The Income Elasticity of Import Demand: Micro Evidence and An Application

David Hummels*
Purdue University, NBER

Kwan Yong Lee†
University of North Dakota

April 2018

* Department of Economics, Krannert School of Management, Purdue University, West Lafayette, IN 49707-1310; email: hummelsd@purdue.edu

† Corresponding Author: Department of Economics and Finance, College of Business and Public Administration, University of North Dakota, ND 58202-8369; Email: kwanyong.lee@business.und.edu; Phone: 1-701-777-2637; Fax: 1-701-777-3365

The Income Elasticity of Import Demand: Micro Evidence and An Application

April 2018

Abstract

We construct a synthetic panel of household expenditures from the Consumer Expenditure Survey (CEX) and use the Quadratic Almost Ideal Demand System to estimate expenditure shares and income elasticities of demand that vary by good-time-income. We show that the size and distribution of income shocks drives expenditure change in a manner that varies profoundly across traded goods. Our estimates of expenditure shares and income elasticities could be useful in many applications that seek to explain changes in trade behavior from the demand side, and indicate the strong sensitivity of trade to changes in the tails of the income distribution. We explore an application involving the Great Trade Collapse. Income-induced expenditure changes are positively correlated with the cross-good pattern of import changes, generating a predicted change as much as 40% as large as the raw variation in import declines.

JEL Classification: D12; D31; F10; F14

Keywords: Income Elasticity; Income Distribution; Import Demand; Great Trade Collapse

1. Introduction

After an extended lull, a recent literature has begun to re-emphasize the importance of non-homothetic preferences for explaining patterns of trade. These papers have tended to emphasize forms of non-homothetic preferences that permit relatively easy aggregation of demands over income levels. As such, their estimation involves minimal data requirements, and they are ideal for incorporating into general equilibrium theories and evaluating welfare consequences of trade.

We pursue a different approach to understanding the role of income effects in import demand, using household expenditure data from the US to estimate a parametrically rich non-homothetic demand system. We recover income elasticities of demand for traded goods that are good-income-time varying. Combining this with information on the share of good expenditures at different income levels, we show that the size and distribution of income shocks drives expenditure change in a manner that varies profoundly across traded goods. These estimates could be useful in many applications that seek to explain changes in trade behavior from the demand side. That is, they provide an extension of the classic demand curve instrument – income – by allowing the distribution of income changes hitting a country to differentially affect consumption and import demand for each good and time period. To show this, we explore an application in which we explain changes in import demand over a period that includes the Great Trade Collapse. Income-induced expenditure changes are positively correlated with the cross-good pattern of import changes, generating a predicted change 40% as large as the raw variation in import declines.

We employ the Quadratic Almost Ideal Demand System (QUAIDS), which allows income elasticities to depend non-linearly on prices and incomes. We estimate key parameters using quarterly data from 1995Q1-2010Q1 taken from the US Consumer Expenditure Survey (CEX). The CEX provides household expenditure data for many traded and nontraded goods. We construct a synthetic panel of 10 income bins corresponding to income deciles in each quarter, and aggregate over households within each bin to create a representative household at each income decile.

This has several advantages. First, while individual household purchases of durables are infrequent, the representative household will have positive expenditures for (nearly) all goods and periods. Second, we can control for key demographic characteristics (family size, age, location) that systematically covary with income and that affect expenditures. Third, the synthetic panel structure allows us to exploit cross-sectional variation across bins in a given period to control for unobservable prices and quality of goods, while making use of variation in income and expenditure both within and across bins. Fourth, and most important, the system allows us to estimate spending shares and income elasticities that vary at the level of good-income-time.

Adding over all goods, the top two deciles are responsible for 49 percent of spending on traded manufactures (excluding food) found in the CEX, while the bottom two deciles are responsible for 3 percent of spending. Of note, the extent to which traded good expenditure is driven by the upper deciles varies tremendously across seemingly similar goods and over time. This is best shown by comparing expenditures for the top decile to the fifth (median) decile. The top decile spends 8.8 times more on "Men's Suits" than does the fifth decile, but only 3.6 times as much for "Men's Uniforms". Similarly, the top decile spends 12.8 times more than the fifth decile for "Winter/Water Sporting Equipment" but only 2.7 times more for "Fishing and Hunting Equipment".

Income elasticities differ from one, vary significantly across good-income-time, and are on average falling with income levels. Moreover, the data clearly reject that the ratio of income elasticities for two goods is constant across income levels – a central prediction of Constant Relative Income Elasticity (CRIE) preferences used in the literature.

The combination of expenditure shares and elasticities varying over good-income-time means that even a uniform income shock will result in large changes in the distribution of expenditures across goods categories. Moreover, income shocks are not uniform, and there are pronounced differences in the distribution of income shocks during recent crisis periods. In the period just before the Dot-Com Crash of 2000-01, higher income households experienced a sharp increase, then a more pronounced slowdown in incomes, while changes for lower income households were more muted. In the period just before the Great Trade Collapse of 2008-9, the rise and fall of expenditures was more pronounced in lower and middle income households.

By combining data on the distribution of shocks with our estimates of income-specific expenditure shares and income elasticities we can construct predicted changes in expenditures specific to each good-income-time period. Aggregating over income bins, we have a measure of predicted expenditure change that is good-time varying, arising only from income shocks.

In a final exercise, we explore whether these predicted expenditure changes can explain panel variation in imports and the pattern of import declines during the recent crashes. We regress changes in imports at the good level on changes in actual expenditures on that good taken from the CEX. Of course, actual expenditures depend on goods prices and quality, and a myriad of other endogenous factors. Accordingly, we use our measure of predicted expenditure change arising from income shocks as an instrument for actual expenditure change. The first stage yields a strong fit, and in our preferred second stage specification we find an elasticity of import change with respect to expenditures of 0.15.

A key to understanding the Great Trade Collapse is that the import change was not uniform, and in fact varied dramatically across goods. Using our estimates for the peak of the GTC, we find that a good with an expenditure change in the 10^{th} percentile (large decreases) had an associated import decline 15 percentage points larger than a good with an expenditure change in the 90^{th} percentile. The actual (10-90) gap in import change was on the order of 51 percentage points, suggesting that expenditure changes arising from the distribution of income shocks played a significant role (roughly one-third) in the overall decline. These results are robust to changes in sample years and width of household income bins used in the estimation, and our point estimates are robust to incorporating other variables emphasized in the Great Trade Collapse literature, including inventories, shocks transmitted through supply chains, and financing constraints.

Our emphasis on non-homothetic demand relates to an older branch of the trade literature that studies per-capita income as a determinant of trade patterns. Linder's (1961) seminal work emphasizes how income affects the composition of the consumption basket, and suggests that more similar countries will have higher bilateral trade volumes. Markusen (1986) and Bergstrand (1989, 1990) formalize these insights using Stone-Geary preferences to generate income effects in models of monopolistic competition and trade. Thursby and Thursby (1987) and Francois and Kaplan (1996) formalize and test the Linder Hypothesis. Hunter and Markusen (1988) and Hunter (1991) show that per-capita income can serve as a basis for interindustry trade, and stress the importance of departures from homotheticity in explaining commodity level import demands. In a part of this literature, authors emphasize explicitly the role of the within-country income distribution. Examples include Flam and Helpman (1987), Dalgin, Trindade and Mitra (2008), and Choi, Hummels, and Xiang (2009).

More recently, Caron et al. (2014) and Fajgelbaum and Khandelwal (2016) estimate gravity equations derived from non-homothetic preferences to generate income elasticity estimates. Caron et al. (2014) use CRIE preferences from Fieler (2011) and focus on explaining home bias and biases in the factor content of trade. Fajgelbaum and Khandelwal (2016) and He (2017) use the

Almost Ideal Demand System and focus on measuring the unequal gains from trade across consumers of different income levels. In both cases, the authors combine non-homothetic demand systems with structural assumptions on the production side of the model to generate trade predictions. They estimate sector-level gravity regressions that exploit cross-country variation in per capita incomes at a point in time to explain the level of expenditures and trade across broad sectors of the economy (including agriculture, manufacturing and services). We focus on the distribution of income and expenditures across households *within* the US and focus on how shocks to the household income distribution drive changes in expenditures and import demand within specific traded manufactured goods over time.

Focusing on the distribution of income shocks within a country is non-trivial. Non-homothetic systems used in the older literature, such as the linear expenditure system (LES) derived from Stone-Geary preferences, allow the level, but not the distribution of income within a country, to affect expenditures (i.e. the LES satisfies Gorman aggregation). Further, the LES system generates identical income elasticities for all non-subsistence goods. More recent innovations, such as CRIE, allow elasticities to vary with income levels and across goods, but constrain the ratio of income elasticities to be the same at all income levels. These more restrictive systems are ideal for use in cases where parameters are identified from the aggregated trade behavior of an entire country.

An intriguing difference in the results generated by these different approaches has to do with the behavior of spending on manufactured goods across different income levels. Fajgelbaum and Khandelwal (2016) provide cross-country evidence that budget shares devoted to manufactures fall with income, and income elasticities for manufactures rise with income. Our within-US household panel evidence suggests exactly the opposite. Expenditure shares devoted to manufactures are only 5 percent at the first decile and sharply rise with income (to 25 percent in the top decile, and 40 percent in the top percentile), and associated income elasticities fall.

Our study tangentially relates to the literature that uses Nielson scanner data (Faber and Fally (2017), Handbury (2013), and Jaravel (2016)) and incorporates income effects in the analysis. Faber and Fally (2017) find that rich and poor households source their consumption from different parts of the firm size distribution, and related, Jaravel (2016) finds that rich household gains more from new and innovative goods. Handbury (2013) assesses biases arising from homotheticity in spatial price indexes across income groups, and find the bias is the largest for high-income households. Our emphasis is on estimating budget shares and income elasticities to generate predicted panel variation in national expenditures and imports across a wider range of traded manufactures that do not appear in the scanner data (Handbury and Weinstein (2014), for example, is focused on extremely detailed food products).

Our final application also relates to the literature on the Great Trade Collapse. In one year, beginning in the fourth quarter of 2008, world trade declined by a third, a drop many times larger than the corresponding decline in incomes or output.[1] A variety of explanations have been offered for this severe downturn. Recent papers on trade finance (Ahn, Amiti and Weinstein (2011), Amiti and Weinstein (2011)) and credit tightening (Chor and Manova (2012)) attribute decreases in trade to the reduction in the availability of external finance during crises. Bems, Johnson and Yi (2010) focus on the transmission of shocks through vertical production linkages. Alessandria, Kaboski and Midrigan (2010) examine whether agents depleted inventories as a substitute to buying more from abroad. On the expenditure side, several authors (Baldwin and Taglioni (2009), Eaton, Kortum, Neiman and Romalis (2016)) examine production composition. If international trade occurs disproportionately in sectors whose domestic demand (or production) collapsed the most, we would expect trade to fall more than GDP. Related, Levchenko, Lewis, and Tesar (2011) argue that a reduction in quality demanded after income losses will result in a contraction in the value (price, rather than quantity) of imports.[2]

We have little to say about the supply side of trade in the recent crisis, though we examine whether our estimates are sensitive to including correlates from this literature. Our work is closely related to the composition effect hypothesis, in that we focus on a systematic decline in expenditure for certain categories of goods. Unlike this literature, we offer a direct test of why particular good categories experienced sharp expenditure contractions as a function of income elasticities and the distribution of income shocks.

Finally, we note at the outset that our approach is deliberately stark. We are not trying to fully explain the Great Trade Collapse or to fully explain what gives rise to expenditure changes

[1] In 2008q3 and 2008q4, world trade flows were 15% below their previous level (Baldwin and Taglioni (2009)). The trade growth rate of 23 OECD countries reached a record negative growth of -37% in April 2009 (Araújo and Oliveira Martins (2011)). Within the US, GDP declined by 3.8% from its peak to the trough, real U.S. imports fell by 21.4% and real exports fell by 18.9% over the same period (Levchenko, Lewis and Tesar (2010)).

[2] Levchenko et al. (2010) test multiple hypotheses, finding support for vertical production linkages and a composition effect, but no support for the credit tightening hypothesis. Haddad, Harrison, and Hausman (2010) provide a simple explanation for this finding: import price in sectors requiring high external finance rose by much more than the prices in other sectors, which offsets the decline in quantities.

on imported goods over time. Rather, we are interested in one aspect of expenditure change arising from the distribution of income shocks and whether that expenditure change can generate some significant portion of the relevant change in trade behavior. The advantage of this approach is that we can identify the relevant income effects from the household data, and a stark specification provides some hope of being able to implement the resulting instrument outside the immediate context. Undoubtedly there are interesting questions about how changes in the availability of household credit, the housing crisis, or an overhang of consumer spending on durables, may have had significant changes in the pattern of expenditures in this period. We put all this to the side to focus on incomes.

The paper is organized as follows. Section 2 develops the methodology for estimating budget shares, income elasticities and expenditure changes. Section 3 describes the CEX data, and construction of the synthetic panel. Section 4 presents stylized facts and key results from estimating the demand system. Section 5 reports results linking expenditure change to the trade decline, along with robustness checks. We conclude with remarks on the broader applicability of our estimates in section 6.

2. Methodology

2.1. Overview

To begin, write imports M of good g at time t as a share of national income Y:

$$\frac{M_{gt}}{Y_t} = \frac{M_{gt}}{E_{gt}} \cdot \frac{E_{gt}}{Y_t} \tag{1}$$

The first term is the share of imports in expenditures E for good g. The second is the expenditure share of good g in national income. Much of the focus of the literature on the Great Trade Collapse is on the first term, explaining why imports as a share of expenditures would decline. Our focus is on the second term, explaining movements in the expenditure shares on good g over time.[3]

The problem is that expenditures are endogenous to many of the supply shocks posited in the literature. For example, if financing constraints raise traded goods prices and demand is price

[3] In section 5 we use a variance decomposition to show that, in our data, 42 percent of the panel variation in equation 1 is driven by variation in the second term. Aggregating over all goods, 56 percent of the time series variation in aggregate imports/GDP is driven by the second term.

elastic, we expect expenditures to decline. Accordingly, we need an instrument for expenditures that is good x time varying and orthogonal to supply shocks. Note that the classic demand instrument, changes in income, provides no cross-good variation if demand for traded goods is homothetic. That is, a 5% fall in income generates an identical 5% reduction in expenditures for all goods. However, cross-good variation in income elasticities arising from non-homothetic demand, combined with a distribution of income shocks, can generate good x time variation in expenditures.

To see how this works, note that the change in aggregate expenditures on a good is a share-weighted aggregation of expenditure change at the household level. To smooth purchases we will focus on bins of similar households (more in the data section below). Denoting traded goods by g, and household bins by b, the change in expenditures over four quarters is:

$$dlnE_{gt} \equiv ln\left(\frac{E_{gt}}{E_{g,t-4}}\right) = ln\left(\sum_{b} S_{gb,t-4} \cdot \frac{E_{gbt}}{E_{gb,t-4}}\right) \tag{2}$$

where the operator dln indicates year-to-year quarterly differences of logarithms, and $S_{gb,t-4}$ is the share of bin b in national expenditures for good g at $(t-4)$. To prevent confusion, note that high income households may devote a relatively small share of their budget to a particular good and yet be responsible for an outsized share of economy-wide spending. We are interested in the predicted change in expenditures. To build this up from the level of household bin, we need to estimate the level of household bin spending on good g and how that spending changes in response to changes in income.

2.2. Expenditure Shares and Income Elasticities in the QUAIDS

The Quadratic Almost Ideal Demand System (QUAIDS) was first introduced by Banks, Blundell, and Lewbel (1997) as an extension of the Almost Ideal Demand System (AIDS). In QUAIDS, budget shares depend not only on the log of real total expenditure but on its square. The quadratic allows more flexibility in expenditure responses while still satisfying theoretical restrictions necessary for well-behaved utility. The QUAIDS in budget share form is:

$$w_g = a_g + \sum\nolimits_k c_{gk}\, ln\, p_k + b_g \ln\left(\frac{y}{P}\right) + \frac{d_g}{\prod_k p_k^{b_k}}\left(ln\left(\frac{y}{P}\right)\right)^2 \tag{3}$$

The household budget share for good g is w_g, y is total expenditure for the household, p_k is the price of a good k, and a_g, b_g, d_g and c_{gk} are parameters.[4] $\ln P$ is a price index defined as $\ln P = a_0 + \sum_g a_g \log p_g + \frac{1}{2}\sum_g \sum_k c_{gk} \log p_g \log p_k$ By setting $d_g = 0$, this system nests the more commonly used AIDS.

Using equation (3) we can calculate the income elasticity for each good g:

$$\eta_g = 1 + \left(b_g + \frac{2d_g}{\prod_k p_k^{b_k}} \ln\left(\frac{y}{P}\right)\right)\frac{1}{w_g} \tag{4}$$

From equation (4), we can infer three properties of the income elasticity:

1. η_g can differ from one.
2. η_g varies across income levels for a particular good.
3. The sign of b_g and d_g determine if a good is income elastic or inelastic at a given income level and price index.

To illustrate these properties, let $p_g = 1\ \forall g$ and set $a_0 = 0$, so that $P = 1$. Then, η_g reduces to:

$$\eta_g = \left(1 + \frac{b_g}{w_g}\right) + \frac{2d_g}{w_g} \ln y$$

It is immediate that if both $b_g > 0$, $d_g > 0$, then $\eta_g > 1$ and is increasing in income at all income levels. Conversely, if both are negative, then $\eta_g < 1$ and is decreasing in income at all income levels. However, if b_g and d_g have opposite signs, goods can switch from income inelastic to income elastic and vice versa as incomes vary. Figure 1 displays these cases.

2.3. Estimation Methodology: Budget Shares

We estimate the relevant parameters of equation (3) using data on income and expenditures from a panel of households. Rewriting (3) to incorporate household bin b and time t variation:

$$w_{gbt} = a_g + \sum_k c_{gk}\, ln\, p_{kt} + b_g \ln\left(\frac{y_{bt}}{P_t}\right) + \frac{d_g}{\prod_k p_{kt}^{b_k}}\left(ln\left(\frac{y_{bt}}{P_t}\right)\right)^2 \tag{5}$$

[4] For well-behaved utility, the following restrictions are necessary: $\sum_g a_g = 1, \sum_g b_g = 0, \sum_g c_{gk} = \sum_k c_{gk} = 0, c_{gk} = c_{kg},\ and\ \sum_g d_g = 0$.

We assume that demand parameters a_g, b_g, d_g and c_{gk} are time invariant. The various price terms pose the main difficulty in estimation because we do not have price data for the specific goods in the CEX.[5]

To resolve the difficulty in estimation, we assume that, after conditioning on location, households of varying income within the US face the same vector of prices at a point in time. This means that the expression $a_g + \sum_k c_{gk}\, ln\, p_{kt}$ can be eliminated by incorporating a good-time fixed effect, α_{gt}. We proxy for the QUAIDS-appropriate price index using the CPI. (We also experiment with allowing the level of prices to vary between urban and rural regions.) Note that the quadratic income term interacts with an aggregated measure of prices that is common across goods but varies over time, $\prod_k p_{kt}^{b_k}$. However, if this price measure takes on the same value for each household at a point in time, we can absorb this variation by interacting the quadratic income with a vector of time dummy $\boldsymbol{T_t}$. To complete the specification, we incorporate a vector of demographic characteristics $\boldsymbol{X_{bt}}$ which may affect expenditures such as age of household head, family size, and location (urban/rural).

$$w_{gbt} = \alpha_{gt} + \beta_g \ln\left(\frac{y_{bt}}{CPI_t}\right) + \delta_{gt} \cdot \boldsymbol{T_t}\left(ln\left(\frac{y_{bt}}{CPI_t}\right)\right)^2 + \boldsymbol{\beta}\, \boldsymbol{X_{bt}} + \varepsilon_{gbt} \tag{6}$$

We estimate equation (6) separately for each good g, exploiting panel variation across household bins and time. Using estimates from equation (6), we obtain predicted budget share, $\widehat{w}_{gbt}$, and income elasticities, $\hat{\eta}_{gbt}$, for a household of income y_{bt} but with otherwise average demographic characteristics:

$$\widehat{w}_{gbt} = \hat{\alpha}_{gt} + \hat{\beta}_g \ln\left(\frac{y_{bt}}{CPI_t}\right) + \hat{\delta}_{gt} \cdot \left(\ln\left(\frac{y_{bt}}{CPI_t}\right)\right)^2 + \widehat{\boldsymbol{\beta}}\, \overline{\boldsymbol{X}}_{\boldsymbol{bt}} \tag{7}$$

$$\hat{\eta}_{gbt} = 1 + \left(\hat{\beta}_g + 2\hat{\delta}_{gt} \ln\left(\frac{y_{bt}}{CPI_t}\right)\right)\frac{1}{\widehat{w}_{gbt}} \tag{8}$$

These elasticities are of independent interest. In addition, they also enable us to implement an instrumenting strategy for changes in expenditures at the good-time level and potentially explain changes in import demand. Recalling equation (2), expenditure change at the national

[5] Broda and Weinstein (2010), Handbury (2013), and Handbury and Weinstein (2014), employ Nielsen scanner data to emphasize differences in the availability and set of prices facing households in the US. We do not employ these data because we do not have access to them, and because these data cover a subset of the goods covered in the CEX.

level is a share-weighted average of expenditure changes happening within each household bin. We want the change in expenditure arising only from changes in income. This is:

$$dlnE'_{gt} \equiv ln\left(\sum_{b} S'_{gb,t-4} \cdot \frac{E'_{gbt}}{E_{gb,t-4}}\right) \tag{9}$$

where $\frac{E'_{gbt}}{E_{gb,t-4}}\left(= exp\left(\eta_{gbt} \cdot ln\left(\frac{y_{bt}}{y_{b,t-4}}\right)\right)\right)$ is the change in expenditure of bin b arising only from change in income for good g, and $S'_{gb,t-4}$ is the share of bin b in national expenditures induced by income change for good g at $(t-4)$.

3. Data

We employ data from the quarterly interview panel survey of the CEX from 1995q1-2010q1. Each consumer unit (CU) in the sample is interviewed once per quarter for five consecutive quarters,[6] and they report expenditures on major items of expense over the preceding quarter. CEX covers a complete range of household expenditures, including services, non-durable and durable goods. The CEX data are organized by universal classification codes (UCC). There are about 320 UCCs, of which 102 we classify as traded goods.

We are interested in examining how changes in income affect expenditures on traded goods, including consumer durables. The short panel dimension of the CEX prevents us from examining within household changes in income. In addition, durable goods purchases at the household level are infrequent and hence households register zero expenditures for many goods in most periods.

To overcome these problems, we create a synthetic panel with households aggregated into decile bins by total expenditure in each quarter (we also experiment with using 20 bins). We use total expenditures in place of income for three reasons. One, reported incomes and total expenditures are very highly correlated. Two, the income field in the CEX has known

[6] The sample design of CEX is a rotating panel survey in which one-fifth of the sample that has completed its final interview is dropped and a new group added in each quarter. Specifically, each quarterly sample is divided into three panels of approximately equal size, each of which is nationally representative. CUs in these panels are interviewed once during the first, second, or third month of each quarter for five consecutive quarters. After CUs have been in the sample for five quarters, they are replaced by new CUs.

measurement problems at the household level. Three, we have nothing to say about savings behavior or how households spend beyond apparent income, this latter issue being especially problematic when fitting expenditures at very low income levels. Henceforth, we will use “income” and “total expenditures” interchangeably.

There are approximately 300 households (CEX Consumer Units) in each bin in each quarter. In the bottom 7 deciles the income range spanned by a bin is $925 on average, though the range of incomes rises sharply in the top two deciles. Within each bin we construct average expenditures across households for each category of purchases within the CEX, including 102 traded goods. We also keep track of household characteristics within each bin. For numerical demographic characteristics such as age and family size, we use averages within bins. For categorical characteristics, we use shares of categories within each bin; for example the share of households living in urban areas.

Following standard practice in the CEX literature, the sample is restricted to improve the measurement of consumption. In particular, households (HH) are dropped from the sample in these cases: multiple consumer units in the HH; HH lives in student housing; the head/spouse of HH is farmer/fisher; the HH does not complete all interviews; HH has incomplete information on income, negative income, or zero income.[7] Additionally, topcoded expenditures are dropped from the sample, and to remove potential outliers we drop the top and bottom 1 percentile of income bins.

For some of our exercises, we will match expenditure data from the CEX to trade data. The CEX data are organized by UCC which we sort into traded goods and non-traded services. We match UCC product descriptions to those found in 10-digit HS import data descriptions, building on a concordance constructed by Ardelean and Lugovskyy (2015). In many cases, there are multiple HS codes corresponding to a single UCC, and we aggregate these HS codes into a single good category. Note that our data cover consumer expenditures, and not expenditures on industrial supplies. In this period we match codes representing 27% of imports by value and will focus primarily on these goods. In some cases we also aggregate similar UCC’s. The list of UCC product descriptions and concordance to HS codes is available on request.

[7] To be clear, we use reported income, not observed expenditure.

4. Expenditure Shares and Changes, and Income Elasticities

From equations (2) and (9) we know that the aggregate response of expenditures to an income shock depends on how spending is initially distributed across households, and the responsiveness of each household to a change in income. In this section, we use the CEX data, and income elasticities estimated from it, to show how profoundly different the effect of an income shock on traded good spending can be depending on where that shock hits.

Figure 2 displays the (over-time) average budget shares for traded manufactures (excluding food) for each of 20 income bins in our data.[8] The share of expenditures devoted to traded goods is less than 5 percent for the bottom decile, rising to more than five times that number for the top decile. In contrast, food and housing comprises half the expenditures of low income households, but only a quarter of spending at the upper end. Repeating this exercise using one percentile increment bins results in a more continuous distribution of spending shares in the upper deciles. In this case, spending on traded manufactures reaches as high as 40 percent of household income in the 99th percentile, and spending on food and housing as low as 20 percent.

Why do these data differ so markedly from the cross-country evidence provided in Fajgelbaum and Khandelwal (2016) (henceforth F-K), whose Figure II shows the share of aggregate expenditures on manufacturing falling in income? First it is notable that housing and food expenditure data displayed here are consistent with micro-household evidence showing income elasticities significantly below one for these categories.[9] In the CEX data, rising expenditure shares on traded manufactures mirror declining expenditure shares on food and housing. Similarly, when F-K use the 2013 CEX data in a robustness check, they estimate a positive income elasticity (.037) for manufacturing consumption. These give us confidence that our calculations properly capture the within country pattern for the US.

[8] Complete information on the share of good *g* spending for each income bin *b* is captured in Appendix Table A.2.

[9] Haurin (1991), Ioannides and Rosenthal (1994), Polinsky (1977), and Zorn (1993) estimate income elasticity of demand for housing, ranging 0.35 to 0.75. Alderman (1986) reports that estimates of the income elasticity of demand for food ranges between 0 and 1 in many countries. Recently, Aguiar and Bils (2015) use the US CEX data to estimate housing and food (at home) expenditure elasticity to be approximately 0.9 and 0.4, respectively.

Why then is the cross-country data taken from national accounts and trade statistics different from household expenditure data? Two possibilities relate to the treatment of housing and intermediate goods. Personal consumption expenditures on an already-built housing stock is meant to be captured in national accounts data. However, in practice this requires imputing the rental value of owner-occupied housing. It is plausible to us that the quality of imputation needed to identify the value of the housing service flow might vary with the sophistication of the national statistical agency, and be missed for lower income countries.

Similarly, spending on intermediate inputs will be included in national accounts and trade statistics but will be omitted from household expenditures. Attempts to split absorption into industrial versus household use necessarily relies on industrial and household survey data. It is difficult to characterize these absorption imputations with great specificity because it requires knowing details about survey methodologies for data construction for many countries. However, it is instructive to focus solely on trade flows themselves and ask: how much of imported manufactures consist of goods for household consumption versus industrial absorption (intermediate and capital goods), and does this share depend on the income of the importer?

To answer this, we concord UN COMTRADE data at the HS6 level to Broad Economic Categories (BEC) data in order to characterize each HS6 flow as household consumption or industrial absorption for the countries used in F-K. Aggregating up to the categories used in F-K allows us to calculate the share of household goods in trade for each sector. When the final consumption share is regressed on the log of GDP per capita with sector fixed effects, the estimated coefficient on the final consumption share is 0.034, remarkably close to the slope of the Engel curve from the CEX data. Figure 3 displays the aggregate relationship between the final consumption share and per capita income. While the aggregate value of manufactured imports is declining in national income, the share of imports going to households (rather than industrial absorption) is sharply increasing in national income. To us this suggests that household survey data is likely a more accurate depiction of the distribution of expenditures across households, or at least that users of cross country data must be especially careful to account for the presence of traded intermediates.

We turn now to our estimation of equation (6) and calculation of corresponding budget shares and income elasticities captured in equations (7) and (8). Since our estimates vary across 102 expenditure categories, 10 income bins and 65 time periods, we have a total of nearly 6500 elasticities and budget shares. To show relevant properties, we report illustrative examples and capture full details in an appendix. For completeness, the appendix also includes expenditure shares and elasticities for expenditure categories (non-traded services, food) that are not incorporated in our import demand estimation.

In Figure 4, we display income elasticities for two specific goods (TV and computer game software; women and girls (W/G) coats) with variation across income deciles at two points in time. While both show elasticities dropping with income, there are quite significant differences in the level of the elasticity, the dispersion across income levels (much greater for software), and in over-time changes in the elasticity (software rises; W/G coats rise at low income levels and vice versa). Of particular note, software is income elastic throughout the income distribution, W/G coats are income inelastic at higher income levels. Picking up this sort of variation is a strength of the highly flexible QUAIDS system.

To show that we have not cherry-picked these examples, we report income elasticities for each decile and good category in Appendix Table A.1. (To simplify presentation, these elasticities are based on a specification where we estimate a single quadratic term for each good g rather than a g-t specific term.) Income elasticities exhibit significant variation across goods within income bins, and across income bins within goods.

It is useful to compare these estimates to the new trade literature that focuses on income effects. A number of authors (Fajgelbaum and Khandelwal (2016) and He (2017)) have employed the AIDS when estimating income effects.[10] Our QUAIDS nests AIDS, and so in Figure 5 we directly compare estimated income elasticities of demand from both systems. The three panels scatter AIDS vs. QUAIDS estimates at the 90th, 50th, and 10th percentiles of income. At the upper end of the income distribution, estimates are very similar (most points lie along the 45-degree line). At the lower end, QUAIDS estimates exhibit much greater dispersion than AIDS. In short, AIDS seems to fit the data quite well at the high end, but misses important variation at low incomes that

[10] Caron et al. (2014) also use the AIDS as a robustness test.

is captured by the quadratic terms. We return to the distinction between AIDS and QUAIDS in our application below.

Several new trade papers (Fieler (2011); Caron et al. (2014)) use non-homothetic CRIE (constant relative income elasticity) preferences. These preferences allow income elasticities to differ from 1 and to differ across income levels. However, they constrain the *relative* income elasticity between two goods to be constant over income levels. We can evaluate whether our estimates support this restriction. We calculate the relative income elasticity for every pair of goods gg', income bin and time, and express them relative to the mean (across bins and time) relative elasticity for gg'. Figure 6 displays the distribution of these values. Were relative elasticities constant, we would find values of 0 throughout but there are clearly large deviations from this baseline.

To be clear, CRIE preferences are a very powerful tool for incorporating non-homotheticities into general equilibrium trade theories and for performing associated welfare calculations. Our point is that a more flexible functional form estimated from household micro data allows us to generate richer variation in these elasticities than are permitted by CRIE, and that this greater variability may be useful for identifying income-induced shocks to good-level import demand.

Recall from equations (2) and (9) that changes in aggregate expenditures for a good are a function of the change in expenditures for each income bin, weighted by the share of that bin in aggregate expenditures. In Appendix Table A.3, we report the share of bin b in aggregate spending on good g. Aggregating over all manufactured goods, the top two deciles are responsible for 49 percent of spending, while the bottom two deciles are responsible for 3 percent of spending.

Of note, the extent to which traded good expenditure is driven by the upper deciles varies tremendously across seemingly similar goods and over time. This is best shown by comparing expenditures for the top decile to the fifth (median) decile. The top decile spends 8.9 times more on “Men’s Suits” than does the fifth decile, but only 3.2 times as much for “Men’s Uniforms”. Similarly, the top decile spends 13.4 times more than the fifth decile for “Winter/Water Sporting Equipment” but only 2.7 times more for “Fishing and Hunting Equipment”. The difference in the high-end spending shares will result in profoundly different changes in aggregate spending in the presence of non-uniform income shocks.

To explore how spending shares and elasticities interact, we use equation (9) to calculate the effect that a 10% rise in income would have on expenditures for the two goods (TV and computer game software; W/G coats) shown in Figure 7. The vertical axis shows aggregate (summed over all households) expenditure change for good g and the horizontal axis shows a series of left and right skewed income shocks that aggregate up to a 10% increase in total incomes.[11]

Starting in the middle of Figure 7, we see that a uniform 10% increase in incomes results in expenditure increases of nearly 10% for W/G coats and 12.5% for software. When we skew these shocks to the left (giving more income to the richest households and less income to the poorest), the expenditure response becomes more highly dispersed, with game software rising by 16% and W/G coats rising by only 8%. When we skew these shocks to the right (giving more income to the poorest households and less to the richest), this pattern reverses. Note that the responsiveness of aggregate expenditures to these distributional changes varies considerably over goods as a function of the relevant income elasticities and the share of each income group in aggregate expenditures. For W/G coats, the former effect dominates – W/G coats are income inelastic at high incomes. For software, the latter effect tends to dominate. Software has high income elasticities throughout the income distribution, and generates large expenditure responses to the income change. But that effect becomes more muted when income is given to poorer households because their baseline expenditures comprise only a small part of overall spending on software. The disparate response across goods, and its dependence on the distribution of income shocks, generates an ideal source of variation for an econometrician, i.e., no two "10% income shocks" are created alike when it comes to their expenditure effects.

This point is especially important when we consider the distribution of income shocks during two recent recessions, the Dot Com Crash (DCC), and the Great Trade Collapse. During the two recessions, expenditure declined throughout the income distribution. However, the distribution of expenditure shocks is distinctly different during the two recessions. During the DCC, the top decile experienced sharp expenditure reductions. During the 2008-2009 recession, the fall of expenditures were more pronounced in the bottom decile and fifth-seventh deciles.

[11] We use income in 2008q4 as a baseline and then shock the income distribution by varying slopes and intercepts in the formula $y_{bt}^{i} = \alpha^{i} + \left(1.1 - \frac{10 \cdot \alpha^{i}}{\sum_b y_{b,t-1}}\right) y_{b,t-1}$ so that aggregate income rises by 10%. Note that a right skewed shock that adds up to an aggregate 10% income rise implies a very large increase in incomes for poor households.

Given our results on expenditure shares and elasticities, the distribution of income shocks in these two episodes should lead to significantly different effects on trade.

Taken together, we have significant variation across income bins in the share of spending on particular goods; the change in aggregate expenditures (income) in particular periods, and the income elasticity of demand for particular goods and income levels. This provides the raw material for an instrument for aggregate expenditure change that might be able to match the variability in expenditures and imports that occur over time.

5. Application: Explaining Import Change During Recent Crises

Recalling equation (1), we can express imports of good g as a share of GDP as a product of the import share of expenditures for good g and good g's share in aggregate expenditures. Taking logs of equation (1) and expressing in first differences yields

$$dln\left(\frac{M_{gt}}{Y_t}\right) = dln\left(\frac{M_{gt}}{E_{gt}}\right) + dln\left(\frac{E_{gt}}{Y_t}\right)$$

where $dln\left(\frac{M_{gt}}{Y_t}\right) \equiv ln\left(\frac{M_{gt}}{Y_t} / \frac{M_{g,t-4}}{Y_{t-4}}\right)$ and similarly for the other two terms. Using actual expenditures from the CEX, a simple variance decomposition of this expression shows that about 42 percent of the total variation in imports/GDP is due to variation in the second term. This is true whether we calculate the variance over g-t unconditionally, or after subtracting time or good means.

We can make a similar point using predicted expenditures for each g-t calculated from equation (9) rather than actual expenditures. For all g-t, we calculate the year-on-year change in imports and predicted expenditures and display the distribution of these changes in Figure 8. The top panel shows the pre-crisis time periods in a histogram. The bottom panel scatters the good-level changes in imports and predicted expenditures during the recent crisis period (including a 45-degree line for reference). These graphs make clear two key points. One, there is tremendous variation across goods in year-on-year changes in imports and expenditures which we can exploit to test the role of income changes. Two, there are periods in which imports grow faster than expenditures (pre-crisis) and periods in which they grow slower (during the crisis). This is not surprising. After all, imports depend both on the trade share and the expenditure share, and we know of many supply side changes in these periods that led to rising, then falling imports. But the order of magnitude of changes is comparable, suggesting that expenditure change can be a quantitatively important part of the story.

We now turn to the estimation of our main equation and begin very simply.

$$dlnIM_{gt} = \beta\ dlnE_{gt} + \theta_g + \rho_t + \varepsilon_{gbt} \tag{11}$$

Conceptually, we can think of equation (11) as rewriting equation (1) by multiplying both sides through by GDP, and assuming that the ratio of imports to expenditures is absorbed into differencing, good fixed effects, or the error term. All variables are in log change over four quarters (sample 1995Q1 through 2010Q1), and g corresponds to a good from the CEX (we have matched HS10 imports data to the 102 traded UCC codes in the CEX and aggregated). In the IV specification, we instrument for actual expenditures using predicted expenditures, $dlnE'_{gt}$, arising from income shocks interacted with good x income bin x time varying income elasticities as described in equation (9) above.

A few notes on threats to identification. By first differencing we eliminate seasonality in the quarterly data and we eliminate level differences across goods in expenditure shares, and in supply characteristics such as price, quality, and variety. By incorporating good fixed effects (θ_g) we allow for good specific time trends in these components, and a year fixed effect (ρ_t) controls for aggregate shocks that affect trade or the macroeconomy and are common to all goods. (Given the year fixed effect, estimates written as shares of GDP and estimates written as equation (11) will be very similar.) For reference we also provide estimates without differencing and show how eliminating sources of variation changes elasticities. In all estimates we cluster standard errors on goods to account for serial correlation in the first differences.[12]

We might be concerned that the income shocks used to construct the instrument are endogenous. This could occur in two ways. One, the income shocks could arise from the trade

[12] This is particularly important in this context because first differencing may inadvertently introduce serial correlation within a good time series. To explain, suppose that IM_{gt} or E_{gt} exhibits an idiosyncratic, one time increase – perhaps a shipment scheduled for January arrives in December, or perhaps the CEX samples a few households with extraordinarily high purchases in a month. Our differencing strategy means that an idiosyncratic increase at time t will correspond to an idiosyncratic decrease at time t+4.

shock itself – in the second stage, declining trade could affect incomes and be correlated with predicted changes in expenditures. However, this should be an aggregate phenomenon, not a product level effect, and absorbed into year fixed effects. Two, both the trade shock and the expenditure shock could arise from a third cause (e.g. financing constraints) that might differentially affect certain goods. We examine this issue in depth below, focusing on supply side explanations for the Great Trade Collapse.

In Table 1 we report OLS and IV estimates of equation (11). For completeness we report estimates with and without good and year fixed effects, and also an IV specification in which log-levels rather than log-differences of variables are employed.

The top panel of Table 1 reports OLS estimates and in each case we find a very small response, an elasticity of about 0.03. IV estimates are considerably different. In the first stage we see that the change in expenditures arising solely from income shocks is a very good predictor of actual expenditure changes. When using log-changes in the variables coefficients are precisely estimated with an elasticity of 0.4 to 0.5, and F-stats are large in specifications with and without fixed effects. When using log-levels we see larger elasticities, especially in specifications that omit product fixed effects. In the second stage, using log-changes, we find a (highly significant) elasticity with the largest coefficients (0.154) in our preferred specifications with saturated fixed effects. When we use log-levels without product fixed effects, elasticities are as high as 0.6.

Several things are notable about these results. First, estimators that make use of the full panel variation exhibit much larger elasticities, while constraining variation only to the within time-series on products are more weakly (but still significantly positively) related to imports. Conceptually, both dimensions of the panel data relate variation in incomes to household expenditures and then to imports. (Recall from the equation (6) and the Figures above that the *level* of national spending on a category, and not just its *changes*, depends critically on the level and distribution of household income.) However, the estimators in differences or with product fixed effects are much more conservative in that they eliminate cross-product variation in the level of spending, whether it depends on income or prices or non-income related budget shares.[13]

[13] For reference, our estimates of budget equation (7) can be performed with and without the income terms, the latter being equivalent to imposing homothetic preferences. Including income terms increases the fit of the regression from 0.3 to 0.4 on average.

Note also that the IV coefficients are much larger than the OLS results. Why? The OLS results could be biased downward if omitted supply side factors in the regression are positively correlated with imports and negatively correlated with expenditures.[14] More likely, the CEX expenditure data are measured with error, either due to household reporting bias or because infrequent purchases of durable goods in smaller samples induce fluctuations in first differenced data. This induces attenuation bias in the OLS regressions, but by projecting raw expenditure data on the instruments we eliminate the error and the attenuation bias.

Above we noted the differences in income elasticities arising from the QUAIDS specification versus the AIDS specification. Do these differences affect our ability to fit changes in imports? In Table 2 we report the second stage of equation (11) estimated in log-changes but where the instruments are constructed in several different ways. The first row is our baseline values for QUAIDS income elasticities where we employ all point estimates regardless of precision. In the second row, we rely only on more precisely estimated coefficients – we replace estimated income elasticities with a value of 1 if the p-value exceeds 0.1 for a joint test of significance for the income terms. Here we see little change in the coefficients, meaning that the estimates are not driven by imprecise elasticity values.

The use of good-time varying quadratic terms can lead to greater imprecision in the estimates. We next constrain the quadratic terms to be good specific, using all point estimates (third row) and again replacing elasticities with a value of 1 when p-values exceed 0.1 (fourth row). There is little effect on the estimation. Finally, we suppress the quadratic term altogether and estimate an AIDS specification (fifth and six rows). Again, this has little effect on our ability to explain import changes.

How do we explain the invariance of our main results to significant changes in the specification of the income elasticities used to construct the instrument? The key can be found in Figure 5, where we showed that the use of quadratic terms is most important for explaining the behavior of very low income households. However, these households are responsible for such a small share of overall traded goods purchases that changing their predicted behavior has little effect on the aggregates. Getting high income households right – which is done equally well by AIDS

[14] While it is easy to think of omitted supply side factors, it is harder to identify any with this particular sign configuration. For example, suppose financing constraints raised import prices relative to domestic purchases. This should be negatively correlated with both imports and expenditures, generating upward bias in the OLS estimates.

or QUAIDS or QUAIDS with various restrictions on quadratic terms – is key to getting trade volumes right. But this also suggests that the aggregate import behavior of a country may be a rather low quality indicator of the behavior of its poorest households, and says nothing about whether AIDS versus QUAIDS is a better guide to the distribution of normative effects.

Robustness

Our specification is intentionally stark, but we briefly describe some robustness checks related to data construction and specifications. (All results available on request.) Our strategy of grouping households into bins is designed to register positive expenditures for all goods categories in all periods. However, by using 10 bins we group together dissimilar households, especially at the high end of the income distribution, and we lose some of the data variation that would be useful in picking up quadratic effects in estimation. We re-repeat all our estimation of budget shares and calculated elasticities using 20 income bins, then repeat our Table 1 specification with this sample. We see little qualitative change in our estimates in either the first or second stage.

Our estimation of budget share equation (6) assumes that all households face the same vector of prices. We do not have price data that would enable us to relax this assumption, but we do have an urban/rural geography indicator for households in the CEX. The vector of prices, especially for housing, is likely to be very different for these sets. We re-estimate budget shares and income elasticities using only urban households, then re-estimate our Table 1 specifications. This slightly increases the point estimate on predicted expenditures, but the change is well within the standard error. Related, our estimates of equation (7) assume that the CPI is an adequate proxy for a more appropriate QUAIDS specific aggregate price index. If this is incorrect, then we would expect estimates of the quadratic term may be subject to attenuation bias. Here we point back to the results from Table 2 and the invariance of our main results to treatment of the quadratic term to suggest that the use of the CPI as a good or bad proxy is unlikely to matter much for the main message.

Our specifications allow for no dynamic responses to income changes, and this may be particularly inappropriate in the case of durable and semi-durable goods. To address this, we tried three robustness checks. First, some of our goods – especially cars, trucks, and motorcycles – exhibit very infrequent purchases at the household level and have expenditure shares highly concentrated in upper income households. We experimented with dropping these goods from the estimation and found no qualitative difference in any results. Second, we tried re-estimating Table 1 using only durable or only semi-durable goods and found no qualitative difference in the results.[15] Finally, we experimented with including both contemporaneous expenditure variables and their first lag. In these cases, the total effect was unchanged, but the estimated effects were split evenly between contemporaneous and first-lagged effects.

We found the invariance of the results to the types of goods somewhat surprising, but believe this is likely the result of aggregating households into a synthetic panel. If we were sampling the same household for long periods of time, presumably past income shocks would have a pronounced effect on past durables purchases and these lagged effects would condition current purchases in an important way. But since we see households for only five quarters and durable purchases are sparse, it is very difficult to identify these sorts of effects in the household data.

Explaining the GTC

As a final exercise, we examine whether income shocks play a quantitatively important role in explaining the Great Trade Collapse. Table 1 uses the full 1995-2010 period, but it might be useful to explore the cross-good variation in trade declines in these periods in isolation. Specifically, the GTC generated an unusually large set of changes in expenditures and imports and we wish to see whether we can fit the relationship excluding it from the sample. In Table 3 we experiment with the sample years, showing only the most conservative specifications with good and time fixed effects. In column 1 we omit the GTC period and find quite similar elasticities to the comparable estimates in Table 1. In columns 2-3 we experiment with including a dummy variable for the GTC period, as well as interacting that dummy with our predicted expenditure variable. The point estimate on the interaction is negative but not significant, suggesting no change in the imports-expenditure pattern in the crisis year. (Even if we ignore the significance and take the point estimates at face value, adding the direct and interaction coefficients implies that the relationship between imports and expenditures is still positive during the GTC).

We next examine whether the key relationships are robust to including other explanations for the Great Trade Collapse. Recalling equation (1), changes in imports for a good can arise either from a change in expenditures for that good relative to GDP or from changes in the share of imports

[15] The import demand elasticity with respect to the instrumented expenditure change is 0.184 for semi-durables and 0.155 for durables.

in expenditures. While our focus is on the former, the trade literature has focused a great deal of attention on supply characteristics that affect the latter. Our specifications incorporate first differences, good and time fixed effects. First differences eliminate any supply factors that are good specific but not changing year to year. Good fixed effects eliminate any good specific trends in these factors, and time effects eliminate any remaining variation that is common to all goods in a year. This makes it challenging to identify sources of exogenous variation in supply factors that are good-time varying and that might be correlated with the income generated good-time varying expenditure shocks of interest to us.

One possibility is to focus on the variables uncovered in the literature on the Great Trade Collapse. Many of these explanations focus on variables or channels that are not relevant to our sample, for example, the effect of the GTC on demand for industrial supplies. More relevant, and more likely to be correlated with household decisions, is a channel operating through the financial markets. Several authors hypothesize that trade was disrupted because firms in the US were less willing to extend trade credit to partners abroad. To consider this trade credit channel, we use two variables, the ratio of account payable to cost of goods sold and the ratio of account receivable to sale as a proxy for intensity of trade credit used and extended, respectively. These two are from Levchenko, Lewis, and Tesar (2010). Since their two proxies are available at more disaggregated level (NAICS 4-6 digit) than our industry level (NAICS 3-4 digit), however, we use average values of the two proxies within industries. These variables are time invariant.

Since these variables are measured at the industry level, we need to match them to product level trade and expenditure data used in the rest of the paper. We add the industry-level supply side controls to our baseline while keeping the import change and the predicted expenditure change at the CEX level.[16] We employ year but not product effects, as product effects would be collinear with the trade credit measures. In addition to these variables, we use an import share weighted sum of distance between the US and source countries as a proxy for trade cost that might potentially affect import change.

[16] We could instead simply aggregate the product data up to the industry level. This approach does not change point estimates for our expenditure variable but the loss of so many observations results in a loss of statistical significance. (Results are available on request).

Table 4 reports the results.[17] In columns 1-2, we experiment with the supply side controls. In columns 3-4, we interact these variables with a dummy variable that takes a value of 1 during the GTC period. This tells us whether goods belonging to industries that are more reliant on trade credit see different changes in their trade behavior during the GTC. We find that the estimated coefficients on the predicted expenditure change have similar size to our baseline results in Table 1. This indicates that the effect of income driven expenditure changes is robust to including other supply side controls.

Since our estimates are robust to the sample period and other supply side correlates of the GTC, we wish to explore quantitatively how important income shocks are to explaining the trade collapse. We begin with a focus on product level changes in trade. Table 5 reports data on the cross-good distribution of expenditure and import changes in each quarter of the GTC. Because of outliers we focus on the goods in the 10th, 50th, and 90th percentiles of predicted expenditure change (where the largest decline is 1st percentile and the largest increase is 99th). For example, in 2008q4, the 10th percentile good had a 59.5% decline in year-on-year expenditures, the median good saw a 16.2% drop, while the 90th percentile good saw a 27.6% increase in year-on-year expenditures. Employing the estimated elasticity of imports with respect to (instrumented) expenditures of 0.154 from Table 1, the associated import changes for 2008q4 are -9.2%, -2.5%, and +4.3%. Expressing that in percentage point differentials, we report in the table that the 10th percentile good had an associated import reduction that is 6.66 percentage points larger than the median good and 13.4 percentage points larger than the good with a 90th percentile change in expenditures. Comparing this to the actual variation in import changes at these percentiles, it appears that the change in expenditures generates a predicted change in imports 25-40 percent as large as the raw variation in import changes we observe in the data.

We believe our estimates are on most solid ground when we remain focused on which products saw the largest trade decreases because the specifications eliminate common shocks that may or may not be related to our main story. More ambitiously, we might aggregate our product

[17] We also experiment with an inventory variable in spirit similar to Alessandria, Kaboski and Midrigan (2010) using data from the US Census Bureau. Specifically, we convert end-of-moth total inventories available at the M3 level to quarterly inventories, and concord M3 to NAICS level. When we include this variable, the point estimate of the predicted expenditure change is unchanged. However, the inventory variable coverage is sparse resulting in a loss of 70% of CEX product observations and a loss of statistical significance in the regression. (Full set of regression results including the inventory measure is available on request).

level predictions $dlnE'_{gt}$ to assess how useful our estimated expenditure shocks are for explaining the aggregate change in imports relative to GDP. Recalling equation (1) recast as an aggregate expression, we can provide a time series variance decomposition of the left and right hand side. That is

$$var\left(ln\frac{M_t}{Y_t}\right) = var\left(ln\frac{M_t}{E'_t}\right) + var\left(ln\frac{E'_t}{Y_t}\right) + 2cov\left(ln\frac{M_t}{E'_t}, \frac{E'_t}{Y_t}\right)$$

Note that this is equivalent to assuming that the coefficient on $ln\frac{E'_t}{Y_t}$ is equal to one. The variance in predicted aggregate expenditures represents 56 percent of the variance in log imports/GDP for this time period.

Next, we plot the predicted changes in expenditure against the actual trade share (manufactured household goods only) in Figure 9. Several things are notable. First, there is clear seasonality in the time series (eliminated in the product level regressions above). Second, the predicted change in expenditures is more variable than the actual trade share (see right and left y-axes). This is consistent with the product level elasticity estimates considerably lower than 1 in Table 1. Third, when we compress this volatility, we see a high degree of covariance between predicted expenditures and actual trade shares from 1998-2006, and again from 2008q4 onward. The failure of income and expenditure changes to match the time series in the trade pattern comes not during the GTC, but in the two years before it. This figure suggests that, relative to predicted household expenditures on traded goods, the puzzle isn't the sharp fall in trade so much as it is the sharp rise that preceded it.

6. Conclusion

We estimate budget shares and income elasticities from household variation in expenditures for the US using QUAIDS, a non-linear non-homothetic demand system. We show that expenditure shares and income elasticities vary dramatically across income levels and violate the assumption of constant relative income elasticities found in several recent papers on non-homothetic demand. Interacting these shares and elasticities with the distribution of income shocks within the US provides an excellent instrument for good x time variation in expenditures on traded goods. Income-induced expenditure shocks are positively correlated with the cross-good pattern of import changes, and during the Great Trade Collapse these shocks generate a predicted change as much as 40% as large as the raw variation in import declines.

While we have provided an application focused on explaining the Great Trade Collapse, our findings could be useful for several additional literatures. First, we show that spending on traded goods is concentrated in upper income households (the top two income deciles are responsible for nearly half of traded goods expenditures), and that expenditures on traded manufactures are rising in income. While this is consistent with other household micro evidence, it seems counterintuitive given the structural change literature and Fajgelbaum and Khandelwal's (2016) recent trade paper. That works suggests that when we look at cross-country data, higher incomes are associated with shifts away from traded manufactures in production. We show that this is largely the result of including intermediate inputs in the cross country data but not the household data.

Our estimates, and the resulting instrument for expenditure change, have application to two additional literatures. The political economy literature shows that optimal tariffs depend on the elasticity of export supply. The existing literature (notably, Soderbery (2010, 2015) and Broda, Limão, Weinstein (2008) based on the technique in Feenstra (1994)) identifies export supply elasticities by assuming that shocks to export supply and import demand are independent. If this identifying assumption does not hold, instruments are necessary for consistent estimation. While instruments for supply are straightforward to construct, instruments for import demand that are good x time varying have, prior to this paper, not appeared in the literature.

Similarly, the role of demand shocks has become prominent in another literature focused on firm-level export data and on export "failures". A number of papers, including Kee and Krishna (2008), Lawless and Whelan (2008), Eaton, Kortum, and Kramarz (2011), Munch and Nguyen (2014), and Nguyen (2012) incorporate demand shocks in a heterogeneous firm framework to reconcile the canonical Melitz (2003) model with the data. Put another way, these authors rely on unobserved demand residuals to fit the data, suggesting that import demands are highly idiosyncratic, varying across countries and time even within narrowly defined product groups. A related literature, following from Besedeš and Prusa (2006), emphasizes the short duration of trade relationships at the country pair x product level. In both cases, the availability of micro-founded demand shocks, rather than residuals, could prove useful in extending our understanding of firm-level trade and trade volatility arising from demand.

Acknowledgements: We thank Thibault Fally, Anson Soderbery, Chong Xiang, Masha Brussevich, Kan Yue and two anonymous referees for helpful comments and suggestions. All errors are our own. Lee gratefully acknowledges research support from the CoBPA Summer Research Grant Program at University of North Dakota.

References

Alderman, H. (1986). “The effect of food price and income changes on the acquisition of food by low-income households.” *Intl Food Policy Res Inst*.

Aguiar, M., and Bils, M. (2015). “Has consumption inequality mirrored income inequality?” *American Economic Review*, 105(9), pp. 2725-2756.

Ahn, J., Amiti, M., and Weinstein, D. (2011). “Trade finance and the great trade collapse.” *American Economic Review*, 101(3), pp.298-302.

Alessandria, G., Kaboski, J. P., and Midrigan, V. (2010). “The Great Trade Collapse of 2008–09: An Inventory Adjustment?” *IMF Economic Review*, *58*(2), pp.254-294.

Amiti, M., and Weinstein, D., (2011). “Exports and Financial Shocks.” *Quarterly Journal of Economics*, 126 (4), pp.1841-1877

Araújo, S., and Oliveira Martins, J. (2011). “The Great Synchronisation: Tracking the Trade Collapse with High-Frequency Data.” Paris Dauphine University.

Ardelean, A and Lugovskyy, V. (2015). “Trade Liberalization and Quality Upgrading.” CAEPR Working Paper No. 004-2015.

Banks, J., Blundell, R., and Lewbel, A. (1997). “Quadratic Engel curves and consumer demand.” *Review of Economics and Statistics*, *79*(4), pp.527-539.

Baldwin, R., and Taglioni, D. (2009). “The great trade collapse: Causes, Consequences and Prospects.” *VoxEU.org*.

Bems, R., Johnson, R., and Yi, K.M. (2010). “The Role of Vertical Linkages in the Propagation of the Global Downturn of 2008," *IMF Economic Review*, 58(2), pp.295-326.

Bergstrand, J. H. (1989). “The Generalized Gravity Equation, Monopolistic Competition, and the Factor-Proportions Theory in International Trade.” *Review of Economics and Statistics* 71(1), pp.143-153.

Bergstrand, J. H. (1990). "The Heckscher-Ohlin-Samuelson Model, The Linder Hypothesis and the Determinants of Bilateral Intra-Industry Trade," *Economic Journal*, 100(403), pp.1216-1229.

Besedeš, T., and Prusa, T. J. (2006). "Ins, outs, and the duration of trade." *Canadian Journal of Economics*, *39*(1), pp.266-295.

Broda, C., Limão, N., and Weinstein, D. (2008). "Optimal Tariffs and Market Power: The Evidence." *American Economic Review*, 98(5), pp. 2032–2065.

Broda, C., and Weinstein, D. (2010). "Product creation and destruction: Evidence and price implications." *American Economic Review*, 100 , pp.691–723.

Caron, J., Fally, T., and Markusen, J. R. (2014). "International trade puzzles: A solution linking production and preferences." *Quarterly Journal of Economics*, *129*(3), pp.1501-1552.

Choi, Y., Hummels, D., and Xiang, C. (2009), "Explaining import quality: the role of the income distribution" *Journal of International Economics* 77, 265-276.

Chor, D., and Manova, K. (2012). "Off the Cliff and Back? Credit Conditions and international trade during the global financial crisis." *Journal of International Economics*, 87(1), pp. 117-133.

Dalgin, M., Trindade, V., and Mitra, D. (2008), "Inequality, nonhomothetic preferences, and trade: a gravity approach" *Southern Economic Journal* 74(3) 747-774.

Eaton, J, Kortum, S, and Kramarz, F. (2011). "An Anatomy of International Trade: Evidence from French Firms." *Econometrica*, 79(5), pp.1453-1498.

Eaton, J., Kortum, S., Neiman, B., and J. Romalis, (2016). "Trade and the Global Recession." *American Economic Review*, 106(11), pp. 3401-3438.

Faber, B., and Fally, T. (2017). "Firm heterogeneity in consumption baskets: Evidence from home and store scanner data." NBER Working Paper (No.23101).

Fajgelbaum, P. D., and Khandelwal, A. K. (2016). "Measuring the unequal gains from trade." *The Quarterly Journal of Economics*, 131 (3), pp.1113-1180

Feenstra, R. C. (1994). "New Product Varieties and the Measurement of International Prices," *American Economic Review*, 84(1), pp. 157-177.

Fieler, A. C. (2011). "Nonhomotheticity and bilateral trade: Evidence and a quantitative explanation." *Econometrica*, *79*(4), pp.1069-1101.

Flam, H. and Helpman, E. (1987). "Vertical product differentiation and north-south trade" *American Economic Review* 77, 810-822.

Francois, J.F., Kaplan, S. (1996). “Aggregate demand shifts, income distribution, and the Linder hypothesis.” *Review of Economics and Statistics* 78 (2), pp. 244–250.

Haddad, M., Harrison, A., and Hausman, C. (2010). “Decomposing the great trade collapse: Products, prices, and quantities in the 2008-2009 crisis.” NBER Working Paper (No.16253).

Handbury, J. (2013). “Are poor cities cheap for everyone? Non-homotheticity and the cost of living across US cities.” *Wharton mimeo*.

Handbury, J., and Weinstein, D. E. (2014). “Goods prices and availability in cities.” *The Review of Economic Studies*, 82 (1), pp.258-296.

Haurin, D. R. (1991). “Income variability, homeownership, and housing demand.” *Journal of Housing Economics*, *1*(1), pp.60-74.

He, Zheli (2017). “Trade and Real Wages of the Rich and Poor: Cross-Country Evidence.” *Columbia University mimeo*.

Hunter, L. and Markusen, J. R. (1988). “Per Capita Income as a Basis for Trade.” in Robert Feenstra, ed., *Empirical Methods for International Trade*, MIT Press, Cambridge MA, London

Hunter, L. (1991). “The contribution of nonhomothetic preferences to trade.” *Journal of International Economics*, 30(3), pp. 345-358.

Ioannides, Y. M., and Rosenthal, S. S. (1994). “Estimating the consumption and investment demands for housing and their effect on housing tenure status.” *Review of Economics and Statistics*, pp.127-141.

Jaravel, X. (2016). “The unequal gains from product innovations.” SIEPR mimeo.

Kee, H. L. and Krishna, K. (2008). “Firm-Level Heterogeneous Productivity and Demand Shocks: Evidence from Bangladesh.” *American Economic Review: Papers and Proceedings*, 98(2), pp. 457-462.

Levchenko, A. A., Lewis, L. T., and Tesar, L. L. (2010). “The collapse of international trade during the 2008–09 crisis: in search of the smoking gun.” *IMF Economic Review*, *58*(2), pp.214-253.

Levchenko, A.A., Lewis,L.T., and Tesar, L. L. (2011). “The "Collapse in Quality" Hypothesis.” *American Economic Review: Papers and Proceedings*, *101*(3), 293-297.

Linder, S. B. (1961). “An Essay on Trade and Transformation,” Almqvist and Wiksells, Uppsala.

Lawless, M., and K. Whelan. (2008). “Where Do Firms Export, How Much, and Why?” Mimeo, University College Dublin.

Markusen, J. R. (1986). “Explaining the Volume of Trade: An Eclectic Approach.” *American Economic Review*, 76(5), pp.1002-1011.

Melitz, Marc (2003). “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity.” *Econometrica*, 71, pp.1695-1725.

Munch, J. R. and Nguyen, D. X. (2014). “Decomposing Firm-level Sales Variation.” *Journal of Economic Behavior & Organization*, *106*, pp.317-334.

Nguyen, D. X. (2012). “Demand uncertainty: Exporting delays and exporting failures.” *Journal of International Economics,* 86, pp.336-344.

Polinsky, A. M. (1977). “The demand for housing: A study in specification and grouping.” *Econometrica*, pp. 447-461.

Soderbery, A. (2010). “Investigating the asymptotic properties of import elasticity estimates.” *Economics Letters*, 109(2), pp.57-62.

Soderbery, A. (2015). “Estimating import supply and demand elasticities: Analysis and implications.” *Journal of International Economics*, 96(1), pp.1-17.

Thursby, J. G. and Thursby, M. C. (1987). “Bilateral Trade Flows, the Linder Hypothesis, and Exchange Risk.” *Review of Economics and Statistics*, 69(3), pp. 488-495.

Zorn, P. M. (1993). “The Impact of Mortgage Qualification Criteria on Households′ Housing Decisions: An Empirical Analysis Using Microeconomic Data.” *Journal of Housing Economics*, *3*(1), pp. 51-75.

<Table 1: OLS and IV Regression Results>

VARIABLES	(1)	(2)	(3)	(4)
	Dep Var: $dlnIM_{gt}$			
$dlnE_{gt}$	0.0328***	0.0301***	0.0307***	0.0273***
	(0.0103)	(0.0103)	(0.00963)	(0.00921)
Observations	5,349	5,349	5,349	5,349
R-squared	0.002	0.051	0.040	0.090
Product FE	No	Yes	No	Yes
Year FE	No	No	Yes	Yes

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Dep Var: $dlnIM_{gt}$				Dep Var: $\ln IM_{gt}$			
$dln\hat{E}_{gt}$	0.098**	0.138***	0.104**	0.154***				
	(0.042)	(0.051)	(0.046)	(0.058)				
$ln\hat{E}_{gt}$					0.608***	0.164***	0.617***	0.201***
					(0.114)	(0.051)	(0.114)	(0.055)
Observations	5,349	5,349	5,349	5,349	5,708	5,708	5,708	5,708
R-squared	0.001	0.051	0.040	0.090	0.182	0.912	0.187	0.915
Product FE	No	Yes	No	Yes	No	Yes	No	Yes
Year FE	No	No	Yes	Yes	No	No	Yes	Yes

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Dep Var: $dlnE_{gt}$				Dep Var: $\ln E_{gt}$			
$dlnE'_{gt}$	0.552***	0.467***	0.491***	0.385***				
	(0.072)	(0.080)	(0.074)	(0.081)				
$\ln E'_{gt}$					0.869***	0.557***	0.868***	0.531***
					(0.016)	(0.038)	(0.017)	(0.037)
F-stat	59.26	34.91	18.36	17.03	2804.68	223.32	454.27	52.54

Note: Clustered robust standard error in parentheses, *** p<0.01, ** p<0.05, * p<0.1. In all regressions, a constant term is estimated but suppressed.

<Table 2: Testing Baseline Specification>

Quadratic Term	QUAIDS	Dep Var: $dlnIM_{gt}$			
δ_{gt} (g-t specific)	All Income Elasticities (Baseline)	0.0981**	0.138***	0.104**	0.154***
	Income elasticities replaced with 1 if p-value of joint test>0.1	0.0929**	0.124**	0.103**	0.149**
δ_g (g-specific)	All Income Elasticities	0.0857**	0.125***	0.0846**	0.135***
	Income elasticities replaced with 1 if p-value of joint test >0.1	0.0862**	0.115**	0.0894**	0.125**
n/a	AIDS	Dep Var: $dlnIM_{gt}$			
	All Income Elasticities	0.0883***	0.124***	0.0893**	0.133***
	Income elasticities replaced with 1 if p-value of β_g >0.1	0.0869**	0.124***	0.0853**	0.130***
Product FE		No	Yes	No	Yes
Year FE		No	No	Yes	Yes

Note: In all regressions, we use clustered robust standard error, and constant terms estimated, but suppressed.
*** p<0.01, ** p<0.05, * p<0.1.

<Table 3: Sensitivity to Sample Period >

VARIABLES	(1)	(2)	(3)
	Dep Var: $dlnIM_{gt}$		
$dln\widehat{E}_{gt}$	0.145**	0.154***	0.160**
	(0.0623)	(0.0583)	(0.0623)
λ_t		-0.00475	-0.0171
		(0.0161)	(0.0206)
$\lambda_t * \ dln\widehat{E}_{gt}$			-0.0947
			(0.207)
Observations	4,539	5,349	5,349
R-squared	0.074	0.090	0.090

Note: Cluster robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. In all regressions, good and year fixed effect are included. Constant terms are estimated but suppressed. In column (1), we omit the GTC period. In column (2)−(3), we include a dummy variable for the GTC period and an interaction between the dummy with the expenditure variable. λ_t is the dummy variable which equals to one if 2008q4≤ t ≤ 2009q2 or zero otherwise.

<Table 4: Sensitivity to Additional GTC Controls>

VARIABLES	(1)	(2)	(3)	(4)
	Dep Var: $dlnIM_{gt}$			
$dln\hat{E}_{gt}$	0.105**	0.099**	0.101**	0.097**
	(0.046)	(0.045)	(0.047)	(0.045)
$(\text{Accounts payable/cost of goods sold})_g$		-0.140*		-0.129*
		(0.075)		(0.076)
$(\text{Accounts receivables/sales})_g$		0.103		0.085
		(0.161)		(0.167)
λ_t			0.012	0.012
			(0.026)	(0.238)
$\lambda_t * \ dln\hat{E}_{gt}$			0.126	0.031
			(0.196)	(0.188)
$\lambda_t * \ (\text{Accounts payable/cost of goods sold})_g$				-0.225
				(0.186)
$\lambda_t * \ (\text{Accounts receivables/sales})_g$				0.207
				(0.432)
$dln(Trade\ cost)_{gt}$	-0.135	-0.137	-0.134	-0.134
	(0.229)	(0.228)	(0.229)	(0.228)
Observations	5,349	5,349	5,349	5,349
R-squared	0.041	0.042	0.041	0.043

Note: Clustered robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. In all regressions, year fixed effect is included. Constant terms are estimated but suppressed. λ_t is the dummy variable which equals to one if 2008q4$\leq$ t $\leq$ 2009q2 or zero otherwise.

<Table 5: Magnitude of GTC Response: Cross Good Variation>

t	$dlnE_{gt}$			Expenditure Effect	Imports Actual	Expenditure Effect	Imports Actual
	10^{th}	50^{th}	90^{th}	$10^{th} - 50^{th}$	$10^{th} - 50^{th}$	$10^{th} - 90^{th}$	$10^{th} - 90^{th}$
2008q4	-0.595	-0.162	0.276	-6.66	-16.81	-13.41	-37.48
2009q1	-0.608	-0.113	0.362	-7.61	-29.87	-14.93	-50.88
2009q2	-0.583	-0.126	0.312	-7.05	-22.52	-13.78	-40.94

Note: To get the expenditure effect, we use the estimated elasticity of 0.154 from Table 1. Change in expenditures generates a predicted change in imports about 25-40 percent as large as the variation in import changes we observe in the data.

Figure 1: Properties of Income Elasticity

$b_g > 0, d_g < 0$:

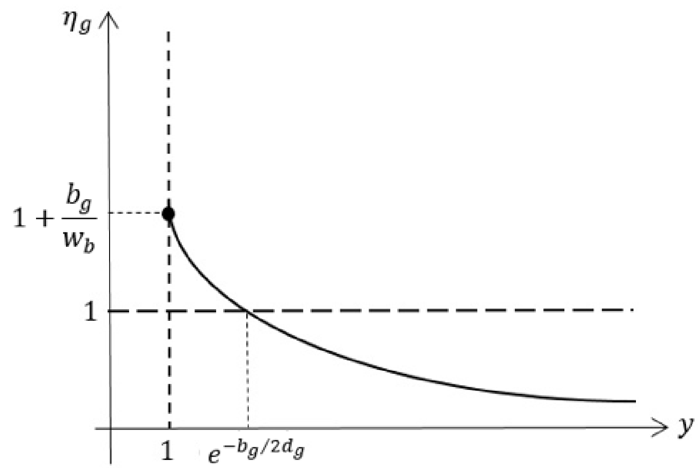

$b_g < 0, d_g > 0$

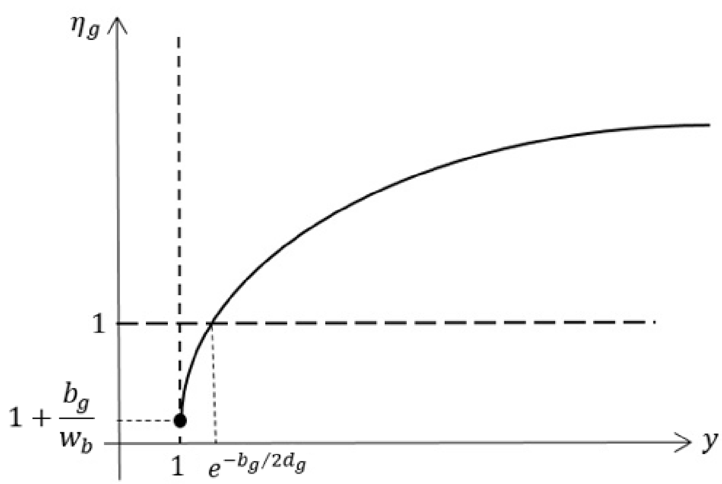

Figure 2: Budget Shares for by Household Income

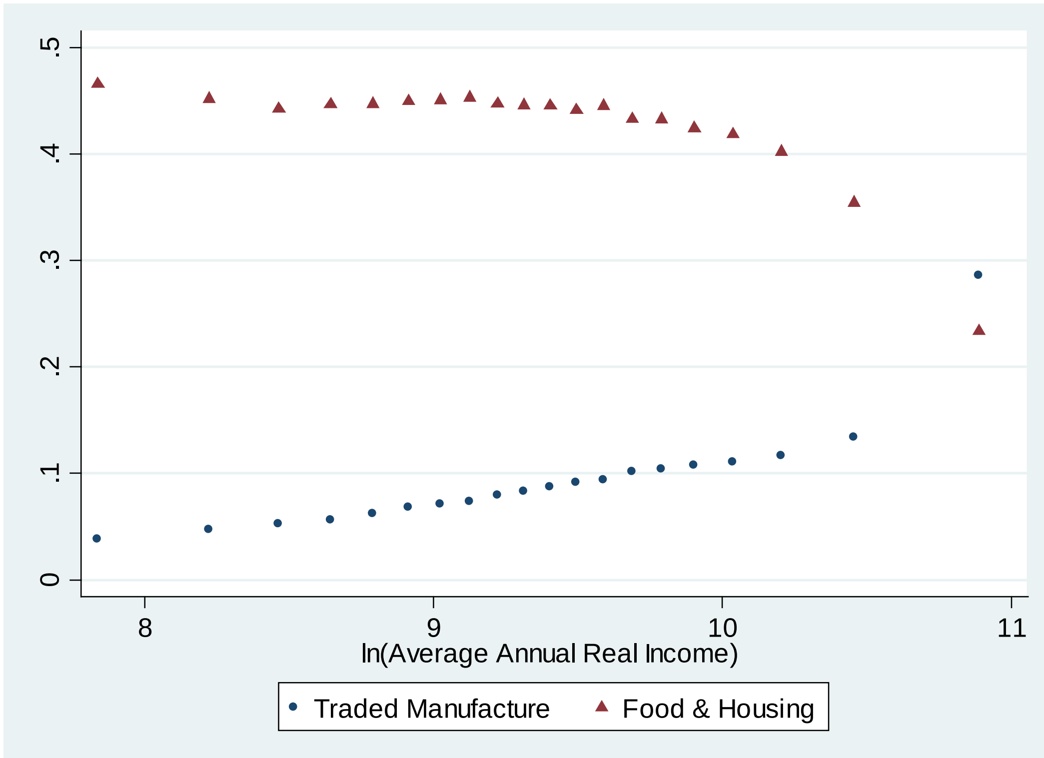

Note: Each point corresponds to the (over-time) average budget shares for traded manufactures (other than food) and food and housing for each of 20 income bins in our data.

Figure 3: Final Consumption Share in Manufactured Imports

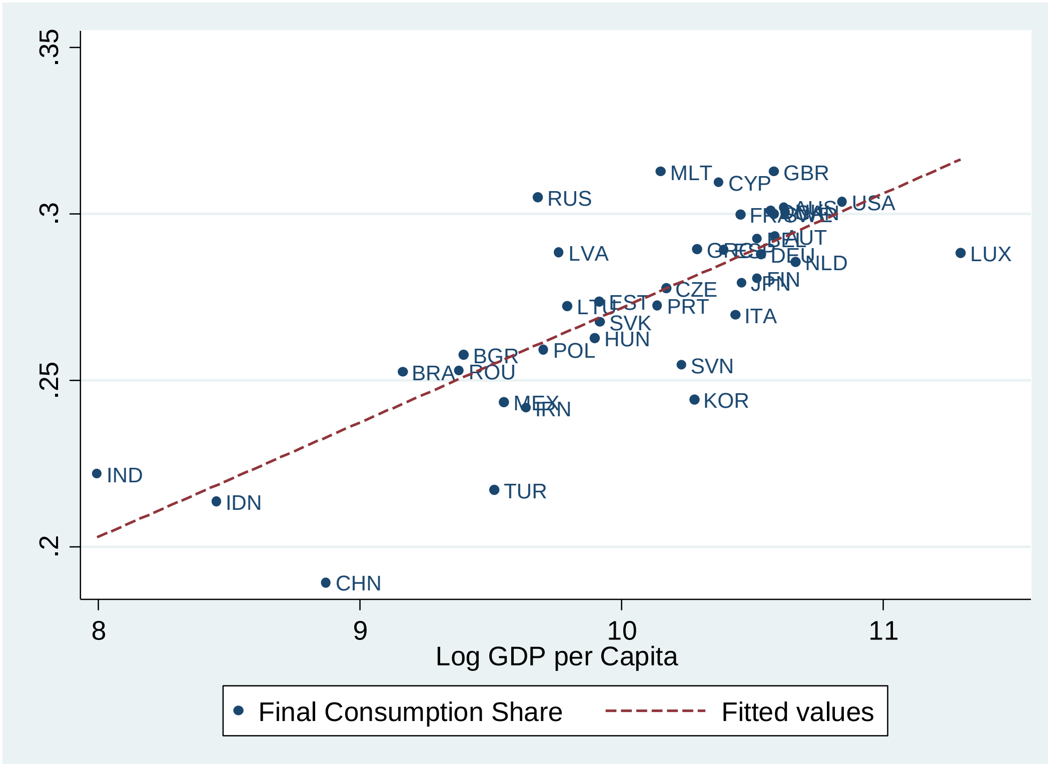

Figure 4: Income Elasticities for Two Goods – Across Incomes and Over Time

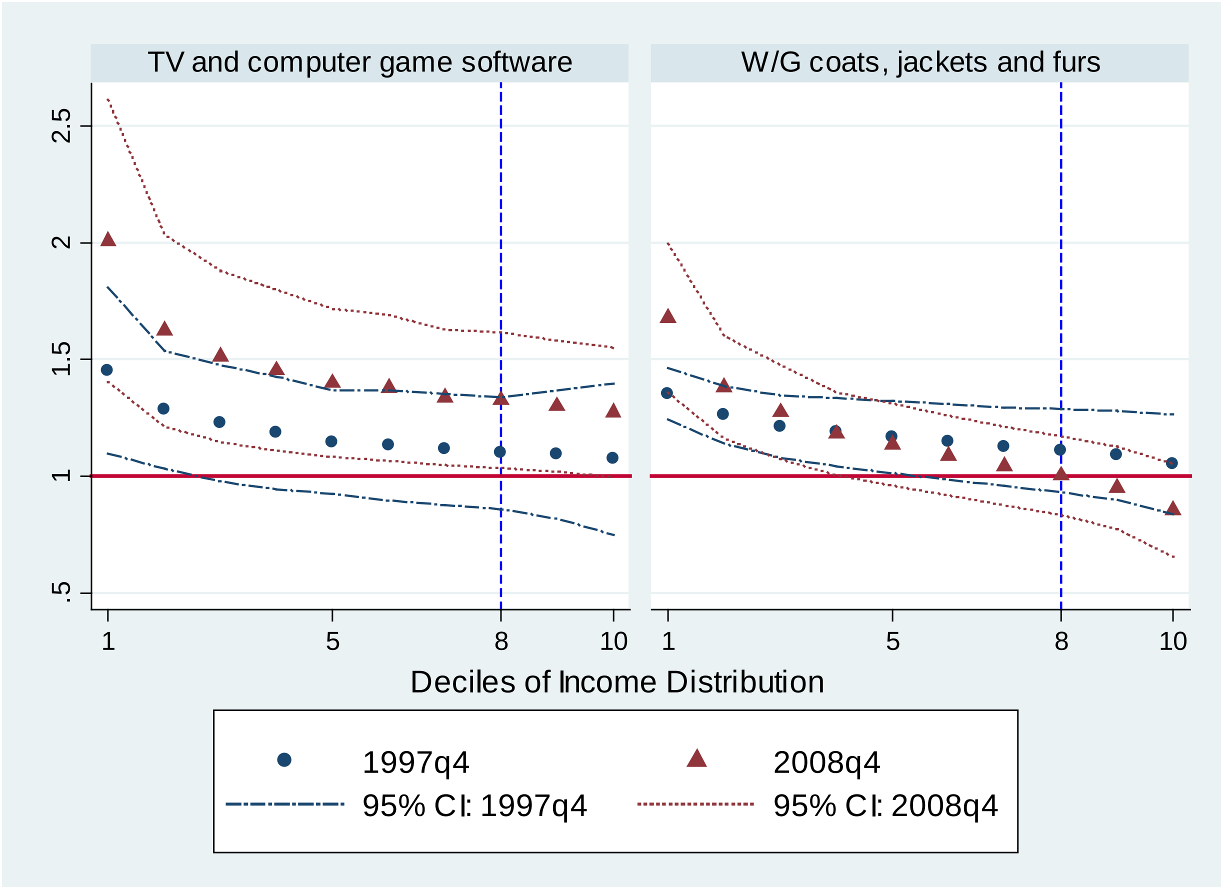

Figure 5: Comparing QUAIDS and AIDS Income Elasticity at Different Income Level

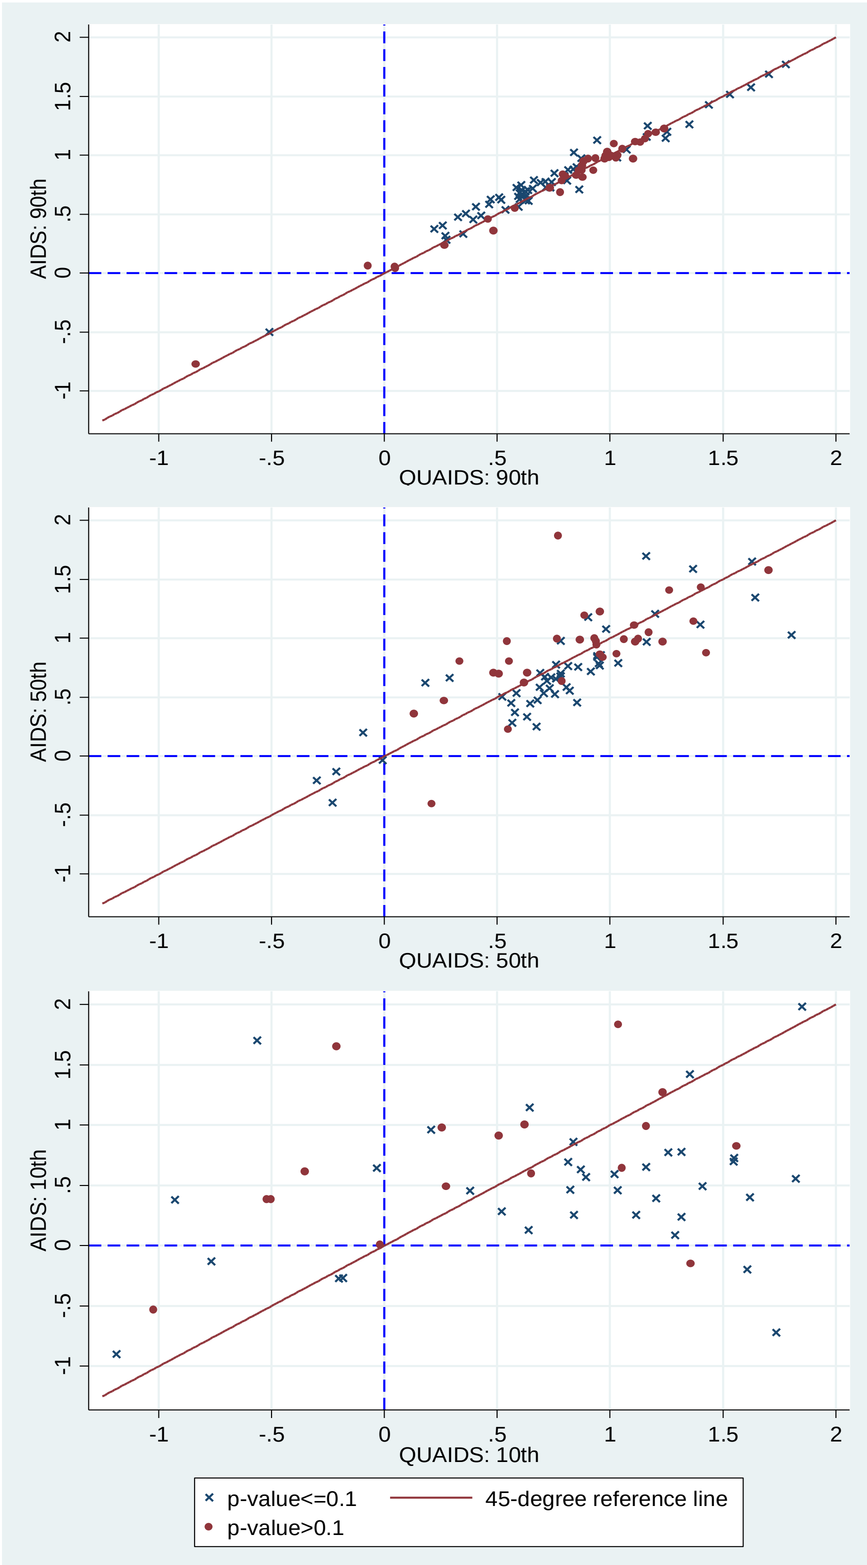

Note: Each data point represents an income elasticity for one CEX product in a single year, 2001q1.

Figure 6: Departure from CRIE Baseline

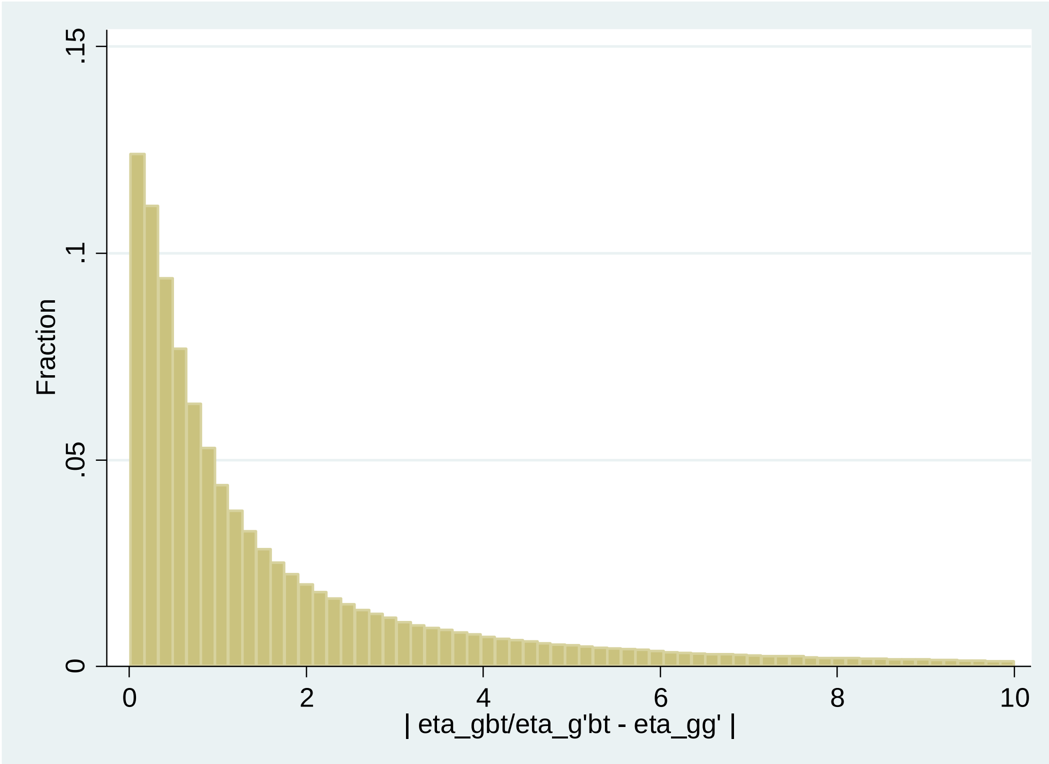

Note: For each good g we calculate the income elasticity for each bin-time and express it relative to the income elasticity for good g' in the corresponding bin-time. Subtracting the mean ratio for gg' generates the distribution of values above. The CRIE baseline is a value of 0.

Figure 7: Mean-Preserving Shocks and Aggregate Expenditure Change for Two Goods

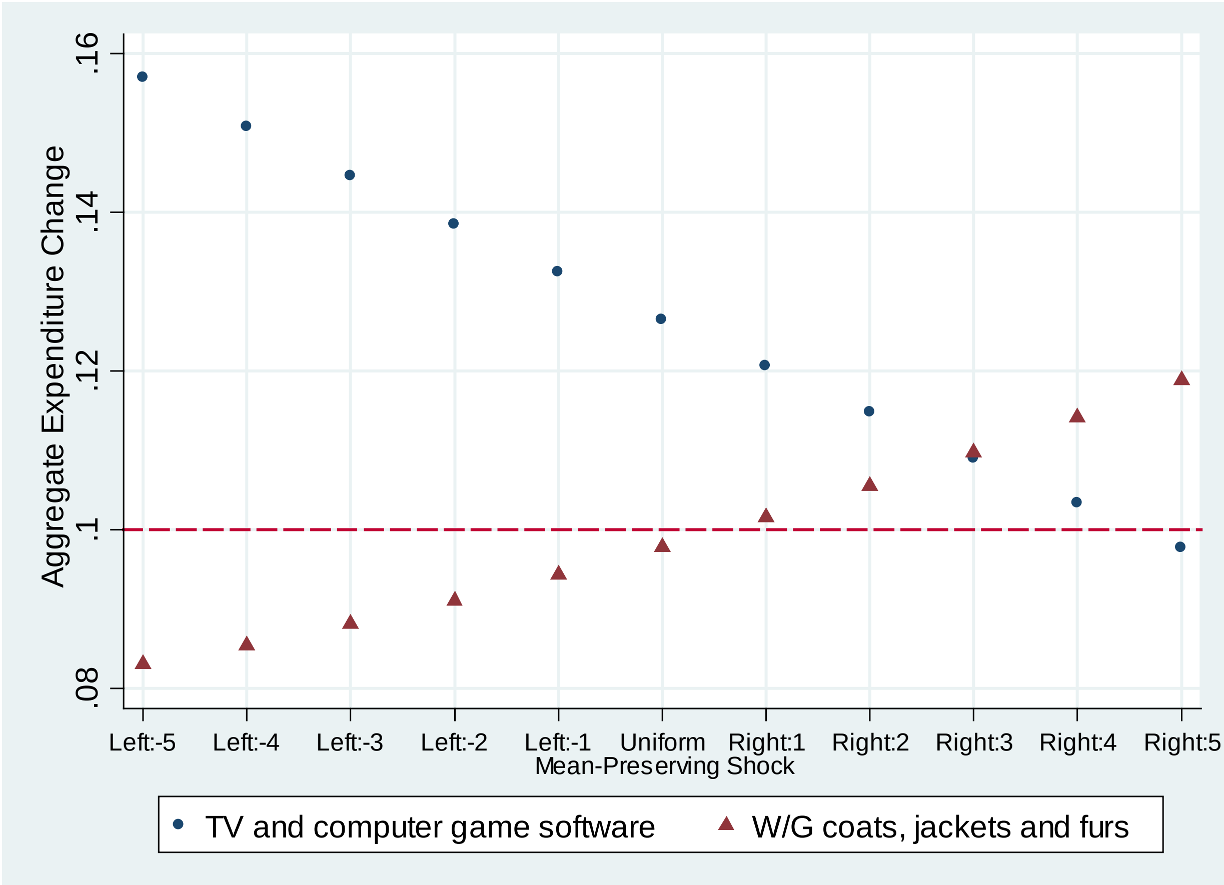

Note: Figure 7 reports aggregate expenditure changes of two expenditure categories under a variety of mean preserving shocks that raise aggregate incomes by 10%. We generate skew using the equation: $y_{bt}^{i} = \alpha^{i} + \left(1.1 - \frac{10 \cdot \alpha^{i}}{\sum_b y_{b,t-1}}\right) y_{b,t-1}$ For right-skewed shocks, we changing α^{i}, intercept values, from 1 to 5 while for the left-skewed shocks, the intercept changes from -5 to -1

Figure 8: The Distribution of Predicted Expenditure Change and Import Change

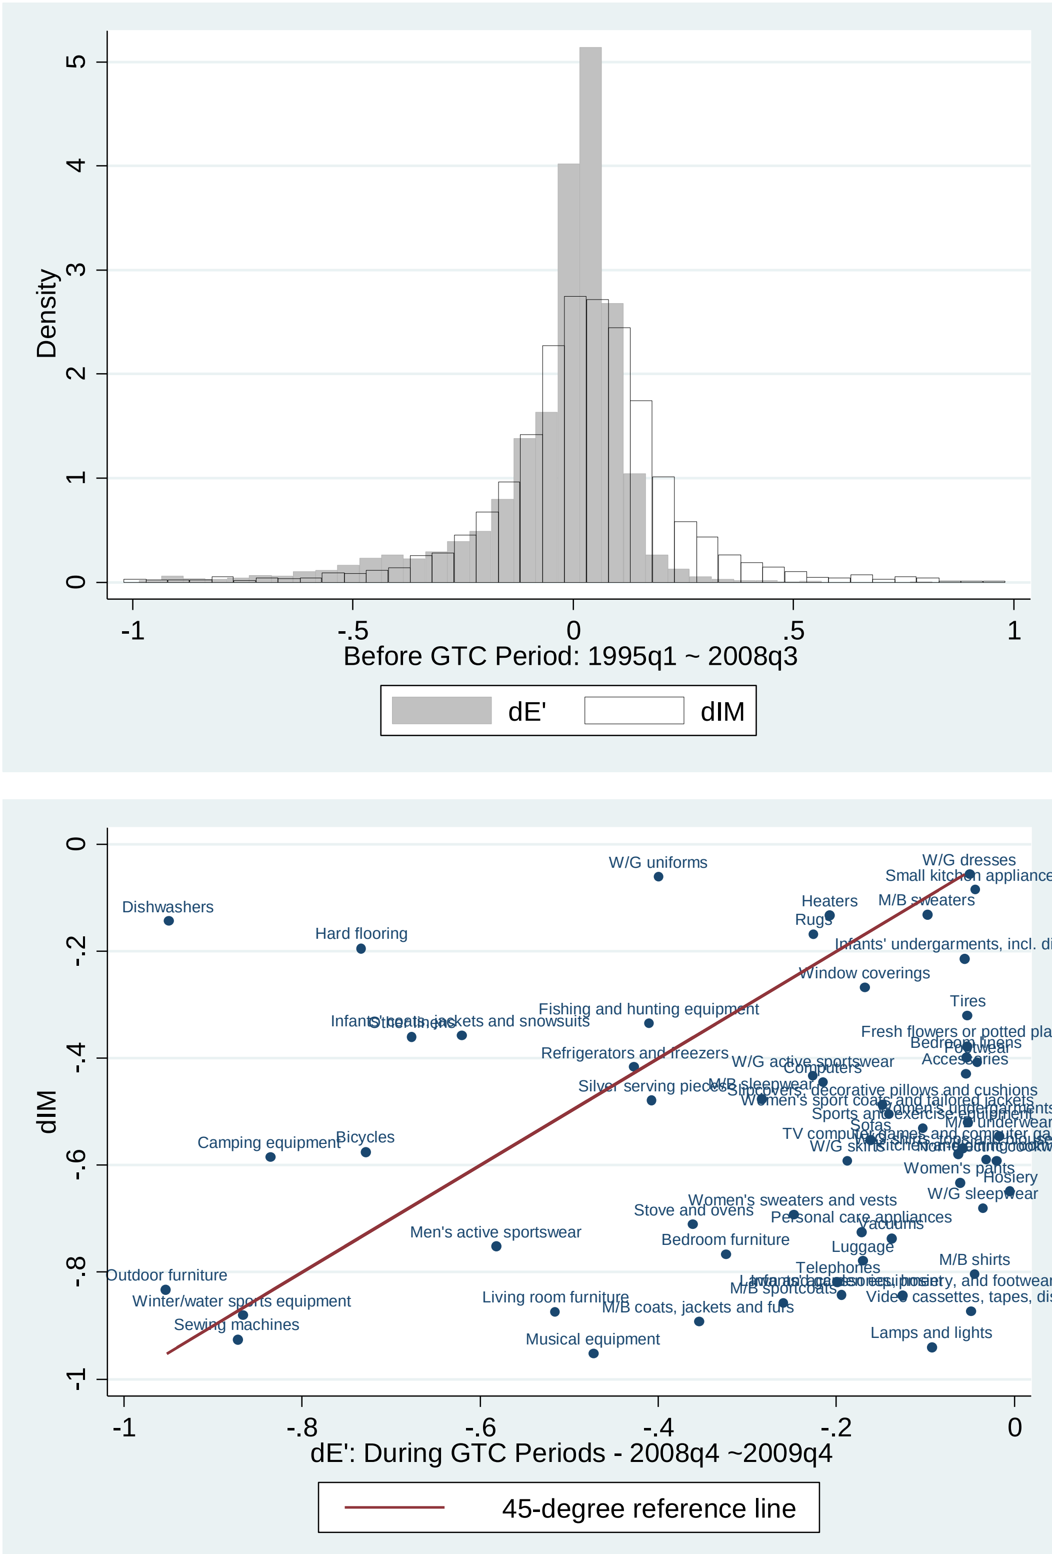

Figure 9: Pass-through of Aggregate Expenditure Share onto Import Share