parameters in this exercise, the only assumption being that they are non-time-varying.

We are interested in computing a Paasche price index that is consistent with our assumptions on the expenditure shares $s^h_{v_g,94}$ before the devaluation and with the observed changes in prices. We proceed in three steps. First, we use observed pre-devaluation prices $p_{v_g,94}$ to infer the taste shifters $a^h_{v_g}$ for each variety v_g of each product category g that are consistent with the assumption that before the devaluation high- (low-) income households put equal weight on varieties priced above (below) the median. These taste shifters are given by:

$$\frac{a^h_{v_g}}{a^h_{v'_g}} = \left[\frac{p_{v_g,94}}{p_{v'_g,94}}\right]^{\sigma_g-1} \frac{s^h_{v_g,94}}{s^h_{v'_g,94}} = \left[\frac{p_{v_g,94}}{p_{v'_g,94}}\right]^{\sigma_g-1}, \tag{15}$$

where the second equality comes from our baseline assumption that the high- (low-) income households consume all varieties v_g above (below) the median price with equal shares in 1994.

Second, we plug in the implied taste shifters and the observed prices in 1996 in equation (14) to obtain the relative shares in 1996:

$$\frac{s^h_{v_g,96}}{s^h_{v'_g,96}} = \left[\frac{p_{v_g,94}}{p_{v'_g,94}}\right]^{\sigma_g-1} \left[\frac{p_{v_g,96}}{p_{v'_g,96}}\right]^{1-\sigma_g}.$$

Using the equation above and noting that shares must add up to one, $\sum s^h_{v_g,96} = 1$, we obtain the expenditure share of each variety in 1996 as a function of the price changes and the elasticity if substitution.

$$s^h_{v_g,96} = \frac{\left[p_{v_g,96}/p_{v_g,94}\right]^{1-\sigma_g}}{\sum_{v_g}\left[p_{v_g,96}/p_{v_g,94}\right]^{1-\sigma_g}}. \tag{16}$$

Third, we use the imputed shares (16) to measure the Within price index using Paasche price indices, which capture substitution away from varieties for which inflation was high following the devaluation. Such substitution is clear in equation (16): when $\sigma_g > 1$, varieties increasing in price in relative terms will see their shares fall. Given the considerable uncertainty regarding the appropriate value of σ_g, we treat it as a free parameter ranging between 0 and 30, and assess the the sensitivity of our results to it.[14]

Appendix Figure A4 presents the results of computing the Within effect with Paasche instead of Laspeyres price indices.[15] It depicts the resulting $\widehat{P}^h_{Within,t}$ for the high- and low-income households as a function of σ_g. Using end-of-period weights unsurprisingly lowers $\widehat{P}^h_{Within,t}$ at high levels of substitution elasticity. This is intuitive: there is substantial dispersion in price changes at the variety level. Allowing agents to substitute towards varieties with the smallest price changes following the devaluation and assuming those varieties are very close substitutes mitigates the welfare impact of the increase in prices. We highlight, however, that the gap between $\widehat{P}^h_{Within,t}$ between high- and low-income households is evident at different values of σ_g. Indeed, the percentage point gap in $\widehat{P}^h_{Within,t}$ between the rich and the poor is about the same under $\sigma_g = 30$ as it is under $\sigma_g = 0$. Note that what is important for the Within effect is not whether agents substituted per se, but rather whether the high- and low-income households had differential substitution possibilities. These possibilities depend on whether the price increases were concentrated in a few varieties or broad-based across all the varieties consumed by each type of household. It turns out that while allowing for substitution between varieties affects the level of $\widehat{P}^h_{Within,t}$, is does not erase the disparity in $\widehat{P}^h_{Within,t}$ between high- and low-income households.[16]

4 Mechanisms

This section evaluates different mechanisms that may be responsible for the relative price changes underlying the indices computed in the previous section. Our analysis follows that in Burstein et al. (2005), who argue that the primary force behind the large drop in real exchange rates after large devaluations is the slow adjustment in the price of non-tradeable goods and services. Our contribution in this section is to provide new

[14] Broda and Weinstein (2010) report elasticities of substitution between product varieties in the range of 7 to 11 in barcode-level data. A potential concern is that the σ_g's may be different for high- and low-income households. There are now several sets of income-specific estimates of σ_g from scanner data that find no difference between high- and low-income households in the average level of σ_g (Handbury, 2013; Argente and Lee, 2015; Faber and Fally, 2016), so we assume that it is the same for all households.

[15] Formally, these Within indices are obtained by using the shares in equation (16) to compute the household-specific price indices $\widehat{P}^h_{g,t}$ defined in equation (6), and using the resulting $\widehat{P}^h_{g,t}$'s for the computation of the Within price indices in equations (9) and (10).

[16] Burstein et al. (2005) and Burstein et al. (2010) show that large devaluations lead to "flight from quality:" substitution from expensive towards cheaper varieties. To the extent that high-income households are better able to substitute towards cheaper varieties following a devaluation (as they start out consuming relatively more of the high-priced varieties), this type of substitution pattern within product categories should if anything amplify the anti-poor welfare effects of a devaluation.

evidence that cross-sectional heterogeneity in these dimensions can also account for differential price changes across goods and varieties, and therefore carries distributional consequences across consumers.

We first show that low-income households spend a higher fraction of their income on tradeable product categories, and among tradeables, on categories with systematically lower non-tradeable component. This together with the changes in the relative price of tradeables to non-tradeables following the devaluation provides an account of the Across effect. We then evaluate whether the leading explanations for incomplete exchange rate pass-through into retail prices are consistent with the relative price changes underlying the Within effect. We discuss the role of local distribution costs, tradeable goods that are locally produced, and variable markups in generating relative price changes within product categories.

4.1 A simple framework for understanding relative price changes

Competitive retailers combine physical goods with distribution services in fixed proportions to sell the goods to consumers. The retail price of variety v_g is given by:

$$P_{v_g,t} \quad = \quad P^T_{v_g,t} + \nu_{v_g} P^D_t, \tag{17}$$

where $P^T_{v_g,t}$, P^D_t and ν_{v_g} denote the price of the physical good, the price of distribution services, and the amount of distribution services required to provide one unit of the retail variety v_g. The proportional price change for retail variety v_g is given by

$$\widehat{P}_{v_g,t} \quad = \quad \eta_{v_g} \widehat{P}^T_{v_g,t} + \left[1 - \eta_{v_g}\right] \widehat{P}^D_t, \tag{18}$$

where $1 - \eta_{v_g} \equiv \nu_{v_g} P^D / P_{v_g}$ is the distribution margin for variety v_g. We are interested in understanding how differences pass-through into retail prices affect consumers differentially across the income distribution. In what follows, we assume that distribution services are purely non-tradeable, so that $\widehat{P}^D_t = \widehat{P}^N_t$, where P^N_t is the price of non-tradeable goods. We also assume that the price of the tradeable goods at the dock or at the factory gate relative to the price of non-tradeables moves in proportion to the exchange rate – $\widehat{P}^T_{v_g,t} - \widehat{P}^N_t = \alpha_{v_g} \widehat{E}_t$, where $\alpha_{v_g} \geq 0$. The parameter α_{v_g} captures in a reduced form the fact that pass-through into prices at the dock can be incomplete and can differ across varieties. We discuss different sources of incomplete pass-through into border prices below. Combining these assumptions, equation (18) becomes:

$$\widehat{P}_{v_g,t} = \widehat{P}^N_t + \eta_{v_g,t-1}\alpha_{v_g}\widehat{E}_t. \tag{19}$$

Aggregating up to the good category, the change in the price index for category g, $\widehat{P}_{g,t} \equiv \frac{1}{V_g}\sum_{v_g \in g}\widehat{P}_{v_g,t}$, is given by:

$$\widehat{P}_{g,t} = \widehat{P}^N_t + \eta_g\alpha_g\widehat{E}_t + cov_v\left(\eta_{v_g}, \alpha_{v_g}\right)\widehat{E}_t; \tag{20}$$

where $1-\eta_g \equiv 1 - \frac{1}{V_g}\sum_{v_g \in g}\eta_{v_g}$ is the average share of distribution services among varieties of g, α_g captures the average pass-though in category g, and $cov_v\left(\eta_{v_g}, \alpha_{v_g}\right)$ is the covariance between the distribution margins and pass-through into border prices within product category g. In what follows, we ignore this covariance and focus on the first order terms.

Equations (19) and (20) relate changes in retail prices following a devaluation to local distribution margins and pass-through into border prices. They state that varieties and product categories for which distribution margins are high and pass-through into border prices is low will experience smaller proportional price changes. To the extent that expenditure patterns across the income distribution are systematically related to these product characteristics, large devaluations will have distributional consequences.

Differences in pass-through into border prices, captured by the parameter α_{v_g}, can be driven by multiple factors, including differences in markup changes across varieties. In what follows, we focus on one dimension of heterogeneity in α_{v_g} across goods: the distinction between goods produced purely for local consumption and goods that are actually traded internationally. We focus on this dimension because it has played a prominent role in the literature on large devaluations and because it is one dimension that we can measure in the data (see, e.g. Burstein et al., 2005). Appendix D lays out a complete accounting framework in which price changes are also affected by changes in markups following Burstein and Gopinath (2015), to illustrate where variable markups can potentially enter, and reviews the available literature on their role. Importantly, the exercises below are still valid in the presence of variable markups.

4.2 Understanding the Across effect

Our explanation for the Across effect relies on two premises: (i) the differences in the nontradeable component of different product categories explain the good-level price changes following the devaluation; and (ii) there is a systematic relationship between the nontradeable component and expenditure shares of high- and low-income households: the poor have higher effective expenditure shares in tradeables. We now provide empirical evidence on each of these in turn.

4.2.1 Distribution margins, local goods, and price changes

This section shows how the observed price changes following the devaluation are related to differences in distribution costs and the share of local goods across product categories. With that in mind, we assume that there are two types of tradeable goods: those that are produced purely for local consumption and those that are actually traded internationally. Under these assumptions, equation (20) can be written as:

$$\widehat{P}_{g,t} = \widehat{P}_t^N + \eta_g \alpha_{loc} \widehat{E}_t + \eta_g \theta_g \left(\alpha_{int} - \alpha_{loc}\right) \widehat{E}_t, \tag{21}$$

where α_{int} and α_{loc} control the pass-through into border prices for internationally traded and local goods respectively, and θ_g is the share of internationally traded goods in product category g. Note that to the extent that $\alpha_{int} > \alpha_{loc}$, pass-through will be higher for internationally traded goods.[17]

Distribution margins and price changes Figure 1 has already documented that the relative price of tradeables to non-tradeables increased following the devaluation. We now show that among the categories classified as tradeables, the prices of goods with higher distribution margins increased by less. To take equation (20) to the data, however, we need to know the distribution margins for disaggregated product categories. Unfortunately, these data are not available for Mexico for a period close to the 1994 devaluation. Thus, we focus on retail margins from the 2004 Mexican Retail Census. The underlying assumption behind the exercise is that the variation in distribution margins across product categories is at least partly technologically determined, and thus the 2004 data are informative of the cross-category variation in distribution margins in 1994. To the extent this measure provides a noisy indicator of Mexican distribution margins in 1994, the noise will likely bias us towards finding no patterns in the data.

We define the retail margin as the ratio of the retail price to the cost of the merchandise that is purchased in order to sell at the retail establishment. The Retail Census reports this information by store types. We match these store categories by hand to the product categories in the Mexican consumer price data. The store types and the resulting matches

[17] Burstein et al. (2005) show that a model in which the prices of local goods move like non-tradeable prices following a large devaluation ($\alpha_{loc} \approx 0$), while pass-through into traded prices is almost complete ($\alpha_{int} \approx 1$) can successfully account for the aggregate pass-through after large devaluations in a panel of countries.

are reported in Appendix Table A12. According to these data, the distribution margins range from about 0.15 to about 0.82 across products, with the mean of 0.45 and the median of 0.44. Appendix Table A13 reports the 5 categories with the lowest and highest distribution margins in our data.

Figure 4 reports the scatterplot of the good-level price changes $\widehat{P}_{g,t}$ following the devaluation (the change from October 1994 to October 1996) against the one minus the distribution margin η_g as in (20). Each dot represents a tradeable product category. There

Figure 4: Price changes and distribution margins

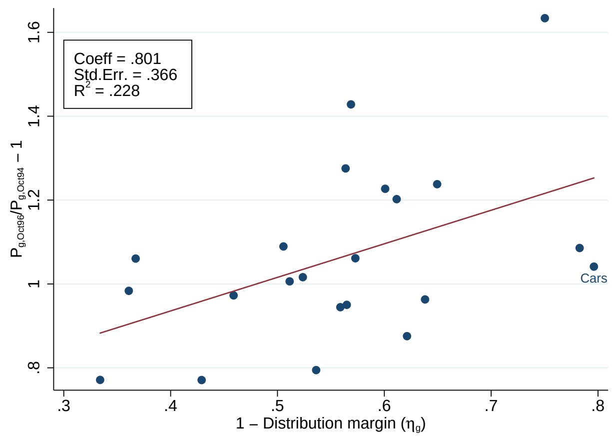

Note: This figure presents the scatterplot of the price change in each good against one minus the distribution margin (η_g) together with an OLS fit following the 1994 Mexican devaluation. The box in the top left corner reports the coefficient, robust standard error, and the R^2 in that bivariate regression.

is a positive and statistically significant relationship between these variables: the product categories with lower distribution margins experienced larger price increases, exactly as implied by (20). In spite of the fact that our data on distribution margins come from the 2004 Census, the relationship is strongly significant, and the R^2 in this bivariate regression is 0.23.

To establish more firmly that this pattern is due to the devaluation, Appendix Figure A5 plots the same relationship in two placebo periods: one immediately pre-devaluation and one in the mid-2000s. The picture is very different, with the point estimates for the slope of the relation negative for the pre-devaluation period, and close to zero and insignificant in the mid-2000s.

Local goods and price changes We now evaluate whether among tradeables, prices of product categories with a higher share of local goods increased by less. It is difficult to quantify the share of local goods in each category g. We use two alternative proxies for the importance of local goods. First, we calculate the import content of absorption in each category g, that is we set $\theta_g = M_g / [Y_g + M_g - X_g]$, where Y_g, M_g, and X_g denote production, imports, and exports in category g respectively. This measure is a lower bound on the share of pure tradeable goods, as it does not count goods produced and consumed in Mexico but that are also exportable. Hence, the second measure is openness at the sector level relative to production and imports, that is: $\theta_g = [M_g + X_g] / [Y_g + M_g]$. Imports, exports, and production data for sufficiently disaggregated sectors that can be mapped intro the DOF categories are not available in input-output matrices. For this reason, we compute proxies for θ_g from the UN Food and Agricultural Organization's FAOSTAT database, that reports imports, exports, and production quantities and values for 60 agricultural products in 1994 in Mexico. Appendix Table A14 reports the matches between Mexican CPI categories and items in FAOSTAT, the two measures of θ_g, and the differences in consumption shares in each category between the top and the bottom income deciles. These categories combined represent nearly 15% of total consumption expenditure in Mexico in 1994.

Figure 5 reports the scatterplot of the product-level price changes $\widehat{P}_{g,t}$ following the devaluation (the change from October 1994 to October 1996) against the one minus the share of purely traded goods, θ_g as in (21). Each dot represents a tradeable product category. There is a positive relation between the share of pure traded goods and the observed price changes during the devaluation. The relationship is strongly significant under our two alternative measures for the share of pure traded goods. Appendix Figure A6 reports the scatterplots for two placebo periods, and shows that the positive relationship does not hold absent a large devaluation.

4.2.2 Distribution margins, local goods and consumption patterns

We now evaluate how expenditure shares across product categories are related to observed distribution margins and the share of local goods in each category. Combining (7) and (21), the Across price index for household h following a devaluation can be written as:

$$\widehat{P}^h_{Across,t} = \widehat{P}^N_t + \omega^h_T \left[\alpha_{loc} \sum_{g \in G} \widetilde{\omega}^h_g \eta_g + [\alpha_{Int} - \alpha_{loc}] \sum_{g \in G} \widetilde{\omega}^h_g \eta_g \theta_g \right] \widehat{E}_t. \qquad (22)$$

Here, $\omega^h_T \equiv \sum_{g \in T} \omega^h_g$ denotes the share of tradeable goods consumed by household h, and $\widetilde{\omega}^h_g \equiv \sum_{g \in T} \frac{\omega^h_g}{\sum_{g \in T} \omega^h_g}$ denotes h's share of spending on tradeable category g in total tradeables expenditure.

Figure 5: Price changes and share of local goods

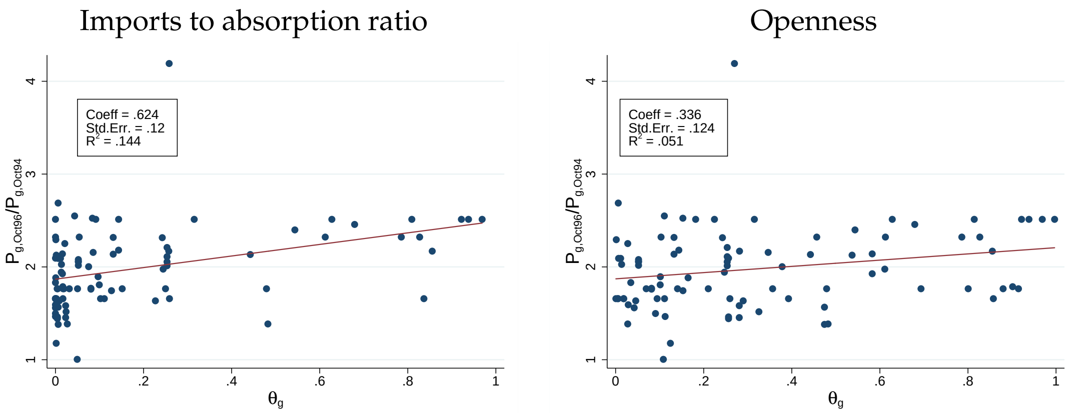

Note: This figure presents the scatterplots of the price change in each good against one minus the share of local goods in each product category (θ_g) together with an OLS fit following the 1994 Mexican devaluation. The box in the top left corner reports the coefficient, robust standard error, and the R^2 in that bivariate regression. 'Imports to absorption ratio' refers to θ_g proxied by $\theta_g = M_g / \left[Y_g + M_g - X_g\right]$. 'Openness' refers to θ_g proxied by $\theta_g = \left[M_g + X_g\right] / \left[Y_g + M_g\right]$.

According to equation (22), differences in the changes in the Across price index across households are driven by: i) the share of expenditure on tradeable product categories, ω_T^h, and ii) expenditure shares across tradeable product categories with different distribution margins and local goods shares $\sum_{g \in G} \widetilde{\omega}_g^h \eta_g$ and $\sum_{g \in T} \widetilde{\omega}_g^h \eta_g \theta_g$. To the extent that the poor consume relatively more of the tradeable categories, $\omega_T^{poor} > \omega_T^{rich}$, the across price index will rise more for the poor. In addition, if the poor consume tradeables with low distribution margins $\sum_{g \in G} \widetilde{\omega}_g^{poor} \eta_g > \sum_{g \in G} \widetilde{\omega}_g^{rich} \eta_g$ and low local goods shares, $\sum_{g \in T} \widetilde{\omega}_g^{poor} \eta_g \theta_g > \sum_{g \in T} \widetilde{\omega}_g^{rich} \eta_g \theta_g$, the Across price index will rise more for the poor.[18] In what follows, we combine the expenditure data from the 1994 Mexican household survey with the sectoral values for η_g and θ_g computed in the previous subsection to study this relation.

First, we show that the poor do indeed have higher expenditure shares on tradeable categories: $\omega_T^{poor} > \omega_T^{rich}$. We sort households into income deciles and compute the expenditure shares of each decile in tradeable and non-tradeable goods.[19] The results are depicted in Figure 6a. Expenditure shares on tradeable goods decrease monotonically as we move up the income distribution. The difference is quantitatively large: the bottom decile's tradeable expenditure share is 0.58, compared to 0.4 for the top decile. Appendix Table A5 reports income-specific expenditure shares across broad consumption

[18]To see this, note that $\alpha_{loc} \geq 0$ and $\alpha_{Int} \geq \alpha_{loc}$ in equation (22).

[19]Appendix Table A4 classifies the consumption categories in the Mexican CPI the into tradeables and non-tradeables (source: Bank of Mexico).

categories. The largest differences are in the Food, Beverages, and Tobacco and Education categories (the expenditure shares of 42% for households at the bottom income decile vs. 11% for households at the top in Food, and of 3% for the bottom decile vs. 15% for the top decile in Education). Higher-income households also have larger expenditure shares in housing, which is partly accounted for by the fact that the imputed expenditure shares in 'owner-occupied housing' are larger for the richer households. Note however that this does not account for the bulk of the expenditure differences across the income distribution.

Figure 6: Expenditures by income decile

(a) Tradeable share of expenditures

(b) Distribution shares

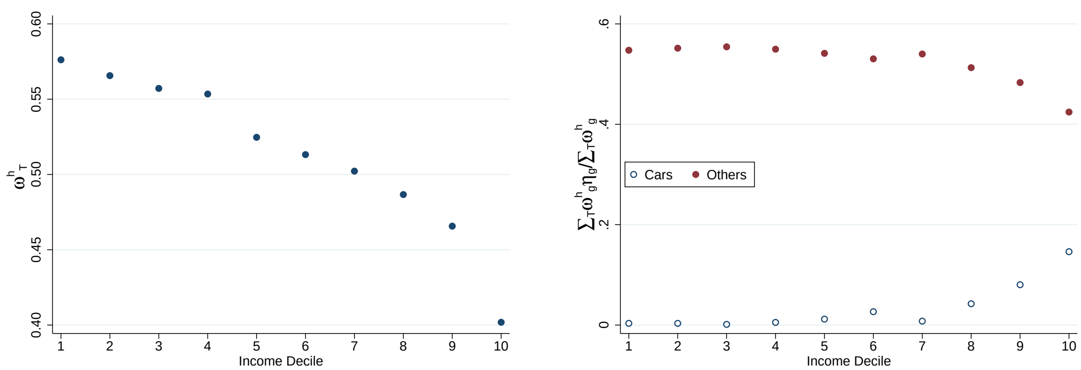

Note: Figure 6a plots the expenditure share of tradeables by income decile in the 1994 ENIGH household survey. Figure 6b plots one minus the distribution margin expenditure share for tradeables, $\sum_{g \in T} \widetilde{\omega}_g^h \eta_g$, by income decile in the 1994 ENIGH household survey.

Second, we establish whether among tradeables, the poor exhibit higher expenditure shares in categories with low distribution margins and a low share of local goods. Because the distribution margins and local goods shares come from different data sources, we cannot compute distribution margins and local goods shares at the same level of disaggregation. To evaluate these two margins in isolation, we proceed in two steps. First, we assume that there are no differences in local goods across product categories ($\theta_g = \bar{\theta}$), and evaluate how $\sum_{g \in G} \widetilde{\omega}_g^h \eta_g$ varies across households. Second, we assume instead that there are no differences in distribution margins across product categories ($\eta_g = \bar{\eta}$), and evaluate how $\sum_{g \in G} \widetilde{\omega}_g^h \theta_g$ varies across households.

Distribution margins and consumption patterns Figure 6b reports one minus the local distribution margin for tradeable expenditure, $\sum_{g \in T} \widetilde{\omega}_g^h \eta_g$, by income decile. In categories other than cars, the pattern is clear. Expenditure-weighted tradeable content falls as income increases. Even restricting attention to tradeables, high-income households have higher effective non-tradeable shares, as they consume more in categories with higher distribution margins. The difference is substantial, falling from about 0.55 to 0.42 between the bottom and top deciles.

Cars is an expenditure category that does not fit this pattern. According to the Retail Census data, cars have a lower than average distribution margin, but are consumed disproportionately more by those at the top of the income distribution. Interestingly, however, Figure 4 shows that for cars the increase in the price was low relative to what would be predicted by their low retail margins. Thus, even though cars are a low-distribution margin good consumed disproportionately more by high-income households, they do not eliminate the substantial Across effect found in the data.

Local goods and consumption patterns We now evaluate how expenditure shares across product categories are related to observed local goods shares. The categories for which θ_g can be computed in FAOSTAT is only a subset of the T tradeable categories. Thus we report results for the weighted share of local goods in the FAOSTAT categories, that is, instead of $\sum_{g \in T} \widetilde{\omega}_g^h \theta_g$ we compute $\sum_{g \in F} \frac{\omega_g^h}{\sum_{g \in F} \omega_g^h} \theta_g$, where F is the set of tradeable goods for which the FAO data are available.

The results are depicted in Figure 7. Expenditure shares on local goods decrease modestly as we move up the income distribution. The bottom decile's expenditure share in pure traded goods is between one and two percentage points higher in the bottom decile than in the top decile. Appendix Table A5 reports the differences in income-specific expenditure shares across broad consumption categories between the top and the bottom income deciles. The largest differences are in the Meat and Milk categories, where the expenditure shares of the top decile are 14 and 7.5 percentage points higher than of the bottom decile, and in Maize and Beans, for which the bottom decile expenditure shares are 11-13 percentage points higher than the top decile shares.

All in all, there is more support in the data for the role of distribution margins than local goods in generating the Across effect. While both the distribution margin and local good differences predict correctly the cross-section of price changes following the devaluation, we find at best weak evidence that consumption baskets of lower-income households are significantly skewed towards categories with more pure traded goods.

Figure 7: Tradeable share of expenditures by income decile

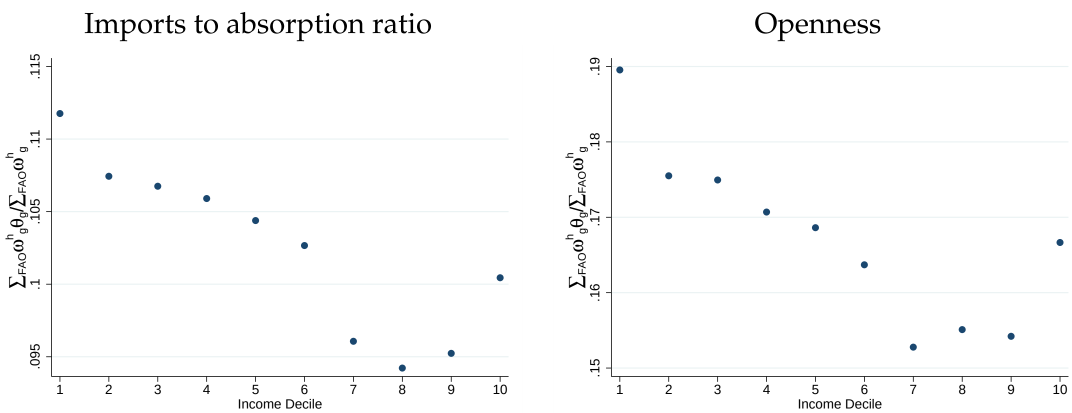

Note: This figure plots the expenditure the share of local goods in each product category (θ_g) by income decile in the 1994 ENIGH household survey. 'Imports to absorption ratio' refers to θ_g proxied by $\theta_g = M_g / \left[Y_g + M_g - X_g\right]$. 'Openness' refers to θ_g proxied by $\theta_g = \left[M_g + X_g\right] / \left[Y_g + M_g\right]$.

4.3 Understanding the Within effect

4.3.1 Distribution margins and the Within effect

Differences in distribution margins within product categories can lead to a Within effect if (i) the relative price of varieties with low distribution margins increased following the devaluation; and (ii) the poor tend to consume varieties with lower distribution margins.

We first assess whether differences in distribution margins can rationalize the observed variation in price changes across varieties within product categories post-devaluation. Equation (19) implies that the difference between the price change of any variety v_g and the change in the average price in category g is given by:

$$\widehat{P}_{v_g,t} - \widehat{P}_{g,t} = \left(\frac{\eta_{v_g,t-1}\frac{\alpha_{v_g}}{\alpha_g} - \eta_{g,t-1}}{\eta_{g,t-1}} \right) \times \eta_{g,t-1}\alpha_g \widehat{E}_t. \tag{23}$$

Equation (23) is the theoretical prediction for variety-level price changes following the devaluation. It states that prices will increase proportionately more for varieties that have low distribution margins (high $\eta_{v_g,t-1}$), and that have higher pass-through into border prices, $\alpha_{v_g} > \alpha_g$. Note that we observe the left-hand side of (23) directly. If we could find proxies for the variation in distribution margins and pass-through into border prices $\left(\frac{\eta_{v_g,t-1}\frac{\alpha_{v_g}}{\alpha_g} - \eta_{g,t-1}}{\bar{\eta}_{g,t-1}} \right)$ and average exchange rate pass-through into retail prices $\eta_{g,t-1}\alpha_g \widehat{E}_t$, we could evaluate this theoretical prediction empirically.

An important challenge in taking (23) to the data is that differences in distribution margins and tradeability across varieties of the same g are not directly observed. We circumvent this challenge by focusing on subsets of products g that are composed of identical physical goods sold in different retail outlets. Restricting attention to identical physical goods implies that their pass-through into border prices is identical: $\alpha_{v_g} = \alpha_g$. As a result, we can then infer differences in distribution margins from differences in their observed prices. To implement this approach, we manually parse verbal product descriptions, and classify goods as being "the same product" if they have an identical verbal description and weight. To ensure that we are grouping identical products, we impose two additional constraints. First, the product description must contain a brand name, and thus we exclude products whose descriptions only contain product characteristics – for instance a type of cut of meat – but do not contain brand names. Second, we limit the sample to goods that have prices quoted in kilos or liters. The resulting sample consists of 1297 products that have identical product descriptions (e.g. "Corn Flour, Maseca, Bag of 1 KG"), spread over 79 product categories (e.g. "Corn Flour").

For this subset of products, (23) simplifies to:

$$\widehat{P}_{v_g,t} - \widehat{\bar{P}}_{g,t} = \left(\frac{\eta_{v_g,t-1} - \bar{\eta}_{g,t-1}}{\bar{\eta}_{g,t-1}}\right) \times \bar{\eta}_{g,t-1}\bar{\alpha}_g \widehat{E}_t, \tag{24}$$

and we can use equation (17) to infer differences in distribution margins from observed price differences:

$$\frac{\eta_{v_g,t-1} - \bar{\eta}_{g,t-1}}{\bar{\eta}_{g,t-1}} = \frac{\bar{P}_{g,t-1} - P_{v_g,t-1}}{P_{v_g,t-1}}. \tag{25}$$

In these expressions, the bars denote the averages among only the identical products within each g.[20]

We then assume that distribution costs and changes in exchange rates do account for observed changes in *average* prices (Burstein et al., 2005), and calibrate $\eta_{g,t-1}\alpha_g \widehat{E}_t$ to match the observed changes in average prices in each category. That is, using equation (18) we match $\eta_{g,t-1}\alpha_g \widehat{E}_t = \widehat{\bar{P}}_{g,t} - \widehat{P}^N_t$.

Based on these two proxies, we compute predicted price changes in the two years following the devaluation for individual varieties using equation (23). The first column of Table 4 reports the results of a linear regression of actual price changes on the predicted price changes. The estimated coefficient is close to 1 and strongly significant. The R^2 is

[20] Appendix D.1 derives the model prediction in the presence of multiplicative retail markups, and shows that our approach of proxying distribution margin differences with proportional price differences is valid when retail markups are the same across varieties, or more generally as long as the differences in retail markups are not too negatively correlated with differences in distribution margins across stores (so that the most expensive stores are not the ones that have lower distribution margins).

equal to 0.135, which means that relying on inferred distribution margins alone we can account for almost one-sixth of the variation in the observed price changes. Appendix Figure A7 plots the observed vs. the predicted price changes across identical products sold in different outlets in the two years following the devaluation. A strong positive relation between the predicted and the observed price changes is evident. We conclude that differences in distribution margins across retailers can indeed explain a significant fraction of the observed variance in price changes following the devaluation.

Finally, the relation between observed price changes and differences in distribution margins is nonexistent in non-devaluation periods. We recompute predicted price changes for two alternative periods in which the nominal exchange rate is roughly constant: i) The January 1994 – October 1994 period, which is the longest time period before the devaluation for which we have variety-level price data, and ii) the January 2004 – January 2006 period. We compare the observed vs. predicted price changes in Appendix Figure A8, and report the estimated coefficients in the last two columns of Table 4. It is clear from the figures that differences in distribution margins do not have explanatory power for differences in price changes in the absence of large exchange rate movements.

Table 4: Predicted vs. observed price changes

	Devaluation: Oct94 – Oct96	Placebo I: Jan94 – Oct94	Placebo II: Jan04 – Jan06
Slope	1.426***	0.161	-0.0865*
	(0.282)	(0.110)	(0.0519)
Observations	5,079	5,084	5,742
R^2	0.135	0.002	0.003

Notes: ***: significant at the 1% level; *: significant at the 10% level. This table reports the results of estimating equation (23) for the devaluation period (first column) and two placebo periods. The prices are for identical goods sold in different stores.

Distribution margins and consumption patterns It remains to link consumption of varieties with different distribution margins to income. Appendix A provides robust empirical evidence that poorer households consume lower-priced varieties. We show above that at least for varieties of identical physical goods, distribution margins are low for the cheaper varieties (see equation 25). Appendix C.2 provides some direct evidence to support this claim based on an alternative data source, the Economist Intelligence Unit

CityData.[21]

A recent paper by Atkin et al. (2016) uses a rich collection of barcode, store, and household-level data in Mexico over 2011-2014 to show that (i) products with identical barcodes are 12% cheaper in foreign-owned stores compared to domestically-owned stores; and (ii) higher-income households spend a higher fraction of their retail expenditure in foreign stores. How are these observations reconciled with the evidence in Table A1 that the poor pay lower prices within product categories? First, Atkin et al. (2016) also show that similar but not identical products are actually more expensive in foreign-owned stores, presumably because they are of higher quality. Since richer households tend to buy higher-quality varieties, this is consistent with the observation that higher-priced varieties are consumed by the high-income households. Second, even for identical (barcode-level) products the analysis in Atkin et al. (2016) does not establish that the poor actually pay more than the rich. Their estimated coefficient reflects the average price difference between all foreign- and non-foreign-owned stores. It does not rule out the possibility that both sets of stores are highly heterogeneous and that the poor shop in particularly cheap domestically-owned stores, and/or that they buy from foreign-owned stores the goods that are cheaper in those stores.

4.3.2 Local goods and other explanations

In contrast to our findings across food categories in FAO data, a common conjecture is that within categories low-income households consume local goods, whereas the high-income households consume imported goods. If the local goods increase in price by less than imported goods following the devaluation, the resulting Within effect will be pro-poor. Note that our Within effect exercise assumes only that the poor consume the lower-priced varieties in each product category. If those lower-priced varieties are also – plausibly – local goods, our Within effect would capture this difference in consumption baskets across the income distribution. The fact that our Within effect is still anti-poor suggests that the imported vs. local goods distinction is not the main driver of the Within effect.

The Within effect establishes that the more expensive varieties within the same product categories experienced smaller price increases following the devaluation. If the more

[21] A recent paper by Jaimovich et al. (2015) shows that in the US low-end retail establishments – where lower-income households are more likely to shop – are less labor-intensive, and thus likely to exhibit relatively lower retail value added. We acknowledge that this US-based evidence is at best suggestive for our purposes. As documented by Lagakos (2016), the retail sector looks very different in Mexico compared to the US. In addition, distribution margins include services of other factors such as capital and materials inputs, and it is not clear how different types of retail outlets differ in their intensity of the use of those other factors.

expensive varieties represent higher quality, an explanation for this fact could be that higher-quality products have lower exchange rate pass-through at the border α_{v_g}. Several recent papers document this type of effect. Auer et al. (2014) propose a model of variable markups in which low exchange rate pass-through into high quality goods arises endogenously as a result of vertical differentiation, and demonstrate that higher-quality products have lower pass-through using detailed data on car sales in several European countries. Antoniades and Zaniboni (2015) use barcode-level data from several retailers in the UAE to show empirically that pass-through into retail prices is indeed lower for high quality goods. Chen and Juvenal (2016) use bottle-level data for Argentina's wine exports to show that pass-through is lower for higher-quality wine. In our own data, exchange rate pass-through following the Mexican devaluation was indeed lower for higher-priced than for lower-priced varieties of the same product (results not reported in order to conserve space, but available upon request). Appendix C.3 provides additional evidence of this finding using price data for several devaluation episodes from the Economist Intelligence Unit.

5 Conclusion

Large exchange rate devaluations affect the prices faced by high- and low-income households differentially. Using the 1994 Mexican peso devaluation, we show that the distributional consequences can be large. In the two years following the devaluation, inflation of the consumption basket of those in the bottom decile of the income distribution was between 32 and 39 percentage points higher than for the basket of those in the top decile. Differences in price changes within narrow product categories account for about half of this difference.

We explore in detail one possible explanation for this result: the poor consume fewer non-tradeable goods. This manifests itself at all levels of product aggregation. Poorer households tend to spend a larger overall share of their income on tradeables. Across tradeable categories, the poor have higher expenditure shares in products with systematically lower distribution margins. Finally, within detailed product categories, the poor consume lower-priced varieties that contain relatively less domestic value added. Correspondingly, prices of goods with a smaller non-tradeable component rise more following a devaluation, leading to anti-poor distributional consequences. Another plausible mechanism that can drive the Within effect is differences in markup elasticities with respect to exchange rate changes between higher- and lower-quality goods. The systematic consumption basket differences we identify are likely to occur in other countries and time periods, and thus the results for Mexico may be informative of the effects of other devaluations.

References

Ahlin, Christian and Mototsugu Shintani, "Menu costs and Markov inflation: A theoretical revision with new evidence," *Journal of Monetary Economics*, April 2007, *54* (3), 753–784.

Almås, Ingvild, "International Income Inequality: Measuring PPP Bias by Estimating Engel Curves for Food," *American Economic Review*, 2012, *102* (2), 1093–1117.

Antoniades, Alexis and Nicola Zaniboni, "Exchange Rate Pass-Through into Retail Prices," June 2015. forthcoming, *International Economic Review*.

Argente, David and Munseob Lee, "Cost of Living Inequality during the Great Recession," June 2015. Kilts Booth Marketing Series Paper 1-032.

Atkin, David, Benjamin Faber, and Marco Gonzalez-Navarro, "Retail Globalization and Household Welfare: Evidence from Mexico," May 2016. forthcoming, *Journal of Political Economy*.

Auer, Raphael, Thomas Chaney, and Philip Sauré, "Quality Pricing-to-Market," May 2014. Mimeo, Swiss National Bank and Toulouse School of Economics.

Berger, David, Jon Faust, John H. Rogers, and Kai Steverson, "Border prices and retail prices," *Journal of International Economics*, 2012, *88* (1), 62–73.

Broda, Christian and David E. Weinstein, "Product Creation and Destruction: Evidence and Price Implications," *American Economic Review*, June 2010, *100* (3), 691–723.

Burstein, Ariel T. and Gita Gopinath, "International Prices and Exchange Rates," in Kenneth Rogoff Elhanan Helpman and Gita Gopinath, eds., *Handbook of International Economics*, Vol. 4, Elsevier, 2015, chapter 7, pp. 391 – 451.

_ , Martin Eichenbaum, and Sergio Rebelo, "Large Devaluations and the Real Exchange Rate," *Journal of Political Economy*, August 2005, *113* (4), 742–784.

_ , Nir Jaimovich, and Pablo Andrés Neumeyer, "Consumer behavior during the 2002 Argentine crisis: a macroeconomic analysis with microeconomic data," September 2010. Mimeo, UCLA, Stanford, and Universidad Torcuato Di Tella.

Chen, Natalie and Luciana Juvenal, "Quality, Trade, and Exchange Rate Pass-Through," *Journal of International Economics*, May 2016, *100*, 61–80.

de Carvalho Filho, Irineu and Marcos Chamon, "A Micro-Empirical Foundation for the Political Economy of Exchange Rate Populism," *IMF Staff Papers*, 2008, *55* (3), 481–510.

Engel, Ernst, "Die Produktions- und Ernteerträge und der Getreidehandel im preussischen Staate," *Zeitschrift des Königlichen preussischen statischen Bureaus*, 1857, *2*, 249–89.

_ , "Das Lebenskosten belgischer Arbeiterfamilien früher und jetzt," *Bulletin de Institut International de Statistique*, 1895, *9*, 1–124.

Faber, Benjamin, "Trade Liberalization, the Price of Quality, and Inequality: Evidence from Mexican Store Prices," July 2014. mimeo, UC Berkeley.

_ **and Thibault Fally**, "Firm Heterogeneity in Consumption Baskets: Evidence from Home and Store Scanner Data," September 2016. mimeo, UC Berkeley.

Fajgelbaum, Pablo D. and Amit K. Khandelwal, "Measuring the Unequal Gains from Trade," *Quarterly Journal of Economics*, August 2016, *131* (3), 1113–1180.

Fajgelbaum, Pablo, Gene M. Grossman, and Elhanan Helpman, "Income Distribution, Product Quality, and International Trade," *Journal of Political Economy*, August 2011, *119* (4), 721–765.

Friedman, Jed and James Levinsohn, "The Distributional Impacts of Indonesia's Financial Crisis on Household Welfare: A 'Rapid Response' Methodology," *World Bank Economic Review*, December 2002, *16* (3), 397–423.

Gagnon, Etienne, "Price Setting During Low and High Inflation: Evidence from Mexico," *The Quarterly Journal of Economics*, August 2009, *124* (3), 1221–1263.

Goldberg, Linda S. and José Manuel Campa, "The Sensitivity of the CPI to Exchange Rates: Distribution Margins, Imported Inputs, and Trade Exposure," *Review of Economics and Statistics*, May 2010, *92* (2), 392–407.

Halac, Marina and Sergio L. Schmukler, "Distributional Effects of Crises: The Financial Channel," *Economia*, Fall 2004, *5* (1), 1–67.

Handbury, Jessie, "Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living Across U.S. Cities," July 2013. mimeo, Wharton.

Hausman, Jerry, "Sources of Bias and Solutions to Bias in the CPI," *Journal of Economic Perspectives*, 2003, *17* (1), 23–44.

Jaimovich, Nir, Sergio Rebelo, and Arlene Wong, "Trading Down and the Business Cycle," September 2015. NBER Working Paper No. 21539.

Kraay, Aart, "The welfare effects of a large depreciation: the case of Egypt, 2000-05," in Hanaa Kheir-El-Din, ed., *The Egyptian Economy: Current Challenges and Future Prospects*, Cairo: The American University in Cairo Press, April 2008.

Lagakos, David, "Explaining Cross-Country Productivity Differences in Retail Trade," *Journal of Political Economy*, April 2016, *124* (2), 579–620.

Levinsohn, James A., Steven T. Berry, and Jed Friedman, "Impacts of the Indonesian Economic Crisis. Price Changes and the Poor," in Michael P. Dooley and Jeffrey A. Frankel, eds., *Managing Currency Crises in Emerging Markets*, NBER Chapters, National Bureau of Economic Research, June 2003, pp. 393–428.

Lopez-Acevedo, Gladys and Angel Salinas, "How Mexico's Financial Crisis Affected Income Distribution," 2000. mimeo, The World Bank.

Maloney, William F., Wendy V. Cunningham, and Mariano Bosch, "The Distribution of Income Shocks during Crises: An Application of Quantile Analysis to Mexico, 1992–95," *World Bank Economic Review*, 2004, *18* (2), 155–174.

Porto, Guido G., "Using survey data to assess the distributional effects of trade policy," *Journal of International Economics*, September 2006, *70* (1), 140–160.

Ravallion, Martin, Shaohua Chen, and Prem Sangraula, "New Evidence on the Urbanization of Global Poverty," *Population and Development Review*, December 2007, *33* (4), 667–701.

Schneider, Friedrich and Dominik Enste, "Shadow Economies: Size, Causes, and Consequences," *Journal of Economic Literature*, March 2000, *38*, 77–114.

Verhoogen, Eric, "Trade, Quality Upgrading and Wage Inequality in the Mexican Manufacturing Sector," *Quarterly Journal of Economics*, May 2008, *123* (2), 489–530.