

NBER WORKING PAPER SERIES

THE INCOME ELASTICITY OF IMPORT DEMAND:
MICRO EVIDENCE AND AN APPLICATION

David Hummels
Kwan Yong Lee

Working Paper 23338
<http://www.nber.org/papers/w23338>

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

We thank Thibault Fally, Anson Soderbery, Chong Xiang, Masha Brussevich and Kan Yue for helpful comments and suggestions. All errors are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by David Hummels and Kwan Yong Lee. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Income Elasticity of Import Demand: Micro Evidence and An Application
David Hummels and Kwan Yong Lee
NBER Working Paper No. 23338
April 2017
JEL No. D12,D31,F10,F14

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-income-time. 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 40% as large as the raw variation in import declines.

David Hummels
Krannert School of Management
403 West State Street
Purdue University
West Lafayette, IN 47907-1310
and NBER
hummelsd@purdue.edu

Kwan Yong Lee
University of North Dakota
293 Centennial Drive Stop 8369
Grand Forks, ND 58202-8369
kwanyong.lee@business.und.edu

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 exploiting 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 8.9 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”.

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 time series 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 good 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 0.17.

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 10th percentile (large decreases) had an associated import decline 16 percentage points larger than a good with an expenditure change in the 90th percentile. The actual (10-90) gap in import change was on the order of 41 percentage points, suggesting that expenditure changes arising from the distribution of income shocks played a significant role 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 Francoise 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.

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 "Constant Relative Income Elasticity" preferences from Fieler (2011) and focus on explaining home bias and biases in the factor content of trade. Fajgelbaum and Khandelwal (2016) 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. Further, the LES system generates identical income elasticities for all non-subsistence goods. (See Appendix A). 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. Fajbelbaum and Khandelwal's (2016) cross-country evidence suggests that budget shares devoted to manufactures fall with income, and income elasticities for manufactures rise with income. This pattern lies at the heart of their conclusion that trade is pro-poor. 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, and associated income elasticities fall. While we do not perform any formal welfare calculations, it is hard to see how trade in manufactures could much benefit poor consumers in the US if they are spending as little as 5 percent of their income on these goods.

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 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 Weistein (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) focuses 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 on imported goods over time. Rather, we are interested in one aspect of expenditure change arising

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.

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 :

$$\frac{M_{gt}}{Y_t} = \frac{M_{gt}}{E_{gt}} \cdot \frac{E_{gt}}{Y_t} \quad (1)$$

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 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

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.

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:

$$dE_{gt} \equiv \ln\left(\frac{E_{gt}}{E_{g,t-4}}\right) = \ln\left(\sum_b S_{gb,t-4} \cdot \frac{E_{gbt}}{E_{gb,t-4}}\right) \quad (2)$$

where $S_{gb,t-4}$ 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:

$$w_g = \alpha_g + \sum_k \gamma_{gk} \ln p_k + \beta_g \ln\left(\frac{y}{P}\right) + \frac{\delta_g}{\prod_k p_k^{\beta_k}} \left(\ln\left(\frac{y}{P}\right)\right)^2 \quad (3)$$

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 α_g , β_g , δ_g and γ_{gk} are parameters.⁴ $\ln P$ is a price index defined as $\ln P =$

⁴ For well-behaved utility, the following restrictions are necessary: $\sum_g \alpha_g = 1$, $\sum_g \beta_g = 0$, $\sum_g \gamma_{gk} = \sum_k \gamma_{gk} = 0$, $\gamma_{gk} = \gamma_{kg}$, and $\sum_g \delta_g = 0$.

$\alpha_0 + \sum_g \alpha_g \log p_g + \frac{1}{2} \sum_g \sum_k \gamma_{gk} \log p_g \log p_k$ By setting $\delta_g = 0$ this system nests the more commonly used AIDS.

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

$$\eta_g = 1 + \left(\beta_g + \frac{2\delta_g}{\prod_k p_k^{\beta_k}} \ln \left(\frac{y}{P} \right) \right) \frac{1}{w_g} \quad (4)$$

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 β_g and δ_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 $\alpha_0 = 0$, so that $P = 1$. Then, η_g reduces to:

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

It is immediate that if both $\beta_g > 0$, $\delta_g > 0$, 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 β_g and δ_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:

$$w_{gbt} = \alpha_g + \sum_k \gamma_{gk} \ln p_{kt} + \beta_g \ln \left(\frac{y_{bt}}{P_t} \right) + \frac{\delta_g}{\prod_k p_{kt}^{\beta_k}} \left(\ln \left(\frac{y_{bt}}{P_t} \right) \right)^2 \quad (5)$$

We assume that demand parameters α_g , β_g , δ_g and γ_{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 However, our goal is to construct an instrument for expenditures that is orthogonal to supply shocks, based on consistently estimated income elasticities. That is to say, even if we had price data we would not want to incorporate it because it would invalidate the IV.

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 $\alpha_g + \sum_k \gamma_{gk} \ln p_{kt}$ can be eliminated by incorporating a good-time fixed effect, a_{gt} . We proxy for the QUAIDS-appropriate price index using the CPI. Note that the quadratic income term interacts with an aggregated measure of prices that is common across goods but varies over time, $\prod_k p_{kt}^{\beta_k}$. 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 time dummy T_t . To complete the specification we incorporate a vector of demographic characteristics \mathbf{X}_{bt} which may affect expenditures such as age of household head, family size, and location (urban/rural).

$$w_{gbt} = a_{gt} + \beta_g \ln \left(\frac{y_{bt}}{CPI_t} \right) + \delta_{gt}(T_t) \left(\ln \left(\frac{y_{bt}}{CPI_t} \right) \right)^2 + \boldsymbol{\beta} \mathbf{X}_{bt} + \varepsilon_{gbt} \quad (6)$$

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 $\widehat{\eta}_{gbt}$ for a household of income y_{bt} but with otherwise average demographic characteristics:

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

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

These elasticities are of independent interest. But 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 level is a

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, because these data cover a subset of the goods covered in the CEX, and because we are interested in variation in expenditures that is exogenous to changes in prices.

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:

$$dE'_{gt} \equiv \ln \left(\sum_b S'_{gb,t-4} \cdot \frac{E'_{gbt}}{E_{gb,t-4}} \right) \quad (9)$$

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 quarters⁶ 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 330 UCC, 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. 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.

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.⁷ 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), who show the share of aggregate expenditures on manufacturing falling in income? First it is notable that housing and food expenditure data displayed here is consistent with micro-household evidence showing income elasticities significantly below one for these categories.⁸ In the CEX data, rising expenditure shares on traded manufactures mirror declining expenditure shares on food and housing. Second, data taken from national accounts and trade statistics may differ in important ways from household expenditure data. This is hard to characterize with great specificity because it requires knowing details about data construction for many countries, but a few key areas seem a plausible source of difference. Spending on intermediate inputs will be included in national accounts and trade statistics but will be omitted from household expenditures, and attempts to split absorption into industrial versus

⁷ Complete information on the share of good g spending for each income bin b is captured in Appendix Table B.2.

⁸ 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.

household use necessarily relies on industrial and household survey data. Related, and particularly important in this context, personal consumption expenditures on an already-built housing stock is meant to be captured in national accounts data. However, implementing this requires sophisticated imputation of the rental value of owner-occupied housing. It is plausible to us that the quality of imputation in separating absorption and in identifying the value of the housing service flow might vary with the sophistication of the national statistical agency, and be missed for lower income countries.

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.

In Figure 3, we display income elasticities for four specific goods (Infants Undergarments; Watches; Bedroom Linens, and Women's Sweaters and Vests) with variation across income deciles at three points in time. While all show elasticities dropping with income, there are quite significant differences in the level of the elasticity (about 50% larger for Watches than for Infants Undergarments), the dispersion across income levels (much greater for Bedroom Linens), and in over-time changes in the elasticity (the watch elasticities rises then falls, while the women's sweater elasticities rise over time). Of particular note, Infants Undergarments and Women's sweaters are income elastic at low income levels, but income inelastic at higher income levels in 1998q1. 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 (over-time mean) income elasticities for each decile and good category in Appendix Table B.1. Income elasticities exhibit significant variation across goods within income bins, and across income bins within goods.

It is useful to compare these estimates with a baseline from an important new literature (Fieler 2011; Caron et al. 2014) in trade using 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. In the top panel of Figure 4 we calculate the relative income elasticity for Watches and Bedroom Linens at each decile and the three points in time. The ratio of elasticities rises with income levels, and at different rates over time. To be more systematic, 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' . The bottom panel of Figure 4 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 B.3, we report the share of bin b in aggregate spending on good g . Aggregating over all goods, the top two deciles are responsible for 49 percent of spending on traded goods, 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”. Even though Men’s Suits and Men’s Uniforms have quite similar income elasticities, and those elasticities vary only a little over the income distribution, 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 four goods shown in Figure 5. 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.⁹

Starting in the middle of Figure 5, we see that a uniform increase in incomes results in expenditure increases ranging from just over 13% (for infant undergarments) to over 18% (for women’s sweaters and vests). 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. When we skew these shocks to the right (giving more income to the poorest households and less to the richest), the expenditure response narrows for all but women’s sweaters and vests. Note that the responsiveness of aggregate expenditures to these distributional changes varies considerably over goods as a function of expenditure shares at each income level and the relevant income elasticities. For example, Watches have high income elasticities throughout the income distribution, and generate large expenditure responses to the income change. But that effect becomes more muted when income is given to poorer households because their baseline expenditure shares comprise only a small part of overall spending on watches. 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 that the distribution of income shocks during two recent recessions. Figure 6 shows changes in total expenditure during the Dot Com Crash, 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 more pronounced in the bottom decile and fifth-seventh deciles. 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

⁹ 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%.

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 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

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

where $d\left(\frac{M_{gt}}{Y_t}\right) \equiv \ln\left(\frac{M_{gt}}{Y_t} / \frac{M_{g,t-4}}{Y_{t-4}}\right)$ 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 7. 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, the magnitude of import changes differs somewhat from expenditure changes, with higher import growth rates pre-crisis, and larger import declines during the crisis. This is not surprisingly. After all, 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.

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

$$d\left(\frac{IM}{GDP}\right)_{gt} = \alpha \cdot d\left(\frac{E}{GDP}\right)_{gt} + d(MDIST)_{gt} + \beta_g + \gamma_t + u_{gbt} \quad (11)$$

$$dIM_{gt} = \delta(dE_{gt}) + d(MDIST)_{gt} + \theta_g + \rho_t + \varepsilon_{gbt} \quad (12)$$

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).

Note that by first differencing we eliminate level differences across goods in expenditure shares, and in supply characteristics such as price, quality, and variety. We incorporate good fixed effects (β_g or θ_g) to allow for good specific time trends in these components, and a year fixed effect (γ_t or ρ_t) to control for aggregate shocks that affect trade or the macroeconomy and are common to all goods. In addition, we use sum of import-share weighted distance between the US and its trading partners as a proxy for trade cost, $d(MDIST)_{gt}$. In all estimates we cluster standard errors on goods to account for serial correlation in the first differences.¹⁰

Conceptually, we can think of equation (11) as rewriting equation (1), and assuming that the ratio of imports to expenditures is absorbed into differencing, good fixed effects, or the error term. We return to this below. Equation (12) can be thought of as rewriting (1) after first multiplying both sides by GDP. The difference in the two functional forms is that when we scale imports and good expenditures by GDP, we absorb any effects of income shocks that reduce trade in proportion to changes in GDP.

We report OLS regression results for equations (11) and (12) in Table 1. For completeness we report estimates with and without good and year fixed effects. In each case we find a very small response, an elasticity of about 0.03.

Next, we instrument for actual expenditures using predicted expenditures arising from income shocks interacted with good x income bin x time varying income elasticities as described above. The top panel of Table 2 reports the second stage, the bottom panel the first stage. In the first stage we see that predicted changes in expenditures arising from income shocks is a very good predictor of actual expenditure changes. Coefficients are precisely estimated with an elasticity of

¹⁰ 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 $t+4$.

0.4 to 0.5, and F-stats are large. We have similar results whether we use the level of expenditure or scale it by GDP, and whether we include or exclude fixed effects. In the second stage we find a (highly significant) elasticity with largest coefficients (0.17) in our preferred specifications with saturated fixed effects.

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.¹¹ 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.

Table 1-2 use 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 specifications with good and time fixed effect. In columns 1, 2 we omit the GTC period and find quite similar elasticities to the comparable estimates in Table 2. In columns 3-6 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).

To explore this a little further, Table 4 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

¹¹ 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.

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 0.174 from Table 4, the associated import changes for 2008q4 are -10.3%; -2.8%; and +4.8%. Expressing that in percentage point differentials, we report in the table that the 10th percentile good had an associated import reduction that is 7.5 percentage points larger than the median good and 15.2 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 30-45 percent as large as the raw variation in import changes we observe in the data. The remainder of the table repeats this calculation for each quarter, and using variables scaled by GDP, with similar results.

Robustness

Our specification is intentionally stark, but we briefly describe some robustness checks related to data construction and other covariates. (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 2 specification with this sample. We see little qualitative change in our estimates in either the first or second stage

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.

Recalling equation (1), changes in imports for a good can arise either from a change in expenditures relative to GDP for that good or from changes in the share of imports in expenditures. While our focus is on the former, the trade literature has focused a great deal of attention on supply characteristics, including good price, quality, variety or trade costs between partners. 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 the covariates found in that literature are collinear with our fixed effects, i.e. our results are already robust to these explanations. Exceptions include trade credit measures from Levchenko et al. (2010), and inventory data in the spirit of Alessandria et al. (2010). These variables are measured at the industry level. Using industry rather than CEX good data required us to aggregate our data considerably, and to eliminate roughly 2/3 of our observations. There are two notable features of these results. First, restricting our sample to match these industry level observations does not affect the point estimates on predicted expenditure in baseline specifications that omit the additional variables. However, losing so many observations results in a loss of statistical significance. Second, incorporating these additional controls has no effect on the point estimates. From this we tentatively conclude that these supply side explanations for trade collapse are orthogonal to the income-induced expenditure change we are focused on.

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 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. It would be intriguing to understand how one can find such different patterns looking across countries as opposed to looking within a given country at a point in time. As we discuss in the text, this could be a simple measurement issue tied to how intermediate inputs and expenditures on the housing stock are captured in the two approaches. Of more economic interest, we speculate that the effect could be related to the price of non-traded goods and services. That is, as incomes rise and push up the price of non-traded goods and services such as housing, these price changes absorb a rising portion of household income for the poor, leaving relatively little to spend on traded manufactures.

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.

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 Bilal, 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.

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.

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.

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.

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 Regression Results>

VARIABLES	(1) dIM_{gt}	(2) dIM_{gt}	(3) dIM_{gt}	(4) dIM_{gt}	(5) $d(IM/GDP)_{gt}$	(6) $d(IM/GDP)_{gt}$	(7) $d(IM/GDP)_{gt}$	(8) $d(IM/GDP)_{gt}$
dE_{gt}	0.0333*** (0.0114)	0.0332*** (0.0116)	0.0320*** (0.0100)	0.0314*** (0.0100)				
$d(E/GDP)_{gt}$					0.0286** (0.0115)	0.0282** (0.0118)	0.0315*** (0.0101)	0.0308*** (0.0101)
$d(MDIST)_{gt}$	-0.145 (0.229)	-0.192 (0.223)	-0.143 (0.229)	-0.188 (0.221)	-0.135 (0.230)	-0.181 (0.224)	-0.142 (0.229)	-0.187 (0.221)
Constant	0.00860 (0.00784)	-0.00113 (0.00236)	-0.00773 (0.0239)	-0.0156 (0.0202)	-0.0118 (0.00780)	-0.0210*** (0.00228)	-0.0331 (0.0240)	-0.0410** (0.0203)
Observations	5,400	5,400	5,400	5,400	5,400	5,400	5,400	5,400
R-squared	0.003	0.057	0.042	0.096	0.003	0.057	0.029	0.084
Product FE	No	Yes	No	Yes	No	Yes	No	Yes
Year FE	No	No	Yes	Yes	No	No	Yes	Yes

Clustered standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

<Table 2: IV Regression Results >

VARIABLES	(1) dIM_{gt}	(2) dIM_{gt}	(3) dIM_{gt}	(4) dIM_{gt}	(5) $d\left(\frac{IM}{GDP}\right)_{gt}$	(6) $d\left(\frac{IM}{GDP}\right)_{gt}$	(7) $d\left(\frac{IM}{GDP}\right)_{gt}$	(8) $d\left(\frac{IM}{GDP}\right)_{gt}$
$d\hat{E}_{gt}$	0.0978** (0.0377)	0.131*** (0.0457)	0.116*** (0.0378)	0.174*** (0.0445)				
$d(\hat{E}/GDP)_{gt}$					0.0772* (0.0407)	0.103** (0.0493)	0.112*** (0.0393)	0.170*** (0.0468)
$d(MDIST)_{gt}$	-0.145 (0.230)	-0.192 (0.224)	-0.145 (0.229)	-0.191 (0.221)	-0.135 (0.230)	-0.181 (0.224)	-0.144 (0.229)	-0.190 (0.221)
Constant	0.0142* (0.00767)	0.000116 (0.00217)	-0.0248 (0.0255)	-0.0536** (0.0230)	-0.00661 (0.00783)	-0.0186*** (0.00204)	-0.0474* (0.0252)	-0.0741*** (0.0227)
Observations	5,400	5,400	5,400	5,400	5,400	5,400	5,400	5,400
R-squared	0.003	0.057	0.042	0.096	0.002	0.057	0.029	0.084
Good FE	No	Yes	No	Yes	No	Yes	No	Yes
Year FE	No	No	Yes	Yes	No	No	Yes	Yes

VARIABLES	dE_{gt}	dE_{gt}	dE_{gt}	dE_{gt}	$d\left(\frac{E}{GDP}\right)_{gt}$	$d\left(\frac{E}{GDP}\right)_{gt}$	$d\left(\frac{E}{GDP}\right)_{gt}$	$d\left(\frac{E}{GDP}\right)_{gt}$
dE'_{gt}	0.527*** (0.0352)	0.477*** (0.0411)	0.458*** (0.0437)	0.379*** (0.0504)				
$d(E'/GDP)_{gt}$					0.518*** (0.0360)	0.469*** (0.0416)	0.452*** (0.0442)	0.375*** (0.0507)
F-stat	223.55	137.51	28.58	26.39	207.05	129.2	25.72	23.35

Clustered standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

<Table 3: Sensitivity to Sample Period >

VARIABLES	(1) dIM_{gt}	(2) $d\left(\frac{IM}{GDP}\right)_{gt}$	(3) dIM_{gt}	(4) dIM_{gt}	(5) $d\left(\frac{IM}{GDP}\right)_{gt}$	(6) $d\left(\frac{IM}{GDP}\right)_{gt}$
$d\hat{E}_{gt}$	0.189*** (0.0482)		0.175*** (0.0445)	0.181*** (0.0467)		
$d(\hat{E}/GDP)_{gt}$		0.181*** (0.0503)			0.170*** (0.0469)	0.176*** (0.0491)
λ_t			0.00507 (0.0186)	-0.00736 (0.0201)	0.00978 (0.0186)	-0.000178 (0.0195)
$\lambda_t * d\hat{E}_{gt}$				-0.0870 (0.0621)		
$\lambda_t * d(\hat{E}/GDP)_{gt}$						-0.0848 (0.0626)
$d(MDIST)_{gt}$	-0.280 (0.211)	-0.278 (0.211)	-0.191 (0.221)	-0.191 (0.221)	-0.190 (0.221)	-0.190 (0.221)
Constant	-0.0706*** (0.0235)	-0.0900*** (0.0231)	-0.0536** (0.0230)	-0.0548** (0.0232)	-0.0742*** (0.0227)	-0.0752*** (0.0228)
Observations	4,580	4,580	5,400	5,400	5,400	5,400
R-squared	0.085	0.082	0.096	0.096	0.084	0.084

Clustered standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. In all regressions, good and year fixed effect are included.

Note: In column (1) and (2), we omit the GTC period. In column (3)–(6), we include a dummy variable for the GTC period as well as an interaction between the dummy with the expenditure variable. λ_t is the dummy variable which equals to one if $2008q4 \leq t \leq 2009q2$ or zero otherwise.

<Table 4: Magnitude of GTC Response: Cross Good Variation>

t	dE_{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	-7.52	-16.81	-15.15	-37.48
2009q1	-0.608	-0.113	0.362	-8.60	-29.87	-16.87	-50.88
2009q2	-0.583	-0.126	0.312	-7.96	-22.52	-15.57	-40.94
2009q3	-0.497	-0.078	0.353	-7.29	-22.00	-14.79	-38.34
2009q4	-0.542	-0.021	0.329	-9.07	-23.37	-15.15	-40.76

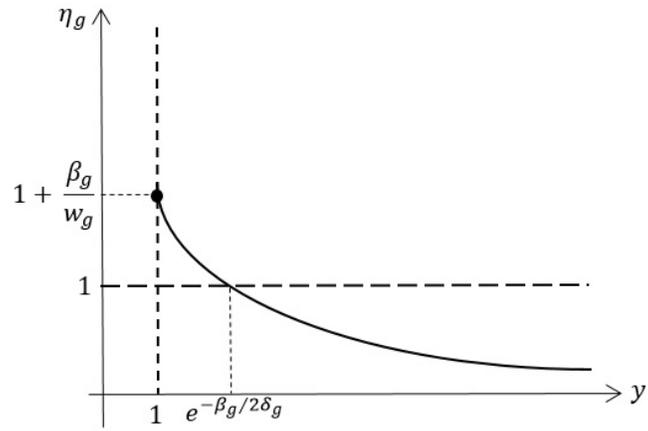
Note: To get the expenditure effect, we use the estimated elasticity of 0.174 from Table 3. Change in expenditures generates a predicted change in imports about 30-45 percent as large as the variation in import changes we observe in the data.

t	$d(E/GDP)_{gt}$			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	-0.569	-0.136	0.302	-7.35	-16.81	-14.80	-37.48
2009q1	-0.588	-0.094	0.381	-8.40	-29.87	-16.48	-50.88
2009q2	-0.562	-0.105	0.333	-7.78	-22.52	-15.22	-40.94
2009q3	-0.482	-0.063	0.368	-7.13	-22.00	-14.45	-38.34
2009q4	-0.529	-0.007	0.342	-8.86	-23.37	-14.80	-40.76

Note: To get the expenditure effect, we use the estimated elasticity of 0.170 from Table 3. When scaled by GDP, a predicted change in imports generated from change in expenditure explains about 28-44 percent variation in import changes observed in the data.

Figure 1: Properties of Income Elasticity

$\beta_g > 0, \delta_g < 0$:



$\beta_g < 0, \delta_g > 0$

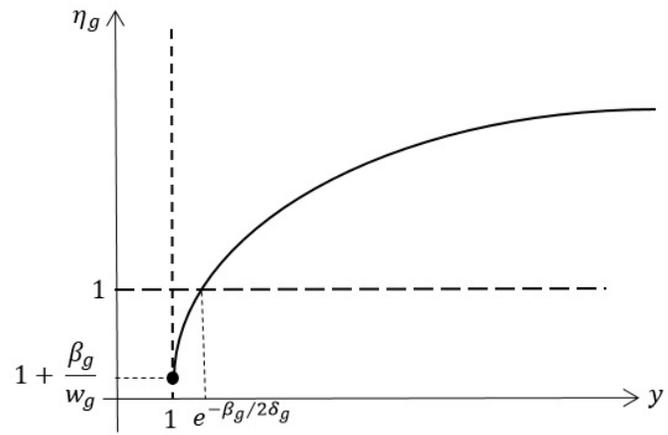
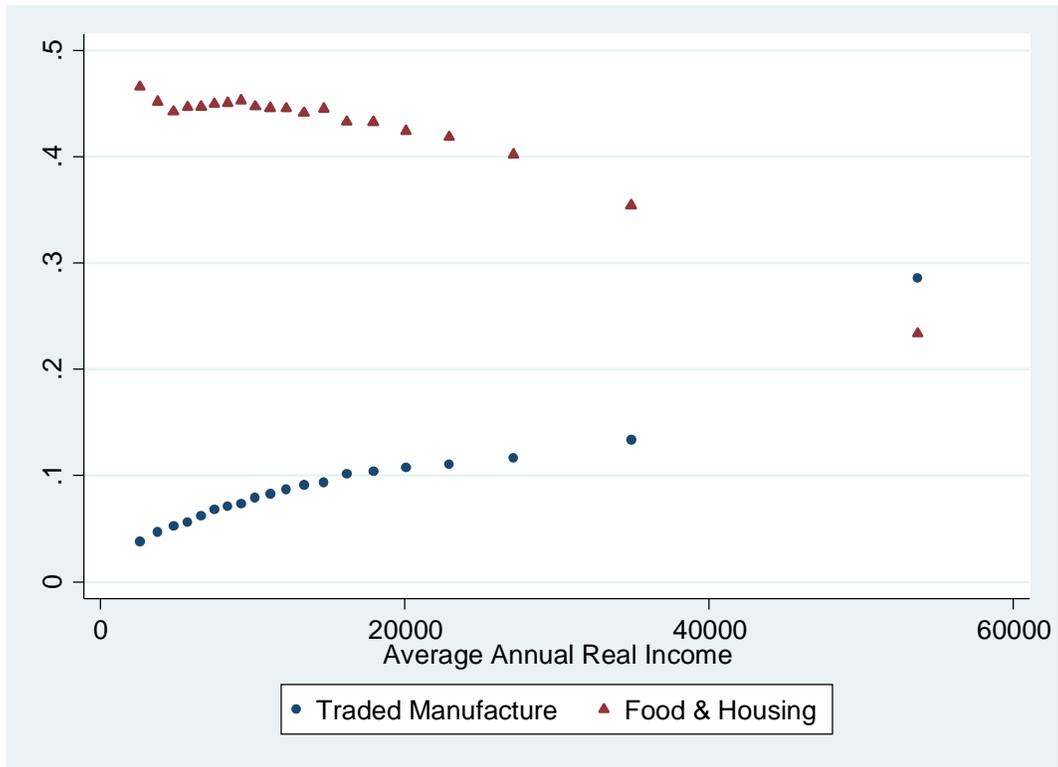


Figure 2: Budget Shares for Traded Manufactures, Food and Housing by Household Income



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: Income Elasticities for Four Goods – Across Incomes and Over Time

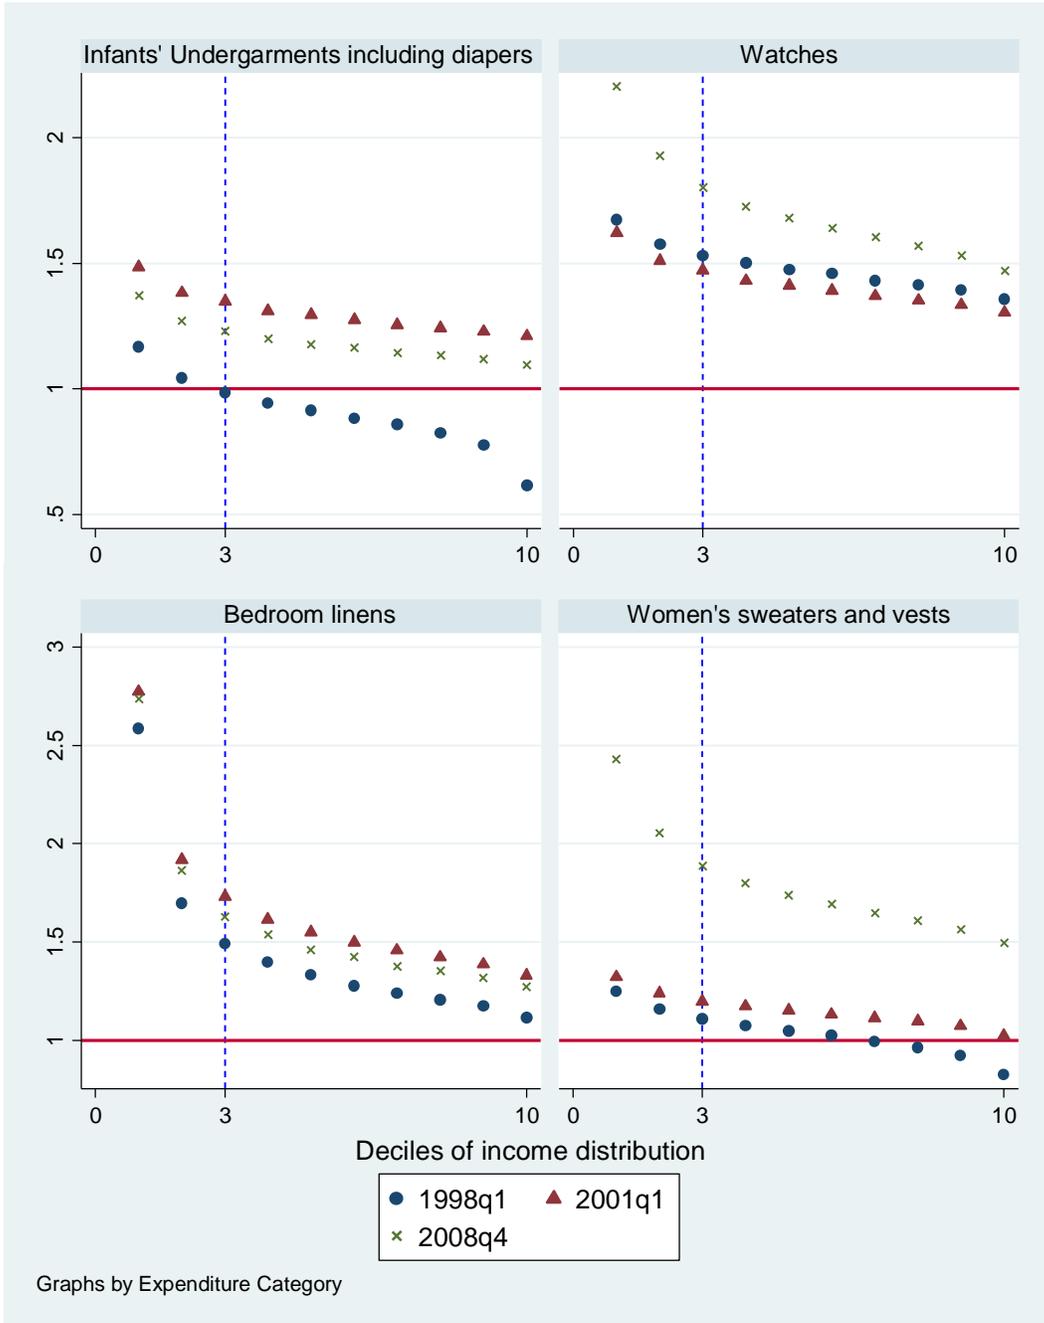
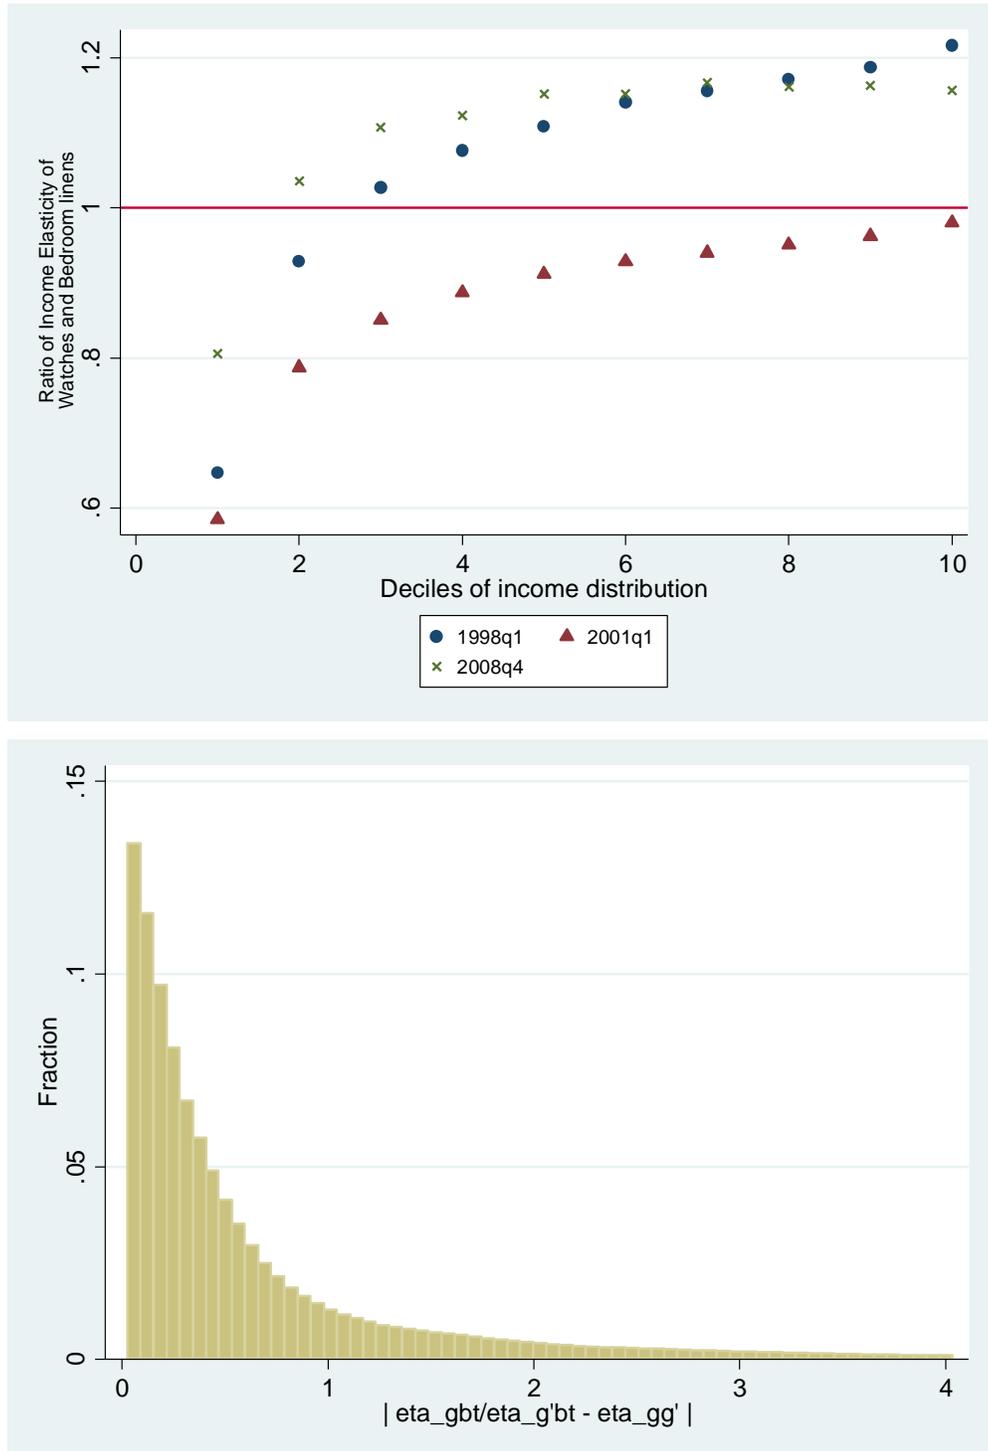
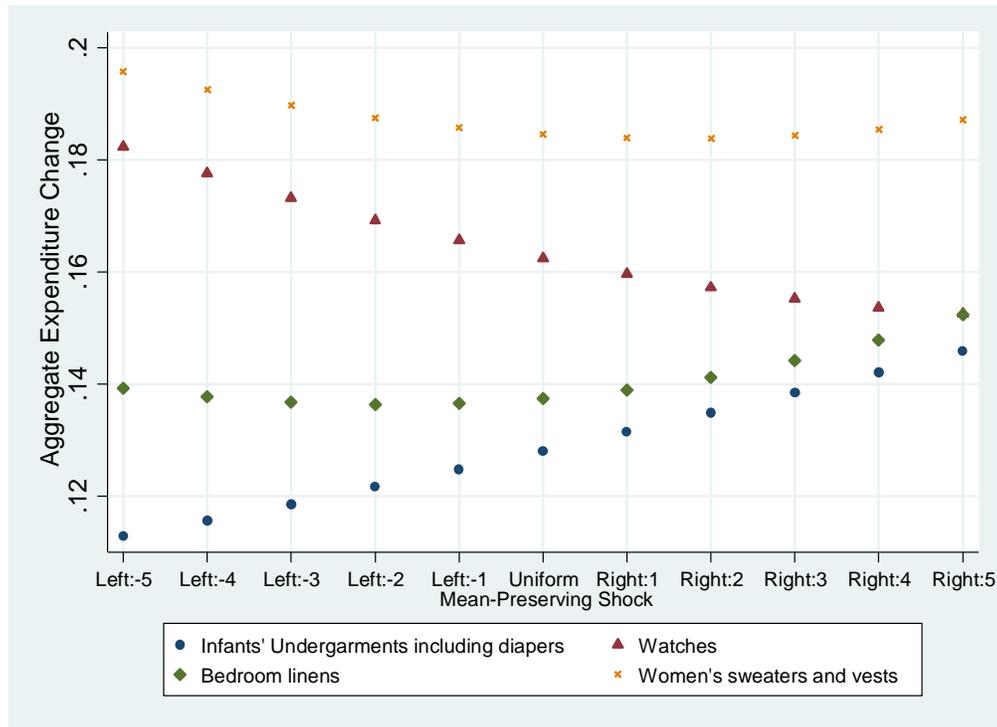


Figure 4: Departure from CRIE Baseline



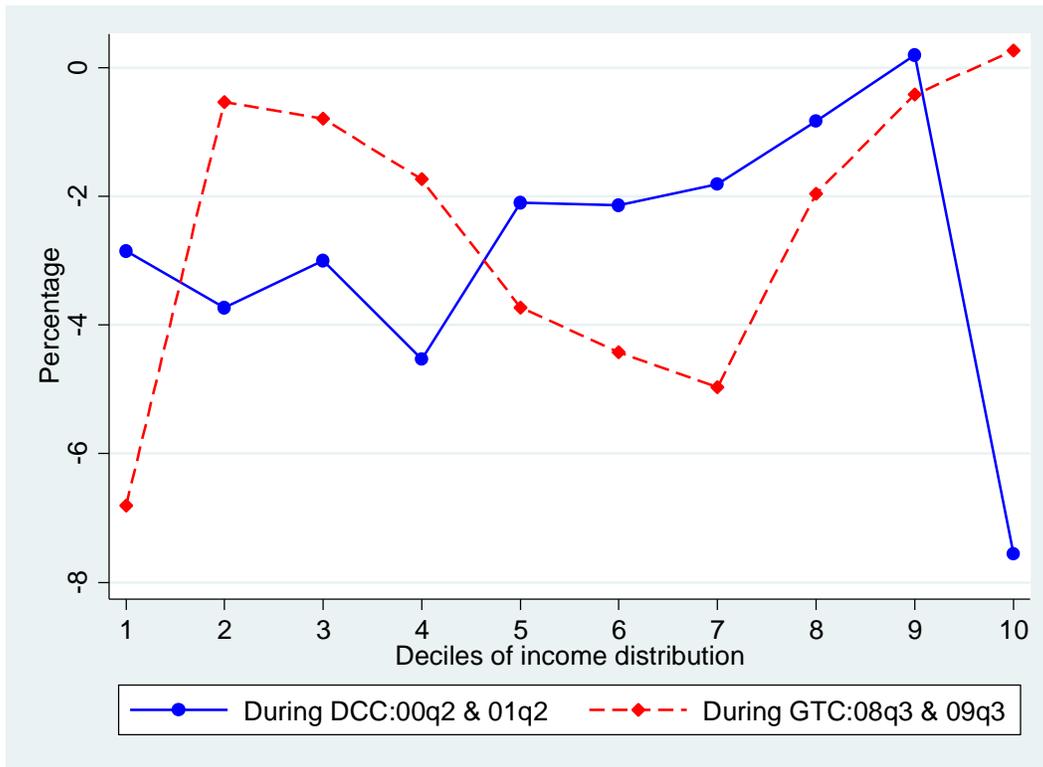
Note: Top panel shows ratios of income elasticity of watches and bedroom linens in 1998q1, 2001q1, and 2008q4 at each of decile points. Bottom panel reports distribution of departures from CRIE baseline.

Figure 5: Mean-Preserving Shocks and Aggregate Expenditure Change for Four Goods



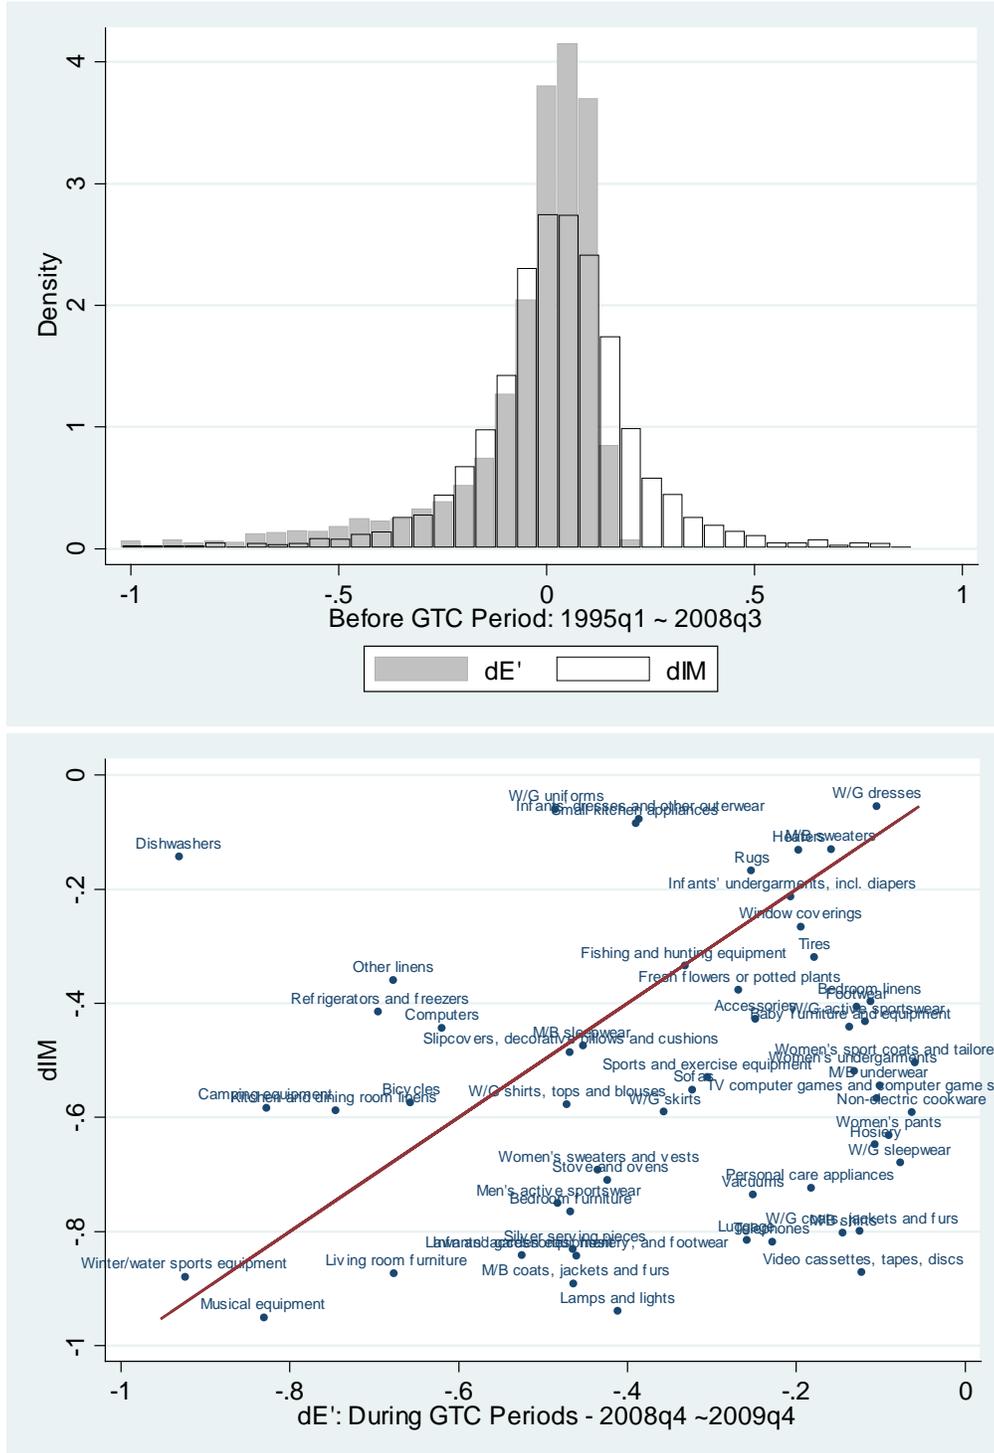
Note: Figure 5 reports aggregate expenditure changes of four expenditure categories under the mean-preserving shocks: the uniform shock, left- and right-skewed shocks. For the uniform shock, the income of all decile points increases by 10%: $y_{bt}^{uniform} = 1.1y_{b,t-1}$ where $t - 1 = 2008q4$ (one of the GTC periods). For the right-skewed shocks, we illustrate changes in skewness by changing α^i , intercept values, from 1 to 5 while for the left-skewed shocks, the intercept changes from -5 to -1: $y_{bt}^i = \alpha^i + \left(1.1 - \frac{10 \cdot \alpha^i}{\sum_b y_{b,t-1}}\right) y_{b,t-1}$ where subscript b and t indicate bins and time, respectively, and superscript i is used to indicate a type of income shock.

Figure 6: Expenditure Change by Deciles during Dot Com Crash and Great Trade Collapse



Note: Figure 6 shows year on year percentage changes in total expenditure during Dot Com Crash (DCC) and Great Trade Collapse (GTC) at each decile point.

Figure 7: The Distribution of Predicted Expenditure Change and Import Change



Note: Figure 7 reports the distribution of year-on-year predicted expenditure change and import change, with the top panel showing variability before the Great Trade Collapse (GTC) period and the bottom panel scattering the good-level changes in imports and predicted expenditures during the period of GTC.