Demographic Transition and Engel's Law across the Development Spectrum

Davide M. Difino*

Goethe University Frankfurt

August 20, 2025

Latest version here

Economic progress brings with it two key patterns. Firstly, we observe the progressive aging of the population. Secondly, as nations' economies grow, the portion of food in total aggregate expenditures tends to decrease. Using country and household-level expenditure data from 20 countries across the entire development spectrum, this work documents that, as the age of household members increases, the proportion of total household expenditures dedicated to food also increases. A large heterogeneity between rich and developing countries emerges when using household-level variables – such as head's and average age – as standard in existing literature. This gap disappears when considering the exact household composition. This finding suggests that – at any development level – the demographic transition leads to a higher overall food share of total expenditures, slowing down structural transformation out of food consumption. I test this hypothesis by constructing a quantitative model that accounts for household demographic composition and documents that the demographic transition is a sizable force that slows down structural transformation. Due to the observed co-movement of demographic transition, structural change, and income growth, not accounting for demography leads to an underestimation of the income effect in almost all countries in the sample.

JEL classification: E17, J11, O57

*Email: difino@wiwi.uni-frankfurt.de

I am grateful to Georg Duernecker, Lidia Cruces, Celina Proffen, and the participants in the RGS Doctoral Conference, EAYE conference, STEG, and FQMG Brown Bags for their helpful comments and suggestions. This document is an output from the research initiative "Structural Transformation and Economic Growth" (STEG), a programme funded by the Foreign, Commonwealth & Development Office (FCDO), contract reference STEG_LOA_512_Difino. The views expressed are not necessarily those of FCDO. All errors are my own.

1 Introduction

Economic progress brings with it two key patterns. Firstly, we observe a movement from a demographic regime marked by high mortality and high fertility rates to one characterized by low mortality and low fertility rates. Referred to as the demographic transition, it portends a slow and gradual aging of society. Secondly, as nations advance economically, the portion of food in total expenditures tends to decrease. This latter trend, known as Engel's Law, is one aspect of the broader structural transformation, wherein expenditures, output, and employment are reallocated across broad economic sectors, from agriculture to manufacturing and services. I shall delve into the interplay between these two phenomena in this work. If demographic factors such as age impact food expenditure decisions – as I will document – then the shifting demographic structure brought on by the demographic transition should shape the sectoral distribution of total expenditures. Is the demographic transition a driving force behind the observed shift away from food expenditures in the structural transformation?

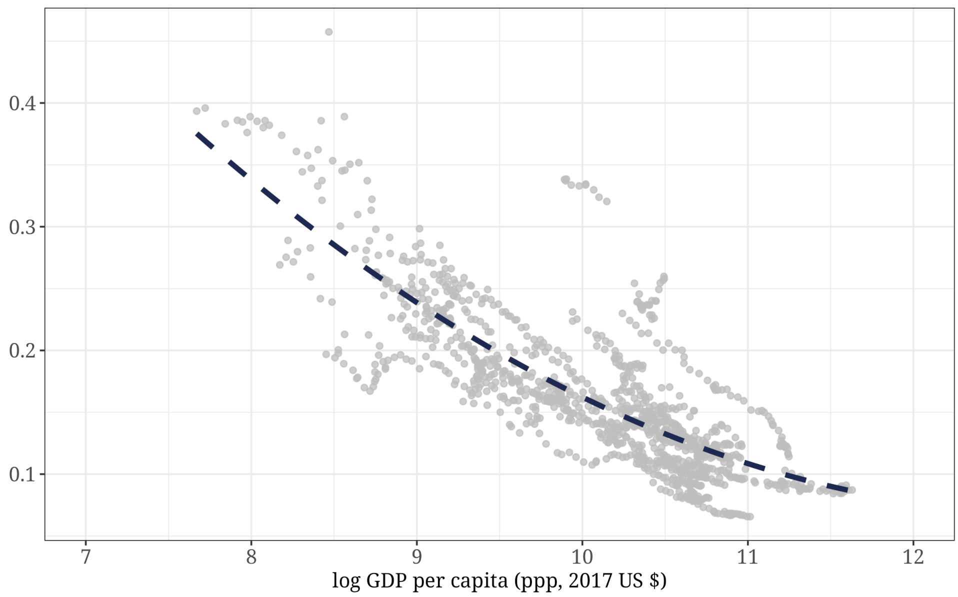

This figure shows the relationship between the food (at home) share of total expenditures and GDP per capita at purchasing power parity for 41 countries in 1959-2018. The blue line represents a quadratic fit. The source of the expenditure data is the Organisation for Economic Co-operation and Development (OECD)'s "Final Consumption Expenditures of Households." Expenditure categories are defined according to use ("Classification of individual consumption by purpose", COICOP), and food (at home) refers to "01 - food and non-alcoholic beverages". Notice that expenditure on food away from home belongs to "11 - Restaurants and hotels". The source of GDP data is the World Bank

Figure 1: food as share of total expenditure and GDP per capita

The reduction in the food share of final consumption and demographic transition are both highly correlated to the rise of real income, the metonym for economic development (see figure 1). The relationship between income and sectoral demand - among the earliest and most persistent contributions in econometrics - goes by the name of Engel's Law: "The poorer is a family, the greater is the proportion of the total [family expenditures] which must be used for food" (Engel, 1895[1]). The idea is simple: once basic needs like food are met, households have more money to spend on non-essential items like manufactured goods and services. More affluent households will spend more on essentials than poorer ones, but the share of total expenditure will decline. This evidence, observed at the household level, has significant implications for the larger macroeconomic structure. Recent research argues that this income effect is an essential driver of long-term structural transformation (Herrendorf et al., 2013, Boppart, 2014, Comin et al., 2021, Alder et al., 2022). Structural transformation itself is a relevant element of modern macroeconomics, affecting growth (Baumol, 2012 and Duernecker et al., 2024) and the transmission of productivity shocks (Baqaee and Farhi, 2019 and Baqaee and Rubbo, 2023).

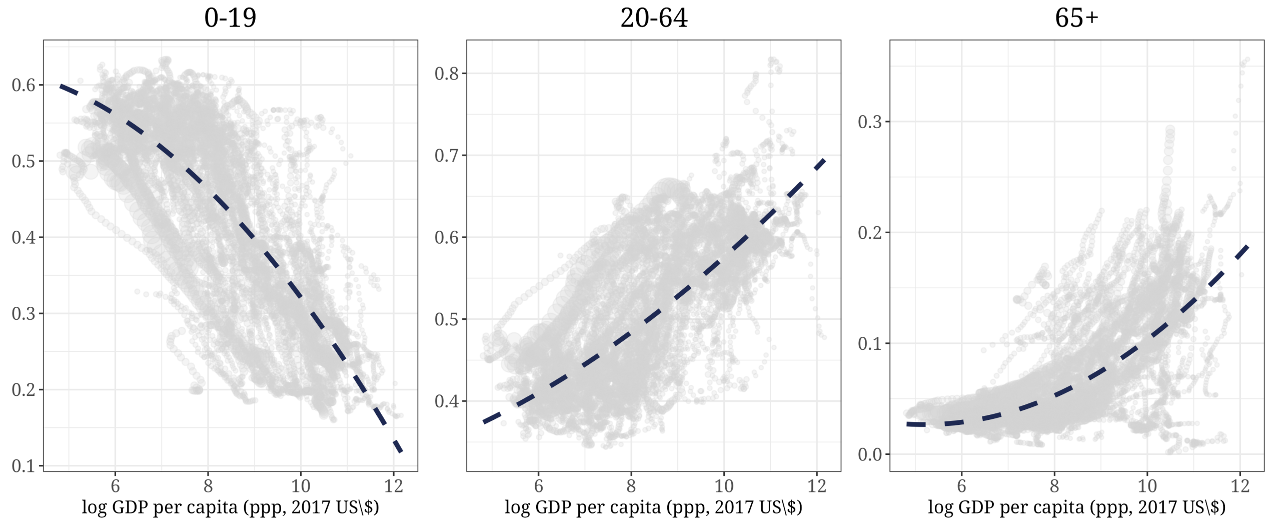

This figure shows the relationship between age structure (across three age groups) and GDP per capita at purchasing power parity. The blue lines represent a quadratic fit. The outliers at the bottom right of panel "60+" are oil-rich Gulf Countries: their GDP per capita is higher than their socio-demographic structure would predict. The source of the demographic data is the World Population Prospects 2024 by the United Nations, while the source of the GDP data is the World Bank.

Figure 2: Age groups shares of total population

The connection between income and demography is a matter that has captured the attention of scholars for centuries, dating back at least to Malthus (1798). Recently, (Doepke et al., 2022) conducted a comprehensive literature review on this topic. In particular, fertility and mortality rates tend to be lower in more prosperous countries than in developing ones (see,

[1]As quoted by Zimmerman (1932), p. 80.

for example Delventhal et al., 2021). As fertility and mortality decline, populations tend to age, resulting in a higher median age in wealthier societies. This relationship is illustrated in Figure 2, where we observe a strong negative correlation between GDP per capita and the share of younger individuals and a positive correlation with older individuals. However, the underlying shift in demographic structure is not monotonic. As shown in Figure 3, the proportion of adults under 70 initially increases before declining as development progresses. The opposite is true for the proportion of working-age individuals, which initially declines before increasing. Finally, the proportion of elderly individuals increases steadily over time and development. The cause of this non-monotonicity lies in the asymmetrical timing of the decline in fertility and mortality: if mortality, and particularly infant mortality, declines before fertility does, then a "baby boom" may occur (i.e., an increase in the proportion of younger individuals in the total population). This phenomenon happened in many developed countries in the 1950s and is currently happening in some African countries (Delventhal et al., 2021).

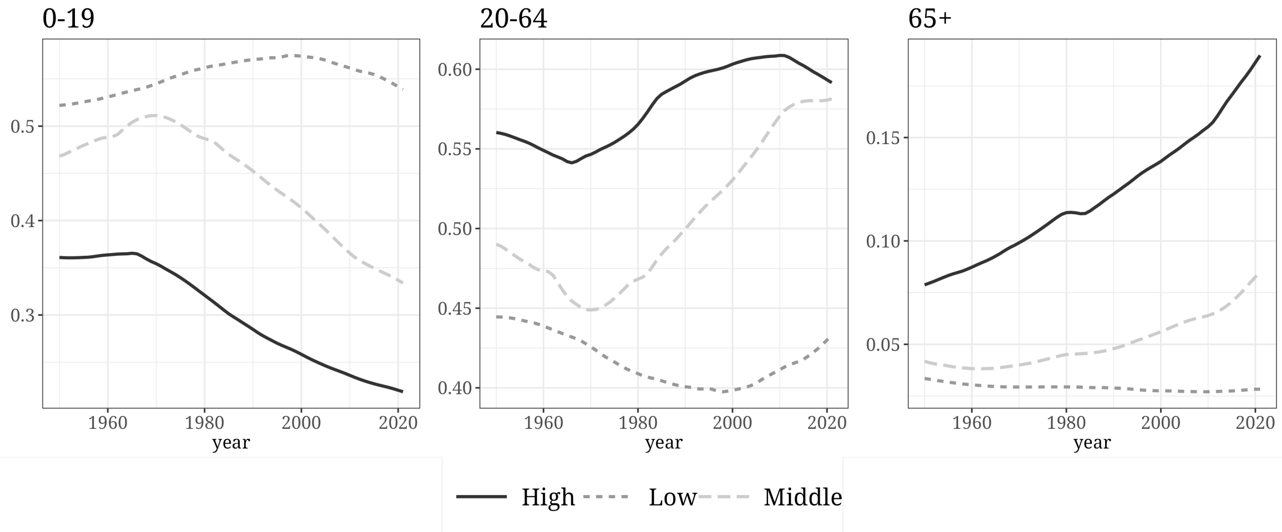

This figure shows the evolution over time of the shares of the total population for three different age groups. Each line represents the population living in countries at different levels of development, the World Bank's country income classification. Notice that "Low" and "Lower-middle income" have been merged into "Low income" to ensure a sufficient number of observations. Therefore, "Middle income" refers to "Upper-middle income" countries according to the World Bank's country income classification. The source of the data is the UN

Figure 3: Age groups shares of total population by income level

The first contribution of this work is to document that age is a significant predictor of higher expenditure shares in food consumption. At the aggregate level, countries with higher median age spend a larger share of their aggregate expenses in food *for home consumption* after controlling for total expenditures. This phenomenon also holds at the household level: using data from over 20 countries across the development spectrum, I document a strong relationship between average age and household food expenditure shares. In particular, the food expenditure share of household total expenditures is increasing in the average age of its members for all countries in the sample. However, there is a significant heterogeneity between developing and developed countries. This heterogeneity disappears when considering the exact composition of the household: the marginal contribution of individual members of different ages does not differ significantly across the development spectrum. The empirical observation that age affects expenditure allocation is consistent with the results of empirical works in the field, such as Aguiar and Hurst (2013), Foster (2015), and Mao and Xu (2014), with the former two using data from the US and the latter from China. The finding that older

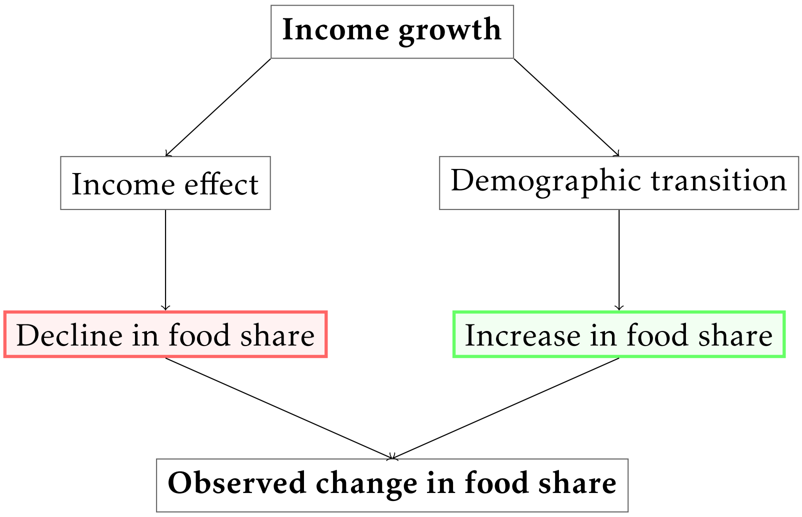

Figure 4

individuals drive higher household expenditure shares in food consumption suggests that the demographic trends might be a slowing force upon structural change out of food consumption. To quantify the size of this mechanism, I build a demand-side, quantitative model centered around the PIGL demand system (Boppart, 2014), a natural choice due to aggregability properties and the ability to capture all critical drivers of structural change. An original contribution is to shape the household taste shifter to map individual member preferences. Calibrated using the same microdata from the empirical exploration, the model confirms that the evolution of age composition in the economy has been a significant slowing down force upon structural change. Furthermore, it confirms the claim that household composition is a driver of the observed heterogeneity between developing levels when using household-level variables. A counterfactual exercise reveals that shutting down the demographic channel would decrease the change rate of food share of aggregated expenditure by between 0.1 to

0.5 percentage yearly for most countries. That is, demographic transition is a sizable force upon structural change.

The observed co-movement of demographic trends, structural change, and income represent a confounding factor to the impact of income on food expenditures at the aggregate level. Income growth drives aging (the demographic transition) and a decline in food expenditures (Engel's Law). However, if aging also affects expenditures, the observed income effect combines two distinct forces - one driven by demography and one driven purely by income (see figure 4). Therefore, a new question arises: to what extent can we attribute the observed (gross) income effect to the demographic transition? To decompose the two forces at play, I run a counterfactual exercise: estimating the income effect without considering the changes in demographic structure leads to an underestimation of the income effect by up to 20%.

This work joins the body of literature that explores the impact of long-run demographic trends on the macro-economy. For example, Aksoy et al. (2019) explores the effects of demographic trends upon macroeconomic indicators such as investment, savings, and hours worked per capita; Jones (2022) inquires on the impact upon the economic growth of taking the demographic transition to its logical extreme. Closely related to this paper, Brembilla (2018) and Cravino et al. (2022) explore the impact of aging upon structural transformation into services in the US. This work expands on this literature in three different directions: first, it focuses on a separate but concomitant facet of structural change - namely, food expenditures. Second, it considers a large set of countries and documents cross-country heterogeneity. Finally, it considers the household's composition as a driver of the observed heterogeneity. However, what sets this work apart from Cravino et al. (2022)'s - a work that shares much of the methodology used here - is the outcome: while they observe that aging drives structural change toward services, I observe that aging hinders structural change out of food expenditures. These results are not conflictual and highlight a complex relationship between the evolution of demographic characteristics and sectorial demand.

2 Cross-country evidence

Let's start by presenting evidence of the relationship between demographic characteristics and food expenditures, defined as food purchased for home consumption. In the next section, I will use country-level data to present some suggestive evidence that a larger share of the elderly population is correlated with a more extensive food share of total expenditures after controlling for the usual drivers, such as income and prices. I will then use microdata from various countries to further explore this phenomenon at the household level and across the development spectrum.

2.1 Country-level evidence

In this section, I will use country-level data to evaluate the correlation between demographic structure and sectorial consumption. For this purpose, I've employed "Final consumption expenditures of households" data from the Organisation for Economic Co-operation and Development (OECD), covering 41 countries from 1959 to 2018. Price data comes from the same source, while demographic variables are obtained from the UN Population Division.

I estimated the following model to explore the relationship between food expenditures and demographic structure.

$$\omega^j_{i,t} = \alpha_i + \beta^j_1 \cdot \text{Median Age}_{i,t} + \beta^j_2 \cdot \log(\text{Exp\_tot}_{i,t}) + \beta^j_3 \cdot \log(\text{Exp\_tot})^2 + \beta^j_4 \cdot \text{Rel\_price}_j + \epsilon^j_{i,t}. \quad (1)$$

In this equation, $\omega^j_{i,t}$ represents the share of total expenditures allocated to sector j in country i at time t. Exp_tot is the private expenditures per capita at purchasing power parity (2017 US dollars). The quadratic term allows for a non-linear relationship between expenditure shares and total expenditures. α_i is a country-level fixed effect[2]. Finally, Rel_price is the ratio between the consumer price index for sector j and the one of the complement sector nj computed using the OECD deflators. Given the price level of sector j and the total Consumer price index, the price index of the complement sector nj is calculated via an

[2] Collinearity with median age does not allow for country-year fixed effects.

(inverse) Tornqvist index[3]:

$$\Delta \log P^{nj}_{i,t} = \frac{\Delta \log P_{i,t} - \left(\frac{\omega^{j}_{i,t}+\omega^{j}_{i,t-1}}{2}\right)\Delta \log P^{j}_{i,t-1}}{\left(\frac{\omega^{nj}_{i,t}+\omega^{nj}_{i,t}}{2}\right)}. \tag{2}$$

Table 1 presents the results of estimating equation 1 for food consumption. We can observe a negative and statistically significant relationship between food consumption shares and total expenditures consistent with Engel's Law. Notably, this relationship appears to be linear. We can also examine the impact of relative prices on food consumption and find that higher relative prices are associated with more substantial expenditures on food. This outcome aligns with expectations based on the gross complementarity of these broad expenditure categories.

Of particular interest, this analysis reveals that countries with higher median age tend to allocate a greater share of their income to food consumption. This result remains robust even after controlling for expenditure levels and relative prices. However, It is worth noting that before controlling for expenditure levels, the unconditional relationship between food shares and age is negative, likely reflecting that richer countries tend to have older populations.

Table 2 shows the result of estimating equation 1 for services. Again, we observe that countries with a larger share of older individuals also spend a larger share of their income on services. The impact of total expenditures and relative prices are aligned with our priors. This result is consistent with the main finding of Cravino et al. (2022).

The cross-country analysis in Cravino et al. (2022) reveals a negative relationship between the share of the population aged 65 or older and food consumption, controlling for GDP per capita. I have replicated their analysis using GDP data from Maddison in Table 3. After reducing the sample to the one they use – focusing on richer countries over a shorter period – the results align with those of Cravino et al. (2022). This finding suggests that the relationship

[3]Equation 2 follows naturally from the definition of the CPI index when using a Tornqvist index:

$$\Delta \log P_{i,t} \equiv \left(\frac{\omega^{j}_{i,t}+\omega^{j}_{i,t-1}}{2}\right)\Delta \log \log P^{j}_{i,t} + \left(\frac{\omega^{nj}_{i,t}+\omega^{nj}_{i,t-1}}{2}\right)\Delta \log \log P^{nj}_{i,t}.$$

Dependent Variable:	Food Share of Household Expenditures						
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Variables							
Median age	-0.0082***	0.0023**	0.0021*	0.0021**	0.0018*	0.0021**	0.0016**
	(0.0014)	(0.0010)	(0.0012)	(0.0008)	(0.0010)	(0.0009)	(0.0006)
log(GDP per capita)		-0.1368***	-0.1720			-0.1108***	
		(0.0153)	(0.1227)			(0.0177)	
log(GDP per capita)^2			0.0019				
			(0.0071)				
log(Total Exp. per capita)				-0.1392***	-0.2176		-0.1101***
				(0.0191)	(0.1737)		(0.0199)
log(Total Exp. per capita)^2					0.0045		
					(0.0099)		
Relative prices						0.0448***	0.0435***
						(0.0093)	(0.0121)
Fixed-effects							
Country	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics							
Observations	1,244	1,166	1,166	1,241	1,241	1,098	1,174
R^2	0.83972	0.93026	0.93039	0.91975	0.92027	0.96884	0.95522
Within R^2	0.49376	0.78131	0.78171	0.74794	0.74957	0.87960	0.82633

Clustered (Country) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

The table presents OLS estimates of the coefficients in equation 1. The dependent variable is the food share of total household expenditure. The source of the expenditure data is the Organisation for Economic Co-operation and Development (OECD)'s "Final Consumption Expenditures of Households." Expenditure categories are defined according to use ("Classification of individual consumption by purpose," COICOP), and food refers to "01 - food and non-alcoholic beverages". Median age data are from the UN, GDP per capita from the World Bank's WDI, and total expenditures per capita are calculated using OECD expenditure and UN population data. GDP per capita and total expenditures per capita are expressed in 2015 US dollars. Relative prices are derived from OECD deflators using equation 2, with 2015 = 1. Standard errors clustered at the country level are in parentheses, with p-values in brackets (*: $p < 0.1$; **: $p < 0.05$; ***: $p < 0.01$).

Table 1

Dependent Variable:	Services Share						
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Variables							
Median age	0.0117***	0.0071***	0.0054***	0.0069***	0.0054***	0.0078***	0.0075***
	(0.0011)	(0.0018)	(0.0019)	(0.0019)	(0.0016)	(0.0012)	(0.0013)
log(GDP per capita)		0.0700***	-0.3019*			0.0333**	
		(0.0197)	(0.1578)			(0.0145)	
log(GDP per capita)^2			0.0203**				
			(0.0080)				
log(Total Exp. per capita)				0.0687**	-0.3886*		0.0254
				(0.0281)	(0.2071)		(0.0217)
log(Total Exp. per capita)^2					0.0262**		
					(0.0109)		
Relative prices						0.1307***	0.1422***
						(0.0296)	(0.0276)
Fixed-effects							
Country	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics							
Observations	1,231	1,151	1,151	1,187	1,187	1,098	1,174
R^2	0.90090	0.92813	0.93633	0.91420	0.92362	0.94768	0.93673
Within R^2	0.65041	0.75039	0.77887	0.69655	0.72985	0.81722	0.77558

Clustered (Country) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

The table presents OLS estimates of the coefficients in equation 1. The dependent variable is the services share of total household expenditure. The source of the expenditure data is the OECD's "Final Consumption Expenditures of Households." Expenditure categories are defined according to use ("Classification of individual consumption by purpose," COICOP), and services refer to the COICOP variable "P314B". Median age data are from the UN, GDP per capita from the World Bank's WDI, and total expenditures per capita are calculated using OECD expenditure and UN population data. GDP per capita and total expenditures per capita are expressed in 2015 US dollars. Relative prices are derived from OECD deflators using equation 2, with 2015 = 1. Standard errors clustered at the country level are in parentheses, with p-values in brackets (*: $p < 0.1$; **: $p < 0.05$; ***: $p < 0.01$).

Table 2

Dependent Variables:	Food Share			Services Share		
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
Share aged 65 or more	-2.519***	-0.1355	-0.3828	3.592***	1.273***	0.8670**
	(0.7512)	(0.2289)	(0.2413)	(0.6833)	(0.3321)	(0.2995)
log(GDP per capita)		-0.1314***	-0.2769***		0.1208***	-0.1534**
		(0.0088)	(0.0399)		(0.0118)	(0.0627)
log(GDP per capita)^2			0.0082***			0.0154***
			(0.0023)			(0.0037)
Fixed-effects						
Country	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	285	285	285	247	247	247
R^2	0.79763	0.97184	0.97424	0.83869	0.95242	0.95865
Within R^2	0.56747	0.93980	0.94495	0.68536	0.90720	0.91935

Clustered (Country) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

The table presents OLS estimates of the coefficients in equation 1. The dependent variable in the first three columns is the food share of total household expenditures, and the services share in the last three. The source of the expenditure data is the OECD's "Final Consumption Expenditures of Households." Expenditure categories are defined according to use ("Classification of individual consumption by purpose," COICOP). The share of individuals aged 65 or more of the total population comes from the UN, while GDP per capita comes from Maddison. GDP per capita is expressed in 2015 US dollars. Relative prices are derived from OECD deflators using equation 2, with 2015 = 1. Standard errors clustered at the country level are in parentheses, with p-values in brackets (*: $p < 0.1$; **: $p < 0.05$; ***: $p < 0.01$).

Table 3

between population age and food expenditures may be highly susceptible to changes in the sample due to country-level heterogeneity. Nonetheless, I interpret the results of this section as providing tentative support for one of the central empirical claims of this study, namely that aging increases food consumption. In the next section, I will use cross-country micro data at the household level to confirm this claim.

3 Household evidence

This section employs cross-sectional household-level data from the Luxembourg Income Study (LIS) Database, an extensive database that collects data from diverse national surveys, such as the CPS (United States) and the German Socio-Economic Panel. It spans five decades and covers about 50 countries. While it focuses on income data, it also provides a range of expenditure variables, including total and food expenditures according to the COICOP classification, the same as those used in OECD expenditure data. Furthermore, this work uses the ERF-LIS database, which combines the Harmonised Household Income and

Expenditure Surveys (HHIES) by the Economic Research Forum (ERF) with the LIS[4]. The two datasets provide a well-assorted subset of 20 countries that report these expenditure variables, covering the entire development spectrum from Mali to Switzerland. The list of the surveyed countries and waves is depicted in table 4.

cname	years
Australia	2004, 2010, 2016
China	2002, 2013, 2018
Egypt	1999, 2004, 2008, 2010, 2012, 2015, 2017
Georgia	2009, 2010, 2011, 2012, 2013, 2014, 2015, 2016, 2017, 2018, 2019
Guatemala	2006, 2014
Hungary	1991, 1994, 1999, 2005, 2007, 2009, 2012, 2015
India	2004, 2011
Iraq	2007, 2012
Israel	1992, 1997, 2001, 2002, 2003, 2004, 2005, 2006, 2007, 2008, 2009, 2010, 2011, 2012, 2013, 2014, 2015, 2016, 2017, 2018, 2019
Italy	1991, 1993, 1995, 1998, 2000, 2002, 2004, 2006, 2008, 2010, 2012, 2014, 2016
Ivory Coast	2002, 2008, 2015
Jordan	2002, 2006, 2008, 2010, 2013
Mali	2011, 2013, 2014, 2015, 2016, 2017, 2018, 2019
Mexico	1992, 1994, 1996, 1998, 2000, 2002, 2004, 2005, 2006, 2008, 2010, 2012, 2014, 2016, 2018
Palestine	1996, 1997, 1998, 2004, 2005, 2006, 2007, 2009, 2010, 2011, 2017
Peru	2004, 2005, 2006, 2007, 2008, 2009, 2010, 2011, 2012, 2013, 2014, 2015, 2016, 2017, 2018, 2019
Poland	1999, 2004, 2005, 2006, 2007, 2008, 2009, 2010, 2011, 2012, 2013, 2014, 2015, 2016, 2017, 2018, 2019
Russia	2000, 2004, 2007, 2010
Serbia	2006, 2007, 2008, 2009, 2010, 2011, 2012, 2013, 2014, 2015, 2016, 2017, 2018, 2019
Slovenia	1997, 1999, 2004, 2007, 2010, 2012, 2015
South Africa	2008, 2010, 2012, 2015, 2017
South Korea	2006, 2008, 2010, 2012, 2014, 2016, 2017, 2018, 2019
Switzerland	2000, 2002, 2004
Tunisia	2005, 2010
United Kingdom	1990, 1991, 1992, 1993
Vietnam	2005, 2007, 2009, 2011, 2013

Table 4: County-Wave coverage of the combined LIS-ERFLIS database

This empirical analysis includes households with strictly positive food consumption (and therefore positive total consumption) and non-negative income, and sample weights are applied. These weights are based on the country's population and enable cross-country comparisons.

The baseline model, inspired by Aguiar and Bils (2015) and Cravino et al. (2022), is estimated as follows:

$$\log\left(\omega_h^f\right) = \alpha_{c,t} + \beta \cdot \log\left(\text{Income}_h\right) + D^{\text{Head Age}} + \gamma \mathbf{X}_h + \xi_h \tag{3}$$

[4]In appendix 5 I replicated this section's analysis using two different expenditure surveys, U.S. Bureau of Labor Statistics (2025)'s for the US and Instituto Nacional de Estadística y Geografá (INEGI)'s ENIGH for Mexico.

Here, ω_h^f represents the food share of total expenditures for household h, and $\alpha_{c,t}$ denotes the country-year fixed effects. The key element of equation 3 is $D^{\text{Head age}}$, which is a boolean variable representing the average group the household belongs to. The socio-demographic controls vector $\mathbf{X}$ includes the number of household members, household type (e.g., single individual, couple, couple with children, couple living with parents), and number of earners. The latter controls for retirement, which literature suggests might affect consumption (Hurst, 2008 and Aguila et al., 2011). Following Aguiar and Hurst (2013), income is instrumented with total expenditures to account for measurement errors. Time indexes are absent due to the cross-sectional nature of the LIS dataset, where each household is observed only once. The logarithmic shape of Equation 3 follows from the typical log-linear shape of the Marshallian demand, and thus the sectoral shares implied by commonly used utility functions such as the PIGL preferences that will be used in section 4. The log-linear formulation allows for easy interpretation of the D_h^{m} coefficients: $\exp\left(D_h^{\mathrm{m}}\right)$ is the ratio between the consumption shares of a household with an average age $= m$ and the consumption shares of an omitted reference age group, in this case, the age group $25-29$.

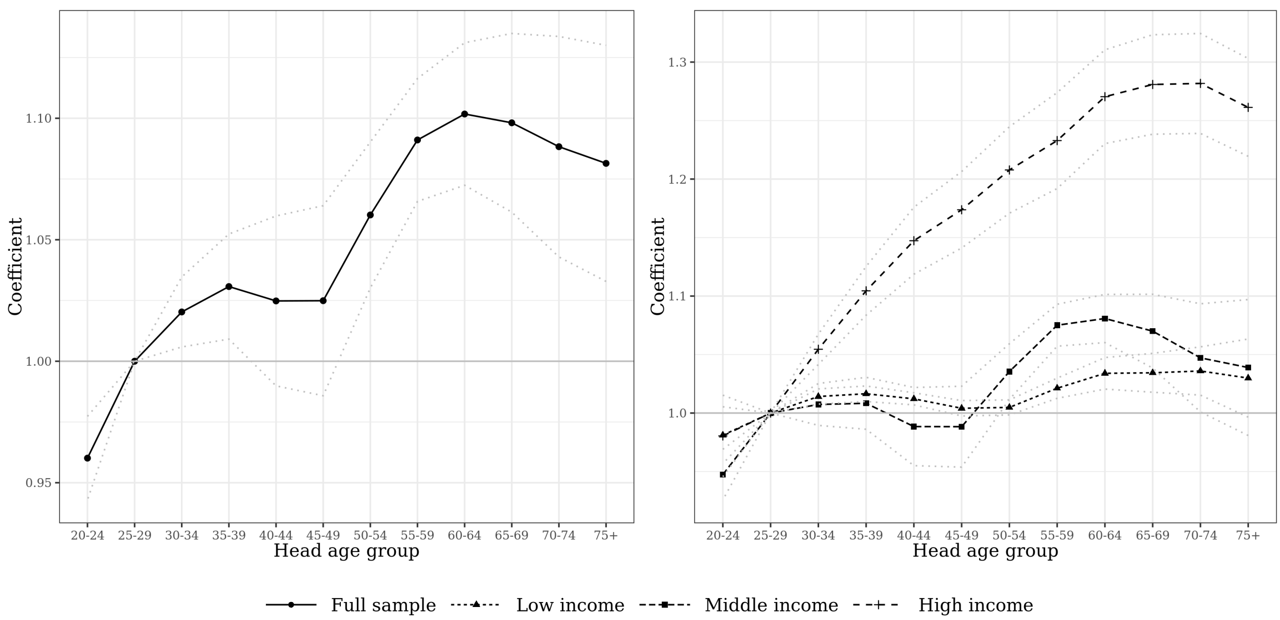

This figure plots the estimated values for $\exp(D_h^{\text{Head Age}})$ from equation 3, but using the head's age instead of the average age. The right-hand panel shows the value estimated across all countries, while the left-hand represents the value with the sample split across the World Bank's income classification. Notice that "Low" and "Lower-middle income" have been merged into" Low income" to ensure a sufficient number of observations. Therefore, "Middle income" refers to "Upper-middle income" countries according to the World Bank's income classification.

Figure 5: Estimated value of $\exp(D_h^{\text{Head Age}})$ by income group

Figure 5 shows the OLS estimates for the age coefficients from equation 3. The left-hand panel presents the estimates using the entire sample: it shows that older households tend to consume a higher share of their total expenditures on food. For example, households aged between 60 and 64 consume, on average, 10% more than households aged between 25 and 29, even after controlling for income and other demographic factors. The right-hand panel of figure 5 shows the estimates by income group. The difference in food expenditures between young and older households appears significantly larger in richer countries. While still statistically significant, the difference between high and low-income countries is almost an order of magnitude: for example, households whose head's age is between 65 and 69 in "high-income" countries consume, on average, 30% more than households whose head's age is between 25 and 29, while in "low income" countries, the difference is approximately 3%.

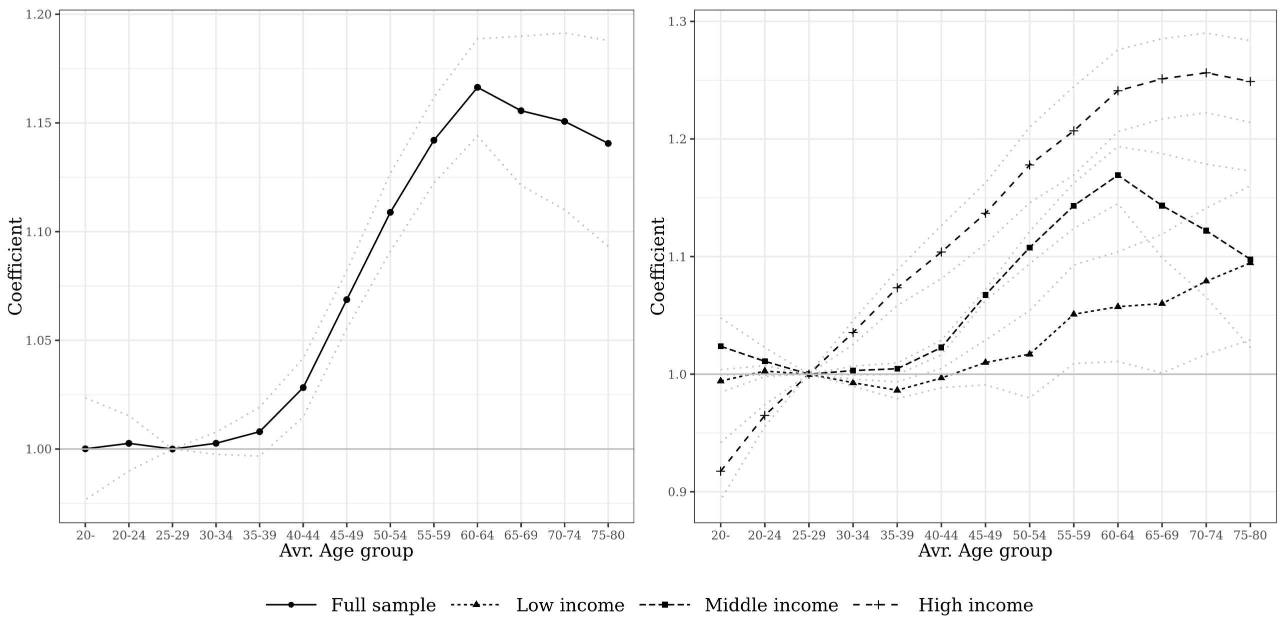

This figure plots the estimated values for $\exp(D_h^{\text{Avr. Age}})$ from equation 3. The right-hand panel shows the value estimated across all countries, while the left-hand represents the value with the sample split across the World Bank's country income classification. Notice that "Low" and "Lower-middle income" have been merged into "Low income" to ensure a sufficient number of observations. Therefore, "Middle income" refers to "Upper-middle income" countries according to the World Bank's country income classification.

Figure 6: Estimated value of $\exp(D_h^{\text{Avr. Age}})$ by income group

Figure 6 shows the coefficient of the same model as before but with average age as the relevant demographic variable instead of the head's age. The results show that the average age of the household is correlated with a higher level of expenditure: households whose head's age is between 60 and 64 consume, on average, 15% more than households whose head's age is between 25 and 29. However, the differences between income groups are not as stark compared to the previous model.

3.1 Accounting for household demographic structure

The models estimated in the previous section suggest that older households with a higher head's age spend more on food than younger households. Significantly, this phenomenon is stronger in richer countries. However, using the household's average age instead of the head's reduces the gap significantly. What can explain this heterogeneity? One possible explanation hinges on the differences in household structure between developed and developing countries, as seen in figure 7. While the regression in the previous section does account for household type, it fails to account that differences in household structure are (also) differences in demographic structure. For example, consider two households with two members and an average age of 25: the first is composed of a couple of individuals aged 25, while the second is composed of a single parent aged 40 and 10-year-old children. In the previous section, the differences in observed food expenditures between these two households would be accounted for by the different "Household types" they belong to ("Couple without children" and "Single parent with children" respectively). However, the difference between these two households is also demographic: arguably, what drives the differences in food expenditure shares is that a 10-year-old individual has different needs and preferences compared to a 25-year-old one.

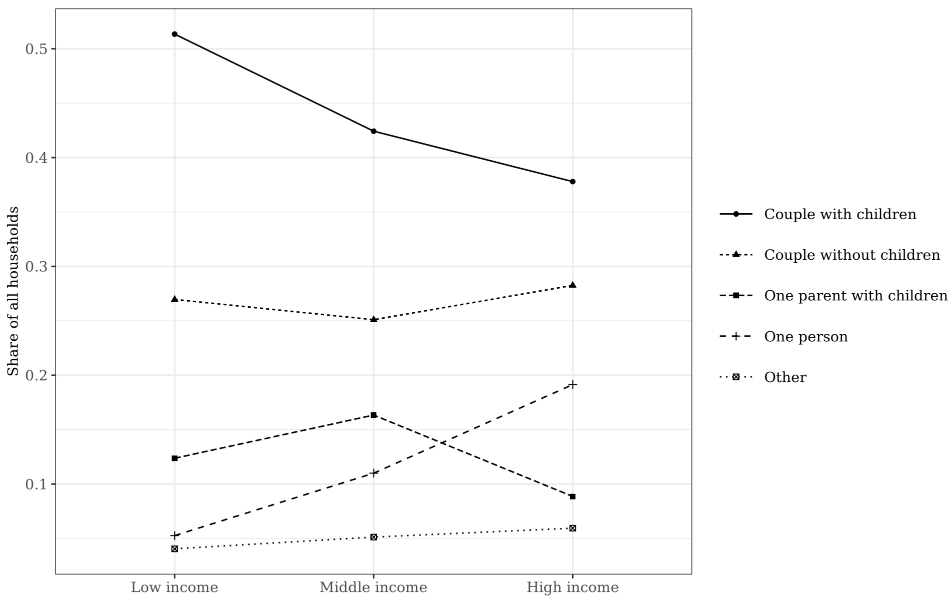

This figure plots the distribution of household type by development level in the LIS database. The hosehold type provided by LIS are: "[100]one person household", "[210]couple without children" ,"[310]couple without children and relatives", "[220]couple with children", "[320]couple with children and relatives", "[230]one parent with children", and"[330]one parent with children and relatives". Notice that "Low" and "Lower-middle income" have been merged into "Low income" to ensure a sufficient number of observations. Therefore, "Middle income" refers to "Upper-middle income" countries according to the World Bank's country income classification.

Figure 7: Distribution of Household types by income group

The mapping between individual and household behavior introduces several complications, such as intra-household bargaining and the presence of public goods, that require more detailed consumption data than what is available in the LIS dataset[5]. One way to circumvent the intra-household distributional issues is to replicate the work in the previous section but only for households composed of a single individual. The outcomes are shown in figure 8: the differences in food shares between single-member households with different ages are significant (up to 30%), and – relevantly – the heterogeneity between low-and high-income countries disappears.

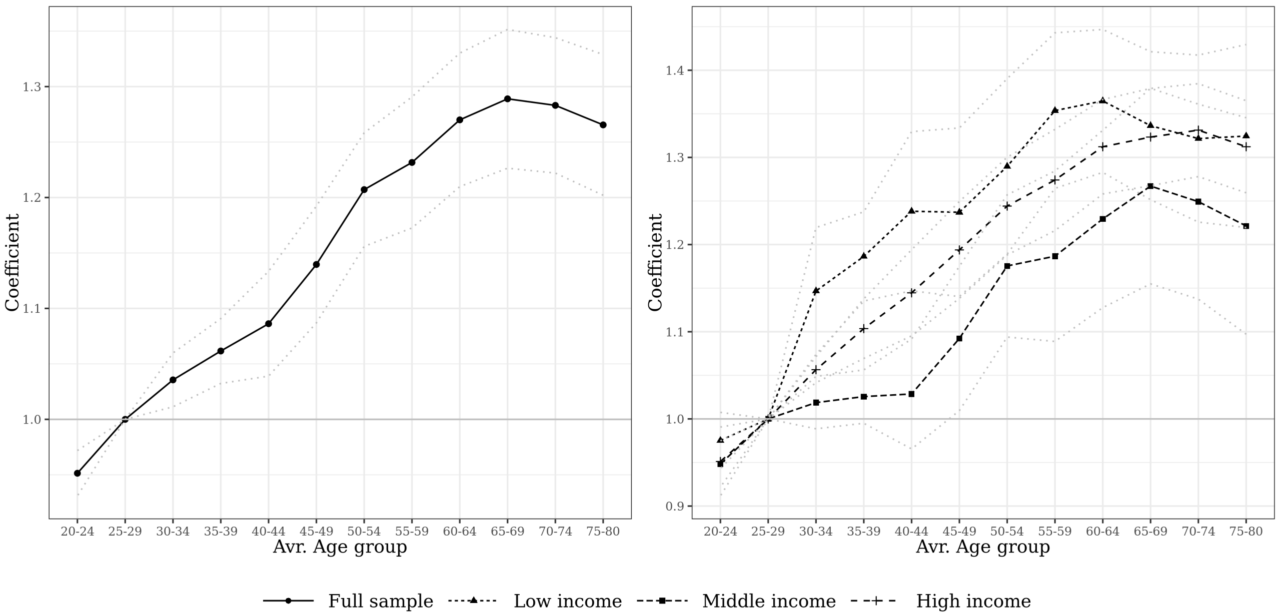

This figure plots the estimated values for $\exp(D_h^{\text{Avr. Age}})$ from equation 3. The sample is reduced to households composed of only a single member. The right-hand panel shows the value estimated across all countries, while the left-hand represents the value with the sample split across the World Bank's country income classification. Notice that "Low" and "Lower-middle income" have been merged into "Low income" to ensure a sufficient number of observations. Therefore, "Middle income" refers to "Upper-middle income" countries according to the World Bank's country income classification.

Figure 8: Estimated value of $\exp(D_h^{\text{Avr. Age}})$ by income group, only single-member households

A different approach is estimating an additional member's marginal impact on food expenditure shares. That is, estimating

$$\omega_h^f = \sum_a \beta_a \cdot N_{a,h}. \tag{4}$$

where $N_{a,h}$ is the number of household members aged a in household h. To gain some intuition of β_a, notice that household h's (total) food expenditures are defined as the sum of the food

[5]For example, Browning et al. (2013) and Lechene et al. (2019) have proposed approaches for empirically decomposing intra-household expenditure allocation. However, these require more detailed data such as imputable consumption categories.

expenditures for each member i:

$$E_h^f \equiv \sum_i E_i^f.$$

Assume now that individuals of the same age are symmetric. We can thus aggregate by age group:

$$E_h^f \equiv \sum_a N_{a,h} \cdot E_{a,h}^f.$$

where $N_{a,h}$ is the number of household members aged a. Thus, household h's food share of total expenditures is

$$\omega_h^f \equiv \frac{E_h^f}{E_h} = \sum_a \phi_{h,a}\omega_{h,a}^f \cdot N_{a,h}.$$

where $\phi_{h,a}$ is the share of total household expenditures that are imputed to each member aged a:

$$\phi_{h,a} \equiv \frac{E_{h,a}}{E_h}, \qquad 0 \leq \phi_{h,a} \leq 1$$

It can be interpreted as a measure of the intra-household bargaining power of individuals aged a. We have also defined

$$\omega_{h,a} \equiv \frac{E_{h,a}^f}{E_{h,a}}, \qquad 0 \leq \omega_{h,a} \leq 1,$$

the food share of individual expenditures for individuals aged a. While the data available do not allow to decompose $\phi_{h,a}$ and ω^f, we can estimate the average, bargain-power adjusted food expenditure share of individual consumption ($\phi_{h,a} \cdot \omega_{h,a}^f$) by estimating equation 4. Figure 9 shows the results, where I used non-linear least squares instead of the more traditional OLS to ensure that the coefficients are positive. As shown, the bargaining-power-adjusted food expenditure share does not differ between income groups, suggesting that accounting for exact household composition closes the gap observed in the previous section.

4 Structural model

The previous section suggested that age affects food expenditure shares. The next objective is assessing how the demographic transition, i.e., the country's age composition shift, affects aggregate food expenditure share. To accomplish this, this section presents a demand-side,

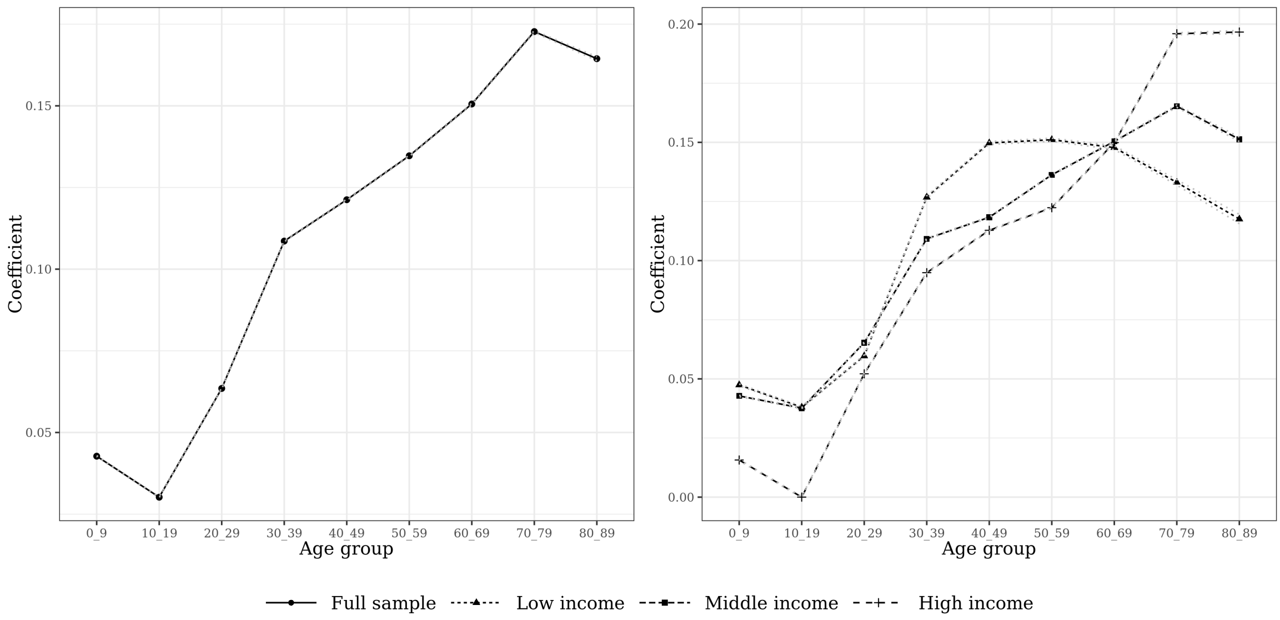

This figure plots the estimated values for β from equation 4. The right-hand panel shows the value estimated across all countries, while the left-hand represents the value with the sample split across the World Bank's country income classification. Notice that "Low" and "Lower-middle income" have been merged into "Low income" to ensure a sufficient number of observations. Therefore, "Middle income" refers to "Upper-middle income" countries according to the World Bank's country income classification.

Figure 9: Estimated value of β by income group

partial equilibrium model using the PIGL class of preferences (Muellbauer, 1975), which have a long history as the basis for empirical work[6]. This class of preferences has been applied to the context of structural change by Boppart (2014) and possesses appealing aggregability properties that will be useful when exploring the impact of demographic trends on the aggregate allocation of food expenditures. Additionally, PIGL preferences are non-homothetic, which makes them capable of capturing the dynamics of the Engel Law.

4.1 Model

In the economy, there are two goods: food ("f") and non-food ("n"). Household h's demographic structure is exogenously given, specifically the distribution of members across M demographic groups. After observing the exogenous price vector $\mathbf{P}$, the household allocates its consumption across the two sectors. Preferences follow a PIGL (Boppart, 2014) functional form:

$$\mathcal{V}^h(\mathbf{P}, E_{h,t}) = \frac{1}{\epsilon}\left[\frac{E_t^h}{P_t^n}\right]^{\epsilon} - \frac{\nu_t^h}{\gamma}\left[\frac{P_t^f}{P_t^n}\right]^{\gamma} - \frac{1}{\epsilon} + \frac{\nu_t^h}{\gamma},$$

[6]The Almost Ideal Demand System is derived from the logarithmic form of PIGL, known as PIGLOG.

where $E_{h,t}$ is the *per capita* expenditure level and P_t^f, P_t^n are the food and non-food prices respectively. Parameters can be interpreted as follows: ϵ is the real expenditures elasticity, while γ represents the marginal impact of a change in relative prices. Finally, ν_t^h is a household-specific taste-shifter. To see its role in the allocation of expenditures across sectors, notice that from Roy's identity, we can derive the household's h food share of total spending:

$$\omega_h^f = \left(\frac{E_{h,t}}{P_t^n}\right)^{-\epsilon} \cdot \left(\frac{P_t^f}{P_t^n}\right)^{\gamma} \cdot \nu_t^h \tag{5}$$

A higher taste-shifter ν_t^h implies a higher food share of total expenditures, given prices and expenditure levels.

In contrast to Cravino et al. (2022), I have defined the complement of the sector that I am interested in (i.e., the non-food) as the "reference sector ." This reference sector serves as the denominator in the indirect utility function. My argument is simple: as shown in equation 5, setting non-food as the reference sector implies log-linear shares of total expenditures for the sector we are interested in. On the other hand, the reference sector's shares are the complement, which is not log-linear. This approach allows for a clearer interpretation of the parameters directly referring to the relevant industry. Additionally, it simplifies the math. For example, we can derive equation 9 directly, while in Boppart (2014) and Cravino et al. (2022) it can be only approximated using two Taylor expansions. As shown in the appendix 5, which sector is used as a reference does qualitatively affect the outcome[7].

Each household is represented by a single agent. However, we know that households comprise diverse individuals with different preferences and abilities that impact the household budget allocation. To account for the household's demographic structure, I propose using a geometric average of individual members' preferences parameters as the household taste-shifter. That is:

$$\nu_t^h \equiv \left(\prod_i^{N_h} \nu_t^i(m)\right)^{\frac{1}{N_h}}$$

In addition to its useful mathematical properties, using a geometric mean aggregator for household taste shifters has strong economic intuition. Let's start by noticing that a geometric

[7]Since the approach by Cravino et al. (2022) requires the use of approximations, the numerical outcomes are slightly different but qualitatively equivalent.

average is mathematically and conceptually equivalent to a Cobb-Douglas aggregator of individual taste-shifters. Therefore, by using a geometric mean aggregator of personal preferences, I assume imperfect substitution between individual consumption within the household[8]. Furthermore, it is assumed implicitly that the elasticity of substitution between individual members is unity. That is, the household cares about welfare distribution across its members and considers all members equally important. I maintain that these assumptions are reasonable in the context of household decision-making.

Following Cravino et al. (2022), the individual[9] taste-shifter takes the form:

$$\nu_t^i(m) = \nu_t \cdot \delta^m \cdot \mu_t^h,$$

That is, the taste-shifter of individual i of household h, belonging to the demographic group m, can be decomposed into an aggregate component (ν_t), a demographic component (δ^m) and, finally, a household-level idiosyncratic one (μ_t^h). Therefore, the household-level taste-shifter can be written as

$$\nu_t^h = \nu_t \cdot \mu_t^h \cdot \underbrace{\exp\left[\sum_m^M s_m^h \cdot \log(\delta^m)\right]}_{\equiv \delta_t^h} \equiv \nu_t \cdot \mu_t^h \cdot \delta_t^h \tag{6}$$

Where $s_m^h \equiv N_h^m / N_h$ is the share of household h members belonging to demographic group m. The demographic component of household preferences, δ_t^h, is the log-linear aggregation of individual member preferences. This structure allows us to incorporate the detailed structure of the household, such as the number and demographic characteristics of individual members.

4.2 Calibration

Combining equations 5 and 6 yields the food share of total expenditures for a generic household h:

$$\omega_h^f = \left(\frac{E_{h,t}}{P_t^n}\right)^{-\epsilon} \cdot \left(\frac{P_t^f}{P_t^n}\right)^{\gamma} \cdot \exp\left[\sum_m^M s_m^h \cdot \log(\delta^m)\right] \cdot \nu_t \cdot \mu_t^h \tag{7}$$

[8]If the household were to maximize any linear combination of member's utility, then we would have perfect substitution across individuals

[9]In Cravino et al. (2022) this is the shape of the *household-level* taste-shifter.

After taking logs, this equation can be estimated empirically by OLS using LIS data. The model that will be fitted is

$$\log(\omega_h^f) = \epsilon \cdot \log\left(\frac{P_t^n}{E_{h,t}}\right) + \sum_m^M s_m^h \cdot \mathrm{DUMMY}_m + \alpha_{c,t} + \xi_h \tag{8}$$

Where α_c is a region-time fixed effect that captures potential local differences in relative prices. Furthermore, it also captures the relative price component in equation 5. The model above allows for a precise estimation of the parameter ϵ while controlling for region-time fixed effect. On the other hand, the estimation of γ using the 2-step approach from Boppart (2014), requires a sufficiently large number of country-year waves to be observed. As is not the case with LIS data, this approach is unfeasible with the data used in this work. However, knowing the value of the γ parameter is not vital for the research question of this paper, as knowing ϵ and the demographic parameters is all is needed to decompose income and demographic effects in equation 9. As standard in the literature, expenditures have been instrumented by total income.

Dependent Variable:	log(food share)			
Income group	Full sample	High income	Low income	Middle income
Model:	(1)	(2)	(3)	(4)
Variables				
ϵ	0.3489***	0.4433***	0.3316***	0.3354***
	(0.0182)	(0.0179)	(0.0141)	(0.0231)
Fixed-effects				
year-Region	Yes	Yes	Yes	Yes
Fit statistics				
Observations	2,413,287	1,037,169	341,903	1,034,215
R^2	0.47805	0.53667	0.47180	0.34816
Within R^2	0.29888	0.33353	0.40070	0.27561

Clustered (Region) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

This table shows the IV estimates for ϵ from equation 8. The estimates for the demographic coefficients are shown in figure 10. Observations with negative income and expenditure levels have been removed. The first column shows the estimates from the entire LIS sample, while the other three columns show the country's income group estimates according to the World Bank's country income classification. Notice that "Low" and "Lower-middle income" have been merged into "Low income" to ensure a sufficient number of observations. Therefore, "Middle income" refers to "Upper-middle income" countries according to the World Bank's country income classification. Standard errors clustered at the region level are in parentheses, with p-values in brackets (*: $p < 0.1$; **: $p < 0.05$; ***: $p < 0.01$).

Table 5

I define the M demographic groups according to age. Specifically, I classify all individual household members according to 9 10-year groups. Table 5 shows the output of the regression. As expected, the coefficient ϵ, ruling the income effect, is positive - i.e., food expenditures share decline with total expenditures. This result is consistent with Engel's Law. Figure 10

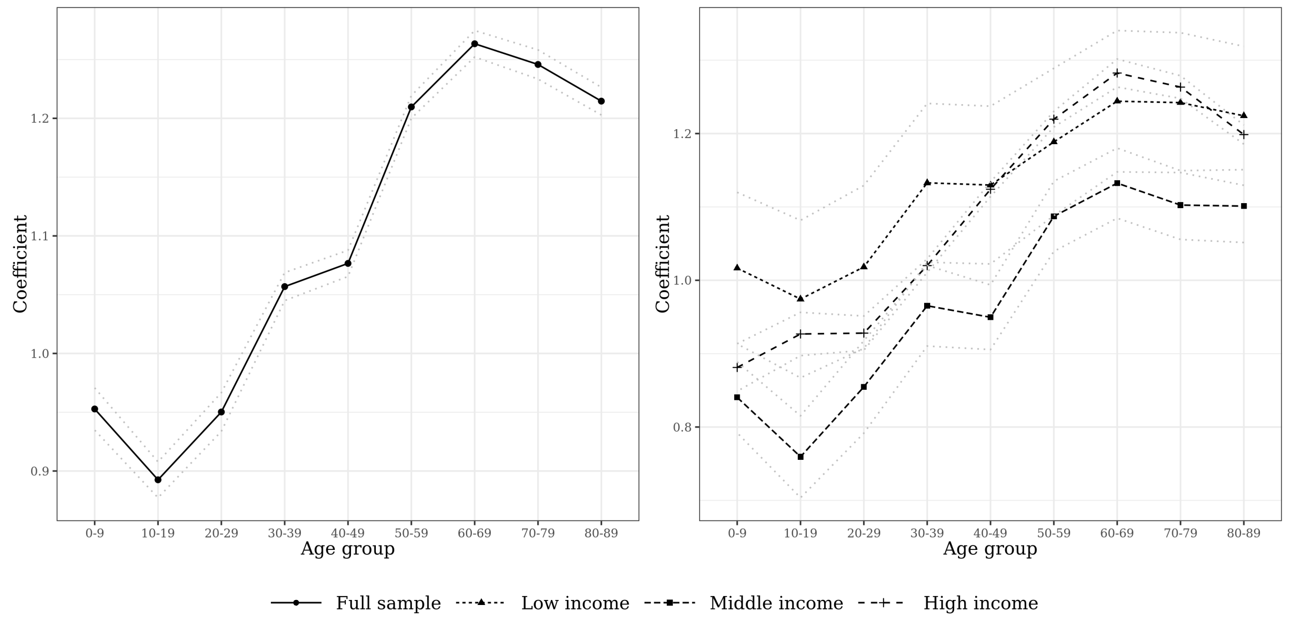

This figure plots the estimated values for δ from equation 8. The right-hand panel shows the value estimated across all countries, while the left-hand represents the value with the sample split across the World Bank's country income classification. Notice that "Low" and "Lower-middle income" have been merged into "Low income" to ensure a sufficient number of observations. Therefore, "Middle income" refers to "Upper-middle income" countries according to the World Bank's country income classification.

Figure 10: Estimated value of δ^m by income group

displays the estimated values for δ^m using the entire LIS dataset and by income group. To understand the figure, consider two households, each with two members. The only difference between these households is the age of one of their members. In the first household, the member is 60 years old, while in the second household, the member is 10. The figure shows their respective δ values are approximately 0.9 and 1.25. Since δ enters multiplicatively in equation 7, the consumption shares of each household must be multiplied by $\sqrt{0.9}$ and $\sqrt{1.25}$ respectively. Note that the square root represents the individual's weight within a two-individual household. Therefore, the household with the 60-year-old member will consume 17%[10] more on food than the household with the 10-year-old member. Since an additional household member's impact on food expenditure shares is weighted by the household size, interpreting the coefficients' numerical implications may be difficult. However, higher coefficients imply higher food expenditure shares *ceteris paribus*. The right-hand side of figure 5 shows the estimated coefficient with the sample split across World Bank's country income classification. The value of the coefficient does not change drastically between income groups, confirming that differences in household structure were the driver of the

[10] $\sqrt{1.25}/\sqrt{0.9} \approx 1.17$. Since we assume that the household is symmetric, the remaining part of equation 7 are the same, and thus cancel over.

heterogeneity observed in the previous section.

Figure 11 and 10 show the results obtained by estimating equation 8 for each county in the LIS dataset. First, we can observe a large heterogenity between the estimates for ϵ across the development spectrum. As figure 11 shows, richer countries have a higher value for ϵ than poorer ones. Therefore, per equation 9, food expenditure shares in richer countries are more sensible to changes in real income. Figure 10 shows the demographic coefficient for each country in the sample: all countires show an upward trend for the coefficients of elderly individuals, suggesting a strong correlation between individuals' age and food expenditure shares.

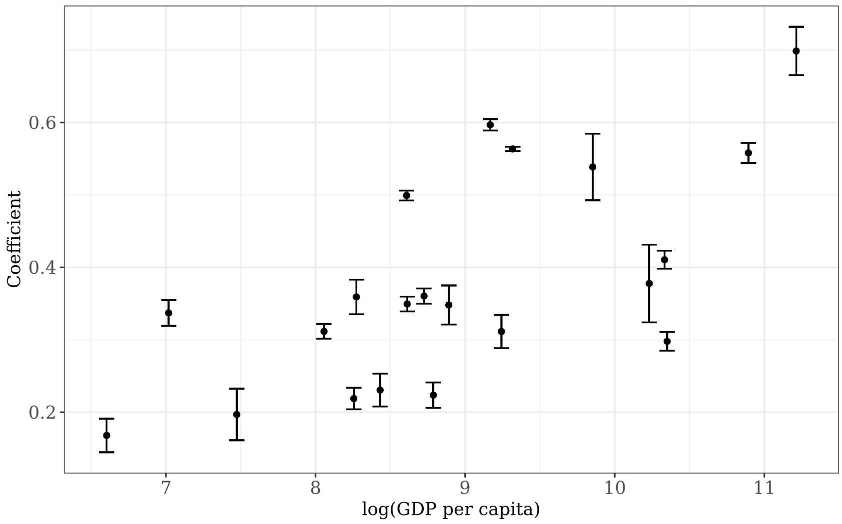

This figure plots the estimated values for ϵ from equation 8 by country. The x-axis represents the average expenditure per capita, weighted using LIS's household-level sample weights.

Figure 11: Estimated values of ϵ by country

4.3 Driver decomposition

One key property of PIGL preferences is that they allow for tractable aggregation. In particular, the aggregated food share of total consumption can be written as

$$\Omega_{\mathrm{f}} \equiv \frac{\sum_h^H E_h^f}{\sum_h E_h} = \left(\frac{P_t^n}{\bar{E}_t}\right)^{\epsilon} \left(\frac{P_t^f}{P_t^n}\right)^{\gamma} \bar{\delta}_t \cdot \theta_t \cdot \nu_t,$$

where

$$\bar{E}_t \equiv \frac{1}{N_t}\sum_h N_{h,t}\cdot E_{h,t} \qquad \text{(Average expenditures per capita)}$$

$$\bar{\delta}_t \equiv \frac{1}{H_t}\sum_h^H \frac{E_{h,t}}{\bar{E}_t}\cdot\delta_t^h \qquad \text{(Expenditure-weighted average of HH demographic shifters)}$$

$$\theta_t \equiv \frac{1}{H_t}\sum_h^H \frac{\delta_t^h}{\bar{\delta}_t}\cdot\left[\frac{E_{h,t}}{\bar{E}_t}\right]^{1-\epsilon} \qquad \text{(Preference-weighted expenditure inequality)}$$

As in Boppart (2014), taking a log change of aggregated share of food consumption from a reference period τ allows us to decompose the different drivers of a change in aggregate food consumption:

$$\hat{\Omega}_t^f = \underbrace{\epsilon(\hat{\mathbf{P}}_t - \hat{E}_t)}_{\text{Income}} + \underbrace{(\gamma - \epsilon\Omega_t^f)(\hat{P}_t^f - \hat{P}_t^n)}_{\text{Substitution}} + \underbrace{\hat{\bar{\delta}}_t}_{\text{Demography}} + \underbrace{\hat{\theta}_t + \hat{\nu}_t}_{\text{Residual}}, \tag{9}$$

where

$$\hat{x}_t \equiv \ln x_t - \ln x_\tau \quad \forall x \qquad \text{(cumulative log change)}$$

$$\hat{\mathbf{P}}_t \equiv (1-\Omega_t^f)\hat{P}_t^n + \Omega_t^f\hat{P}_t^f \qquad \text{(log change in the aggregate price index)}$$

The concept behind equation 9 is easy to understand: the income effect is defined as the effect of a change in *real* expenditures ($\hat{\mathbf{P}}_t - \hat{E}_t$), while the substitution effect is defined as the log change in relative price time the price elasticity ($\gamma - \epsilon\Omega_t^f$). Finally, the demographic effect is the change in the expenditure-weighted average demographic taste-shifters.

Figure 12 illustrates the breakdown of the income and demographic drivers for all countries in the sample. As we can see, the income effect is a critical mechanism that reduces aggregate food expenditure shares, consistent with Engel's Law and unsurprising. However, the most interesting finding is the effect of changes in demographic structure. As shown, demographic changes are a positive driver of food consumption in all countries except Guatemala, where the changes in demography structure resulted in a decline in the food shares of total consumption. Interestingly, a few countries observed increased aggregate food consumption: Hungary, Mali, Guatemala, Vietnam, and Jordan.

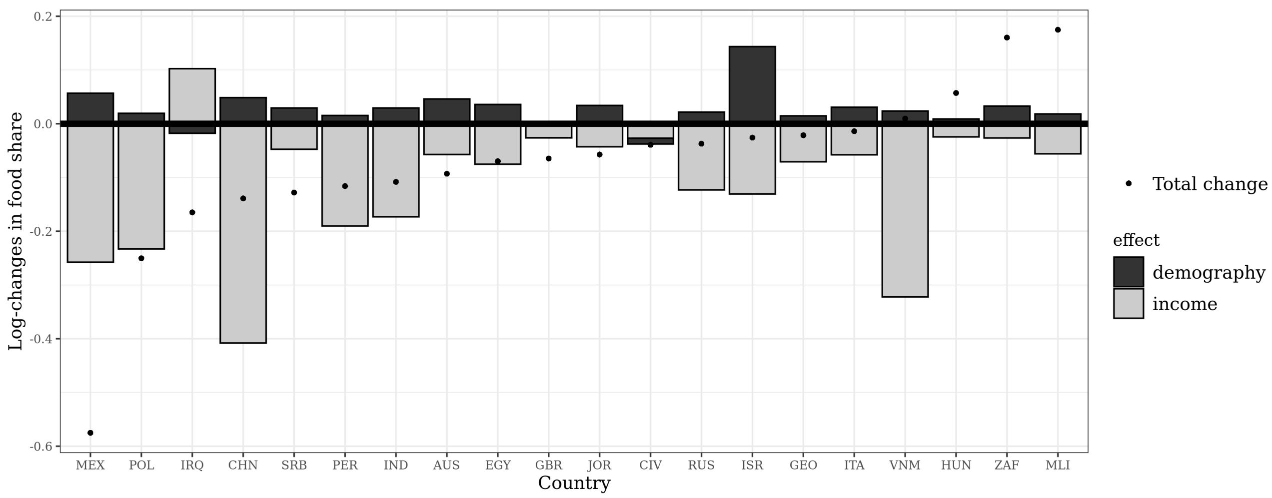

This figure shows the value of the different drivers from equation 9. Countries are observed over a distinct period (see table 6), and the black dot shows the total change over said period. Values represent log changes of aggregated food expenditures.

Figure 12: Decomposition of income and demographic drivers of structural transformation

Country	Interval	$\Delta\Omega_f$	$\Delta\Omega_f$ (CAGR)	$\Delta\Omega_f$ (counterfactual)	$\Delta\Omega_f$ (counterfactual, CAGR)	Δ	Δ (yearly)
Mali	2011-2019	10.45	2.21	9.28	1.98	1.17	0.23
India	2004-2011	-4.54	-1.53	-5.69	-1.95	1.15	0.42
Ivory Coast	2002-2015	-1.71	-0.30	-1.25	-0.22	-0.46	-0.08
Vietnam	2005-2013	0.37	0.12	-0.53	-0.17	0.90	0.29
Egypt	1999-2017	-2.85	-0.39	-4.24	-0.59	1.39	0.20
Jordan	2002-2013	-2.16	-0.52	-3.38	-0.83	1.22	0.31
Georgia	2009-2019	-0.75	-0.21	-1.26	-0.36	0.51	0.15
Iraq	2007-2012	-6.90	-3.24	-6.21	-2.90	-0.69	-0.34
South Africa	2008-2017	3.70	1.80	2.90	1.43	0.80	0.37
Serbia	2006-2019	-4.66	-0.98	-5.64	-1.20	0.98	0.22
Peru	2004-2019	-3.77	-0.77	-4.23	-0.87	0.46	0.10
China	2002-2018	-4.22	-0.86	-5.56	-1.17	1.34	0.31
Russia	2000-2010	-1.27	-0.37	-1.98	-0.59	0.71	0.22
Mexico	1996-2018	-20.10	-2.58	-21.53	-2.83	1.43	0.25
Hungary	1999-2015	2.04	0.36	1.71	0.30	0.33	0.06
Poland	1999-2019	-7.43	-1.24	-7.93	-1.34	0.50	0.10
Italy	1998-2016	-0.49	-0.08	-1.55	-0.25	1.06	0.17
Israel	2001-2019	-0.47	-0.14	-2.87	-0.94	2.40	0.80
United Kingdom	1990-1993	-1.30	-2.13	-1.30	-2.13	0.00	0.00
Australia	2004-2016	-1.25	-0.77	-1.82	-1.15	0.57	0.38

This table shows the outcome of a counterfactual exercise where the demographic driver from equation 9 has been shut down. The column $\Delta\Omega_f$ shows the observed change in the aggregate food expenditure share in the interval shown in the second column. Column $\Delta\Omega_f$ (CAGR) shows the observed yearly cumulated growth rate. The following two columns show the counterfactual change and the implied yearly growth rate after shutting down the demographic channel. Column Δ shows the differences between the observed and the counterfactual changes: positive values imply that the food shares are lower when the demographic channel is turned off. That is, demographic changes led to an increase in food expenditure shares.

Table 6: Counterfactual exercise

To quantify the impact of the demography on food consumption, I conducted a counterfactual exercise. Table 6 demonstrates the observed change in food consumption shares ($\Delta\Omega_f$ in percentile points) and the implied compounded change rate over the observed period (in % points). The following two columns show how food consumption would be if we shut down the demographic effect and the compounded change rate. Finally, the last two columns display the difference between the observed baseline and the counterfactual. Changes in demographic structure significantly affect the annual growth rates. For example, we observed that food consumption in Mali increased by 2.21% annually in the period 2011-2019, of which 0.23 percentage points ($\approx$ 10%) can be attributed to the demographic transition. However, there is strong cross-country heterogeneity: most high-income countries observe a small impact, while some countries, such as the United Kingdom, report no effect.

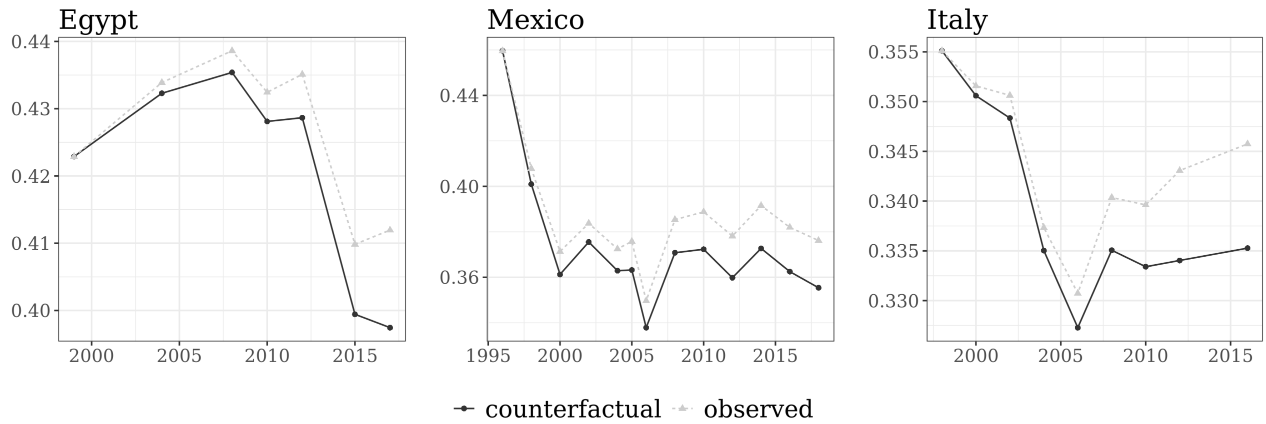

This figure shows the evolution of the aggregate food expenditure shares (at fixed prices) as observed ("observed") and keeping the demographic composition fixed as the first year observed ("counterfactual") for Egypt, Mexico, and Italy. For the latter country, the 2014 observation has been omitted due to the 2013 European sovereign debt crisis: as average spending declined, it resulted in a huge spike in food expenditure shares. Since this work focuses on long-run trends, the observation has been removed for visual clarity.

Figure 13: Evolution over time of aggregate food shares with and without demographic controls

Figure 13 shows the evolution over time of aggregate food expenditure shares at fixed prices with and without demographic controls. The three countries have the longest periods observed by LIS data for each World Bank country's income classification. For all three countries, the demographic transition represents a strong slowing force upon structural transformation: the counterfactual expenditure shares lie consistently below the observed shares, suggesting that keeping the demographic (and household composition) constant would increase food expenditure shares.

So far, I have documented that income - proxied by total expenditures - and demography affect food consumption. At the aggregate level, income growth decreases food consumption while the demographic transition increases it. However, as I argued in the introduction, demographic transition and income are strongly correlated; a richer country is usually ahead in the demographic transition. Therefore, estimating the income effect alone without considering demographic trends would underestimate the income effect, as the estimation would incorporate the negative demographic component. To test this hypothesis, I conducted another counterfactual exercise where I estimated an equation without demographic controls:

$$\log(\omega_h^f) = \epsilon \cdot \left(\frac{P_t^n}{E_{h,t}}\right) + \alpha_{c,t} + \xi_h. \tag{10}$$

This allows me to estimate the income effect from equation 9 without considering the demographic driver. However, the fixed effect $\alpha_{c,t}$ still accounts for a time trend. Figure 14 shows the ratio between the income effect computed using the coefficients estimated using equation 10 and the one calculated using the parameters from equation 8: accounting for the demographic transition for almost all countries increases the magnitude of the income effect by up to 25%.

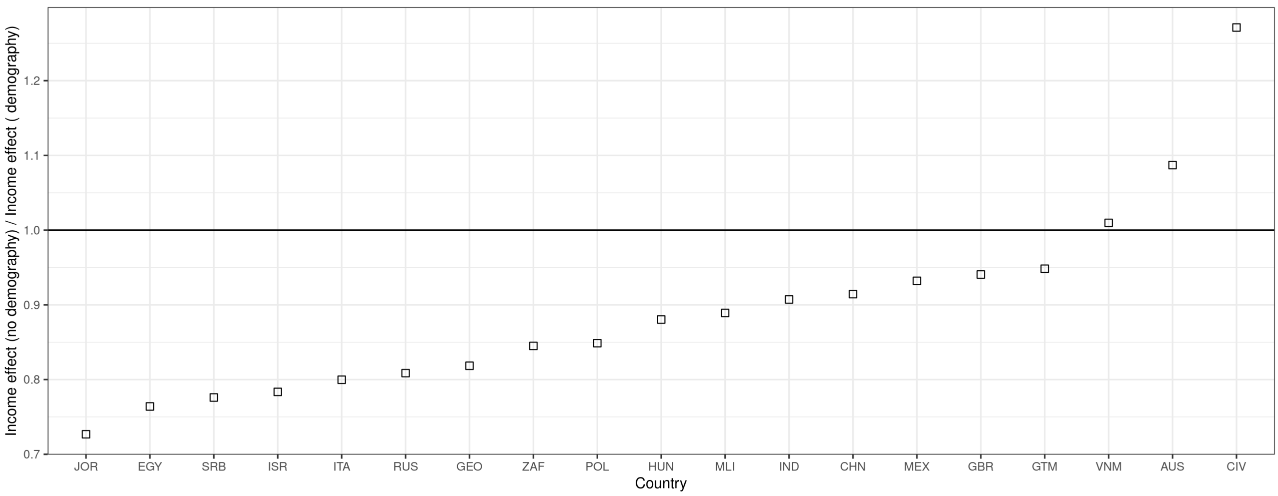

This figure shows the ratio between the income effect from equation 9 computed using the coefficients estimated using equation 10 and the one computed using the parameters from equation 8.

Figure 14: Ratio between income effects without and with demographic controls

5 Conclusion

This study examines the role of demographic transition in influencing structural transformation, particularly regarding household food expenditures. It analyzes aggregated and household-level consumption data from 20 countries across the development spectrum and demonstrates that individuals' age characteristics affect the proportion of total expenditures allocated to food. This relationship holds consistently across both developed and developing nations when the exact household composition is considered, indicating a shared pattern where older individuals drive higher food expenditure shares at the household level.

The quantitative model constructed in this work, based on the PIGL demand system, confirms the hypothesis that demographic transitions slow down structural transformation by increasing the food share of total expenditures. The model's calibration and counterfactual exercises show that shutting down the demographic channel would decrease the change rate of the food share of aggregated expenditure by between 0.1 to 0.8 percentage points yearly for most countries. Furthermore, ignoring demographic factors results in a substantial underestimation (up to 25%) of the income effect on food expenditures.

References

Mark Aguiar and Mark Bils. Has consumption inequality mirrored income inequality? *American Economic Review*, 105(9):2725–56, 2015.

Mark Aguiar and Erik Hurst. Deconstructing life cycle expenditure. *Journal of Political Economy*, 121(3):437–492, 2013.

Emma Aguila, Orazio Attanasio, and Costas Meghir. Changes in consumption at retirement: evidence from panel data. *Review of Economics and Statistics*, 93(3):1094–1099, 2011.

Yunus Aksoy, Henrique S Basso, Ron P Smith, and Tobias Grasl. Demographic structure and macroeconomic trends. *American Economic Journal: Macroeconomics*, 11(1):193–222, 2019.

Simon Alder, Timo Boppart, and Andreas Muller. A theory of structural change that can fit the data. *American Economic Journal: Macroeconomics*, 14(2):160–206, 2022.

David Baqaee and Elisa Rubbo. Micro propagation and macro aggregation. *Annual Review of Economics*, 15(1):91–123, 2023.

David Rezza Baqaee and Emmanuel Farhi. The macroeconomic impact of microeconomic shocks: Beyond hulten's theorem. *Econometrica*, 87(4):1155–1203, 2019.

William J Baumol. *The cost disease: Why computers get cheaper and health care doesn't*. Yale university press, 2012.

Timo Boppart. Structural change and the kaldor facts in a growth model with relative price effects and non-gorman preferences. *Econometrica*, 82(6):2167–2196, 2014.

Laurent Brembilla. The demographics of structural change. In *24th Conference of the Society of Computational Economics and Finance*, 2018.

Martin Browning, Pierre-Andre Chiappori, and Arthur Lewbel. Estimating consumption economies of scale, adult equivalence scales, and household bargaining power. *Review of Economic Studies*, 80(4):1267–1303, 2013.

Diego Comin, Danial Lashkari, and Martí Mestieri. Structural change with long-run income and price effects. *Econometrica*, 89(1):311–374, 2021.

Javier Cravino, Andrei A Levchenko, and Marco Rojas. Population aging and structural transformation. Technical report, National Bureau of Economic Research, 2022.

Matthew J Delventhal, Jesús Fernández-Villaverde, and Nezih Guner. Demographic transitions across time and space. Technical report, National Bureau of Economic Research, 2021.

Matthias Doepke, Anne Hannusch, Fabian Kindermann, and Michèle Tertilt. The economics of fertility: A new era. Technical report, National Bureau of Economic Research, 2022.

Georg Duernecker, Berthold Herrendorf, and Akos Valentinyi. Structural change within the services sector and the future of cost disease. *Journal of the European Economic Association*, 22(1):428–473, 2024.

E Engel. Die productions-und consumtionsverhältnisse des königreichs sachsen (1857). wiederabgedr. *Derselbe, Die Lebenskosten belgischer Arbeiter-Familien früher und jetzt. Ermittelt aus Familien-Haushaltrechnungen und vergleichend zusammengestellt. Dresden*, 1895.

AC Foster. Consumer expenditures vary by age. beyond the numbers: Prices and spending, 4 (14). us bureau of labor statistics, 2015. URL https://www.bls.gov/opub/btn/volume-4/consumer-expenditures-vary-by-age.htm.

Berthold Herrendorf, Richard Rogerson, and Akos Valentinyi. Two perspectives on preferences and structural transformation. *American Economic Review*, 103(7):2752–89, 2013.

Erik Hurst. The retirement of a consumption puzzle. Technical report, National Bureau of Economic Research, 2008.

Instituto Nacional de Estadística y Geografá (INEGI). Encuesta nacional de ingresos y gastos de los hogares (enigh). URL https://www.inegi.org.mx/programas/enigh/.

Charles I Jones. The end of economic growth? unintended consequences of a declining population. *American Economic Review*, 112(11):3489–3527, 2022.

Valérie Lechene, Krishna Pendakur, and Alex Wolf. Ols estimation of the intra-household distribution of consumption. Technical report, IFS Working Papers, 2019.

Luxembourg Income Study (LIS) Database. Luxembourg income study (lis) database. URL https://www.lisdatacenter.org.

Thomas Malthus. An essay on the principle of population. printed for j. johnson. *St. Pauls church-yard, London*, pages 1–126, 1798.

Rui Mao and Jianwei Xu. Population aging, consumption budget allocation and sectoral growth. *China Economic Review*, 30:44–65, 2014.

John Muellbauer. Aggregation, income distribution and consumer demand. *The Review of Economic Studies*, 42(4):525–543, 1975.

Organisation for Economic Co-operation and Development (OECD). Final consumption expenditure of households. URL https://stats.oecd.org/.

United Nations. World population prospects 2024 revision. URL https://population.un.org/wpp/.

U.S. Bureau of Labor Statistics. Consumer expenditure surveys (cex), 2025. URL https://www.bls.gov/cex/.

World Bank. World development indicators. URL https://databank.worldbank.org/source/world-development-indicators.

Carle C Zimmerman. Ernst engel's law of expenditures for food. *The Quarterly Journal of Economics*, 47(1):78–101, 1932.

Impact of a change in the reference sector

Assume that the preferences take the form

$$\mathcal{V}^h(\mathbf{P}, E_{h,t}) = \frac{1}{\epsilon}\left[\frac{E_t^h}{P_t^f}\right]^{\epsilon} - \frac{\nu_t^h}{\gamma}\left[\frac{P_t^n}{P_t^f}\right]^{\gamma} - \frac{1}{\epsilon} + \frac{\nu_t^h}{\gamma}.$$

We used the food price (P_t^f) as the reference sector. Then, the food share of total consumption is

$$\omega_t^f = 1 - \underbrace{\left(\frac{E_{h,t}}{P_t^f}\right)^{-\epsilon} \cdot \left(\frac{P_t^n}{P_t^f}\right)^{\gamma} \cdot \nu_t^h}_{\equiv \omega_t^n} \tag{11}$$

Notice that equation 11 is non-log-linear and thus cannot be estimated directly by OLS. However, we can log-linearize the *non-food shares* ω_t^n. That is, under the same assumptions as in the main model,

$$\log(\omega_t^n) = \epsilon \cdot \log\left(\frac{P_t^f}{E_{h,t}}\right) + \sum_m^M s_m^h \cdot \text{DUMMY}_m + \alpha_{c,t} + \epsilon_t \tag{12}$$

Notice that compared to equation 8, the parameters γ and the demographic dummies δ_m have different values and interpretations as they refer to the complement sector. Again, following Boppart (2014), we can aggregate consumption as

$$\Omega_t^{\mathrm{n}} \equiv \frac{\sum_t^H E_t^n}{E_t} = \left(\frac{P_t^f}{E_t}\right)^{\epsilon}\left(\frac{P_t^n}{P_t^f}\right)^{\gamma} \bar{\delta}_t \cdot \theta_t \cdot \nu_t, \tag{13}$$

Since $\Omega_t^f = 1 - \Omega_t^n$, the log difference between time t and T is:

$$\begin{aligned}
\hat{\Omega}_t^n \equiv& \log\Omega_t^n - \log\Omega_T^n = \epsilon(\hat{P}_t^f - \hat{E}_t) + \gamma(\hat{P}_t^n - \hat{P}_t^f) + \hat{\bar{\delta}}_t + \hat{\theta}_t + \hat{\nu}_t \\
=& \epsilon(\hat{P}_t - \hat{E}_t) + (\gamma - \epsilon\Omega_t^n)(\hat{P}_t^n - \hat{P}_t^f) + \hat{\bar{\delta}}_t + \hat{\theta}_t + \hat{\nu}_t \\
\hat{\Omega}_t^f \equiv& \log\Omega_t^f - \log\Omega_T^f = \log(1 - \Omega_t^n) - \log(1 - \Omega_T^n)
\end{aligned}$$

First order Taylor approximation around time T

$$\log\Omega_t^f \equiv \log(1-\Omega_t^n) \approx \log(1-\Omega_T^n) - \frac{1}{1-\Omega_T^n}(\Omega_t^n - \Omega_T^n) = \log(\Omega_T^f) - \frac{1}{\Omega_T^f}(\Omega_t^n - \Omega_T^n)$$

and

$$\log\Omega_t^n \approx \log\Omega_T^n + \frac{1}{\Omega_T^n}(\Omega_t^n - \Omega_T^n) = \log\Omega_T^n - \left(1 - \frac{\Omega_t^n}{\Omega_T^n}\right)$$

substituting yields

$$\hat{\Omega}_t^f \approx -\frac{\Omega_T^n}{\Omega_T^f}\left(1-\frac{\Omega_t^n}{\Omega_T^n}\right)$$
$$\hat{\Omega}_t^n \approx \left(1-\frac{\Omega_t^n}{\Omega_T^n}\right)$$

Therefore

$$\hat{\Omega}_t^f \approx -\frac{\Omega_T^n}{\Omega_T^f}\cdot\hat{\Omega}_t^n = -\frac{\Omega_T^n}{\Omega_T^f}\left[\underbrace{\epsilon(\hat{P}_t - \hat{E}_t)}_{income} + \underbrace{(\gamma - \epsilon\Omega_t^n)(\hat{P}_t^n - \hat{P}_t^f)}_{substitution} + \underbrace{\hat{\hat{\delta}}_t}_{demography} + \underbrace{\hat{\theta}_t + \hat{\nu}_t}_{residual}\right] \tag{14}$$

Figure 15 shows the drivers from equation 14 computed using the parameters coming from the OLS estimations of 12. The resulting drivers are quantitatively similar to the one estimated before. The only major difference is a flip in Australia's demographic driver sign.

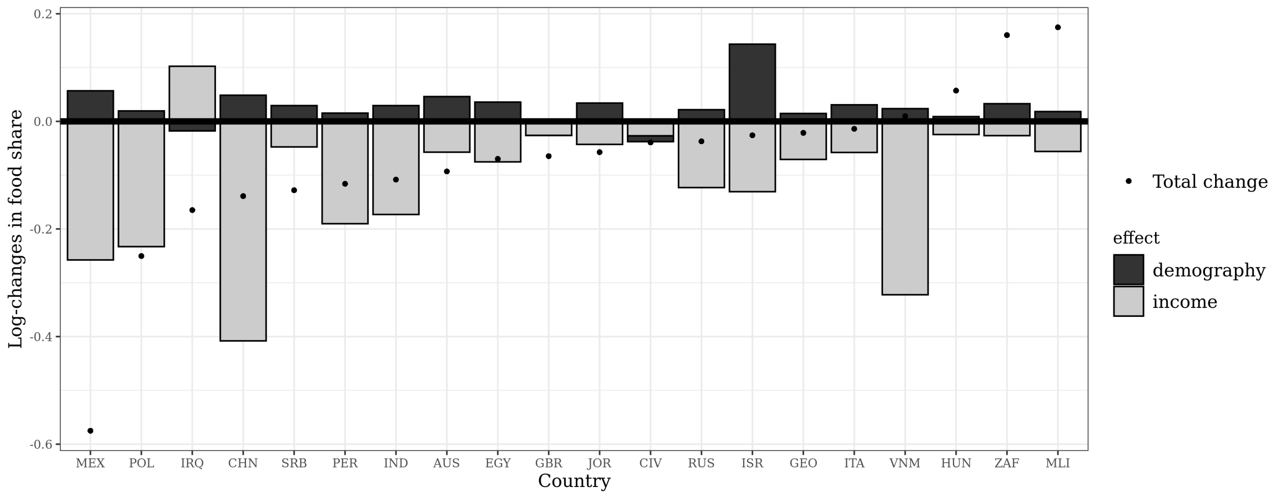

This figure shows the value of the different drivers from equation 9. Countries are observed over a distinct period (see table 6), and the black dot shows the total change over said period. Values represent log changes of aggregated food expenditures.

Figure 15: Decomposition of income and demographic drivers of structural transformation (alternate reference sector)

Replication with external datasets

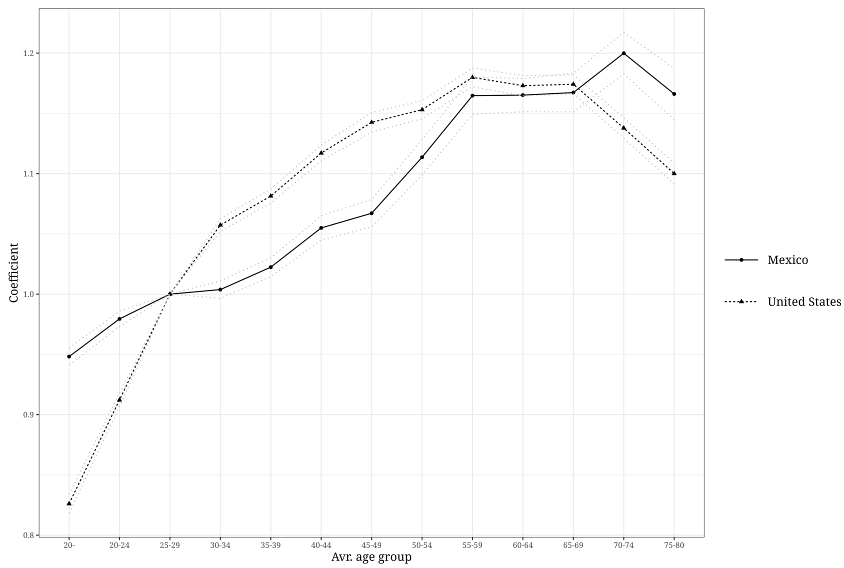

This figure plots the estimated values for $\exp(D_h^{\text{Avr. Age}})$ from equation 3 using data from the United States (CE) and Mexico (ENIGH).

Figure 16: Estimated value of $\exp(D_h^{\text{Avr. Age}})$ for US and Mexico.

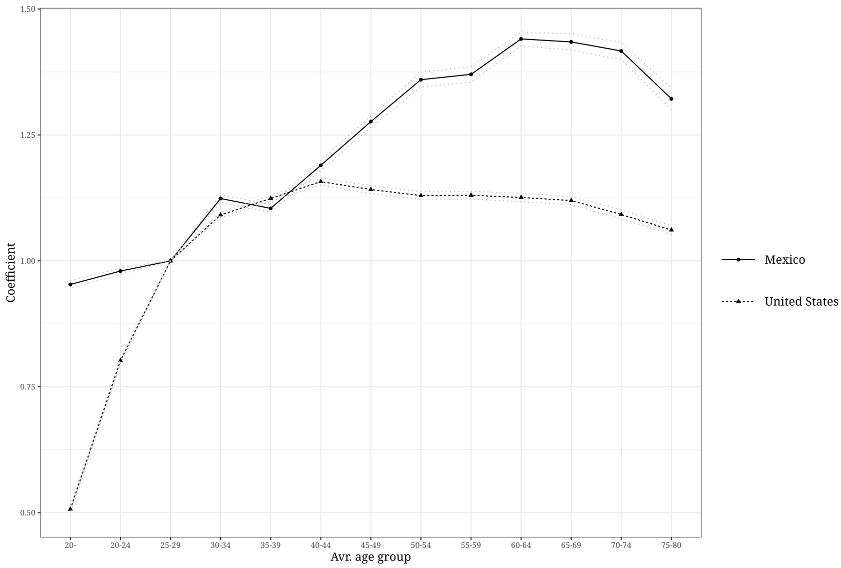

This figure plots the estimated values for $\exp(D_h^{\text{Avr. Age}})$ from equation 3 using data from the United States (CE) and Mexico (ENIGH). The sample has been reduced to households composed of only a single member.

Figure 17: Estimated value of $\exp(D_h^{\text{Avr. Age}})$ for US and Mexico (one-member household)

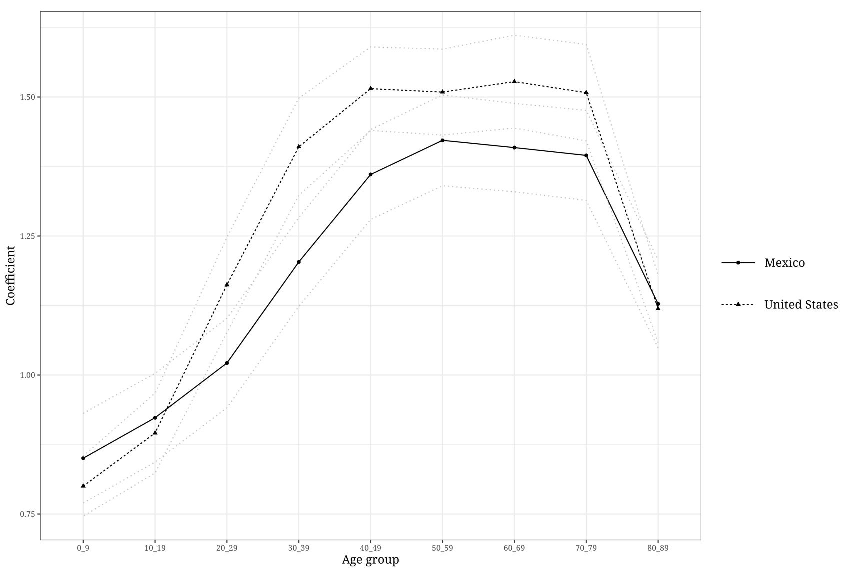

This figure plots the estimated values for δ from equation 8 using data from the United States (CE) and Mexico (ENIGH).

Figure 18: Estimated value of δ_m for US and Mexico