

A New Engel on the Gains from Trade*

David Atkin,[†] Benjamin Faber,[‡] Thibault Fally[§] and Marco Gonzalez-Navarro[¶]

July 2018

-Preliminary and Incomplete-

Abstract

Measuring the gains from trade and their distribution is challenging. Recent empirical contributions have addressed this challenge by drawing on rich and newly available sources of microdata to measure changes in household nominal incomes and price indices. While such data have become available for some components of household welfare, and for some locations and periods, they are typically not available for the entire consumption basket. In this paper, we propose and implement an alternative approach that uses rich, but widely available, expenditure survey microdata to estimate theory-consistent changes in income-group specific price indices and welfare. Our approach builds on existing work that uses linear Engel curves and changes in expenditure on income-elastic goods to infer unobserved real incomes. A major shortcoming of this approach is that while based on non-homothetic preferences, the price indices it recovers are homothetic and hence are neither theory consistent nor suitable for distributional analysis when relative prices are changing. To make progress, we show that we can recover changes in income-specific price indices and welfare from horizontal shifts in Engel curves if preferences are quasi-separable (Gorman, 1970; 1976) and we focus on what we term “relative Engel curves”. Our approach is flexible enough to allow for the highly non-linear Engel curves we document in the data, and for non-parametric estimation at each point of the income distribution. We first implement this approach to estimate changes in cost of living and household welfare using Indian microdata. We then revisit the impacts of India’s trade reforms across regions.

Keywords: Gains from trade, household welfare, price indices, non-homothetic preferences, Engel curves.

JEL Classification: F63, O12, D12.

*David Atkin and Ben Faber thank the Stanford Economics Department and SIEPR for their hospitality while writing this paper. We thank participants at seminars and conferences for helpful comments. Amir Hafiz Amir Hisham, Tasnia Hussain, Yulu Tang, and Guan Yi Lin provided excellent research assistance.

[†]Department of Economics, MIT and NBER.

[‡]Department of Economics, UC Berkeley and NBER.

[§]Department of Agricultural and Resource Economics, UC Berkeley and NBER.

[¶]Department of Agricultural and Resource Economics, UC Berkeley.

1 Introduction

Much of the recent trade literature has been preoccupied with estimating the gains generated by international trade. Given recent political upheaval in the US and elsewhere, there is an increased urgency to accurately measure the distribution of these gains. Measuring the effects of trade on household incomes and price indices, however, requires extremely detailed microdata that are seldom, if ever, available. In particular, capturing changes in household cost of living in the denominator of real incomes requires, at minimum, knowledge of price changes for every item in household expenditures down to the variety level. If we take seriously the fact that products change in quality, new products arrive and old products exit, and new shopping options open and close, then even knowledge of all price changes is not sufficient (e.g. Hausman 1996; Feenstra 1994; Redding and Weinstein, 2017). As the existing literature has shown, shocks to trade openness with other countries, or across regions within a country, can affect each of these dimensions.

Recent papers estimating the gains from trade have addressed this challenge by bringing to bear newly available and very rich sources of microdata on consumption prices and quantities (e.g. Atkin, Faber and Gonzalez-Navarro, 2018; Borusyak and Jaravel, 2018; Hottman and Monarch, 2018). While such data have become available to estimate accurate price indices for some components of household welfare—e.g. developed-country retail consumption using scanner microdata covering roughly 20 percent of consumption, or developing-country expenditure surveys on well-measured basic foodstuffs and fuel covering 50 percent or more of consumption—these types of data are infeasible to collect for the entire consumption basket. Accurately measuring prices and quantities for services is particularly fraught with difficulty. Furthermore, even in the richest data environments evaluating changes in welfare from observed price data still requires strong functional form assumptions (e.g. quantifying the gains from variety).

In this paper, we instead propose and implement a new approach that uses rich, but widely available, expenditure survey microdata—and in particular does not require observing reliable price data for all consumption categories¹—to estimate theory-consistent changes in exact household price indices and welfare at every point of the household income distribution. We then implement this approach to quantify changes in household welfare across Indian regions and over and time, revisit the impacts of India’s 1991 trade reforms studied by Topalova (2010), and investigate the gains from trade in a cross-country setting.

Our approach is related to a longstanding literature using Engel curves and expenditure changes on income-elastic goods to recover unobserved changes in real income (e.g. Hamilton, 2001; Costa, 2001; Young, 2012; Nakamura et al., 2016). The initial goal of this literature was to correct biases in US consumer price index (CPI) measures due to difficulties in measuring prices, quantities and quality changes for consumption categories such as services. For example, Hamilton (2001) uses observed changes in food consumption shares and estimates of food Engel curves to correct US CPI estimates. This exercise has since been repeated for many developed and developing countries (see references contained in Nakamura et al., 2016). Almås (2012) applies

¹As we discuss below, while our method recovers the full price index from expenditure data on a subset of consumption and total outlays alone, we also make extensive use of available price data from well-measured consumption categories to test the preference restrictions and identifying assumptions behind our methodology, and to compute correction terms if necessary.

this approach to correct for unobserved biases in purchasing power parity comparisons across countries, while Young (2012) estimates real income growth in sub-Saharan Africa.

The bulk of this literature estimates linear Engel curves generated by the Almost Ideal Demand System (AIDS).² While this approach leans heavily on non-homotheticity—if demand is homothetic Engel curves are horizontal and so changes in expenditure shares are uninformative about changes in welfare—we show that existing applications only correctly recover changes in the price index under a specific realization of unobserved price changes, such that changes in the price index are uniform across households at different income levels (i.e. by assuming away the income-group specific price indices generated by a non-homothetic demand system).³

To make progress on these challenges, our analysis proceeds in four steps. In the first step, we document two motivating facts using the Indian expenditure survey microdata. First, we document that Engel curves in the data are non-linear: we formally reject linearity in the relationship between budget shares (y-axis) and log total outlays per capita (x-axis)—what we call “textbook” Engel curves following Working’s now standard formulation (1943)—for about 90 percent of goods and services at a 95 percent confidence level. Second, we show that Engel curves shift over time within a given market, and across markets within the same period, and that those horizontal shifts are not uniform across households of different income levels.

In the second step, we propose a novel methodology that addresses the drawbacks of the existing approaches using Engel curves for welfare estimation, and is consistent with the two motivating facts. In particular, our method uses observed horizontal shifts in what we term “relative Engel curves” across time or space to recover theory-consistent changes in exact price indices and household welfare at each point of the income distribution. To fix ideas, consider the textbook Engel curve for food at two different points in time in the same market. The horizontal distance between curves at any point in the distribution of log nominal outlays per capita (the variable on the x-axis) reveals the change in log nominal outlays that holds the food share constant across the two sets of prices. First, Lemma 1 shows in a very general setting (for any rational utility function), that this horizontal distance in the log income space (x-axis) recovers the change in the price index at any point in the income distribution, but only under the assumption of constant relative prices across the two periods. However, if there are no relative price changes, shifts in Engel curves must be parallel—in violation of the second motivating fact in the data—and changes in price indices must be uniform across the income distribution. Unfortunately, we also prove in Lemma 2 that when relative prices are allowed to change in ways consistent with the second motivating fact above, then horizontal distances between textbook Engel curves do not in general recover changes in price indices.

To make progress, we then add additional structure to the very general preferences above in order to relax the restrictions on unobserved price realizations that preclude price index changes from being income-group specific. Our approach focuses on the broad class of quasi-separable

²Some studies, such as Almås (2012) and Almås, Beatty and Crosely (2018), have also used quadratic Engel curves under QUAIDS preferences.

³We review the existing approach in Section 2 below. Almås, Beatty and Crosely (2018) also make note of this shortcoming. As we show below, allowing for non-uniform price index changes under the existing AIDS methodology re-introduces the need for observing the full vector of consumption prices. Their paper addresses this challenge either by using available regional price information, bounding the estimates or imposing additional structure on relative price effects. In a related paper, Almås and Kjelsrud (2016) apply these approaches to measuring inequality in India.

preferences (following Gorman (1970; 1976))⁴ and what we term “relative Engel curves”, that describe how relative expenditure shares within any given subset of goods or services G (i.e. spending on $i \in G$ as a share of total spending on all $i \in G$) vary with log total household outlays per capita. If preferences are quasi-separable, we prove that as long as relative prices remain constant within group G —an assumption we can relax by replacing it with an orthogonality condition or correction term when we turn to estimation below—horizontal shifts in the relative Engel curves of the goods within G reveal changes in exact income-group specific price indices (i.e. the price index for the full consumption basket). It is then straightforward to recover the change in welfare for households at any point in the income distribution from the distance in outlays to go between period 0 and period 1 relative expenditure shares along either period 0’s relative Engel curve (to recover the equivalent variation (EV)) or period 1’s relative Engel curve (to recover the compensating variation (CV)).

This approach has several benefits. We are able to estimate both income-group specific price indices and theory-consistent welfare measures using only commonly available household expenditure survey data. By not placing restrictions on the realization of relative price changes outside of the subset G (or relative price changes between G and other product groups), price indices can differ arbitrarily across income groups, and we can accommodate rich patterns in the data that are consistent with the second motivating fact. Moreover, the approach is flexible enough to allow for the highly non-linear Engel curves we document in the first motivating fact (quasi-separable demand systems can be of any rank, in the sense of Lewbel (1991)), as well as allowing for non-parametric estimation at each point of the income distribution.

An obvious question to ask is how general is the quasi-separable class of preferences? Quasi-separability requires that subsets of goods or services are separable in the expenditure function (not the utility function), so that relative budget shares within a subset G of goods are functions of relative prices within G and household utility. This is less restrictive than the more common assumption of direct separability across goods in the utility function (hence the term “quasi”). The relative expenditures of goods and services within subset G are still a function of all prices in the rest of the economy, but prices outside the subset can only affect the relative expenditures within G through their effect on utility. Examples of preferences in this class are several variants of PIGLOG preferences (e.g. Deaton and Muellbauer 1980), non-homothetic CES preferences as in Gorman (1965), Hanoch (1975) and more recently Comin et al. (2015), and a general class of Gorman preferences discussed in Fally (2018). Relative to these special cases, our approach is more flexible and can, for example, allow for arbitrary own and cross-price effects within and outside of the subset G .⁵

In the third step, we form a bridge between the theoretical results above and the empirical implementation described below. Our estimation approach follows directly from our main theoretical proposition and uses expenditure survey microdata to separately estimate non-parametric relative Engel curves for every period and every good in group G . The horizontal difference

⁴Deaton and Muellbauer (1980) also refer to quasi-separability as implicit separability. Blackorby et al. (1991) distinguish quasi and implicit separability. We describe this class of preferences in more detail just below.

⁵For example, the preferences used in recent work by Comin et al. (2015) satisfy quasi-separability for any possible grouping of goods or services with respect to the rest of household consumption, with a single parameter determining all price effects. Relative to these special cases, our approach is more flexible and can, for example, allow for arbitrary own and cross-price parameters within and outside of the subset G .

between curves across time or across space at various points of the household income distribution reveals changes in the price index for households at those points. To take this approach to the data, we derive four corollaries to our theoretical results that define a set of conditions that need to hold for unique and unbiased identification. First, we show that in order to invert relative Engel curves and obtain unique welfare estimates, we require a unique mapping between household utility and the vector of relative budget shares within G . A testable and sufficient condition is that at least one good or service within G has a relative Engel curve that is monotonic with respect to outlays. Second, we show that we can replace the assumption that relative prices are fixed within group G with an orthogonality condition between changes in relative prices (or taste shocks) and the local slopes of relative Engel curves at a given point in the distribution of household nominal outlays. If this orthogonality condition is violated, the condition naturally generates a first-order correction term that ensures estimates remain unbiased. With these results in hand, we propose restricting our estimation to subsets of well-measured goods for which we have reliable price data—recall that only a subset of goods are required to obtain the complete price index—so that we can explicitly test the orthogonality condition and correct any bias if the test is violated. Third, we address sample selection issues that can arise from the fact that levels of household utility observed in one period or location are not necessarily observed in another. For example, when evaluating price index changes for poor households in period 0, there may be no equally poor households in period 1 if there is real income growth. This means that relative Engel curves may not overlap across the range of observed budget shares in both periods, and we derive suitable sample selection correction methods to address this possibility. Fourth, we derive a set of tests for quasi-separability that again can be implemented using the available expenditure and reliable price information in the expenditure surveys.

In the final step, we implement our methodology in two applications. First, we draw on Indian expenditure survey microdata to quantify changes in welfare over time at different points in the income distribution.⁶ Relative to existing CPI estimates for rural India 1987/88-1999/2000 (Deaton, 2003b) that are based on calculating standard price index numbers using changes in prices of those food and fuel products with reliable quantity information and no evidence of multiple varieties within a given market, we find that our estimates based on relative Engel curves yield broadly comparable levels of consumer price inflation among poorer deciles of the income distribution. Given that these food and fuel products cover around 80 percent of total outlays for poorer income deciles in rural India over this period, it is reassuring our estimates are very similar. But our estimation also brings to light that price inflation has been far from uniform across the income distribution, with significantly higher inflation rates for poorer households—something that is not apparent from calculating standard price indices even, as we show, when using income-group specific expenditure weights. While estimates based on a single price index suggest that there has been significant convergence between poor and rich household income deciles in India over time, we find that cost of living inflation has been substantially lower among richer Indian households compared to the poor, significantly flattening real income convergence. Given that existing estimates have been mainly based on Laspeyres price indices using changes in local unit values for food and fuel consumption, the most likely explanation for these

⁶In future drafts, we will also quantify welfare differences within India across space.

findings is that higher-income Indian households disproportionately benefited from previously omitted components of consumer inflation over this period in India. For example, our findings are consistent with lower price growth among non-food consumption (manufacturing and services) that richer households spend larger portions of their budget on, including previously unaccounted for changes in product quality and variety between 1987-2000.

In the second application, we use this machinery to revisit the gains from trade within countries. To do so, we revisit Topalova's (2010) analysis of the regional impacts of India's 1991 trade reforms. The paper uses district employment shares to calculate the local exposure to import tariff cuts and then regresses poverty rates on these trade shocks. Her main finding is that rural poverty rates (the fraction of households below the poverty line) increase as a consequence of local labor market shocks due to import competition. Revisiting this analysis using our estimated welfare changes across the income distribution, we present two main findings. First, while Topalova (2010) highlights effects on poverty rates, our approach uncovers adverse effects of import competition across the full distribution of household income, including at the very top of the income distribution. Second, we find that the adverse effects on household nominal expenditures are amplified when taking into account household cost of living inflation. That is, areas adversely affected by import competition experienced higher local price inflation compared to less exposed regions of India, and these effects are most pronounced at the tails of the local income distribution (especially among the richest households). We verify that this somewhat surprising finding is confirmed by raw price information from the food and fuel product groups carefully selected by Deaton (2003b), and further investigate potential mechanisms.

In work in progress, we plan to use the harmonized cross-national expenditure data assembled by the World Bank to estimate both purchasing power parity across countries and the gains from trade at a national level. For the latter, we plan to revisit the seminal contribution of Frankel and Romer (1999). While this database has already been publicly released by the World Bank, the data at the income percentile level is still pending. Hence, these applications are still in progress at the time writing.⁷

In addition to the literatures that we mention above, this paper relates to a large literature on the structure of demand and household preferences (e.g. Gorman, 1995; Blackorby et al., 1978; Lewbel and Pendakur, 2009; Ligon, 2017), and provides several new results and proofs. Recent work by Ligon (2017) explores an alternative approach based on Frisch demand that assumes preferences are directly separable and isoelastic in own prices to recover a measure related to household welfare. While this approach does not recover equivalent or compensating variations, it can be used to recover differences in the marginal utility of money ("neediness") between households. Since quasi-separable preferences are not directly separable and vice-versa, the two approaches complement one another. Redding and Weinstein (2018) show how to use CES preferences to aggregate up from detailed scanner microdata to estimate welfare from US retail consumption, while taking into account changes in product quality and variety.

The rich distributional analysis that is made possible by the non-homotheticity of preferences also generates parallels with the growing literature on non-homothetic preferences and the gains from trade (e.g. Fajgelbaum et al. (2011); Amiti et al. (2018); Berlingeri et al. (2018); Faber and

⁷We have been assured by the World Bank team in charge that this release is high on their list of priorities.

Fally (2017)). Within this small literature, Fajgelbaum and Khandelwal (2016) is most closely related and pursue a more structural approach. They use an AIDS preference structure and compute welfare changes across the entire income distribution based on observed changes in country-level trade and expenditure shares combined with estimates of income and price elasticities for all goods. A key distinction between their approach and ours is that we do not require estimates of all own and cross-price elasticities—which typically require information on prices that we argue above are rarely available—to compute welfare changes.⁸

Finally, this paper is related to a recent and fast-growing literature that combines machine learning tools with “big data” such as satellite images or cellphone records to infer unobserved changes in economic development (e.g. Blumenstock, 2016; Jean et al., 2016). While this method can be trained to yield relatively precise predictions of observable correlates of welfare, e.g. asset ownership, the lack of theoretical foundations in such “big data” approaches may limit their ability to be used to make predictions in other contexts. Perhaps more importantly, variation in observable correlates of welfare, such as nominal incomes or assets, only crudely translate into measures of household welfare. In this sense, by providing theory-consistent estimates of welfare, our methodology complements recent approaches using machine learning tools by providing improved measures of the target attribute that the algorithm tries to predict. More generally, our methodology is widely applicable in the many contexts where expenditure survey data is available, even if reliable price data are not, and researchers want to understand the welfare effects of policies or shocks and the distribution of those effects.

The remainder of the paper is structured as follows. Section 2 provides a brief review of the existing Engel approach to price index estimation under AIDS preferences. Section 3 describes the data and presents a number of stylized facts that motivate the theoretical framework. Section 4 presents the theory. Section 5 derives a number of corollary results for unique and unbiased identification and presents our two estimation strategies. Section 6 applies our methodology in the three applications described above. Section 7 concludes.

2 Review of Existing Engel Approach

In order to clarify our contribution, in this section we briefly recap existing approaches that use Engel relationships to uncover changes in real income (e.g. Nakamura, 1996; Costa, 2001; Hamilton, 2001; Almås, 2012; Nakamura et al., 2016). These papers estimate linear Engel curves using AIDS and use the recovered income elasticity to infer changes in real income from changes in budget shares on income elastic goods.

To be more precise, under AIDS, Engel curves take the following form:

$$\frac{x_{ih}}{y_h} = \alpha_i + \sum_j \gamma_{ij} \log p_j + \beta_i \log(y_h / \Lambda(p)) \quad (1)$$

where y_h is total nominal outlays per capita for household h , $\frac{x_{ih}}{y_h}$ is the household’s budget share spent on good i as part of total consumption, $\sum_j \gamma_{ij} \log p_j$ are own and cross-price effects and Λ is

⁸Fajgelbaum and Khandelwal (2016) show that with knowledge of the full set of preference parameters (income elasticities and matrix of price elasticities), observed variation in expenditure shares can be used to recover unobserved price realizations. To obtain price elasticities in the absence of reliable price data, they assume a single price elasticity for each sector and estimate this parameter from the ratio of the distance coefficient in the gravity equation to the elasticity of trade costs to distance from Novy (2013).

a price aggregator (not a price index) defined by:

$$\log \Lambda(p) = \alpha_0 + \sum_i \alpha_i \log p_j + \sum_{i,j} \gamma_{ij} \log p_i \log p_j$$

with $\sum_j \gamma_{ij} = 0$ for all i . Hence the literature estimates time series regressions of the form:

$$\frac{x_{iht}}{y_{ht}} = \alpha_{it} + \sum_j \gamma_{ij} \log p_{jt} + \beta_i \log y_{ht} + \epsilon_{cit}$$

where the constant is allowed to vary by time period. Then it is straightforward to see that changes in the intercept over time scaled by income elasticities reveal the change in the price aggregator:

$$\frac{d\alpha_{it}}{\beta_i} = -d \log \Lambda(p_t)$$

with $d \log(y_{ht}/\Lambda(p_t)) = d \log y_{ht} + \frac{d\alpha_{it}}{\beta_i}$.^{9,10} If the constant is allowed to vary by location as well as time, this same method provides price aggregator estimates by location. If the constant is only allowed to vary by location, the method can correct for PPP bias across countries as in Almås (2012).

There are several drawbacks to this approach. While total expenditures divided by a price aggregator is an appealing measure of “real income”, it is not a theory-consistent welfare measure since it does not correspond to welfare in this demand system. AIDS is non-homothetic and, thus, allows for income-specific changes in cost of living over time. More precisely, under the preferences represented by the AIDS expenditure function, the proportional change in household welfare is not $d \log(y_{ht}/\Lambda(p_t))$, but:

$$d \log U_{ht} = d \frac{\log(y_{ht}/\Lambda(p_t))}{\prod_j p_{jt}^{\beta_j}} \quad (2)$$

Essentially there are two price aggregators under AIDS, $\Lambda(p_t)$ and $\prod_j p_{jt}^{\beta_j}$, the combination of which generates income-group specific price indices.

Only in the special case where price realizations are such that $\prod_j (p_{jt})^{\beta_j}$ is unchanged over time is $d \log(y_{ht}/\Lambda(p_t))$ proportional to $d \log U_{ht}$. This special case implies that the true change in price indices is uniform across all incomes (i.e. that it is reflected by a single price aggregator across all income groups). This homotheticity-like restriction is unsatisfactory for a method that relies on Engel curves being non-homothetic, a point also noted by Almås, Beatty and Crossley (2018). It is also unsuitable for evaluating distributional effects since price indices are not allowed to change across the income distribution.

In the next section, we will show that both this restriction and the linearity inherent in AIDS Engel curves are inconsistent with empirical evidence, before proposing a new approach that overcomes these shortcomings.

⁹Papers concerned with CPI bias will typically replace y_{ht} with income divided by the CPI price index, with $-\frac{d\alpha_{it}}{\beta_i}$ equal to the error in the CPI rate.

¹⁰As discussed above, notable exceptions are recent papers by Almås and coauthors (Almås and Kjelsrud (2016), Almås, Beatty and Crossley (2018)).

3 Data and Motivating Facts

3.1 Data

Indian NSS Expenditure Surveys

Following Topalova (2010), we draw on two of India’s “thick” rounds of NSS survey data collection for 1987/88 (43rd round) and 1999/2000 (55th round). These surveys provide us with detailed expenditure data on approximately 120,000 households each round residing in more than 400 Indian districts. Districts are further divided into urban and rural areas. Households are asked about their expenditures on 310 goods and services in each survey round. Examples include rice, firewood, washing soap and diesel. Deaton (2003a) and Deaton and Tarozzi (2000), carefully analyze these NSS expenditure survey rounds to identify 136 food and fuel products (from the 310) for which quantities are recorded and prices (obtained from expenditures divided by quantities) are robust to concerns about unobserved product quality or variety.¹¹ These 136 goods cover on average roughly 60 percent of total household consumption in 1999/2000. As we discuss in Sections 5 and 6, most of our welfare analysis will be based on these products (aggregated to 35 product groups such as wheat, leafy greens, and firewood), as they allow us to directly test the preference restrictions and identifying assumptions, and to compute correction terms if necessary.

In addition to total monthly expenditure for the household and the expenditures across products, we observe survey weights to make the data nationally representative, detailed household demographics, household district of residence, whether the household lives in a rural or urban area, and household religion and caste identifiers.¹² In the empirical analysis below, we define a “market” as a combination of the district and rural/urban identifiers, and focus attention on rural India to match the primary analysis of Topalova (2010) that we revisit in the applications.

As the NSS surveys are repeated cross-sections not panels, we cannot track specific households across rounds. We can however track percentiles of the local income distribution across rounds as we discuss in the subsequent sections. Although we do not observe income in the NSS data but only total household outlays, given limited saving in India these will be very similar. For readability, in what follows we occasionally interchange the word outlays with income.

Finally, it is important to note that in the 55th round, the surveys included an additional recall period for several consumption categories. Most consequentially, a 7-day recall period was added for all food products in addition to the standard 30-day recall period asked across rounds. While we only use the responses for the consistently-recorded 30-day recall period, Deaton (2003a,2003b) and others show that households inflated their 30-day reports to be consistent with their 7-day reports, which increased food expenditure relative to non-food. In standard approaches to measuring changes in Indian real incomes, this “recall bias” raises reported household total nominal outlays (used in the numerator for evaluating changes in real incomes)

¹¹In particular, Deaton (2003a) and Deaton and Tarozzi (2000) discard product categories that are likely to contain multiple varieties or quality levels based on both the name of the category (e.g. “other milk products”) or for which they observe bi-modal distributions of prices within the category (e.g. “liquid petroleum”).

¹²The 9 religious categories are Buddhism, Christianity, Hinduism, Islam, Jainism, Sikhism, Zoroastrianism and Other. The 4 castes are Upper Caste, Scheduled Caste, Scheduled Tribe, and in round 55 the category Other Disadvantaged Caste.

even when consistently using the 30-day recall data. Fortunately, our “New Engel” approach is relatively robust to these concerns. As we further discuss in Section 5, since our approach focuses on relative expenditures *within* groups such as food, the recall bias must be systematically related to the income elasticities of products within such groups to bias our welfare estimates. Consistent with this claim, we find qualitatively very similar findings when exploring welfare changes between 1987/88 and 1994/95 when the surveys were unchanged.

Harmonized Cross-Country Expenditure Surveys

For the cross-country empirical exercise, we plan to make use of the World Bank’s Global Consumption Database (GCD) which contains harmonized expenditures data for 92 countries.¹³ The data is drawn from nationally representative household expenditure surveys conducted mostly between the years 2007 and 2010. The number of households per country per year varies from a low of 2,000 observations to a maximum of 100,000 observations. Importantly for us, consumption categories in each country are mapped into a common classification of 110 goods and services taken from the International Comparison Program (ICP). For example, product categories include rice, bread, pasta products, poultry, pork and sugar. In cases in which multiple ICP goods corresponded to a single national survey item, a split was made using national accounts data.¹⁴ To harmonize consumption data across countries, nominal expenditures in local currency are annualized, home-produced goods are valued at local (farm-gate) prices and budget share outliers are winsorized separately by rural/urban and by decile of income separately for each country. Finally, to have the database refer to 2010 for all countries, if the national survey was fielded for example in 2009, household consumption values were multiplied by one plus the growth rate of national consumption per capita.

3.2 Motivating Facts

In this subsection, we use the Indian NSS data described above to establish two motivating facts that guide our theoretical framework in Section 4. We start by plotting Engel curves following Working’s (1943) now standard formulation (e.g. Banks, Blundell, and Lewbell, 1997; Lewbell, 2008) where budget shares are on the y-axis and log total household outlay per capita is on the x-axis. Figure 1 shows non parametric Engel curves for salt consumption, a good that is consumed widely across India and across all income groups, using kernel-weighted local polynomials.¹⁵ As salt is an inferior good, Engel curves are downward sloping with richer households spending a smaller share of their budget on the good. Panel A shows the Engel curves for the largest rural market, Midnapur, estimated separately for each survey round. Panel B shows the Engel curves in the latest survey round 1999/2000 for the largest market in the North, East, South and West of India. Two motivating facts are apparent from looking at these figures.

Motivating Fact 1: Engel Curves Are Non-Linear

From the plots for salt in Figure 1, it is apparent that textbook Engel curves (budget shares on the y-axis and log nominal outlays on the x-axis) are not in general linear. To test this observation

¹³<http://datatopics.worldbank.org/consumption/>

¹⁴For example, the ICP classifies gas and electricity separately, whereas a national survey may ask for expenditures on these categories jointly. In this case, expenditures are split into the two ICP categories using shares from that country’s national accounts.

¹⁵We use the Epanechnikov kernel with the “rule-of-thumb” bandwidth estimator.

more formally, we estimate regressions of the following form separately for each of the 38 goods and service categories included in total household consumption expenditure:

$$\frac{x_{hit}}{y_{ht}} = \theta_{imt} + \beta_{ki}F(\log(y_{hmt})) + \epsilon_{himit} \quad (3)$$

where y_{ht} is total nominal outlays per capita for household h residing in market m during survey round t .¹⁶ $\frac{x_{hit}}{y_{ht}}$ is the household's budget share of good i , θ_{imt} are good-by-market-by-period fixed effects and $F(\cdot)$ is a vector of polynomial terms with the order indexed by k . The error term ϵ_{himit} is clustered at the market-level. Panel A in Table 1 presents the fraction of goods and services for which the data formally reject the null hypothesis that the second-order and above polynomial terms are jointly equal to zero (i.e. that Engel curves are linear). We reject linearity at the 5 percent level of significance for 90 percent of goods and services. As noted previously by e.g. Banks, Blundell and Lewbell (1997), Engel curves appear to be non-linear in the data.

Motivating Fact 2: Shifts in Engel Curves Over Time and Across Space Are Not Parallel

Both panels in Figure 1 illustrate a second fact. As is apparent from the plots for a single market across different time periods (Panel A), and for markets in different parts of the country in the same time period (Panel B), it is apparent that Engel curves shift over time and across space, and that those shifts are not parallel horizontally—i.e. the horizontal shift in Engel curves is not uniform across the income distribution.

To move beyond these illustrative examples, we provide a more formal test of whether shifts in Engel curves across different markets (time or space) are parallel. We focus on rounds 43 and 55 and restrict attention to larger markets, those with at least 100 households in both rounds. For each of these combinations, we flip the axes and run the following regression:

$$\log(y_{hmt}) = \theta_{im} + \delta_{im}Post_t + \beta_{kim}F\left(\frac{x_{hit}}{y_{ht}}\right) + \gamma_{kim}Post_t \times F\left(\frac{x_{hit}}{y_{ht}}\right) + \epsilon_{himit}$$

where θ_{im} are market-by-good fixed effects, $Post_t$ is an indicator for the more recent round 55, and $F\left(\frac{x_{hit}}{y_{ht}}\right)$ is a vector of polynomial terms (up to the $k = 4$ th order) of the budget share of good or service i .¹⁷ To test for parallel horizontal shifts of i 's Engel curve in market m across the two rounds, we test the hypothesis that the four γ_{kim} interaction terms between the polynomials of the budget share and $Post_t$ are jointly equal to zero (i.e. that the only shift in the Engel curve is parallel and captured by δ_{im}). As reported in the second column in Table 1, we formally reject the null of a uniform shift in the Engel curve across the distribution of nominal outlays for 69 percent of the market-by-good cells (at a 95 percent confidence level).

Summary

This section documents two stylized facts using Indian expenditure microdata that motivate the new approach we describe in the next section. First, Engel curves tend to be non-linear in the data. And second, Engel curves shift across markets and over time, and those shifts are frequently non-uniform across the household income distribution.

¹⁶Results are not sensitive to including a full vector of household characteristic controls on the right hand side of equation (3).

¹⁷Again, results are not sensitive to the inclusion of household characteristic controls.

4 Theory

In this section we develop a new approach to estimating changes in household welfare and price indices that i) addresses the drawbacks of the existing Engel approach, and ii) is consistent with the two motivating facts we have documented in the data. In particular, our approach uses observed horizontal shifts in Engel curves over time or across space to recover changes in unobserved household price indices and welfare at each point of the income distribution.

We proceed in two steps. First, we introduce in a very general setting (for any rational utility function) the idea that horizontal shifts in Engel curves are related to changes in household price indices and hence welfare. While appealing, we prove that such an approach uncovers theory-consistent price indices only with the restriction that all relative prices remain unchanged (Lemma 1), but not in general with arbitrary price realizations (Lemma 2)—a restriction that precludes income-group specific changes in price indices and necessarily violates Motivating Fact 2 above.

In the second step, we then show how we can relax these assumptions on realizations of relative prices in a way that allows for income-specific price index changes. To make progress, we focus on a broad class of quasi-separable preferences following Gorman (1970; 1976) and on “relative Engel curves” that describe how relative expenditure shares within any given subset of goods or services G (i.e. spending on $i \in G$ as a share of total spending on all $i \in G$) vary with log total household outlays per capita. We prove that if, and only if, preferences are quasi-separable, horizontal shifts in relative Engel curves within G recover exact price indices when relative prices are fixed within G .¹⁸ By allowing for arbitrary price realizations outside of G , price indices can vary arbitrarily with income and can be estimated at each point of the income distribution, accommodating the rich patterns in the data that are consistent with Motivating Fact 2. Moreover, quasi-separable preferences can be of any rank so that our approach is flexible enough to allow for highly non-linear Engel curves consistent with Motivating Fact 1.

4.1 Using Shifts in Engel Curves to Infer Changes in Price Indices and Welfare

We first introduce the idea that shifts in Engel curves over time and space can potentially uncover unobserved changes in price indices and welfare. Consider comparing an Engel curve, for example food budget shares plotted against log total nominal outlays, at two different points in time or space. The horizontal distance between curves at any point in the income distribution reveals the change in log nominal outlays which holds the food share constant. The close link between this distance and price indices is obvious in the case where there are no changes in relative prices. Then, as long as demand is homogeneous of degree zero in total outlays and prices, a uniform price increase is equivalent to an equally sized fall in outlays. Hence, between points in time or space the price index change expressed in units of log income is exactly equal to the size of the horizontal shift. More generally, this will not be the case when relative prices are changing. This subsection makes this statement precise, and the next two subsections explore the restrictions to preferences required for shifts in Engel curves to reveal changes in price indices.

For ease of exposition, we focus the discussion below on inferring price index changes for

¹⁸In Section 5, we relax the assumption that relative prices with G are constant by replacing it with an orthogonality condition and, if necessary, a bias correction term.

households at a given percentile of the income distribution within a market location over time.¹⁹ However, the period superscripts ($t0$ and $t1$ below) could be analogously replaced with location identifiers (e.g. location j and k) in order to infer price index changes over locations.

As part of this exposition, we introduce notation and a number of key definitions. In what follows, the subscript i indexes goods and services in household consumption (for readability we will typically refer to them simply as goods), h indexes households, and superscript $t0$ and $t1$ indicate time periods 0 and 1 respectively. We denote functions of Engel curves with budget shares on the y-axis and log outlays on the x-axis by $\left(\frac{x_{hi}}{y_h} = E_i(p, y_h)\right)$, where p is the full vector of consumption prices, y_h is household nominal outlays per capita and x_{hi} is household expenditure on good or service i .

We define $P^{t1}(p^{t0}, p^{t1}, y_h^{t1})$ as the exact price index change from period 0 to period 1 prices, at period 1's level of nominal income y_h^{t1} , such that $U_h^{t1} = V(p^{t1}, y_h^{t1}) = V(p^{t0}, \frac{y_h^{t1}}{P^{t1}(y_h^{t1})})$ where V is the indirect utility function—for the sake of exposition, we also use the more compact notations $P^{t1}(y_h^{t1})$ and P^{t1} to refer to this price index change.²⁰ In other words, the price index P^{t1} converts the nominal incomes observed in the period 1 to hypothetical levels of incomes in period 0, holding period 1 utility U_h^{t1} constant. Symmetrically, we define $P^{t0}(p^{t0}, p^{t1}, y_h^{t0})$ as the exact price index change from period 1 to period 0 prices, at period 0's level of nominal income y_h^{t0} , such that $U_h^{t0} = V(p^{t0}, y_h^{t0}) = V(p^{t1}, \frac{y_h^{t0}}{P^{t0}(y_h^{t0})})$. In turn, the price index $P^{t0}(y_h^{t0})$ converts the nominal incomes observed in the period 0 to hypothetical levels of incomes in period 1, holding period 0 utility U_h^{t0} constant.²¹

These two price indices are intimately related to equivalent and compensating variation. $EV_h = e(p^{t0}, u_h^{t1}) - e(p^{t0}, u_h^{t0}) = \frac{y_h^{t1}}{P^{t1}(y_h^{t1})} - y_h^{t0}$ is the amount of money that would bring a household in period 0 to their period 1 utility and $CV_h = e(p^{t1}, u_h^{t1}) - e(p^{t1}, u_h^{t0}) = y_h^{t1} - \frac{y_h^{t0}}{P^{t0}(y_h^{t0})}$ is the amount of money that would need to be taken away from a period 1 household to bring them to their period 0 utility.

Lemma 1 *Assume that relative prices remain unchanged, i.e. $p_i^{t1} = \lambda p_i^{t0}$ for all i and some $\lambda > 0$.*

- i) The price index change for a given income level in period 1, $\log P^{t1}(y_h^{t1}) = \log \lambda$, or a given income level in period 0, $\log P^{t0}(y_h^{t0}) = -\log \lambda$, is equal to the horizontal shift in the Engel curve of any good i at that income level, such that*

$$E_i(p^{t1}, y_h^{t1}) = E_i(p^{t0}, \frac{y_h^{t1}}{P^{t1}(y_h^{t1})}) \quad \text{and} \quad E_i(p^{t0}, y_h^{t0}) = E_i(p^{t1}, \frac{y_h^{t0}}{P^{t0}(y_h^{t0})}).$$

- *Compensating variation and equivalent variation for a given income level are equal to the horizontal distance along Engel curves at $t0$ and $t1$ respectively, between the new and old expenditure share, such that:*

$$\frac{x_{hi}^{t0}}{y_h^{t0}} = E_i(p^{t1}, y_h^{t1} + CV_h) \quad \text{and} \quad \frac{x_{hi}^{t1}}{y_h^{t1}} = E_i(p^{t0}, y_h^{t0} + EV_h).$$

¹⁹If household panel data is available, this approach can be implemented at the level of individual households over time.

²⁰We note that, unlike typical expositions where exact price indices are functions of prices and utility, here we define our price index as a function of prices and nominal income.

²¹Note that the two price indexes are closely related: $y_h^{t1} = y_h^{t0} / P^{t0}(y_h^{t0})$ implies $y_h^{t0} = y_h^{t1} / P^{t1}(y_h^{t1})$.

Figure 2 plots a textbook Engel curve with the budget share on the y-axis and nominal income on the x-axis in order to graphically illustrate Lemma 1. Take as an example a household with an initial nominal income of y_h^{t0} (the bottom-left dot in the figure). The price index change P^{t0} for this household is equal to the horizontal distance (in $\log y_h$ space) between their initial budget share on the period 0 Engel curve and that same budget share on the period 1 Engel curve. As described at the start of this section, the intuition is that since relative prices are not changing, households with the same budget shares must be equally well off and so the nominal income difference between the two points must be equal to the change in the price index. The compensating variation for this household is then revealed by the additional distance that must be traveled in $\log y_h$ space to go from the crossing point on the period 1 Engel curve to the actual budget share of that household in period 1 (the upper-right dot). The same movements in reverse reveal P^{t1} and equivalent variation.

The proof is simple and relies on homogeneity of degree zero of the Marshallian demand and the indirect utility functions: i.e. the lack of money illusion. This ensures that $E_i(p^{t1}, y) = E_i(\lambda p^{t0}, y) = E_i(p^{t0}, y/\lambda)$ when all prices change with a common scalar λ . In our simple example, this common scalar coincides with the price index change. In terms of EV, we can see that $y_h^{t0} + EV_h = y_h^{t1}/\lambda$. Hence we can use Engel curves to infer EV_h : $E_1(p^{t1}, y_h^{t1}) = E_i(p^{t0}, y_h^{t1}/\lambda) = E_i(p^{t0}, y_h^{t0} + EV_h)$. In terms of CV_h , we can see that $y_h^{t1} + CV_h = y_h^{t0}\lambda$ and CV_h can be inferred from $E_i(p^{t0}, y_h^{t0}) = E_i(p^{t1}, y_h^{t0}\lambda) = E_i(p^{t1}, y_h^{t1} + CV_h)$.

Lemma 1 shows that shifts in Engel curves reveal changes in price indices when price changes are uniform across goods. However, if relative prices are unchanged, shifts in Engel curves must be parallel (and price index changes must be identical for households across the income distribution). Motivating Fact 2 above clearly shows this is not the case in the Indian context, and is unlikely to be true in other contexts.

To allow for Engel curves consistent with Motivating Fact 2 (i.e. changing slopes over time implying income-group specific price indices), we must allow relative prices to change. However—as we show in Lemma 2 below—if relative prices are allowed to change arbitrarily, Engel curves will not in general reveal changes in price indices and hence welfare.

Lemma 2 *Horizontal shifts in any good i 's Engel curve do not recover changes in the price index under arbitrary changes in the price of good i relative to other goods, or groups of goods.*

Shifts in the Engel curve for good i reflect both changes in utility and changes in relative prices if the price of i is allowed to change relative to other goods. The proof of Lemma 2 in Appendix B shows that only when demands are Cobb-Douglas will expenditure shares not depend on relative prices, even if in principle we allow such shares to depend on utility. Thus, only when Engel curves are horizontal (the Cobb-Douglas case) will relative price changes not confound shifts in Engel curves, and in that case the flatness of the Engel curve clearly precludes us from identifying price index changes from horizontal shifts in Engel curves. Hence, to be able to relax the assumption of constant relative prices (and thus uniform price index changes across the income distribution), we must impose additional structure on the very general preferences considered in the two Lemmas, and depart from textbook Engel curves that relate expenditure shares to total outlays. The following subsection shows how to make progress.

4.2 Relative Engel Curves and Quasi-Separable Preferences

Lemmas 1 and 2 show that while an appealing concept, shifts in textbook Engel curves will not in general recover changes in price indices when relative prices are changing. We now present an approach using what we term “relative Engel curves”, defined as expenditure shares on a subset of goods plotted against total outlays per capita. We show that as long as preferences are quasi-separable in the parlance of Gorman (1970, 1976), shifts in relative Engel curves do reveal changes in income-group specific price indices under a rich set of relative price changes consistent with Motivating Fact 2. In particular, while we restrict relative price changes on a subset of goods—an assumption we replace with an orthogonality condition or bias-correction term in Section 5—this approach accommodates price indices changing differentially across the income distribution, a possibility that needed to be assumed away to recover theory-consistent price indices in the previous two approaches. Furthermore, the class of household preferences we propose below is flexible enough to allow for the highly non-linear relative Engel curves shown in Motivating Fact 1.

In the following proposition, we introduce “relative Engel curves” which we define as follows:

Definition *Relative Engel curves, denoted by the function $E_{iG}(p, y_h) = \frac{x_{hi}}{x_{hG}}$, describe how relative expenditure shares within a subset of goods G (i.e. spending on $i \in G$ as a share of total spending on all $i \in G$) vary with log total household outlays per capita.*

We will also refer to quasi-separable preferences which we define as follows:

Definition *Preferences are quasi-separable in group G of goods if a household’s expenditure function can be written as:*

$$e(p, U_h) = \tilde{e}(\tilde{P}_G(p_G, U_h), p_{NG}, U_h)$$

where $\tilde{P}_G(p_G, U_h)$ is a scalar function of utility U_h and the vector of the prices p_G of goods $i \in G$, and is homogeneous of degree 1 in prices p_G .

With these two definitions in hand, we turn to the key proposition behind our approach.

Proposition 1 *Suppose that G is a subset of goods. The following three properties hold for any realization of prices leaving relative prices within G unchanged (i.e. $p_i^{t1} = \lambda_G p_i^{t0}$ for all $i \in G$ and for some $\lambda_G > 0$) if, and only if, preferences are quasi-separable in the subset G of goods:*

- i) *The price index change for a given income level in period 0, $\log P^{t0}(y^{t0})$, or in period 1, $\log P^{t1}(y^{t1})$, is equal to the horizontal shift in the relative Engel curve of any good $i \in G$ at that income level, such that*

$$E_{iG}(p^{t1}, y_h^{t1}) = E_{iG}(p^{t0}, \frac{y_h^{t1}}{P^{t1}(y_h^{t1})}) \quad \text{and} \quad E_{iG}(p^{t0}, y_h^{t0}) = E_{iG}(p^{t1}, \frac{y_h^{t0}}{P^{t0}(y_h^{t0})}).$$

- ii) *The log compensating variation for household h , $\log(1 + \frac{CV_h}{y_h^{t1}})$, is equal to the horizontal distance between the new and old expenditure share along the period 1 relative Engel curve, such that*

$$E_{iG}(p^{t1}, y_h^{t1} + CV_h) = \frac{x_{hi}^{t0}}{x_{hG}^{t0}}.$$

- iii) *The log equivalent variation for household h , $\log(1 + \frac{EV_h}{y_h^{t0}})$, is equal to the horizontal distance between the new and old expenditure share along the period 0 relative Engel curve, such that*

$$E_{iG}(p^{t0}, y_h^{t0} + EV_h) = \frac{x_h^{t1}}{x_h^{t0}}.$$

Figure 3 illustrates Proposition 1 graphically, and Appendix B provides the proof.²² The figure is similar to that for Lemma 1, except now the x-axis is the relative budget share within G and the two curves are no longer parallel (generated by relative price changes outside of group G). And a similar logic applies, with the horizontal distance from the initial budget share of household h on the period 0 relative Engel curve to the same budget share on the period 1 relative Engel curve revealing the change in the price index P^{t0} . As before, the additional horizontal distance traveled from that crossing point to the new budget share of that household along the period 1 Engel curve reveals the compensating variation CV_h . However, unlike in Lemma 1, since the Engel curves are no longer parallel the change in the price index P^{t0} and the compensating variation CV_h is different depending on household h 's position in the income distribution. If we go in the other direction for the same household h , we obtain the change in price index P^{t1} and equivalent variation EV_h : i.e. the horizontal distance between h 's budget share in period 1 and that same budget share in period 0 uncovers the price index, and the additional horizontal distance along the period 0 Engel curve required to reach h 's actual budget share in period 0 reveals EV_h . The fact that the Engel curves are no longer parallel also means that the price index estimates P^{t1} and P^{t0} (and welfare estimates EV_h and CV_h) are not equal as they were in Lemma 1.

It is perhaps easiest to understand Proposition 1 by highlighting several steps of its proof. Under quasi-separability, relative expenditure in good i within group G can be written as a compensated function $H_{iG}(p_G, U)$ of utility and relative prices within group G only (see Lemma 3 below). This ensures that, holding relative prices within G constant, expenditure shares within group G solely depend on household utility.

The second step is the link from this compensated demand function to observable relative Engel curves which are functions of total household outlays (nominal income) rather than utility. This is done by substituting the indirect utility function $V(p, y)$ that links outlays and utility given the full vector of prices into the Hicksian demand function for a particular time period t :

$$E_{iG}(p^t, y_h^t) = H_{iG}(p_G^t, U_h^t) = H_{iG}(p_G^t, V(p^t, y_h^t)). \quad (4)$$

In the third step, we show how horizontal shifts in Engel curves in $\log y_h$ space identify changes in price indices. For example, to obtain the price index change at income level y_h^{t1} , $P^{t1}(p^{t0}, p^{t1}, y_h^{t1})$, start with the relative budget share on household h on the relative Engel curve in period 1, $E_{iG}(p^{t1}, y_h^{t1})$. Moving to its Hicksian representation, we obtain:

$$\begin{aligned} E_{iG}(p^{t1}, y_h^{t1}) &= H_{iG}(p_G^{t1}, V(p^{t1}, y_h^{t1})) \\ &= H_{iG}(p_G^{t0}, V(p^{t1}, y_h^{t1})) \\ &= H_{iG}(p_G^{t0}, V(p^{t0}, y_h^{t1} / P^{t1}(p^{t0}, p^{t1}, y_h^{t1}))) \\ &= E_{iG}(p^{t0}, y_h^{t1} / P^{t1}(p^{t0}, p^{t1}, y_h^{t1})). \end{aligned}$$

Equality between the first two lines above is an implication of the homogeneous price change

²²Proposition 1 can also apply to goods that are themselves aggregates of different varieties or quality levels. For example, in case each of the $i \in G$ is itself a group of many barcodes, quasi-separability can accommodate nests of subgroups within group G . In this case, prices across the $i \in G$ are themselves price indices across individual varieties within each of the subgroups.

$p_i^{t1} = \lambda_G p_i^{t0}$ within group G (note that H_{iG} is homogeneous of degree zero in prices p_G^t within group G). Equality between the second and third lines follows from the implicit definition of the price index $P^{t1}(p^{t0}, p^{t1}, y_h^{t1})$ above. The final line simply moves back to relative Engel curves. The distance in nominal income space between household h 's budget share in period 1 and the nominal income of a household in period 0 spending the same budget share reveals the price index change $P^{t1}(p^{t0}, p^{t1}, y_h^{t1})$. Swapping $t0$ and $t1$, we obtain the equivalent result for $P^{t0}(p^{t0}, p^{t1}, y_h^{t0})$.

Proposition 1 is a strong result. It states that, in theory, we can infer an exact measure of unobserved changes in household welfare at any given point of the initial or future distribution of household incomes by observing i) relative expenditure shares across some subset G of goods, and ii) total household outlays. It also states that this is true if, and only if, household preferences fall in the class of quasi-separable demand, and if relative prices are unchanged within the subset of goods G .

By focusing on relative Engel curves, we are able to make progress where Lemma 2 showed we could not do so using traditional Engel curves. Furthermore, Proposition 1 shows that only under quasi-separable preferences will shifts in relative Engel curves reveal changes in price indices and welfare. In particular, by not placing restrictions on relative prices outside of set G , income-group specific price indices can diverge and rich patterns consistent with Motivating Fact 2 (non-parallel shifts in Engel curves) can be easily accommodated.

Naturally, the question that then arises is how restrictive are the conditions on preferences and prices in Proposition 1? The first condition, quasi-separability in the sense of Gorman, is less restrictive than the usual use of the term separability. It implies that relative budget shares within the subset G are a function of relative prices within G and household welfare. That is, instead of imposing separability across goods in the utility function, we impose it in the expenditure function. The difference is that instead of shutting down all relative price effects between goods in G and the rest of consumption, as under separability, quasi-separability allows for prices outside of G to affect both the average and the relative consumption within G through household utility U_h . Furthermore, under quasi-separability we can remain fully flexible on both own and cross-price effects within both G and non- G as we discuss further below.

The second condition in Proposition 1 fixes relative price changes within G , which alongside quasi separability ensures that price changes do not shift relative Engel curves independently of changes in household welfare. As touched upon above, this condition on relative prices does not restrict shifts in relative Engel curves to be parallel. Thus, unlike the two special cases highlighted in Section 2 and Lemma 1, our approach allows for non uniform price index changes across the income distribution. In the next section, when we move from theory to estimation, we further show that this restriction on within- G prices can be replaced by an orthogonality condition between relative price changes within G and the slopes of relative Engel curves across $i \in G$, and if this condition is violated we provide a bias-correction term.

To further explore what structure household utility has to possess to satisfy quasi-separability, and discuss which preferences used in the literature fall within the quasi-separable class, we turn to a final lemma:

Lemma 3 *Preferences are quasi-separable if and only if:*

- i) *Relative compensated demand for any good or service i within group G only depends on utility U_h*

and the relative prices within G :

$$\frac{x_{hi}}{x_{hG}} = \frac{p_i h_i(p, U_h)}{\sum_{j \in G} p_j h_j(p, U_h)} = H_{iG}(p_G, U_h)$$

for some function $H_{iG}(p_G, U_h)$ of utility and the vector of prices p_G of goods $i \in G$.

ii) Utility is implicitly defined by:

$$K(F_G(q_G, U_h), q_{NG}, U_h) = 1$$

where q_G and q_{NG} denote consumption of goods in G and outside G respectively, for some functions $K(F_G(q_G, U_h), q_{NG}, U_h)$ and $F_G(q_G, U_h)$, where $F_G(q_G, U_h)$ is homogeneous of degree 1 in q_G .

This lemma draws on existing results. The equivalence between quasi-separability and condition i) is shown in Blackorby, Primont and Russel (1978),²³ and the equivalence with ii) has been proved in McFadden (1978) and Deaton and Muellbauer (1980).

The equivalence in condition i) of Lemma 3 sheds light on why quasi-separability is necessary for the analysis based on relative Engel curves (as was made clear from the sketch of the proof of Proposition 1 above). Condition i) is required to ensure that relative Engel curves only reflect changes in utility when relative prices outside G remain unobserved and vary arbitrarily, as long as relative prices within G remain constant. This property is used again in the following section to discuss orthogonality conditions when relative prices vary within G .

Both conditions i) and ii) of Lemma 3 can be used to characterize the type of preferences that satisfy quasi-separability. With condition ii), one can see that the preferences used in Comin et al. (2015) and Matsuyama (2015),²⁴ in which utility is implicitly defined by:

$$\sum_i^N \left(\frac{q_i}{g_i(U)} \right)^{\frac{\sigma-1}{\sigma}} = 1 \quad (5)$$

are quasi-separable in any subset of goods. Using condition i), we can also see that Translog (in expenditure functions) and PIGLOG demand systems satisfy quasi-separability in a group of goods G if there are no cross-price effects between goods within G and goods outside of G . Beyond these examples, condition ii) indicates that we can construct highly flexible demand systems that allow for flexible substitution effects within group G (captured by function F_G) and between goods within G and outside G (function K). In particular, Allen-Uzawa price elasticities do not have to be constant across goods within G as in Comin et al (2016).

The properties of quasi-separable preferences mimic those of direct separability in the dual—condition i) is similar to Sono-Leontief characterization of direct separability (i.e. separability in the direct utility function). However, directly-separable preferences are in general not quasi-separable, and vice versa. Finally, note also that quasi-separable demand systems can have any rank in the sense of Lewbel (1991) and so accommodate highly non linear Engel curves. In fact, in equation 5, the rank can be equal to $N - 1$ if none of the $g_i(U)$ are collinear with each other, which is theory consistent as long as each is monotonic in U .

²³In Appendix we also provide a new and simpler proof of this result.

²⁴As well as in Fally (2018) where the price elasticity $\sigma(U)$ varies with utility.

5 From Theory to Estimation

In this section, we build on the theoretical results above to derive an empirical methodology for estimating exact price indices and changes in household welfare using expenditure survey microdata. We then turn to identification and derive four corollaries to our theoretical results above that define the conditions for unique and unbiased identification in the data, and discuss the testable implications for a set of additional robustness checks.

5.1 Implementation

Suppose that we want to estimate the welfare change between two periods for a household with income y_h^{t0} in the reference period $t0$ and y_h^{t1} in the new period $t1$. The graphical exposition of Proposition 1 in Figure 3 provides a simple method of estimating changes in exact household price indices, $\log P^{t0}(p^{t0}, p^{t1}, y_h^{t0})$ and $\log P^{t1}(p^{t0}, p^{t1}, y_h^{t1})$, and household welfare, EV_h and CV_h . Essentially, we use non-parametric methods to estimate very flexible relative Engel curves in both periods, and can then recover changes in income-group specific price indices as well as household welfare from the size of the shift at different points of the income distribution. Repeating this procedure for multiple goods generates multiple price index and welfare estimates that can be combined to increase precision (and allow for unobserved good-specific taste and price shocks as we discuss in the following subsection). We refer to this approach as good-by-good estimation and describe the approach in more detail below.²⁵

Good-by-Good Estimation We first use the expenditure survey microdata to separately estimate non-parametric relative Engel curves for every good $i \in G$ and for each period $t0$ and $t1$. To do so, we use kernel-weighted local polynomial regressions of relative expenditure shares, x_{ihm}^t / x_{Ghm}^t , on total outlays.²⁶ This provides estimates of $\widehat{x_{ihm}^t / x_{Ghm}^t}$ for households indexed by h at every point in the income distribution. In practice, since we do not have true panel data in our applications below, we use h to index the percentile of the distribution and explore price index changes and welfare changes at different percentiles of the income distribution. Accordingly, we estimate Engel curves at 101 points corresponding to percentiles 0 to 100 of the local income distribution.²⁷

With these relative Engel curves in hand, consider estimating the exact price index change for a household at a particular percentile h in period 0, $\log P^{t0}(p^{t0}, p^{t1}, y_h^{t0})$ (i.e. the exact price index change from period 1 to period 0 prices, fixing the household outlays at period 0's level of nominal outlays). The relative Engel curve for period 0 provides a point estimate of relative expenditures for households at this percentile of the initial income distribution, $x_{ihm}^{t0} / x_{Ghm}^{t0}$.

The next step is to estimate the income level $\widehat{E}_{iG}^{-1}(p^{t1}, x_{ihm}^{t0} / x_{Ghm}^{t0})$ associated with this rel-

²⁵We are currently exploring alternative procedures that impose cross-good and cross-period restrictions to improve efficiency and handle non-monotonic Engel curves.

²⁶As noted in the Motivating Facts Section 3, the estimated Engel curves are mostly insensitive to the inclusion of a full vector of household characteristic controls, greatly facilitating the non-parametric estimation. In future versions of this draft, we plan to include the full vector of household controls using the `npregress` and `margins` commands in Stata.

²⁷We first smooth the distribution of income using a local polynomial regression of nominal income on rank in the income distribution divided by the number of households observed in market m at time t to obtain predictions of income at every percentile. For the Engel curves, we use an Epanechnikov kernel with a bandwidth equal to one quarter of the range of the income distribution in a given market. We also present results across a range of alternative bandwidth choices.

ative expenditure share from the crossing point on the relative Engel curve in period 1, where $\widehat{E}_{iG}^{-1}(p^{t1}, \cdot)$ denotes the inverse of the relative Engel curve at period 1 prices. To do so we find the crossing point $\widehat{x_{ih'm}^{t1}/x_{Gh'm}^{t1}}$ and take the corresponding income at this point $\widehat{\log y_{h'}^{t1}}$.²⁸

With estimates of the period $t1$ income level $\widehat{\log y_{h'}^{t1}} = \widehat{\log E_{iG}^{-1}(p^{t1}, x_{ih'm}^{t0}/x_{Ghm}^{t0})}$ in hand, we can calculate the change in the the price index and welfare for a household h that lies at a particular point in the initial distribution. The income-group specific price index change, $\log P^{t0}(p^{t0}, p^{t1}, y_h^{t0})$, is equal to the difference between $\widehat{\log y_h^{t0}}$ (the initial level of household income for h) and the estimate of $\widehat{\log y_{h'}^{t1}}$ —this is the horizontal shift labeled $\log P^{t0}$ in Figure 3. The welfare change for household h , as measured by compensating variation, is recovered from the relationship $\log(1 + CV_h/y_h^{t1}) = \widehat{\log y_{h'}^{t1}} - \log y_h^{t1}$, where y_h^{t1} is the observed period $t1$ level of income for a household at percentile h of the distribution. This expression recovers welfare changes for a hypothetical household that stays at the same point of the income distribution in both periods from the movement along the period $t1$ relative Engel curve. If panel data are available, we could of course recover welfare changes for a specific household h using this methodology.

Finally, consider estimating the exact price index change for households at a particular percentile h in period 1 fixing the household outlays at period 1's level of nominal outlays y_h^{t1} : $\log P^{t1}(p^{t0}, p^{t1}, y_h^{t1})$ (i.e. the exact price index change from period 0 to period 1 prices). We follow the same procedure as above but going in the other direction. We start by finding the period 0 income corresponding to household h 's period 1 relative expenditure by reading off level of total outlays corresponding to where $\widehat{x_{ih'm}^{t1}/x_{Ghm}^{t1}}$ crosses the period 0 Engel curve. We then recover $\log P^{t1}(p^{t0}, p^{t1}, y_h^{t1})$ from the difference between $\widehat{\log y_{h'}^{t1}}$ and this estimate of total outlays $\widehat{\log y_{h'}^{t0}} = \widehat{E_{iG}^{-1}(p^{t0}, x_{ih'm}^{t1}/x_{Ghm}^{t1})}$ —the horizontal shift labeled $\log P^{t1}$ in Figure 3. The welfare change for household h , here measured by equivalent variation, is recovered from the relationship $\log(1 + EV_h/y_h^{t0}) = \widehat{\log y_{h'}^{t0}} - \log y_h^{t0}$, i.e. from the movement along the period $t0$ relative Engel curve. Each good $i \in G$ proves a separate estimate for $\log P^{t0}$, $\log P^{t1}$, CV_h and EV_h . We aggregate this information through a simple average, as well as implementing a selection correction procedure for missing good-level estimates that we describe below.

5.2 Identification

In this subsection, we derive four corollaries related to unique and unbiased identification when taking Proposition 1 to the data.

Invertibility of Relative Engel Curves

The first result derives necessary and sufficient conditions for being able to invert relative Engel curve functions. The following corollary formally defines necessary and sufficient conditions under which our methodology is identified when using expenditure survey microdata.

²⁸In principle you could estimate Engel curves at many thousands of points of the income distribution and find a close to exact match. In practice, we take the two closest percentiles from our 101 points and linearly interpolate nominal outlays between these adjacent percentiles to obtain $\widehat{\log y_{h'}^{t1}}$. As we discuss in Section 5.2, because of issues related to uniqueness we restrict attention to monotonic Engel curves. If there is no interior crossing point we record the price index estimate as missing, and similarly if the slope at $\widehat{x_{ih'm}^{t0}/x_{Ghm}^{t0}}$ and $\widehat{x_{ih'm}^{t1}/x_{Ghm}^{t1}}$ takes opposite signs. See section 5.2 for a description of a selection correction to avoid potential biases from missing estimates.

Corollary 1 *Under the same conditions as Proposition 1:*

- i) *The necessary condition to recover unique estimates of changes in exact price indices and household welfare is that different levels of household utility map into unique vectors of relative budget shares within the subset of goods and services G at any given set of prices.*
- ii) *A sufficient condition for i) to hold is that the relative Engel curve $E_{iG}(p, y_h)$ is monotonic for at least one good or service $i \in G$.*

The first condition is weaker than the second condition. However, the second condition is readily verifiable in the data, and turns out to be true empirically for all markets and time periods that we consider in our empirical specification.

For the good-by-good estimation approach we discuss above, neither i) nor ii) are sufficient to ensure invertibility of all $E_{iG}(p, y_h)$ for $i \in G$. Estimates of $\widehat{E}_{iG}^{-1}(p^{t1}, x_{ihm}^{t0}/x_{Ghm}^{t0})$ or $\widehat{E}_{iG}^{-1}(p^{t0}, x_{ihm}^{t1}/x_{Ghm}^{t1})$ will not be unique if a given i 's relative Engel curve is flat or non-monotonic so we restrict attention to monotonic Engel curves in the current analysis.²⁹ This restriction can be weakened substantially, and we plan to do so in future draft.

Orthogonality Conditions

In the theory section above, we impose the strong restriction that relative prices within the subset of quasi-separable goods G remain unchanged over time. With this restriction, shifts in the relative Engel curve of any single good $i \in G$ reveal the price index. We have also, thus far, abstracted from changes in tastes over time (or across space when comparing markets in the cross-section). In this section, we now formally characterize a set of orthogonality conditions—using estimates from many goods $i \in G$ —that allow us to relax these two strong assumptions in the empirical estimation. The orthogonality conditions naturally lead to a set of bias-correction terms should they be violated.

Under quasi-separable preferences, part i) of Lemma 3 indicates that we can write relative expenditures within G as a function of utility and relative prices within G , with $\frac{x_{ih}^t}{x_{Gh}^t} = H_{iG}(p_G^t, U_h^t)$ and H_{iG} homogeneous of degree zero in prices p_G^t within G . Hence, as long as relative prices within G remain constant, i.e. $p_i^{t1} = \lambda_G p_i^{t0}$, the income level in period 0 that yields the same expenditure share would also be associated with the same utility level, and we can retrieve the exact price index change by inverting the relative Engel curve from period 0 at that expenditure share: $\log\left(\frac{y_h^{t1}}{p^{t1}}\right) = \log E_{iG}^{-1}\left(p^{t0}, x_{ih}^{t1}/x_{Gh}^{t1}\right)$. When relative prices with G are no longer constant, this relationship no longer holds exactly for each good in particular, $H_{iG}(p_G^{t0}, U) \neq H_{iG}(p_G^{t1}, U)$, but we can combine estimates across several goods to obtain an unbiased estimate of the price index if instead we impose conditions on the distribution of price changes within G .

To make progress, we use a first-order approximation to account for the effect of price and taste shocks on expenditure shares (where we denote taste shocks by $\Delta \log \alpha_{ih}$, the change in ex-

²⁹In particular, as non-parametrically estimated Engel curves are often noisy at the extreme tails where there are few households across large ranges of outlays, we restrict attention to good-market combinations where Engel curves in both periods are monotonic between percentiles 5 and 95 and drop relative expenditure share estimates beyond those percentiles in cases where those portions are non-monotonic (replacing those values with a linear extrapolation from the monotonic portion of the curve).

penditure shares from period 0 to period 1 that are due to changes in preferences).³⁰ Relative to initial prices, we obtain the following first-order approximation of $\log \left(\frac{x_{ih}^{t1}}{x_{Gh}^{t1}} \right)$ for any unspecified function H_{iG} :

$$\log \left(\frac{x_{ih}^{t1}}{x_{Gh}^{t1}} \right) = \log H_{iG}(p_G^{t1}, U_h^{t1}) + \Delta \log \alpha_{ih} \approx \log H_{iG}(p_G^{t0}, U_h^{t1}) + \sum_{j \in G} \sigma_{ijh} \Delta \log p_j + \Delta \log \alpha_{ih} \quad (6)$$

where $\sigma_{ijh} = \frac{\partial \log H_{iG}}{\partial \log p_j}$ is the compensated price elasticity of relative consumption of i w.r.t price j and $\Delta \log p_j = \log p_j^{t1} - \log p_j^{t0}$ is the difference in the price of good j from the base period $t0$. The first term on the right-hand side is the movement in relative budget shares that would have occurred purely due to changes in household welfare. The final two terms are potentially confounding factors due to unobserved relative price movements within G and taste shocks within G .

When evaluated under the original Engel curve, the term $\log H_{iG}(p_G^{t0}, U_h^{t1})$ coincides with $\log E_{iG}(p^{t0}, \frac{y_h^{t1}}{P^{t1}})$. Using equation (4) and the definition of the price index, $U_h^{t1} = V(p_{t0}, y_h^{t1}/P^{t1})$, this can be described by the following equalities:

$$H_{iG}(p_G^{t0}, U_h^{t1}) = H_{iG}(p_G^{t0}, V(p^{t0}, y_h^{t1}/P^{t1})) = E_{iG} \left(p^{t0}, \frac{y_h^{t1}}{P^{t1}} \right)$$

Similarly, the term $H_{iG}(p_G^{t1}, \alpha^{t1}, U_h^{t0})$ is equal to $E_{iG}(p^{t1}, \frac{y_h^{t0}}{P^{t0}})$. Assuming that the observed Engel curves E_{iG} are invertible (Corollary 1 above), this implies that our empirical estimates of the horizontal shifts in relative Engel curves (i.e. counterfactual levels of nominal income that hold constant either U_h^{t0} or U_h^{t1}), can be written as a first-order Taylor expansion around the exact price index $\log \left(\frac{y_h^{t1}}{P^{t1}} \right)$:³¹

$$\begin{aligned} \log E_{iG}^{-1} \left(p^{t0}, \frac{x_{ih}^{t1}}{x_{Gh}^{t1}} \right) &= \log E_{iG}^{-1} \left(p^{t0}, \log H_{iG}(p_G^{t0}, \alpha^{t0}, U_h^{t1}) + \sum_{j \in G} \sigma_{ijh} (\Delta \log p_j - \overline{\Delta \log p_G}) + \Delta \log \alpha_{ih} \right) \\ &= \log E_{iG}^{-1} \left(p^{t0}, \log E_{iG} \left(p^{t0}, \frac{y_h^{t1}}{P^{t1}} \right) + \sum_{j \in G} \sigma_{ijh} (\Delta \log p_j - \overline{\Delta \log p_G}) + \Delta \log \alpha_{ih} \right) \\ &\approx \log \left(\frac{y_h^{t1}}{P^{t1}} \right) + (\beta_{ih}^{t0})^{-1} \sum_{j \in G} \sigma_{ijh} (\Delta \log p_j - \overline{\Delta \log p_G}) + (\beta_{ih}^{t0})^{-1} \Delta \log \alpha_{ih} \end{aligned} \quad (7)$$

and equivalently for $P^{t0} = P^{t0}(p^{t0}, p^{t1}, y_h^{t0})$:

$$\log E_{iG}^{-1} \left(p^{t1}, \frac{x_{ih}^{t1}}{x_{Gh}^{t1}} \right) \approx \log \left(\frac{y_h^{t0}}{P^{t0}} \right) + (\beta_{ih}^{t0})^{-1} \sum_{j \in G} \sigma_{ijh} (\Delta \log p_j - \overline{\Delta \log p_G}) - (\beta_{ih}^{t0})^{-1} \Delta \log \alpha_{ih} \quad (8)$$

where $\beta_{ihm}^{t0} = \frac{\partial \log E_{iG}}{\partial \log y_h}$ denotes the slope of the Engel curve (income elasticity) evaluated at income y_h^{t1}/P^{t1} at the initial set of prices p^{t0} , and $\beta_{ihm}^{t1} = \frac{\partial \log E_{iG}}{\partial \log y_h}$ evaluated at income y_h^{t0}/P^{t0} at the set of prices p^{t1} .³² The first term on the right-hand side of both (7) and (8) is the object that we are trying

³⁰Note that instead of a first-order adjustment for some arbitrary demand function within G , an alternative would be to fully specify H_{iG} and compute an exact correction under this specification.

³¹Note that $\sum_{j \in G} \sigma_{ijht} = 0$ due to homogeneity of degree zero, which implies that the price term $\sum_{j \in G} \sigma_{ijh} \Delta \log p_j$ can also be written as: $\sum_{j \in G} \sigma_{ijh} (\Delta \log p_j - \overline{\Delta \log p_G})$, where $\overline{\Delta \log p_G}$ is the average of price changes within G .

³²The β_{ih} are here defined in terms of elasticities (using log shares). The covariance terms that we derive below would be identical if we use shares instead of log shares as long as we weight the covariance term by expenditure

to estimate. The second and third terms are potential confounders as if they are non zero—i.e. relative prices or tastes within G change across periods—the shift in relative Engel curves across periods is not only driven by changes in household welfare over time, but also by unobserved relative price changes or taste shocks.

As described above, we do not use the implied price index or welfare change from a single good i but from an average of many goods within G . While the bias above may be large for a specific good, when averaging over many goods the problem may be less severe. Solving for the counterfactual incomes that hold utility constant across price environments in (7) and (8) above and then averaging across $i \in G$, we obtain:

$$\log \left(\frac{y_h^{t1}}{P^{t1}} \right) \approx \frac{1}{G} \sum_{i \in G} \log \widehat{E}_{iG}^{-1} \left(p^{t0}, \frac{x_{ih}^{t1}}{x_{Gh}^{t1}} \right) - \frac{1}{G} \sum_{i,j \in G} (\beta_{ih}^{t0})^{-1} \sigma_{ijh} (\Delta \log p_j - \overline{\Delta \log p_t}) - \frac{1}{G} \sum_{i \in G} (\beta_{ih}^{t0})^{-1} \Delta \log \alpha_{ih} \quad (9)$$

and

$$\log \left(\frac{y_h^{t0}}{P^{t0}} \right) \approx \frac{1}{G} \sum_{i \in G} \log \widehat{E}_{iG}^{-1} \left(p^{t1}, \frac{x_{ih}^{t0}}{x_{Gh}^{t0}} \right) - \frac{1}{G} \sum_{i,j \in G} (\beta_{ih}^{t1})^{-1} \sigma_{ijh} (\Delta \log p_j - \overline{\Delta \log p_t}) + \frac{1}{G} \sum_{i \in G} (\beta_{ih}^{t1})^{-1} \Delta \log \alpha_{ih}$$

We are now in a position to define the following two orthogonality conditions that yield unbiased estimates of changes in household prices indices and welfare even in contexts where relative prices or tastes within G are changing:

Corollary 2 *Assuming quasi-separability of subset G :*

- i) *To identify $\log P^{t0}(p^{t0}, p^{t1}, y_h^{t0})$, unobserved relative price changes and taste shocks across i within subset G must be orthogonal to the local slope of the relative Engel curve in period $t1$.*
- ii) *To identify $\log P^{t1}(p^{t0}, p^{t1}, y_h^{t1})$, unobserved relative price changes and taste shocks across i within subset G must be orthogonal to the local slope of the relative Engel curve in period $t0$.*

These orthogonality conditions ensure that the second and third terms in equations (9) average out to zero in expectation and so our estimates of $\log \left(\frac{y_h^{t0}}{P^{t0}} \right)$ and $\log \left(\frac{y_h^{t1}}{P^{t1}} \right)$ are unbiased.³³

Corollary 2 is testable in the survey data for products groups where we have reliable price data. As discussed in Section 3, for this reason we focus on the 136 food and fuel goods Deaton (2003b) and Deaton and Tarozzi (2000) identify as having reliable price data. Thus, when restricting estimation to food and fuel product groups G with reliable information on price changes across the $i \in G$ —recall that our methodology does not require estimating relative Engel curves for all goods in order to recover the complete price index—we can empirically test whether observed price changes are systematically correlated with the slopes of the relative Engel curves within G in either period 0 or period 1 at a given income percentile.

shares.

³³Using the expressions in (9), it is straight-forward to see to what extent measurement error in household expenditures (e.g. due to additional recall periods) could give rise to biased estimates. Similar to the omitted relative price effects in (9), such measurement error would have to be systematically related to the slopes of the relative Engel curves within the group G across the products $i \in G$. For example, the addition of a second recall period for goods in G in period $t1$ would give rise to concerns if i) reported relative outlays within a given product group G are affected differentially (not across food and non-food as has been the primary concern in the existing debate), and ii) such differential changes in reported outlays vary systematically between high and low income elastic products within the group G .

If the orthogonality test in Corollary 2 is rejected for these groups G with reliable price data, then our price index estimates will be biased. However, the derivation above also provides a simple first-order correction term that corrects for the bias coming from the confounding price effects in (9). We simply need to calculate the relative-Engel-slope-price-change terms in equation (9) and add them to our price index estimates, which is straightforward to do since we focus on the groups G for which reliable price and expenditure data are available. In our applications below, we further assume a constant elasticity of substitution σ_G within group G (e.g. implicitly as in Comin et al. 2015), so that the bias correction terms for each market-by-decile for P^{t1} is:

$$\frac{1}{G} \sum_{i \in G} (\beta_{ih}^{t_0})^{-1} \sigma_G (\Delta \log p_i - \overline{\Delta \log p_t})$$

where $\beta_{ih}^{t_0}$ is the slope of the relative Engel curve at the crossing point of good i for households in income decile h at time t_0 , and where the term in parenthesis is the log price change for good i relative to the average of log price changes among all goods within G . These bias correction terms are then added to our market-by-decile estimates of the price index to account for the first-order bias coming from confounding relative price changes within G .

A second type of robustness exercise following from Corollary 2 does not depend on using the subset of goods where prices can be observed from survey data. Instead, in many empirical applications, we are interested in the impacts of shocks or policies that generate variation across product groups. For example, in the trade application that we discuss below, tariff changes vary across product groups in a way we can observe in the absence of good price data. If the concern is that, for example, a trade shock leads to relative price changes that are systematically correlated with the slopes of relative Engel curves within G , then an alternative robustness check is to formally test whether variation in the underlying economic shock across goods and services (tariffs in this example) is significantly related to the Engel slope parameters.

Unobserved Welfare Changes (Sample Selection)

Not all levels of household utility in period 0 are necessarily observed in period 1 and vice versa (and similarly not all levels observed in market 0 are observed in market 1 if comparing markets across space). For example, when evaluating price index changes for poor households in period 0, there may be no equally poor households in period 1 if there is real income growth. This means that Engel curves may not always overlap in budget share space for all income percentiles, and this gives rise to sample selection concerns.

Corollary 3 *i) When estimating welfare changes over time for a particular percentile of the income distribution in a given market, sample selection concerns arise across goods when not all good's relative Engel curves overlap across periods.*

ii) When estimating mean welfare changes over time for a particular percentile across all markets, sample selection concerns arise when not all markets have that percentile's welfare levels observed in both periods.

The good-by-good estimation approach described above recovers a mean price index estimate across the $i \in G$ goods for which we can calculate the horizontal shift in relative Engel curves

across periods at a particular percentile. Suppose there is no true overlap across periods because that level of household welfare was simply not observed in the other period. Then it is likely that a few goods experience relative price or taste shocks within G such that their relative Engel curves do overlap across periods. In this scenario, a simple average would yield biased estimates of the price index change and welfare, as the orthogonality condition in Corollary 2 would not hold across the selected subset of goods for which we can measure the horizontal shift in the data. When estimating P^{t0} for poor households or P^{t1} for rich households and there is no true overlap in utility, the cases where we are able to measure horizontal shifts will tend to provide estimates that are biased upwards (as price or taste shocks have vertically shifted the relative Engel curves to be farther apart from one another, leading to more overlap across the income distribution). Conversely, when estimating P^{t0} for rich households or P^{t1} for poor households it can still be the case that there is no true overlap in utility, in which case we have a downward selection bias in the price index estimate.

Similar selection concerns arise also in cases where there is true overlap in terms of utility but not all goods have overlapping Engel curves. In this case, some goods may experience relative price or taste shocks (vertical shifts) that drive the two relative Engel curves closer together in vertical space, such that there is no overlap at one of the tails of the income distribution. By averaging over only the subset of goods for which there is overlap generates biases in the same directions as in the case of no true utility overlap above.

To address such sample selection concerns, we exploit the fact that we observe whether or not a given good or service has missing overlap at a given income percentile (for both P^{t0} and P^{t1}) and whether this good is censored from above or from below at this percentile. Using this information, we can make the identifying assumption that the distribution of price index estimates across the different goods and services i within G is symmetric for a given income percentile in question. Since we know the ordering of the observed and unobserved price index estimates for this household group across all goods and services i within G , at least all those with monotonic Engel curves, the symmetry assumption allows us to consistently estimate the price index change at this point from the median (which is an unbiased estimate of the mean).³⁴ However, the median will not always be observed, for example if most goods are censored from above. In these cases, we can impose a stronger distributional assumption that the price index estimates for a given percentile of the income distribution follow a uniform distribution (see Sarhan (1955)). That allows us to solve for the mean as long as at least two goods overlap. In the following applications section, we report and compare results imposing just the symmetry assumption, as well additionally imposing uniformity when the median is not observed.

A different type of sample selection arises if we don't observe any relative Engel curves that overlap for a given percentile (part (ii) of Corollary 3 above). In this case, we cannot apply the strategies above. Instead, we are faced with a market-level sample selection concern when aggregating across markets (for example, to summarize welfare growth across markets at a give percentile of the income distribution, or to estimate the average effect of trade on welfare at a particular income percentile). In particular, there will be missing markets among poor percentiles

³⁴We rank estimates, placing unobserved estimates above the highest or below the lowest estimate depending on whether they were censored from above or below.

for P^{t0} and there will be missing markets for rich percentiles for P^{t1} . These missing markets are the markets that have seen the largest welfare growth (and the smallest price index growth) at a given percentile, as those welfare levels are the ones for which overlap will not be observed in the data. This type of market-level missing observation is a standard sample selection issue in a regression setting: we only observe the outcomes for a selected subset of markets, and those missing are not drawn at random but instead the probability of being observed is a function of price or income growth.

Applying a standard two-step Heckman selection method to the market-level selection concern in our setting would require somewhat heroic assumptions: the selection process in our applications is directly tied to the outcome (inflation or welfare) that we are interested in, making plausibly exogenous variation that affects selection, but not treatment, extremely unlikely. Instead, we can make use of the facts that i) we are in a setting where the outcome is censored from above or below (ordered censoring), and ii) we observe the distribution of price index or welfare estimates for most markets in our sample at every percentile. Using these two features of our setting, we can apply sensible distributional assumptions to correct for missing markets. In our estimation that we present below, we show that almost no markets remain missing after we implement the good-level selection correction we discuss above (i.e. we observe overlap in relative Engel curves for some goods at a given decile in almost every market in our sample). Therefore, the good-level selection correction is sufficient in our context to solve market-level selection issues.

Tests of Quasi-Separability

Using the results of Lemma 3 in the theory section, we can derive a test of quasi-separability using the expenditure survey data.

Corollary 4(a) *If, and only if, preferences are quasi-separable in group G , then the price elasticity of the uncompensated expenditure share $x_{iG} \equiv \frac{x_i}{x_G}$ (i.e. holding income y fixed) in the price of any good $j \notin G$ is equal to the slope of the relative Engel curve multiplied by good j 's overall budget share:*

$$\left. \frac{\partial \log x_{iG}}{\partial \log p_j} \right|_y = - \frac{p_j q_j}{y} \frac{\partial \log x_{iG}}{\partial \log y}$$

Using this formulation of the cross-price elasticity between outside goods and relative budget shares within G , we can again make use of observed price information in the expenditure survey data to test these restrictions for products and consumption categories that report both consistent quantity measures in addition to expenditure data (see the discussion of Corollary 2 above).

The test above can be reformulated as a test directly involving the price indices P^{t0} and P^{t1} . In the test above, $\left. \frac{\partial \log x_{iG}}{\partial \log p_j} \right|_y$ corresponds to the vertical shift of the relative Engel curve induced by the marginal change in the price of good j . Alternatively, we can explore the horizontal shift induced by this price change which is equal to the ratio $\left. \frac{\partial \log x_{iG}}{\partial \log p_j} \right|_y / \frac{\partial \log x_{iG}}{\partial \log y}$. Under quasi-separability, this ratio coincides with the marginal effect of the change in the price of a good j on the price indices $P^{t0} = P^{t0}(p^{t0}, p^{t1}, y_h^{t0})$ and $P^{t1} = P^{t1}(p^{t0}, p^{t1}, y_h^{t1})$. This observation generates a second test:

Corollary 4(b) *The elasticity of the exact price index P^t , $t \in \{t_0, t_1\}$ with respect to the price of any good j equals the overall expenditure share of good j :*

$$\frac{\partial \log P^t}{\partial \log p_j^t} = \frac{p_j^t q_{jh}^t}{y_h^t}$$

This equality is simply Shephard’s Lemma applied to our price indices. However, since we do not use prices from non- G goods to estimate our price indexes, our estimation strategy does not guarantee that this equality holds. Our estimated price indices fully capture both G and non- G prices only if quasi-separability holds

Finally, note that we can extend and jointly test Corollary 2, 4(a) and 4(b) by exploiting price index changes estimated from multiple G groups. In particular, if more than one group G are used in the estimation, and the observed product aggregation is fine enough to have a large number of individual products i within each of the G (such that the orthogonality conditions across the $i \in G$ are credible), then the price index and welfare results are over-identified. If the orthogonality conditions within each group of G as well as the quasi-separability assumption hold for all groups G , then the estimates based on each G should yield point estimates that are statistically indistinguishable.

6 Applications

In the final section we apply the new methodology introduced in the previous sections. We explore both changes in Indian welfare across time and the welfare impacts of India’s 1991 trade reforms.

[This section is still preliminary. In future drafts, we plan to explore the results in more detail and present a variety of additional robustness exercises. Pending updated data from the World Bank’s Global Consumption Database, we also intend to calculate purchasing power parity numbers and explore the gains from trade across countries.]

6.1 Changes in Indian Price Indices and Welfare Over Time

To implement our new methodology, we use the Indian NSS microdata described in Section 3 to estimate changes in household price indices and welfare over time between 1987-87 and 1999-2000. We do this for 9 income deciles (10-90) in each rural Indian district. We then investigate how our approach based on relative Engel curves compares to the existing Indian CPI statistics that are based on changes in well-measured prices of food and fuel products (see e.g. Deaton, 2003b).

The estimation approach requires quasi-separable groups G . Following Corollaries 2 and 4 of the previous section, we restrict the estimation to the 136 products with reliable price information that we then aggregate to 35 product groups.³⁵ These 35 product groups are part of three broader consumption categories: raw food products (e.g. rice, leafy vegetables), other food products (e.g. milk, edible oils) and fuels (e.g. firewood, kerosene). In our baseline estimation, we assume these three broad groups each form a G product group, with the remaining 174 products (e.g. processed food, durables and services) excluded as part of the NG group, where relative price changes are

³⁵We use a superlative price index (Fisher) to aggregate the observed price changes of the 136 products to the level of 35 product groups (using the relative budget shares within each group in both periods to compute weights).

unrestricted. As we discuss as part of the robustness analysis below, we also report results across alternative groupings and aggregations for both the G s and the $i \in G$. Given the need to estimate Engel curves at many points in the income distribution, we also restrict attention to markets where we observe 100 or more households in each survey round.³⁶

Figure 4 begins by plotting the mean growth rates in nominal household outlays per capita for each decile of the local income distribution across all rural Indian districts (weighting districts by population). In terms of income growth, there is a clear and strong pattern of convergence over this 13-year period with incomes for the poor growing substantially faster. Clearly, any single price index applied to all households in a market would yield the same conclusion, that there has been substantial convergence of real incomes.

The left panel of Figure 5 presents the price index estimation results of the New Engel approach using the uniform sample selection correction described in the discussion of Corollary 3 in Section 5 (again, using survey weighted means across markets by decile). Alongside these estimates, we present price index changes using existing CPI estimates (Laspeyres and Paasche price indices) that follow the methodology of Deaton 2003b and are based on changes in observed prices for food and fuels where price data were deemed reliable. The left panel shows standard CPI estimates that use average expenditure shares across all households in the market to weight price changes (i.e. democratic price index weights, not plutocratic). Mechanically, these do not vary across the income distribution. The middle panel modifies Deaton’s approach and presents decile-specific Laspeyres and Paasche price indices, in which price changes are weighted by the average expenditure shares of households in a given decile. Figure 6 plots the resulting welfare changes from the decile-specific CPI and New Engel approaches.

Two main findings emerge. First, we find that the New Engel approach generates broadly similar estimates of Indian consumer price inflation among low-income deciles compared to existing estimates that are based on changes in observed prices for food and fuel. Since these product groups represent a sizable fraction of rural household consumption for poor households in India (around 80 percent for the poorest decile in 1999/2000, falling to 60 percent for the average household), this finding is reassuring— particularly since no price data was used in the New Engel approach. Second, as discussed above, estimates based on a single price index suggest that there has been significant convergence between poor and rich household income deciles over time, a result that is exacerbated using decile-specific Laspeyres and Paasche price indices. In contrast, our approach finds that cost of living inflation has been substantially lower among richer Indian households compared to the poor, substantially reducing the degree of real income convergence in rural India over this 13-year period.

Given that existing estimates have been mainly based on homothetic Laspeyres price indices using price changes for food and fuel, the most likely explanation for these findings is that higher-income Indian households disproportionately benefited from previously omitted components of inflation over this period in India. Most intuitively, since the rich spend a larger share of their budget on services, durables, and hard to measure non-food non-durables such as clothing and

³⁶For our baseline estimates presented below, we have not yet implemented a bootstrap procedure to obtain standard errors that take into account sampling variation underlying each market’s price index and welfare estimates. In future versions of this draft, we plan to bootstrap the estimation procedure described in the previous section by randomly re-sampling (with replacement) at the level of individual households.

personal care items, , lower price growth in these excluded product groups , including improvements in product quality and variety, would generate just such a finding.

[In work in progress, we plan to use the same approach to evaluate spatial inequality across rural districts in India.]

6.1.1 Selection Corrections, Tests of Assumptions and Robustness Checks

In this subsection, we implement and report a number of tests discussed as part of Corollaries 1-4 in Section 5.

Sample Selection Issues

In Appendix Figures A.1-A.4 we report additional estimation results to illustrate and assess sample selection issues. Figures A.1-A.3 report the estimation results on rural Indian price index changes across different approaches for dealing with the good-level selection correction that we discuss as part of Corollary 3 above. The left panel of Figure A.1 presents the price index estimates that do not correct for the fact that not all $i \in G$ have Engel curves that overlap across the two periods for all deciles, and that these missing estimates are not randomly selected. In particular, missing price index estimates across the $i \in G$ are driven by relative price or taste shocks (vertical shifts) that drive the two relative Engel curves closer together in vertical space, reducing the region of overlap horizontally. When not accounting for such missing estimates (i.e. averaging across the observed estimates $i \in G$), the orthogonality condition under Corollary 2 would likely be violated, typically yielding upward-biased estimates of price index changes (as the missing goods experience price or taste shocks that result in smaller horizontal shifts in relative Engel curves).

This pattern is clearly apparent in the left panel in Figure A.1. For most deciles, inflation estimates are much larger than those in our baseline estimation depicted in the right panel that applies a selection correction. The biggest discrepancies between the two figures occur for P^{t0} among the poorest deciles and P^{t1} among the richest deciles. It is exactly these households for which overlap issues are most severe since with economic growth, the welfare levels of the poorest households in period 0 are typically not observed in period 1, and similarly the welfare levels of the richest households in period 1 are typically not observed in period 0. In contrast, among the richest deciles for P^{t0} and the poorest deciles for P^{t1} , a less severe selection bias operates in the opposite direction. Taken together, the bias from the uncorrected estimates is large enough that the two price indices P^{t0} and P^{t1} would yield opposite conclusions about whether inflation was pro-rich.

The next two panels of Figure A.1 sequentially implement the sample selection corrections we discuss in Section 5. The middle panel only imposes symmetry in the distribution of the price index estimates across all goods $i \in G$ (including those that are unobserved due to a lack of overlap—which recall we can still rank in order to take medians which are unbiased estimates of the mean if the distribution is symmetric). The right panel (our baseline estimation) further imposes the assumption of a uniform distribution, which following Sarhan (1955) allows us to estimate the mean of the distribution even in cases where the median is unobserved (as long as we observe at least two goods with overlap). As is apparent in the Figure, the two approaches yield almost identical results when averaged across markets.

The advantage of the baseline approach is clear from Figure A.2 which depicts the number of markets by decile for which we obtain price index estimates for each of these three approaches to good-level selection issues. While market-level selection issues (missing markets when averaging across markets by decile) arise when just imposing symmetry (middle panel), they are effectively eliminated by imposing the uniformity assumption when the median is missing (right panel). That is, we observe enough overlapping goods $i \in G$ for each market and decile such that market-level selection issues do not arise in our empirical application.

Figure A.3 drills further into sources of missing good-level overlap across the two periods. The figure makes use of the fact that we know for each of the $i \in G$ from which direction the relative Engel curves are censored (i.e. due to missing overlap from above or below) that we use to rank the missing estimates and calculate medians. As discussed above, we find that missing good-level estimates of the price index change tend to be concentrated among poor households for P^{t0} and for rich households for P^{t1} . Finally, Figure A.4 reports the estimation results without restriction attention to markets with at least 100 household observations in both survey rounds. Reassuringly, there does not appear to be a systematic difference in our estimates (even though sampling error is of course much higher among sparsely covered markets).

Orthogonality and Quasi Separability

In Appendix Figures A.5 and A.6 we report additional results that correct for potentially confounding relative price changes across $i \in G$ and test the sensitivity of our estimates to alternative groupings and levels of aggregation across both is and Gs .

Figure A.5 depicts our baseline inflation estimates alongside estimates that apply the correction term for confounding relative price changes within G that we discuss as part of Corollary 2 in Section 5.³⁷ Recall that we focus our analysis on product groups G for which we observe reliable price data following Deaton (2003b). The advantage of this focus is that we can directly compute a first-order approximation of the bias due to any systematic correlation between relative price changes within the Gs and slopes of relative Engel curves. Reassuringly, we find that inflation rates change little with this correction. In other words, we do not find evidence that relative price changes within our three food and fuel product groups are systematically different across more or less income elastic goods in our setting.

Figure A.6 reports estimates from alternative aggregations of goods i and groups G . The left panel depicts our baseline approach with three G groups and 35 i goods. The middle panel keeps the same 35 i goods, but assigns them to finer product groups within food and fuels (six Gs instead of three). The right panel keeps the initial three product groups G , but increases the level of aggregation of the is (25 products instead of 35). Reassuringly, we find very similar results across these alternative product aggregations and choices of product groups G .

In work in progress, we are implementing the formal test for quasi-separability described in Corollary 4(a) above. In particular, we estimate the matrix of cross-price elasticities for price changes of j goods outside of G on expenditures for $i \in G$ and test for what fraction of cells in this matrix (for each of the three groups G), we can reject the hypothesis that the observed cross-price effect is consistent with Corollary 4(a).

³⁷We assume $\sigma = 4$ to calculate the correction term. In future drafts we plan to estimate group and decile-specific elasticities.

Here we present instead the test described in Corollary 4(b). Given that we have reliable price data for only food and fuel products, we implement our test by re-estimating price indices from food expenditures only (by district and by decile), and regressing these price indices on fuel price changes interacted with fuel expenditure shares. Assuming that fuel price changes across districts are orthogonal to other price changes, Corollary 4(b) predicts a coefficient equal to unity. We show the results of this test in Table 2. In the first and third columns, we find that the price index increase across markets is proportional to the fuel price increase and the share of fuels in overall consumption, with a coefficient close to unity. In the second and fourth columns, we exploit within-market variations across deciles to show that our estimated price indices have increased relatively more for households with larger expenditure shares in fuel prices, in line with Corollary 4(b). These results indicate that our estimated price indices, computed from changes in relative expenditure shares within one set of goods (raw foods and other foods), successfully capture price changes in goods outside this group (fuels).

Additional Robustness Checks

Finally, in Appendix Figures A.7-A.9 we report a number of additional robustness checks. Figure A.7 presents results across alternative bandwidth choices and alternative ways of dealing with noise at the tails of the distribution for the non-parametric estimation of relative Engel curves. Reassuringly, these choices do not yield qualitatively different results compared to our preferred baseline estimation discussed above.

Finally, we turn to issues related to recall bias. Figures A.8 and A.9 revisit the estimation after focusing on price index or welfare changes between the 43rd and 50th round (1987/88 to 1994/995) instead of between rounds 43 and 55 as in Topalova. As rounds 43 and 50 both only included a 30-day recall period for foods, the recall bias issues discussed in Section 3 do not arise. Recall from Sections 3 and 5 above that to bias our welfare results, mis-reported outlays due to recall bias would have to i) affect not just total outlays, but relative outlays within groups G , and ii) any such mis-reporting of relative budget shares within G would have to be systematically different across more or less income-elastic goods.

The fact that Figures A.8 and A.9 show very similar price index and welfare results (i.e. pro-rich inflation leading to a substantial flattening of the convergence seen in nominal incomes) provides some reassurance that recall bias is not affecting our baseline welfare estimates. In ongoing work in progress, we make use of the fact that different recall periods were randomly assigned as part of the NSS surveys in “thin” rounds between 50 and 55 to directly test whether both the conditions i) and ii) above are present in the data (by comparing relative outlays within G across randomly assigned recall periods, and by testing whether any recall bias we do find in relative shares is related to income elasticities).

6.2 Effect of Indian Trade Reforms Across Districts

In this section, we revisit the impact of India’s 1991 trade reforms on the welfare of rural households in India. We closely follow the seminal analysis of Topalova (2010) which regresses poverty rates in rural districts on local industry composition-weighted sums of tariff changes due to the 1991 Indian trade reforms. This paper pioneered the use of a shift-share instrument to identify the impacts of trade shocks, an identification strategy that has widespread in the trade

literature in recent years (e.g. Kovak, (2013); Autor, Dorn and Hanson, (2013)).

Her most robust specification regresses the the head count ratio for the 1988/89 and the 1999/2000 NSS rounds on district fixed effects, time fixed effects, a list of time-changing district controls (or levels interacted with the time fixed effects) and the key regressor of interest: district-level exposure to import tariff cuts, measured as the weighted average tariff cut, with weights proportional to the first-period sectoral employment shares in the local market. This explanatory variable is instrumented with the two IVs. The first is the same measure of import competition, but estimated only using tradable industries (taking out confounding effects of differences in industrial employment shares across districts). The second IV is the employment share-weighted initial average level of import tariffs (rather than changes), applying the argument that initially more protected sectors experienced a larger unexpected tariff reduction due to Indian market reforms in the early 1990s.

We revisit this specification of Topalova, and replace the outcome (poverty rates) with our estimated welfare metrics.³⁸ Importantly, our method allows us to calculate impacts at each decile of the income distribution. For exposition, we focus here on the estimated welfare effects as measured by the log of the equivalent variation as a fraction of initial income (the equivalent variation welfare metric presented in Proposition 1).³⁹ Figure 7 presents the point estimates for the New Engel approach and compares them to the effects on nominal outlays across deciles. Figure 8 presents the point estimates of the effect of import competition on a conventional decile-specific Laspeyres price index that is only based on observed price changes (again for food and fuels following Deaton).

Two main findings emerge. First, while existing work on the effects of Indian trade reforms across Indian districts has focused on the effect on poverty rates (Topalova, 2010), our estimation reveals that the adverse effects of import competition on local labor markets are borne by households across the income distribution, including by households in the richest income deciles of rural India. Second, we find that the adverse effects of import competition on local nominal outcomes are amplified when taking into account the effects on household price indices. We find that import competition leads to higher local price inflation in particular among the richer tail of the income distribution. Reassuringly, as becomes apparent from Figure ??, we find that this somewhat surprising finding is also supported by the raw price information from the subset of goods we observe in the data.

[In progress: discussion and additional analysis of potential underlying mechanisms.]

6.3 The Gains from Trade Across Countries

[Work in progress. We plan to use the Global Consumption Database produced by the World Bank and described in Section 3, that has harmonized data from expenditure surveys containing expenditure shares for 92 countries and 107 goods and services. With these data, we intend to calculate purchasing power parity exchange rates across countries and revisit Frankel and Romer's (1999) classic analysis of trade and growth in the cross-section of countries.]

³⁸We find broadly similar findings using Topalova's other specifications, although just as in Topalova, results become less significant. Note that Topalova does not restrict attention to markets with more than 100 households observed. Restricting Topalova's sample to be identical to ours makes her effect size larger.

³⁹Due to the fact that log nominal outlays are right skewed, we have more overlapping Engel curves when calculating P^{f1} and equivalent variation than for compensating variation and so prefer this measure.

7 Conclusion

Measuring changes in household welfare and the distribution of those changes is challenging and requires a combination of detailed microdata that are seldomly, if ever, available to the researcher. In this paper, we propose and implement a new approach to estimate changes in household price indices and welfare across the income distribution from horizontal shifts in what we term relative Engel curves. We prove that if preferences fall within the class of quasi-separable preferences, such an approach uncovers theory consistent and exact price indices only drawing on widely available expenditure survey microdata. In particular, our approach does not require accurate measures of price changes for the whole consumption basket. However, focusing on subsets of goods for which we have reliable price data allows us to recover the full price index, but also to directly test the preference restrictions and identifying assumptions required for unbiased estimation, and to compute correction terms if necessary.

The methodology we present is widely applicable in the many contexts where expenditure survey data is available and researchers want to understand the welfare effects of policies or shocks, and particularly the distribution of those effects. Given the increasing availability of expenditure survey data over time and across space, and the increased interest in distributional analysis, the usefulness of such an approach is likely to grow.

We apply this new machinery to measure changes in Indian household welfare and re-visit the effects of trade across Indian regions. We have three preliminary findings. First, we find that Indian consumer price inflation has been higher for poor households than rich, a finding that is missed by calculating standard price indices from the subset of consumption where prices are observable, even when these price indices uses income-group specific product weights. This fact strongly reduces convergence between rich and poor households. Second, while existing work on the effects of Indian trade reforms across Indian districts has focused on the effect of poverty rates (Topalova, 2010), our estimation reveals that the adverse effects of import competition on local labor markets are borne by households across the entire income distribution, including the richest income deciles. Third, we find that the adverse effects of import competition on local nominal outcomes are amplified when taking into account the effects on household price indices with import competition raising inflation particularly for the rich.

References

- ALMÅS, I. (2012): "International income inequality: Measuring PPP bias by estimating Engel curves for food," *American Economic Review*, 102(2), 1093–1117.
- ALMÅS, I., T. BEATTY, AND T. CROSSLEY (2018): "Lost in translation: What do Engel curves tell us about the cost of living?," *IFS working paper W18/04*.
- ALMÅS, I., AND A. KJELSRUD (2016): "Rags and riches: Why consumption patterns matter for inequality," *Working paper*.
- AMITI, M., M. DAI, R. C. FEENSTRA, AND J. ROMALIS (2018): "How did China's WTO entry benefit US consumers?," Discussion paper, National Bureau of Economic Research.
- ATKIN, D., B. FABER, AND M. GONZALEZ-NAVARRO (2018): "Retail globalization and household welfare: Evidence from Mexico," *Journal of Political Economy*, 126(1), 1–73.
- AUTOR, D. H., D. DORN, AND G. H. HANSON (2013): "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 103(6), 2121–68.
- BANKS, J., R. BLUNDELL, AND A. LEWBEL (1997): "Quadratic Engel curves and consumer demand," *Review of Economics and statistics*, 79(4), 527–539.
- BERLINGIERI, G., H. BREINLICH, AND S. DHINGRA (2018): "The Impact of Trade Agreements on Consumer Welfare," *Journal of the European Economic Association*.
- BLACKORBY, C., R. DAVIDSON, AND W. SCHWORM (1991): "Implicit separability: Characterisation and implications for consumer demands," *Journal of Economic Theory*, 55(2), 364–399.
- BLACKORBY, C., D. PRIMONT, AND R. R. RUSSELL (1978): *Duality, separability, and functional structure: Theory and economic applications*, vol. 2. Elsevier Science Ltd.
- BLUMENSTOCK, J. E. (2016): "Fighting poverty with data," *Science*, 353(6301), 753–754.
- BORUSYAK, K., AND X. JARAVEL (2018): "The Distributional Effects of Trade: Theory and Evidence from the United States," *Harvard University Working Paper*.
- COMIN, D. A., D. LASHKARI, AND M. MESTIERI (2015): "Structural change with long-run income and price effects," *National Bureau of Economic Research Working Paper*.
- COSTA, D. L. (2001): "Estimating real income in the United States from 1888 to 1994: Correcting CPI bias using Engel curves," *Journal of political economy*, 109(6), 1288–1310.
- DEATON, A. (2003a): "Adjusted Indian poverty estimates for 1999-2000," *Economic and Political Weekly*, pp. 322–326.
- (2003b): "Prices and poverty in India, 1987-2000," *Economic and Political Weekly*, pp. 362–368.
- DEATON, A., AND J. MUELLBAUER (1980): "An almost ideal demand system," *American economic review*, 70(3), 312–326.
- DEATON, A., AND A. TAROZZI (2000): "Prices and poverty in India," *Princeton University Working Paper*.
- FABER, B., AND T. FALLY (2017): "Firm heterogeneity in consumption baskets: Evidence from home and store scanner data," *NBER Working Paper 23101*.
- FAJGELBAUM, P., G. M. GROSSMAN, AND E. HELPMAN (2011): "Income distribution, product quality, and international trade," *Journal of Political Economy*, 119(4), 721–765.
- FAJGELBAUM, P., AND A. KHANDELWAL (2016): "Measuring the unequal gains from trade," *The*

- Quarterly Journal of Economics*, 131(3), 1113–1180.
- FALLY, T. (2018): “Integrability and generalized separability,” *CEPR Discussion Paper* 12667.
- FEENSTRA, R. C. (1994): “New product varieties and the measurement of international prices,” *American Economic Review*, 84(1), 155–57.
- FRANKEL, J. A., AND D. H. ROMER (1999): “Does trade cause growth?,” *American economic review*, 89(3), 379–399.
- GORMAN, W. M. (1965): “Production functions in which the elasticities of substitution stand in fixed proportions to each other,” *The Review of Economic Studies*, 32(3), 217–224.
- GORMAN, W. M. (1970): “Quasi-separable preferences, costs, and technologies,” *University of North Carolina, Chapel Hill, mimeo*.
- GORMAN, W. M. (1976): “Tricks with utility functions,” *Essays in Economic Analysis: The Proceedings of the Association of University Teachers of Economics, Sheffield 1975*.
- GORMAN, W. M. (1995): *Collected Works of W.M. Gorman: Separability and aggregation*, vol. 1. Oxford University Press.
- HAMILTON, B. W. (2001): “Using Engel’s Law to estimate CPI bias,” *American Economic Review*, 91(3), 619–630.
- HANOCH, G. (1975): “Production and demand models with direct or indirect implicit additivity,” *Econometrica*, pp. 395–419.
- HAUSMAN, J. A. (1996): “Valuation of new goods under perfect and imperfect competition,” in *The economics of new goods*, pp. 207–248. University of Chicago Press.
- HOTTMAN, C., R. MONARCH, ET AL. (2018): “Estimating Unequal Gains across US Consumers with Supplier Trade Data,” *Federal Reserve Board Working Paper*.
- JEAN, N., M. BURKE, M. XIE, W. M. DAVIS, D. B. LOBELL, AND S. ERMON (2016): “Combining satellite imagery and machine learning to predict poverty,” *Science*, 353(6301), 790–794.
- KOVAK, B. K. (2013): “Regional Effects of Trade Reform: What Is the Correct Measure of Liberalization?,” *American Economic Review*, 103(5), 1960–76.
- LEWBEL, A. (1991): “The rank of demand systems: Theory and nonparametric estimation,” *Econometrica*, pp. 711–730.
- (2008): “Engel curve,” in *The New Palgrave Dictionary of Economics, 2nd edition*, ed. by S. N. Durlauf, and L. E. Blume. Palgrave Macmillan.
- LEWBEL, A., AND K. PENDAKUR (2009): “Tricks with Hicks: The EASI demand system,” *American Economic Review*, 99(3), 827–863.
- LIGON, E. (2017): “Estimating household welfare from disaggregate expenditures,” *Working paper*.
- MATSUYAMA, K. (2015): “The home market effect and patterns of trade between rich and poor countries,” *Northwestern University manuscript*.
- MCFADDEN, D. (1978): “Cost, revenue, and profit functions,” in *Production Economics: A Dual Approach to Theory and Applications*, ed. by M. U. A. for the History of Economic Thought, vol. 1.
- NAKAMURA, E., J. STEINSSON, AND M. LIU (2016): “Are Chinese growth and inflation too smooth? Evidence from Engel curves,” *American Economic Journal: Macroeconomics*, 8(3), 113–44.
- NAKAMURA, L. (1996): “Is U.S. economic performance really that bad?,” *Federal Reserve Bank of*

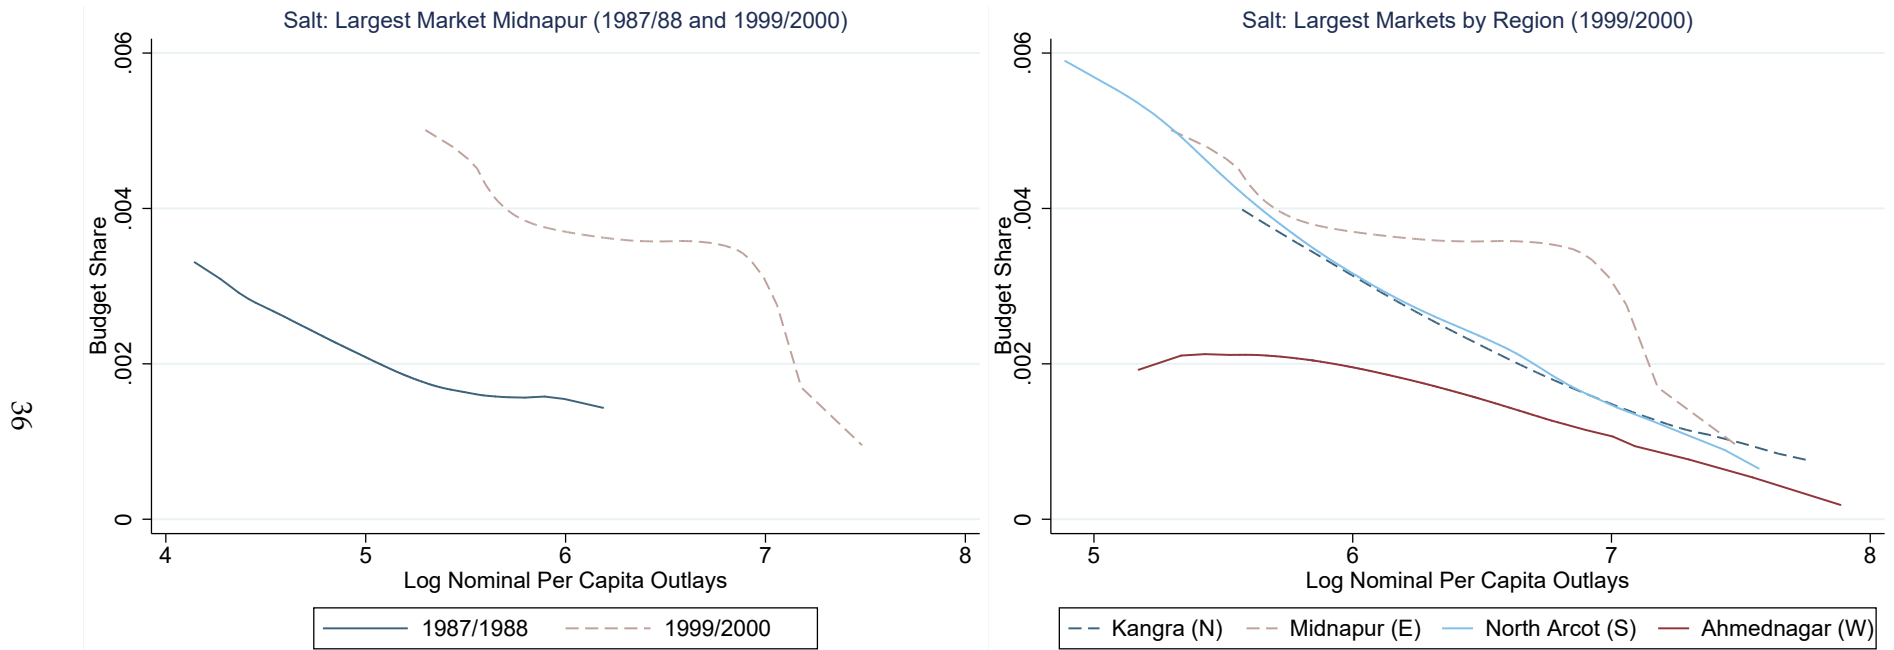
Philadelphia Working Paper No. 95-21/R.

- NOVY, D. (2013): "International trade without CES: Estimating translog gravity," *Journal of International Economics*, 89(2), 271 – 282.
- REDDING, S., AND D. WEINSTEIN (2018): "Measuring Aggregate Price Indexes with Demand Shocks: Theory and Evidence for CES Preferences," *National Bureau of Economic Research Working paper*.
- REDDING, S., AND D. E. WEINSTEIN (2017): "A unified approach to estimating demand and welfare," *NBER Working paper 22479*.
- TOPALOVA, P. (2010): "Factor immobility and regional impacts of trade liberalization: Evidence on poverty from India," *American Economic Journal: Applied Economics*, 2(4), 1–41.
- WORKING, H. (1943): "Statistical laws of family expenditures," *Journal of the American Statistical Association*, 38, 43–56.
- YOUNG, A. (2012): "The African growth miracle," *Journal of Political Economy*, 120(4), 696–739.

8 Figures and Tables

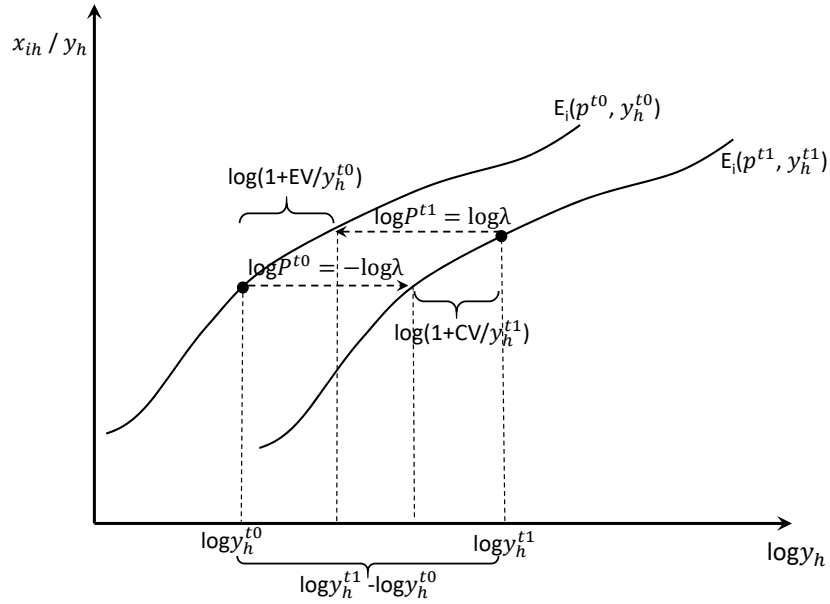
8.1 Figures

Figure 1: Shifts in Engel Curves Over Time and Across Space



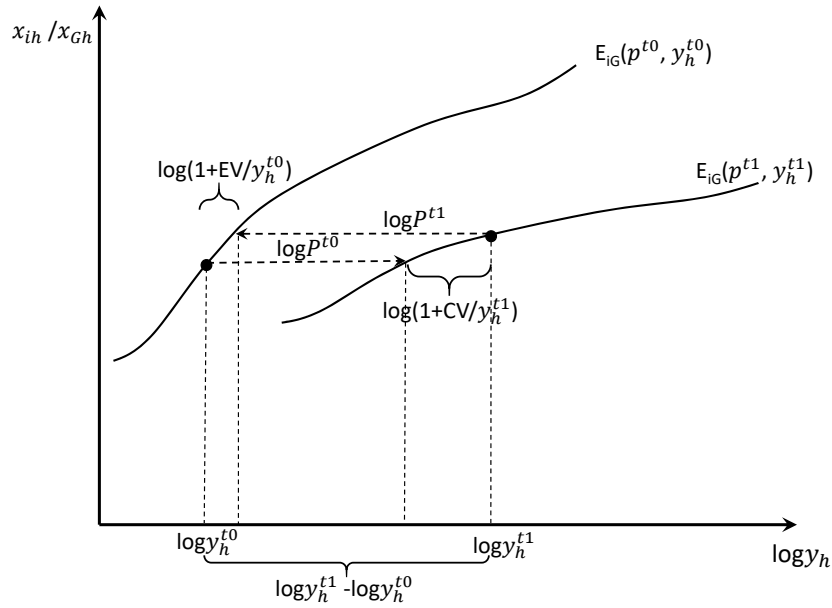
Notes: Figures plot Engel curves for salt over time (NSS 43rd Round 1987-1988 to NSS 55th round 1999-2000) for the largest rural market (Midnapur), and over space for the largest markets in the four broad region of India in terms of numbers of households surveyed . A market is defined as the rural area of an Indian district. See Section 3 for further discussion.

Figure 2: Illustration of Lemma 1



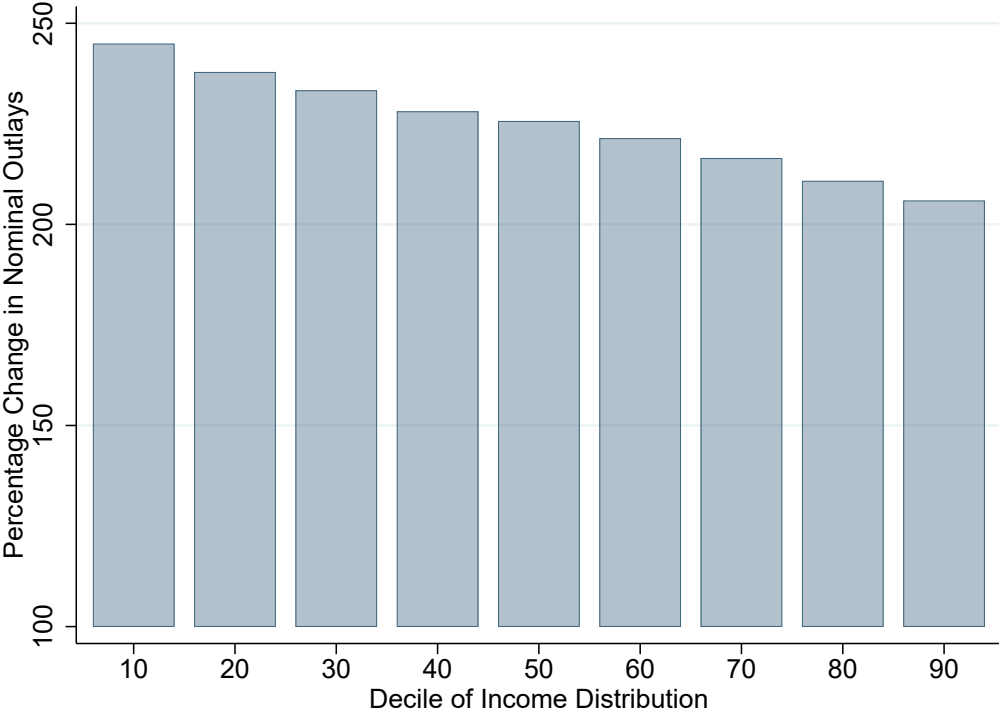
Notes: See Section 4 for discussion.

Figure 3: Illustration of Proposition 1



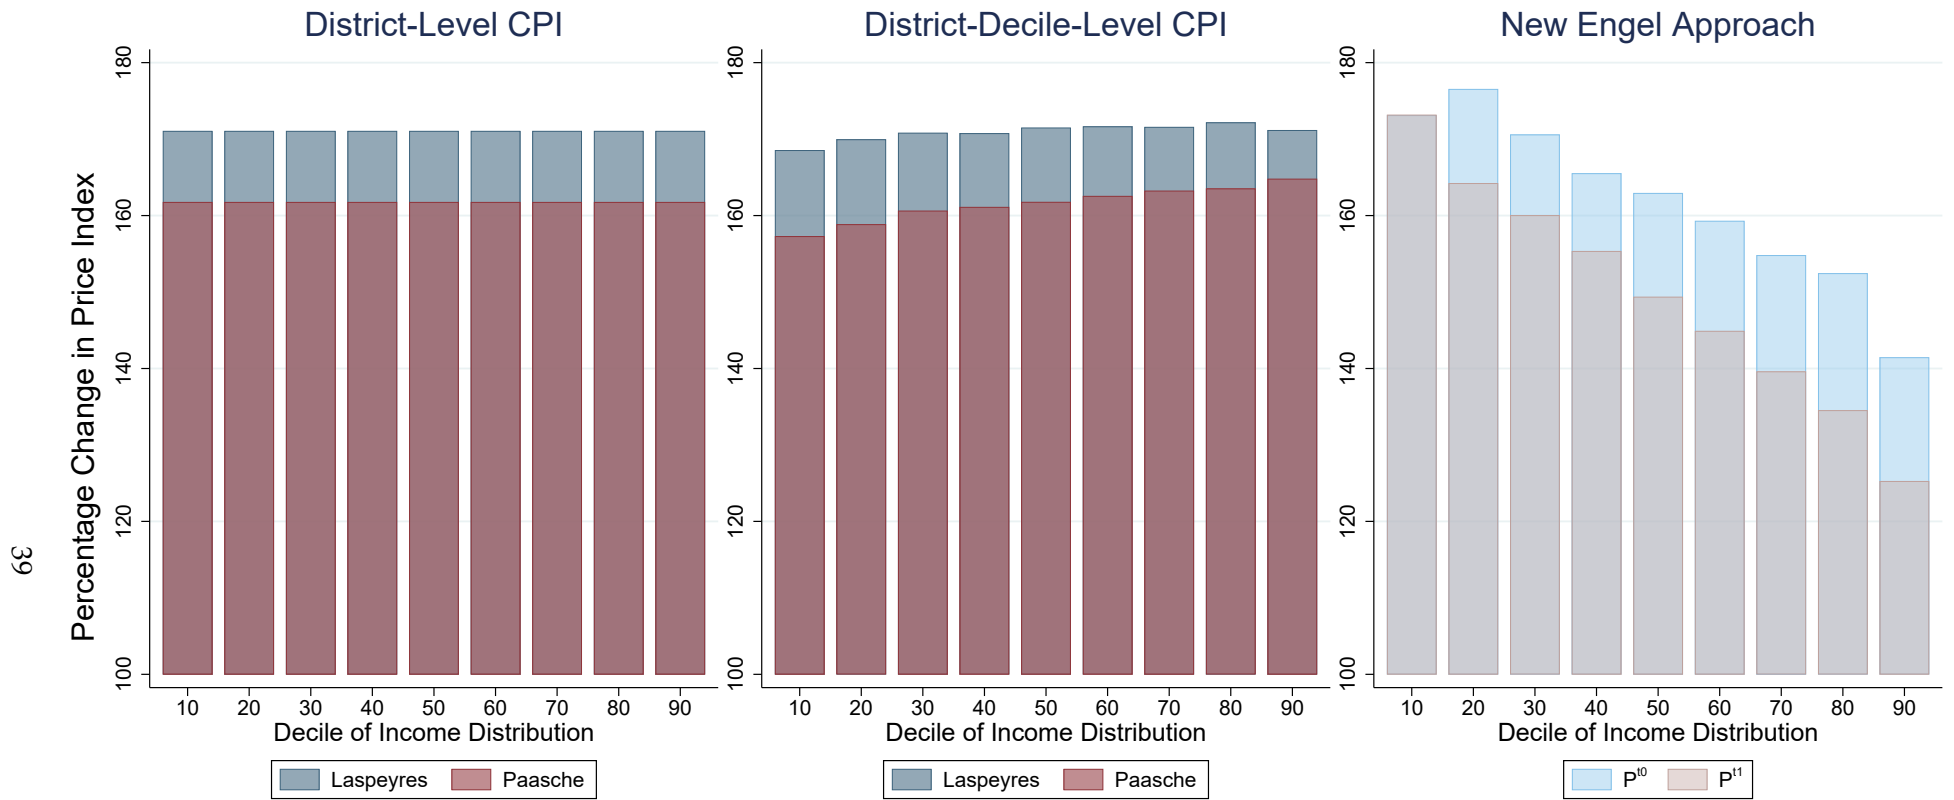
Notes: See Section 4 for discussion.

Figure 4: Indian Growth in Nominal Outlays 1987/88-1999/2000



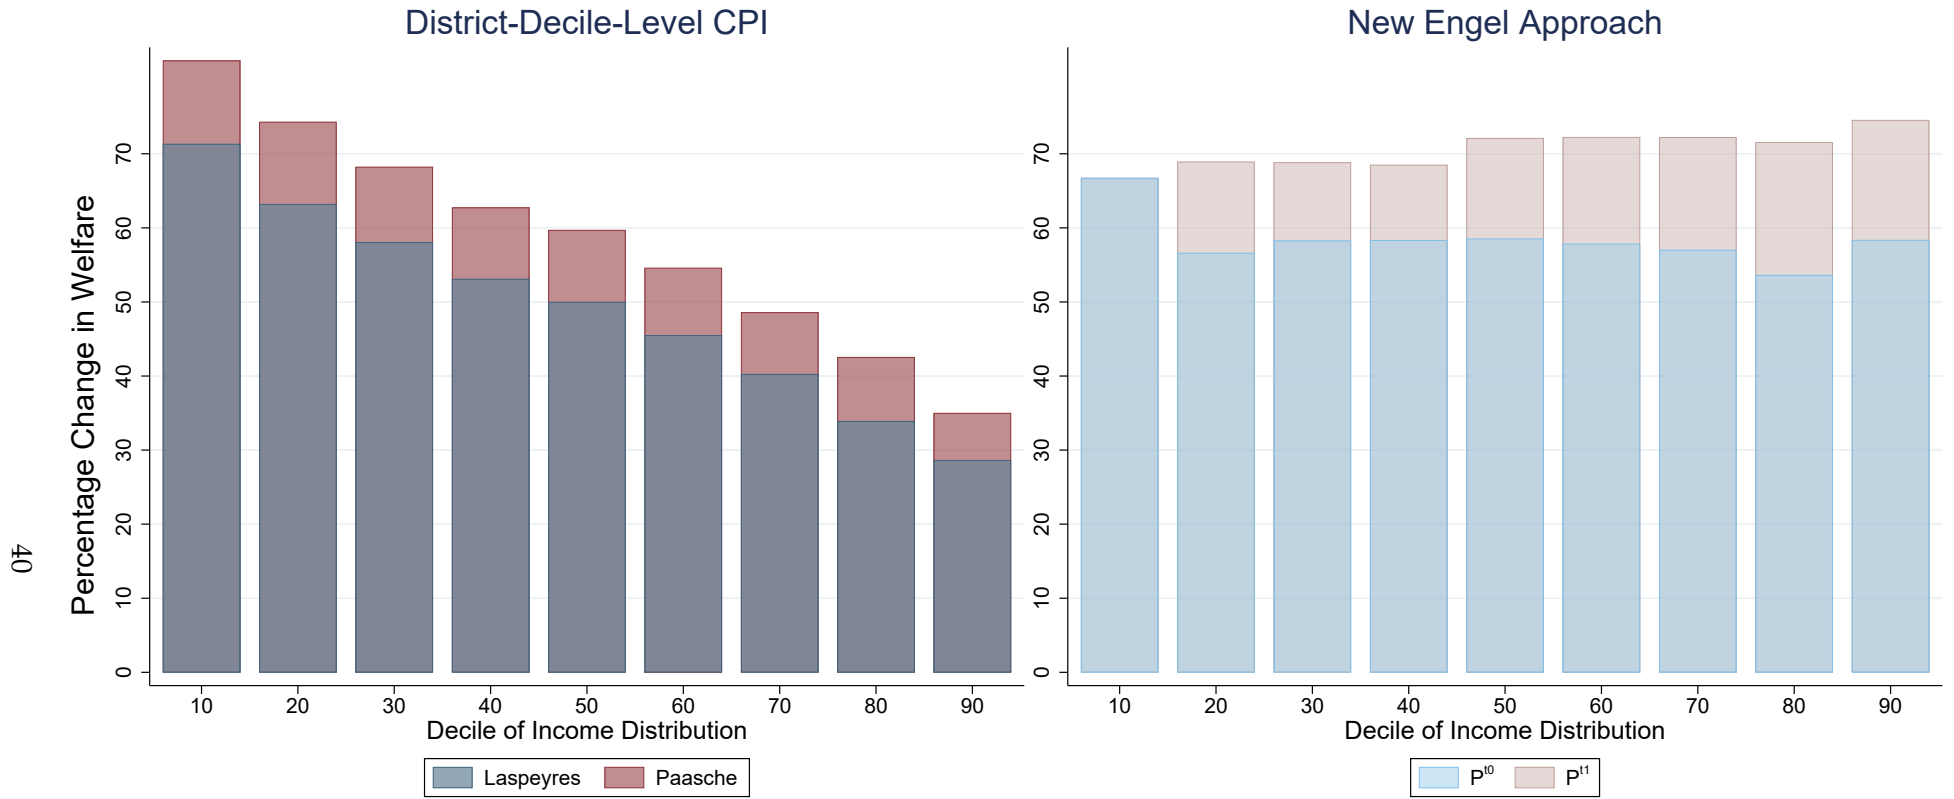
Notes: See Section 6 for discussion.

Figure 5: Indian Cost of Living Inflation 1987/88-1999/2000: New Engel Approach Compared to Existing CPI Estimates



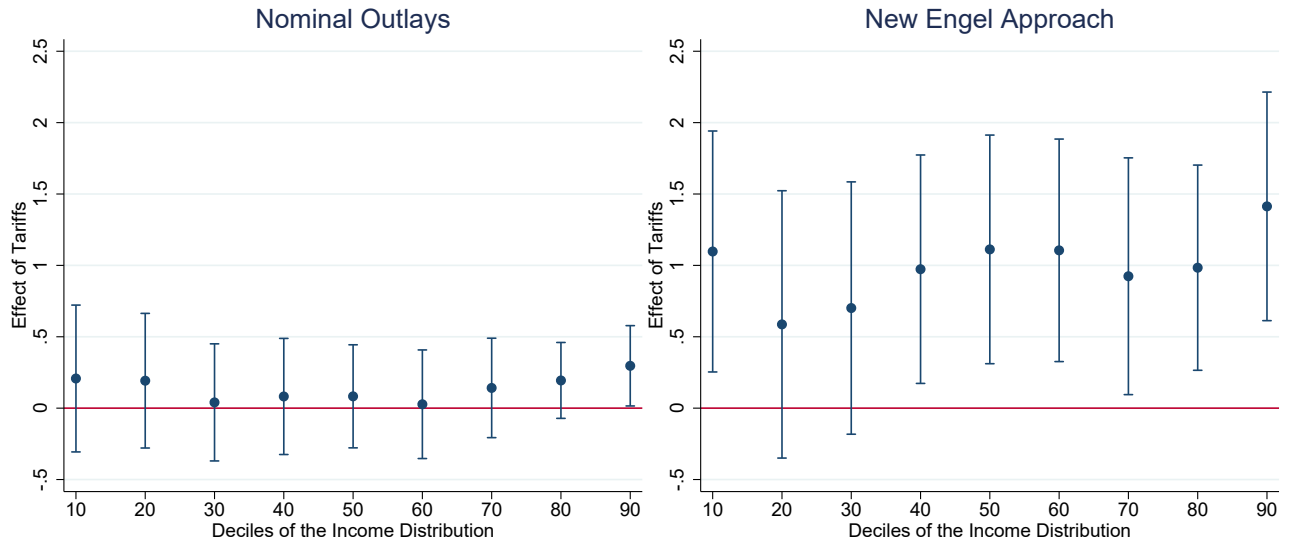
Notes: See Section 6 for discussion.

Figure 6: Indian Welfare Growth 1987/88-1999/2000



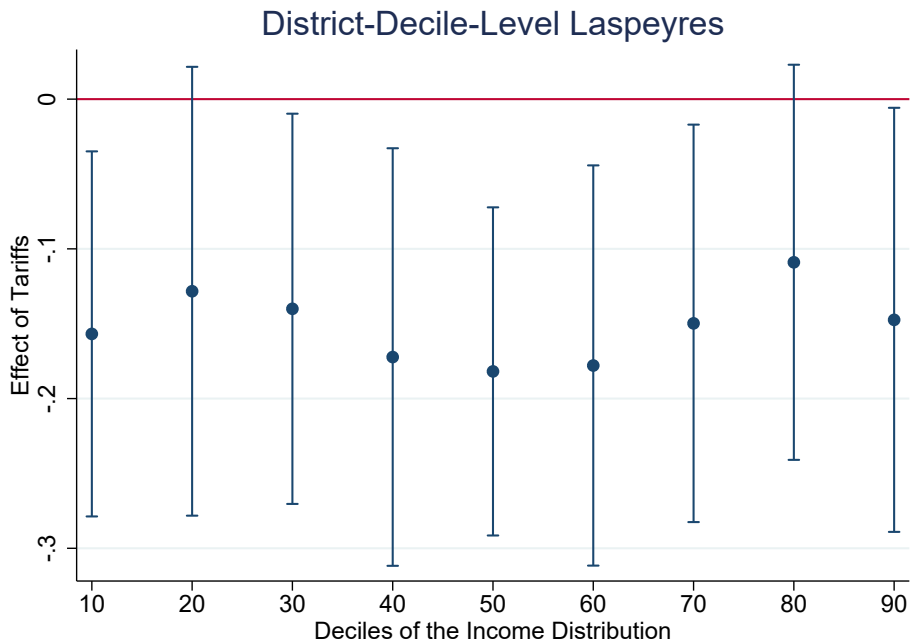
Notes: See Section 6 for discussion.

Figure 7: Effect of Import Competition on Household Nominal Outlays and Welfare



Notes: See Section 6 for discussion. The outcome variable is changes in log welfare using equivalent variation as the welfare metric. Positive point estimates indicate negative effects of import tariff changes.

Figure 8: Effect of Import Competition on Laspeyres Price Index (Only Using Reliable Price Data)



Notes: See Section 6 for discussion. Positive point estimates indicate negative effects of import tariff changes.

8.2 Tables

Table 1: Stylized Facts

	Motivating Fact 1	Motivating Fact 2
	Fraction of Goods and Services Rejecting Linear Engel Curves	Fraction of Market-by-Product Cells Rejecting Uniform Shift in Engel Curve Across Years
	(Out of All Goods and Services)	(Out of All Market-Product Cells)
Fraction Rejecting at the 99% Confidence Level	0.815	0.567
Fraction Rejecting at the 95% Confidence Level	0.899	0.686
Fraction Rejecting at the 90% Confidence Level	0.905	0.745
Controls for Household Characteristics	✓	✓

Notes: See Section 3 for discussion. The first column pools all markets in both periods to estimate the Engel curve for all goods and services in the Indian microdata. Engel curves are estimated by stacking within-market-by-period variation for each good or service across all markets and periods. We reject linearity if the joint test of all 2nd order or higher polynomial terms of log household outlay per capita is significantly different from zero. The second column presents information from market-by-good cells covering all markets with at least 100 households in both survey rounds. We reject a uniform horizontal shift in the Engel curve if the shift in log nominal outlays, moving from Round 43 to 55, is not uniform for different levels of budget shares.

Table 2: Quasi-separability test

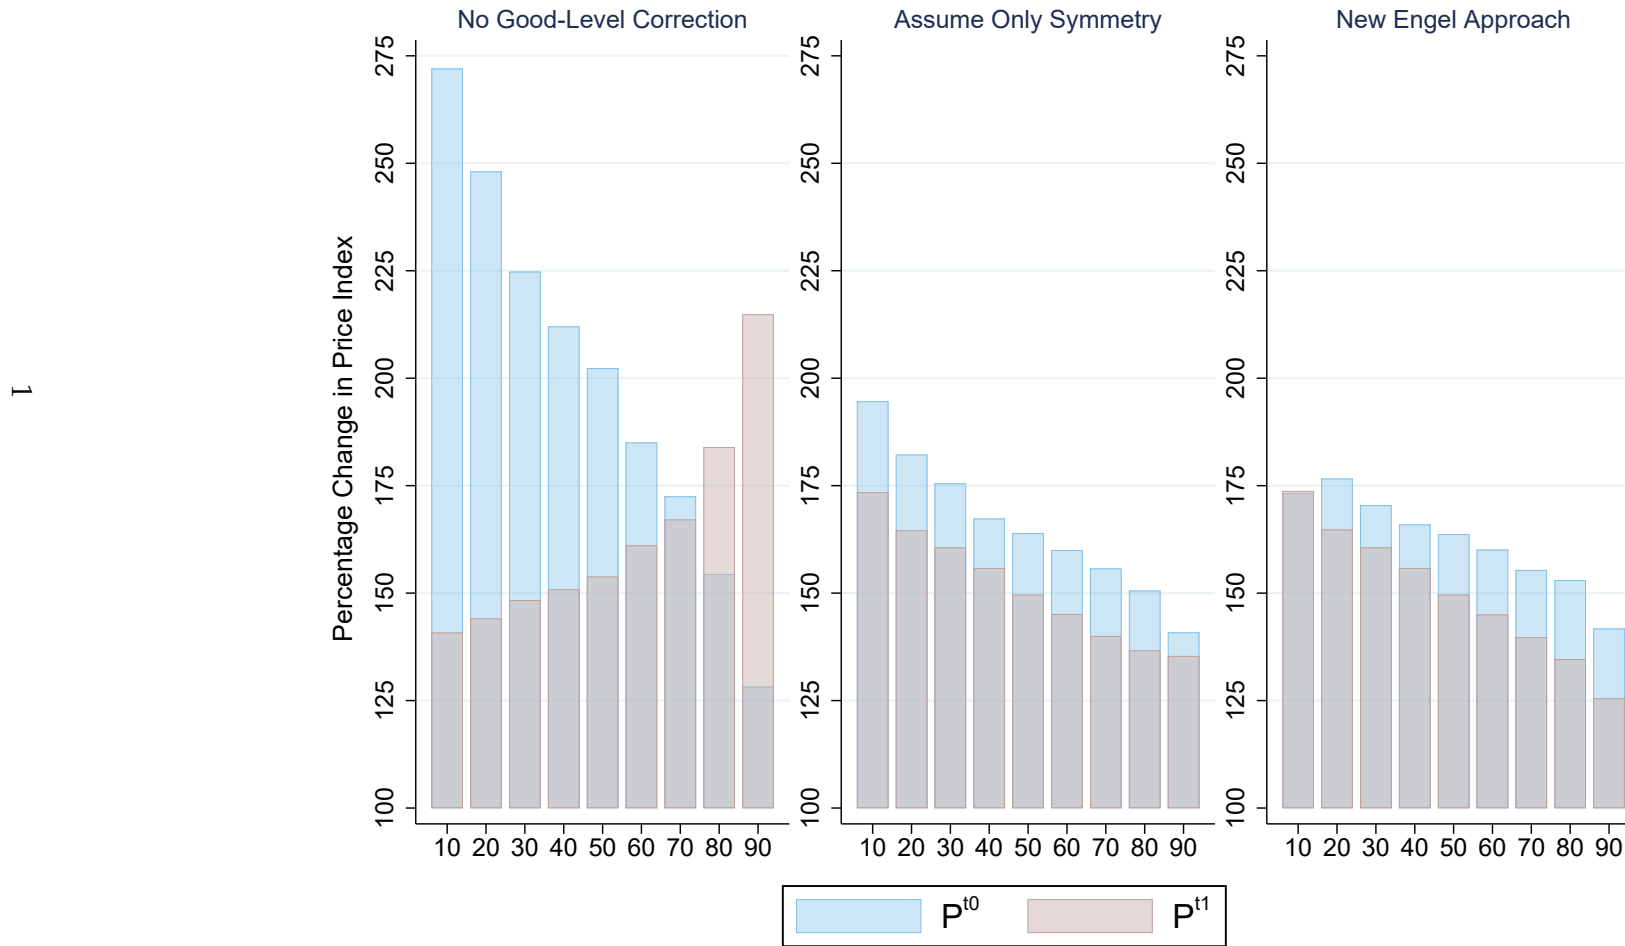
	P0	P0	P1	P1
	Dep. var: $\Delta \log$ Engel Price Index (bias corrected)			
$\Delta \log(\text{Price_Fuel}) * \text{Exp_Share_Fuel}$	1.062*** (0.306)	0.971* (0.531)	1.186*** (0.308)	0.603 (0.518)
ρ -value test $\beta=1$	0.84	0.96	0.55	0.44
Obs	2319	2319	2379	2379
R2	0.015	0.60	0.021	0.68
District FE	No	Yes	No	Yes
Income decile FE	Yes	Yes	Yes	Yes

Notes: The dependent variable in the table is the log change in the bias-corrected new Engel price index estimate obtained using food items only. The explanatory variable is the log change in the price of fuels (obtained using mean district-level expenditure shares) multiplied by the district by decile expenditure share on fuels. The first row of the bottom panel reports the p -value on the test of the coefficient of interest being equal to 1, as required by the quasi-separability test described in subsection 5.2. Regressions are weighted using district weights. Robust standard errors in parenthesis.

Appendices

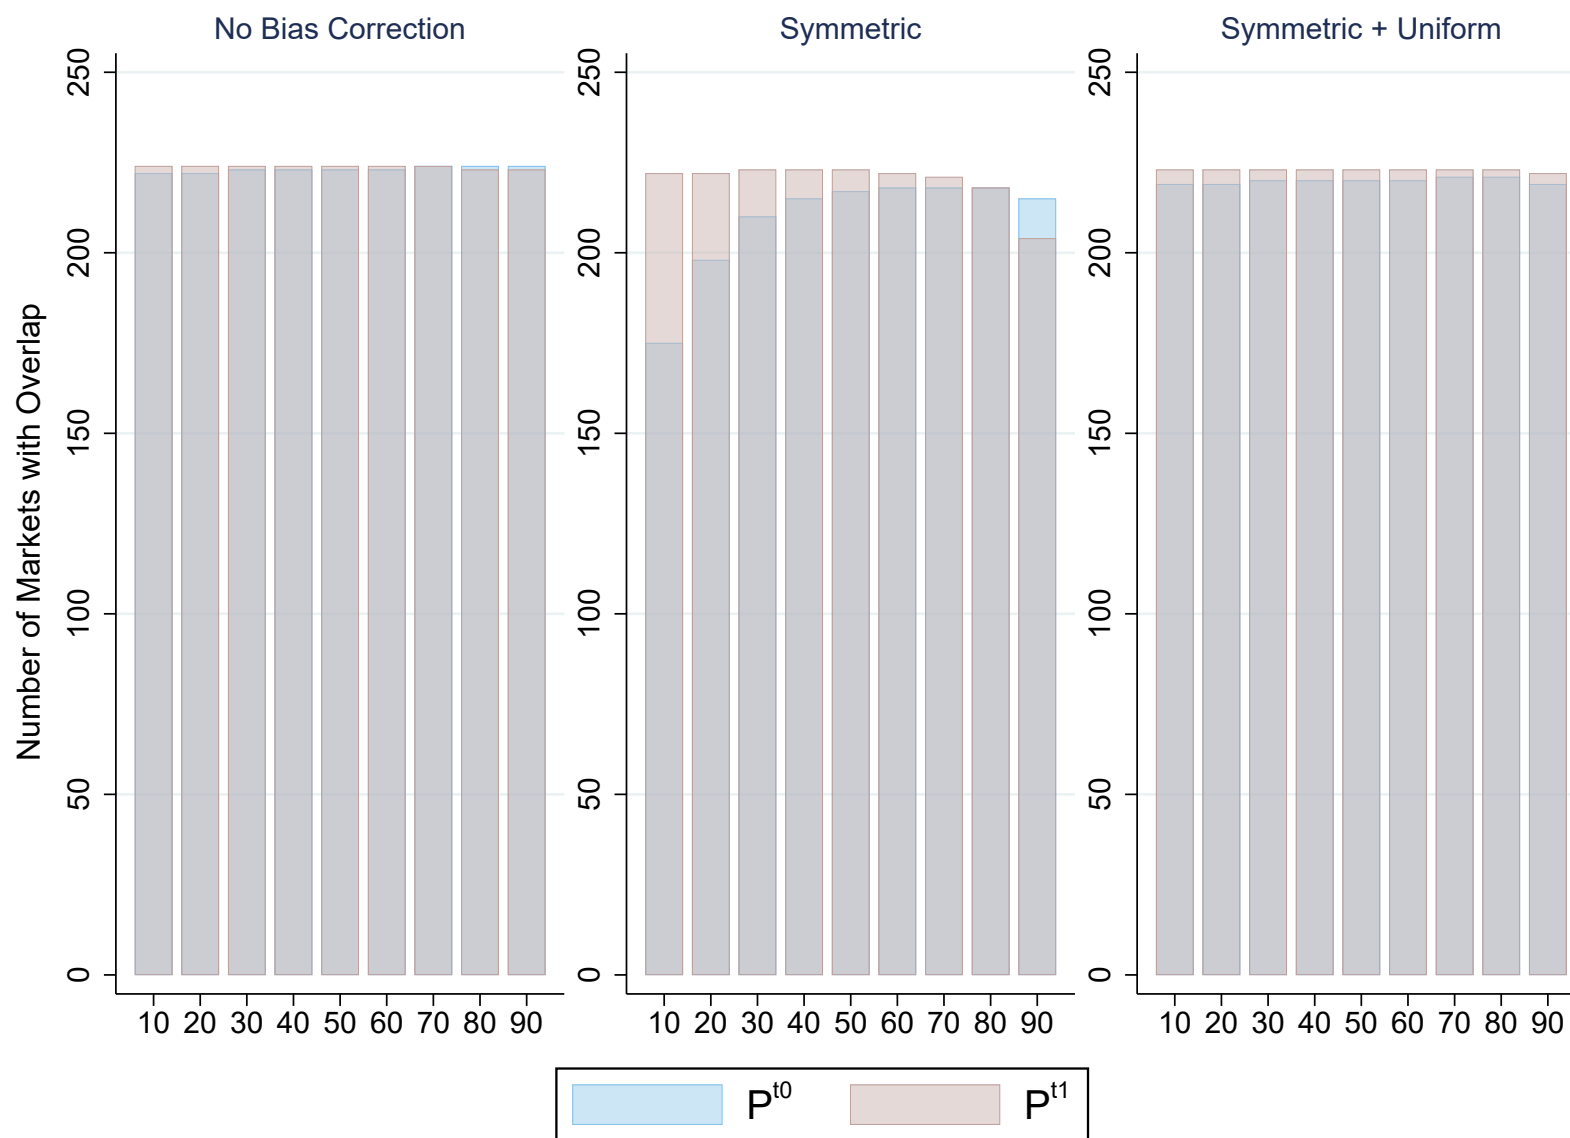
A Additional Figures and Tables

Figure A.1: Good-Level Selection Corrections (1): Price Index Changes With and Without Bias Correction



Notes: See Section 6 for discussion.

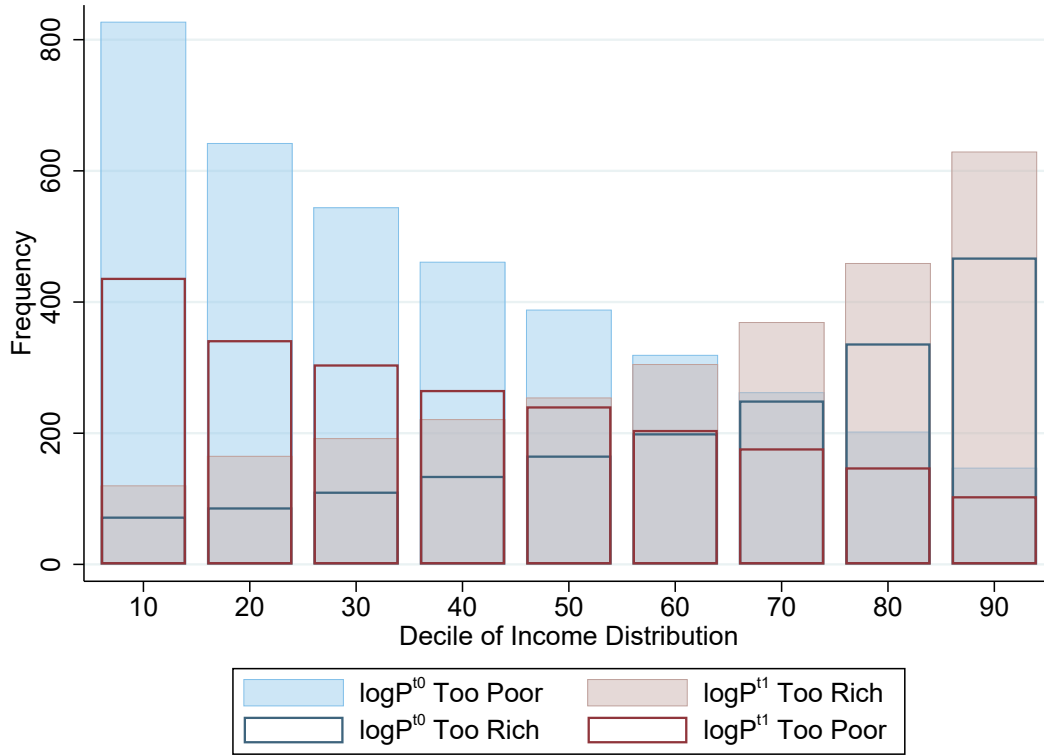
Figure A.2: Good-Level Selection Corrections (2): Number of Markets With and Without Bias Correction



2

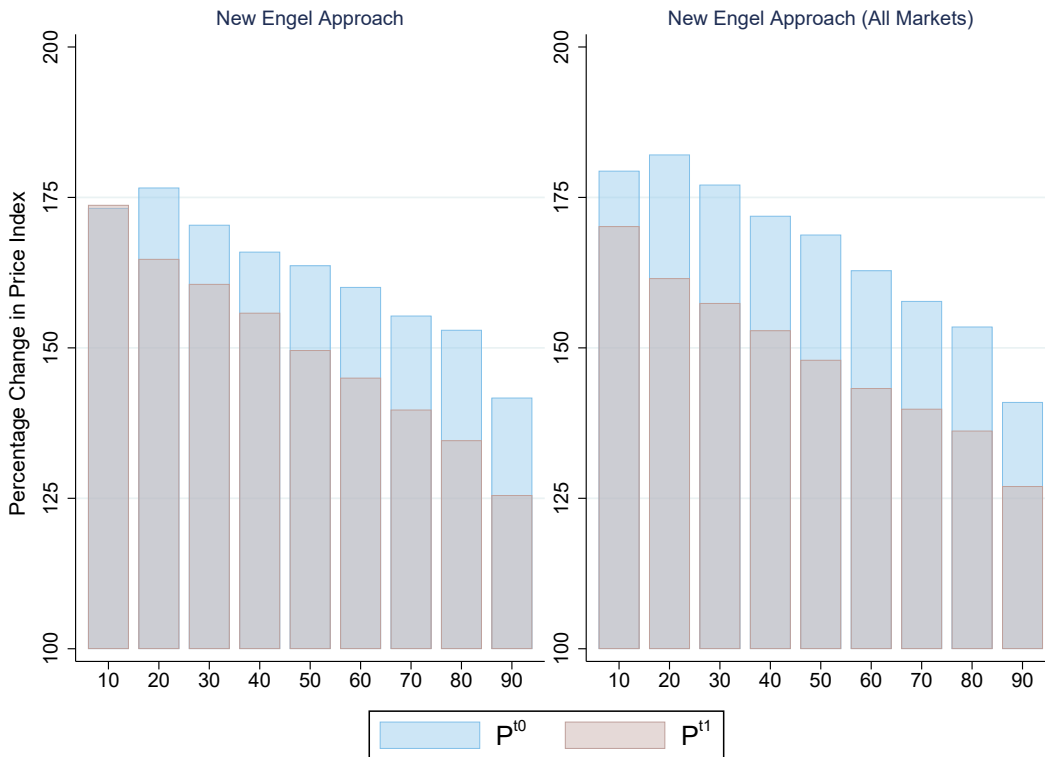
Notes: See Section 6 for discussion.

Figure A.3: Good-Level Selection Corrections (3): Reasons for Non-Overlap



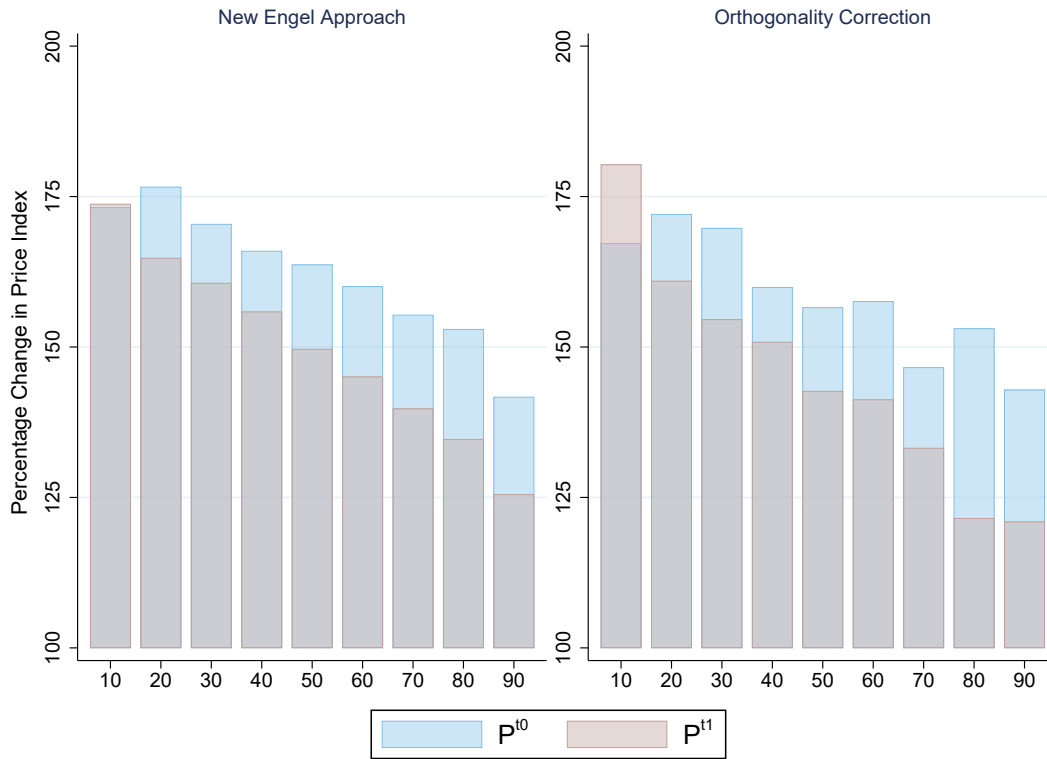
Notes: See Section 6 for discussion.

Figure A.4: Using All Markets (Including Markets <100 Hholds)



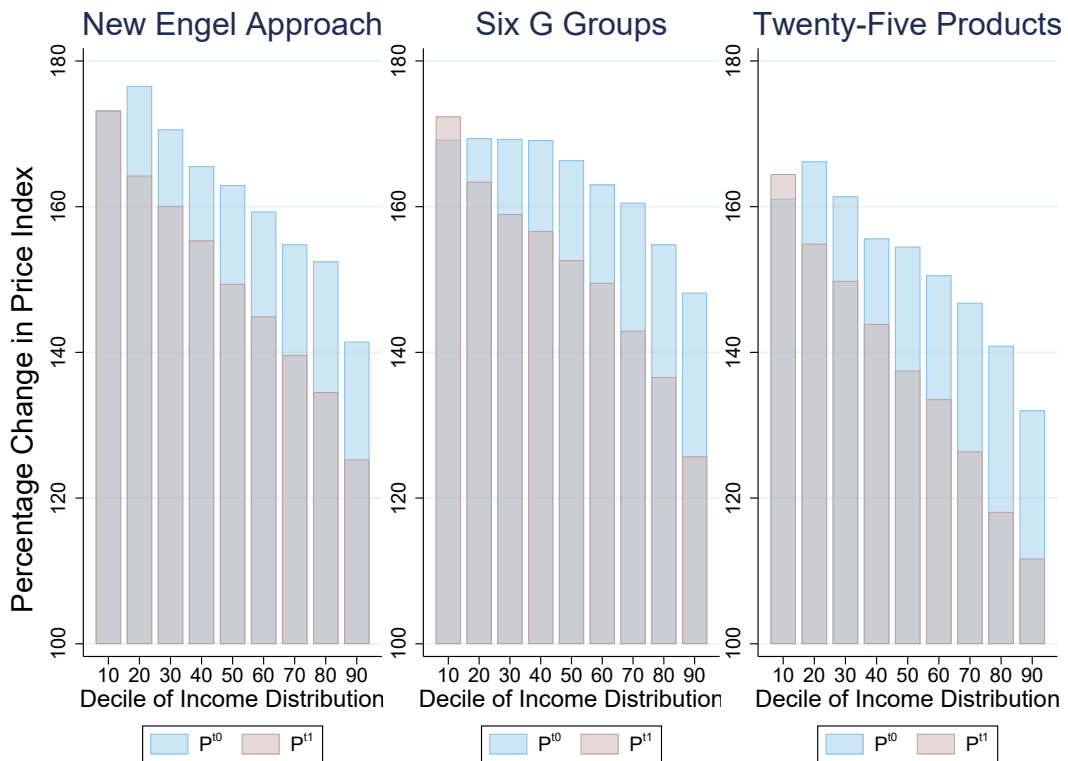
Notes: See Section 6 for discussion.

Figure A.5: Applying Orthogonality Correction for Confounding Price Effects ($\sigma = 4$)



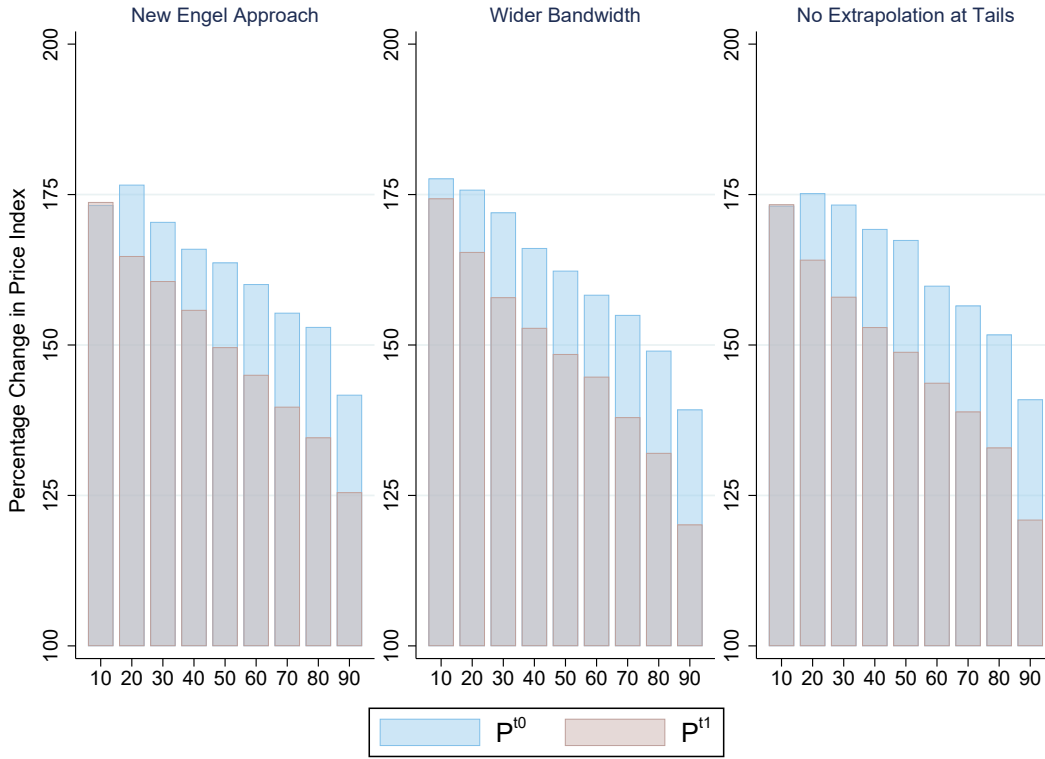
Notes: See Section 6 for discussion.

Figure A.6: Alternative i and G groups



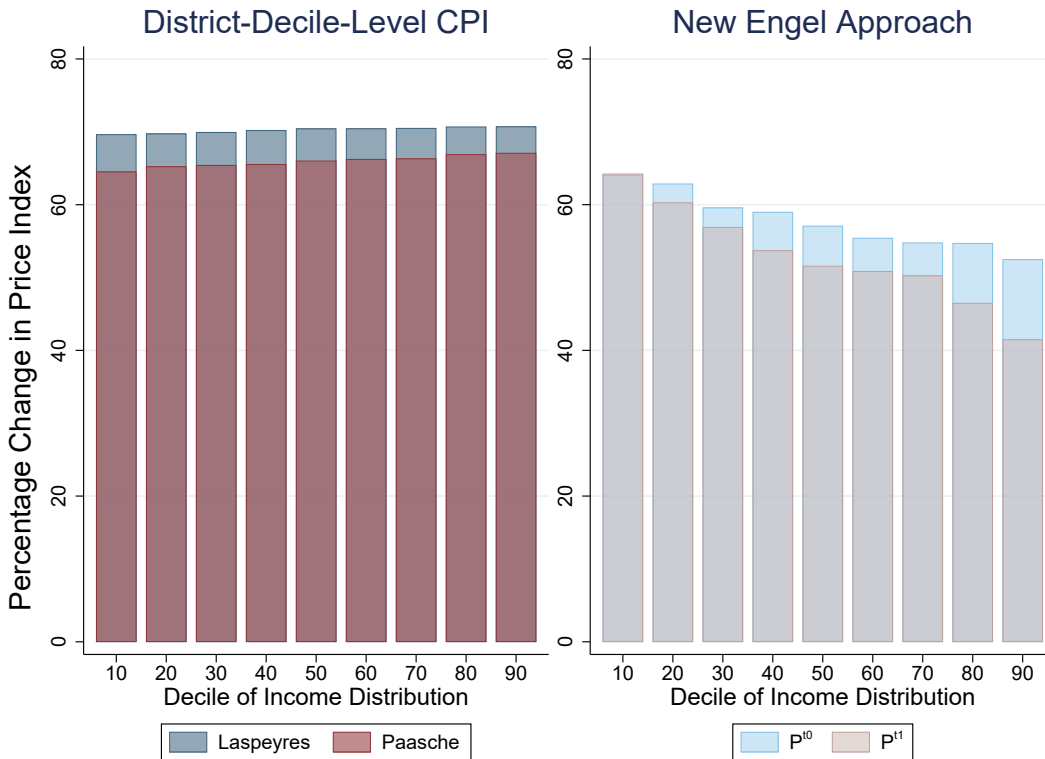
Notes: See Section 6 for discussion.

Figure A.7: Alternative Estimates of Relative Engel Curves



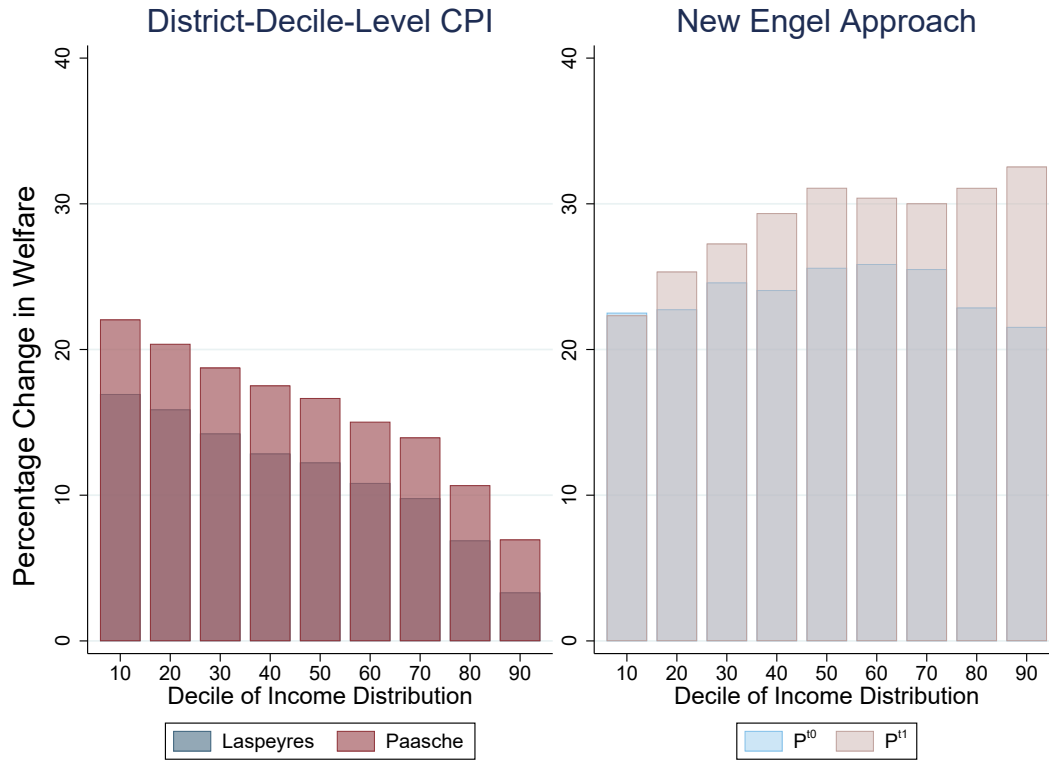
Notes: See Section 6 for discussion.

Figure A.8: Recall Bias: Inflation 1987/88-1994/95



Notes: See Section 6 for discussion.

Figure A.9: Recall Bias: Welfare 1987/88-1994/95



Notes: See Section 6 for discussion.

B Theory Appendix

B.1 Proofs

Proof of Lemma 1

Denote $q_i(p^t, y_h^t)$ the Marshallian demand for good i , function of prices p^t at time t and household h income y_h^t . Denote $E_i(p^t, y) = p_i q_i(p, y) / y$ the Engel curve for good i as a function of income y for a given set of prices p_t , and denote $V(p^t, y_h^t)$ the indirect utility function. In Proposition 1, the key property that we exploit is that q_i , E_i and V are all homogeneous of degree zero in p, y .

- i) Define the price index relative to prices in period $t0$ implicitly as $P^{t1}(p^{t0}, p^{t1}, y^{t1})$ such that: $V(p^{t1}, y_h^{t1}) = V(p^{t0}, \frac{y_h^{t1}}{P^{t1}})$. With the homogeneous change in prices $p^{t1} = \lambda p^{t0}$, it is immediate to verify that $P^{t1} = \lambda$ given that indirect utility is homogeneous of degree zero:

$$V(p^{t1}, y_h^{t1}) = V(\lambda p^{t0}, y_h^{t1}) = V(p^{t0}, \frac{y_h^{t1}}{\lambda})$$

Similarly, define the price index relative to prices in period $t1$ implicitly as $P^{t0}(p^{t0}, p^{t1}, y^{t0})$ such that: $V(p^{t0}, y_h^{t0}) = V(p^{t1}, \frac{y_h^{t0}}{P^{t0}})$. With the homogeneous change in prices $p^{t1} = \lambda p^{t0}$, it is again immediate to verify that $P^{t0} = 1/\lambda$.

Next, using $P^{t1} = \lambda$, we can then check that:

$$E_i(p^{t0}, \frac{y_h^{t1}}{P^{t1}}) = \frac{p_i^{t0} q_i(p^{t0}, y_h^{t1} / P^{t1})}{(y_h^{t1} / P^{t1})} = \frac{\lambda p_i^{t0} q_i(p^{t0}, y^{t1} / \lambda)}{y^{t1}} = \frac{p_i^{t1} q_i(\lambda p^{t0}, y^{t1})}{y^{t1}} = \frac{p_i^{t1} q_i(p^{t1}, y_h^{t1})}{y_h^{t1}} = E_i(p^{t1}, y^{t1})$$

hence the shift (in log) of the Engel curve from period $t0$ to period $t1$ corresponds to the price index change $\log P^{t1}$. Conversely, by switching $t0$ and $t1$, we can prove in the same manner that the shift (in log) of the Engel curve from period $t1$ to period $t0$ corresponds to the price index change $\log P^{t0}$. This proves assertion i).

- ii) Compensating variations CV_h are implicitly defined as $V(p^{t1}, y^{t1} + CV_h) = V(p^{t0}, y_h^{t0}) = U_h^{t0}$. With the homogeneous change in prices $p^{t1} = \lambda p^{t0}$, we can verify that compensating variations CV_h are such that $y_h^{t1} + CV_h = \lambda y^{t0}$:

$$V(p^{t1}, y^{t1} + CV_h) = V(p^{t0}, y^{t0}) = V(p^{t1} / \lambda, y^{t0}) = V(p^{t1}, \lambda y_h^{t0})$$

Next, we can then check that:

$$E_i(p^{t1}, y_h^{t1} + CV_h) = \frac{p_i^{t1} q_i(p^{t1}, y_h^{t1} + CV_h)}{y_h^{t1} + CV_h} = \frac{p_i^{t1} q_i(p^{t1}, \lambda y_h^{t0})}{\lambda y_h^{t0}} = \frac{p_i^{t0} q_i(p^{t1} / \lambda, y_h^{t0})}{y_h^{t0}} = \frac{p_i^{t0} q_i(p^{t0}, y_h^{t0})}{y_h^{t0}} = E_i(p^{t0}, y_h^{t0})$$

hence the initial observed expenditure share $p_i^{t0} q_{ih}^{t1} / y_h^{t1}$ of good i in period $t0$ corresponds to the counterfactual expenditure share of good i at new prices and total outlays $y_h^{t1} + EV_h$. This is assertion ii).

- iii) Equivalent variations EV_h are implicitly defined as $V(p^{t0}, y^{t0} + EV_h) = V(p^{t1}, y^{t1}) = U_h^{t1}$. For EV_h the proof proceeds the same way as for CV_h just by swapping periods $t0$ and $t1$ (and $1/\lambda$ instead of λ).

With the homogeneous change in prices $p^{t1} = \lambda p^{t0}$, we can verify that equivalent variations EV_h are such that $y_h^{t0} + EV_h = y^{t1} / \lambda$:

$$V(p^{t0}, y^{t0} + EV_h) = V(p^{t1}, y^{t1}) = V(\lambda p^{t0}, y^{t1}) = V(p^{t0}, y_h^{t1} / \lambda)$$

Again we can then check that:

$$E_i(p^{t0}, y_h^{t0} + EV_h) = \frac{p_i^{t0} q_i(p^{t0}, y_h^{t0} + EV_h)}{y_h^{t0} + EV_h} = \frac{\lambda p_i^{t0} q_i(p^{t0}, y^{t1} / \lambda)}{y^{t1}} = \frac{p_i^{t1} q_i(\lambda p, y^{t1})}{y^{t1}} = \frac{p_i^{t1} q_i(p^{t1}, y_h^{t1})}{y_h^{t1}} = E_i(p^{t1}, y^{t1})$$

hence the new observed expenditure share $p_i^{t1} q_{ih}^{t1} / y_h^{t1}$ of good i corresponds to the counterfactual expenditure share of good i at former prices at $y_h^{t0} + EV_h$. This proves assertion iii).

Proof of Lemma 2

Suppose that for a certain good i the shift of the Engel curve $E_i(p^t, y_h^t)$ (expenditure share x_{ih}^t / y_h^t plotted against total outlays y_h^t) reflects the price index change for any realization of price changes across periods and any y , i.e. $E_i(p^{t1}, y) = E_i(p^{t0}, y / P^{t1}(y))$. We know already from Proposition 1 that this is true for any preferences if we impose the price changes to be uniform across goods: $p^{t1} = \lambda p^{t0}$. Here we show that:

- **Step 1:** the expenditure share x_{ih} / y_h does not depend on prices, conditional on utility.

- **Step 2:** this expenditure share x_{ih} / y_h does not depend on utility either (i.e. the utility function has a Cobb-Douglas upper tier in i vs. non- i).

Step 1. Stating that the shifts in the Engel curve reflect the price index change means more formally that for any income level y_h^{t1} :

$$E_i(p^{t1}, y_h^{t1}) = E_i(p^{t0}, y_h^{t1} / P^{t1}(y_h^{t1})) \quad (\text{A.1})$$

where $P^{t1}(y_h^{t1})$ is the price index change transforming income at period $t1$ prices to income in $t0$ prices. By definition, the price index change P^{t1} is such that $V(p^{t1}, y_h^{t1}) = V(p^{t0}, y_h^{t1} / P^{t1})$ where V denotes the indirect utility function. An equivalent characterization of the price index is:

$$\frac{y_h^{t1}}{P^{t1}(y_h^{t1})} = e(V(p^{t1}, y_h^{t1}), p^{t0}) = e(U_h^{t1}, p^{t0})$$

using the expenditure function e , denoting utility in period $t1$ by U_h^{t1} . Looking at the share good i in total expenditures and imposing that Engel curves satisfy condition A.1, we can see that it no longer depends on prices p^{t1} once we condition on utility U_h^{t1} :

$$\frac{x_{ih}}{y_h} = E_i(p^{t1}, y_h^{t1}) = E_i\left(p^{t0}, \frac{y_h^{t1}}{P^{t1}(y_h^{t1})}\right) = E_i(p^{t0}, e(U_h^{t1}, p^{t0}))$$

(note that the expenditure share at time $t1$ is independent of prices p^{t0} in another period).

Step 2. So from now on, denote by $w_i(U)$ the expenditure share of good i as a function of utility. Let us also drop the time superscripts for the sake of exposition. Here in step 2 we show that w_i must be constant for demand to be rational.

Suppose that relative prices remain unchanged among other goods $j \neq i$ – but relative prices still vary across i and others. Using the composite commodity theorem (applied to non- i goods), the corresponding demand for i vs. non- i goods should correspond to a rational demand system in two goods. Hence we will do as if there is only one good $j \neq i$ aside from i . We will denote by p_j the price of this other good composite j .

A key (although trivial) implication of adding up properties is that the share of good j in expenditure is given by $1 - w_i(U)$ and only depends on utility. Denote by $e(p, U)$ the aggregate expenditure function. Shephard's Lemma implies:

$$\frac{\partial \log e(p, U)}{\partial \log p_i} = w_i(U) \quad , \quad \frac{\partial \log e(p, U)}{\partial \log p_j} = 1 - w_i(U)$$

Hence, conditional on utility U , the expenditure function is log-linear in log prices. Integrating, we get:

$$\begin{aligned} \log e(p, U) &= \log e_0(U) + w_i(U) \log p_i + (1 - w_i(U)) \log p_j \\ &= \log e_0(U) + w_i(U) \log(p_i / p_j) + \log p_j \end{aligned}$$

This must hold for any relative prices. Yet, the expenditure function must also increase with utility, conditional on any prices. Suppose by contradiction that there exist $U' > U$ such that $w_i(U') \neq w_i(U)$. We can then find $\log(p_i / p_j)$ such that:

$$\log e_0(U) - \log e_0(U') > [w_i(U') - w_i(U)] \log(p_i / p_j)$$

which implies:

$$\log e_0(U) + w_i(U) \log(p_i/p_j) > \log e_0(U') + w_i(U') \log(p_i/p_j)$$

which contradicts the monotonicity of the expenditure function in U . Hence w_i is constant and we have a Cobb-Douglas expenditure function with a constant exponent, leading to a flat Engel curve for good i .

Proof of Proposition 1

Proposition 4 states that quasi-separability in group G is a necessary and sufficient condition for the shifts in within- G Engel curves to exactly reflect price index changes when relative prices do not change within group G . The proof that quasi-separability is a necessary condition relies on part i) of Proposition 5 that we prove in the next section.

Quasi-separability as a sufficient condition. In brief, the intuition is that, thanks to the quasi-separability assumption, relative expenditures in i within group G only depend on the level of utility and within-group relative prices (we hold the latter constant). The first step is to show that quasi-separability implies a relationship as in condition i) of Proposition 5.

Quasi-separability in G implies that the expenditure function can be written:

$$e(p, U) = \tilde{e}(\tilde{P}_G(p_G, U), p_{NG}, U)$$

using Shephard's Lemma we obtain that compensated (Hicksian) demand for two goods $i \in G$ is:

$$h_i(p, U) = \frac{\partial e(p, U)}{\partial p_i} = \frac{\partial \tilde{e}(p, U)}{\partial P_G} \frac{\partial \tilde{P}_G(p_G, U)}{\partial p_i}$$

Taking the sum across goods in G , multiplying by prices and using the assumption that P_G is homogeneous of degree one: $\tilde{P}_G = \sum_i p_i \frac{\partial \tilde{P}_G(p_G, U)}{\partial p_i}$ (Euler's identity), we obtain:

$$\sum_{i \in G} p_i h_i(p, U) = \frac{\partial \tilde{e}(p, U)}{\partial P_G} \sum_i p_i \frac{\partial \tilde{P}_G(p_G, U)}{\partial p_i} = \frac{\partial \tilde{e}(p, U)}{\partial P_G} \tilde{P}_G$$

Looking at relative expenditures in i within group G , we get:

$$\frac{x_i}{x_G} = \frac{p_i h_i(p, U)}{\sum_{j \in G} p_j h_j(p, U)} = \frac{\partial \log \tilde{P}_G(p_G, U)}{\partial \log p_i} \equiv H_{iG}(p_G, U) \quad (\text{A.2})$$

i.e. the expenditure share of good 1 within G depends only on utility u and the vector of prices p_G of goods that belong to group G . Note that compensated demand is homogeneous of degree zero in prices. Hence, using our assumption that relative prices remain constant: $p'_G = \lambda_G p_G$ across the goods of group G , we obtain:

$$H_{iG}(p'_G, U') = H_{iG}(p_G, U')$$

For a consumer at initial utility u , income y and price p , notice that:

$$E_{iG}(p, y) = H_{iG}(p_G, U)$$

Denoting indirect utility by $V(p, y)$, we obtain the key identity behind Proposition 1:

$$H_{iG}(p_G, V(p, y)) = E_{iG}(p, y) \quad (\text{A.3})$$

which holds for any income y (and also any price p and subvector p_G).

Using this equality, we can now obtain each subpart i), ii) and iii) of Proposition 4 on Engel curves:

- i) For part i), define $P^{t1}(p^{t0}, p^{t1}, y_h^{t1})$ the exact price index change at income y_h^{t1} for household h , implicitly defined such that $V(p^{t0}, y^{t1}/P^{t1}) = V(p^{t1}, y_h^{t1})$ where V is the indirect utility function. Using equality (A.3) and the assumption that relative prices remain constant within G : $p_G^{t1} = \lambda_G p_G^{t0}$, We

obtain:

$$\begin{aligned}
E_{iG}(p^{t0}, y^{t1} / P^{t1}(p^{t0}, p^{t1}, y_h^{t1})) &= H_{iG}(p_G^{t0}, V(p^{t0}, y_h^{t1} / P^{t1}(p^{t1}, p^{t0}, y_h^{t1}))) \\
&= H_{iG}(p_G^{t0}, V(p^{t1}, y^{t1})) \\
&= H_{iG}(p_G^{t1}, V(p^{t1}, y^{t1})) \\
&= E_{iG}(p^{t1}, y_h^{t1})
\end{aligned}$$

where we go from the second to third line by noticing that H_{iG} is homogeneous of degree zero in prices (and $p_G^{t1} = \lambda_G p_G^{t0}$). By switching time superscripts $t1$ and $t0$, we prove a similar equality using the other price index $P^{t0}(p^{t0}, p^{t1}, y_h^{t0})$:

$$E_{iG}(p^{t1}, y^{t0} / P^{t0}(p^{t0}, p^{t1}, y_h^{t0})) = E_{iG}(p^{t0}, y_h^{t0})$$

The shift from one to the other Engel curve is given by each price index (which may vary across income levels y_h), from period $t0$ to $t1$ and from $t1$ to $t0$.

ii) By definition, compensating variations CV_h satisfy:

$$V(p^{t1}, y_h^{t1} + CV_h) = V(p^{t0}, y^{t0}) = U_h^{t0}$$

where U_h^{t0} denotes the utility level of household h in period $t0$. With the definition of CV_h and the homogeneity of function H_{iG} described above, we obtain that CV_h satisfies:

$$\begin{aligned}
E_{iG}(p^{t1}, y_h^{t1} + CV_h) &= H_{iG}(p_G^{t1}, V(p^{t1}, y_h^{t1} + CV_h)) \\
&= H_{iG}(p_G^{t1}, U_h^{t0}) \\
&= H_{iG}(p_G^{t0}, U_h^{t0}) \\
&= x_{ih}^{t0} / x_{iG}^{t1}
\end{aligned}$$

where the last term refers to the within-group G expenditure share of good i in period $t0$. This proves part ii) of Proposition 4.

iii) Similarly, by definition, equivalent variations EV satisfy:

$$V(p^{t0}, y_h^{t0} + EV_h) = V(p^{t1}, y_h^{t1}) = U_h^{t1}$$

where U_h^{t1} denotes to the period $t1$ utility level of household h .

With the definition of EV_h and the homogeneity of function H_{iG} , we obtain that EV_h satisfies:

$$\begin{aligned}
E_{iG}(p^{t0}, y_h^{t0} + EV_h) &= H_{iG}(p_G^{t0}, V(p^{t0}, y_h^{t0} + EV_h)) \\
&= H_{iG}(p_G^{t0}, U_h^{t1}) \\
&= H_{iG}(p_G^{t1}, U_h^{t1}) \\
&= x_{hi}^{t1} / x_{Gh}^{t1}
\end{aligned}$$

where the last term refers to the within-group G expenditure share of good i in period $t1$.

Quasi-separability as a necessary condition. The proof starts with the same argument as in Lemma 2: for the shifts in Engel curves to reflect the changes in price indexes, we need within- G expenditure shares to depend only on utility and relative prices within group G . In a second step, we use part i) of Lemma 3 (proven in the following appendix section) to obtain that quasi-separability is required.

Stating that the shifts in relative Engel curve reflect the price index change means more formally that for any income level y_h^{t1} :

$$E_{iG}(p^{t1}, y_h^{t1}) = E_{iG}(p^{t0}, y_h^{t1} / P^{t1}(y_h^{t1})) \quad (\text{A.4})$$

where $P^{t1}(y_h^{t1})$ is the price index change transforming income at period $t1$ prices to income in $t0$ prices. By definition of the price index, P^{t1} is such that $V(p^{t1}, y_h^{t1}) = V(p^{t0}, y_h^{t1} / P^{t1})$ where V denotes the indirect

utility function. Or equivalently:

$$\frac{y_h^{t1}}{P^{t1}(y_h^{t1})} = e(V(p^{t1}, y_h^{t1}), p^{t0}) = e(U_h^{t1}, p^{t0})$$

using the expenditure function e , where we denote utility in period $t1$ by U_h^{t1} . Looking at the share good i in expenditures within group G , and imposing that Engel curves satisfy condition A.4, we can see that it no longer depends on prices p^{t1} once we condition on utility U_h^{t1} :

$$\frac{x_{ih}}{y_h} = E_i(p^{t1}, y_h^{t1}) = E_i\left(p^{t0}, \frac{y_h^{t1}}{P^{t1}(y_h^{t1})}\right) = E_G(p^{t0}, e(U_h^{t1}, p^{t0}))$$

Note that the expenditure share at time $t1$ is independent of prices p^{t0} in another period. Hence there exists a function H_{iG} of within-G relative prices and utility such that:

$$\frac{x_{ih}}{y_h} = H_{iG}(p_G, U_h)$$

This is condition i) of Proposition 5. This condition implies quasi-separability in G , as shown below in the proof of Proposition 5. Hence quasi-separability in G is required if we want the shifts in relative Engel curves to reflect the changes in price indexes.

Proof of Lemma 3

Gorman (1970) and Deaton and Muellbauer (1980) have already provided a proof of the equivalence between quasi-separability and ii), using the distance function. Here for convenience we provide a proof without referring to the distance function.

Blackorby, Primont and Russell (1978), theorem 3.4) show the equivalent between quasi-separability (which they refer to as separability in the cost function) and i). The proof that we provide here is more simple and relies on similar argument as Goldman and Uzawa (1964) about the separability of the utility function.

In the proofs below, we drop the household subscripts and time superscripts to lighten the notation.

• **Quasi-separability implies i).** Actually we have already shown that quasi-separability implies i). In the proof of Proposition 4 above, we have shown in equation (A.2) that we have:

$$\frac{x_i}{x_G} = H_{iG}(p_G, U) = \frac{\partial \log \tilde{P}_G}{\partial \log p_i}$$

if the expenditure function can be written as $e(p, U) = \tilde{e}(\tilde{P}_G(p_G, U), p_{NG}, U)$ where $\tilde{P}_G(p_G, U)$ is homogeneous of degree one in the prices p_G of goods in G .

The most difficult part of the proof Proposition 5 is to show that condition i) leads to quasi-separability:

• **i) implies quasi-separability.**

Let us assume (condition i) that the within-group expenditure share of each good $i \in G$ does not depend on the price of non-G goods:

$$\frac{p_i h_i(p, U)}{x_G(p, U)} = H_{iG}(p_G, U)$$

where $h_i(p, U)$ is the compensated demand and $x_G(p, U) = \sum_{j \in G} p_j h_j(p, U)$ is total expenditure in goods of groups G . As a first step, we would like to construct a scalar function $\tilde{P}_G(p_G, U)$ such that:

$$\frac{\partial \log \tilde{P}_G}{\partial p_i} = \frac{1}{p_i} H_{iG}(p_G, U) \tag{A.5}$$

for each i , and $\tilde{P}_G(p_{G0}, U) = 1$ for some reference set of prices p_{G0} . Thanks to the Frobenius Theorem used notably for the integrability theorem of Hurwicz and Uzawa (1971), we know that such problem admits a solution if $\frac{\partial(H_i/p_i)}{\partial p_j} = \frac{\partial(H_j/p_j)}{\partial p_i}$. We need to check that this term is indeed symmetric for any two goods i

and j in group G :

$$\begin{aligned}
\frac{\partial(H_i/p_i)}{\partial p_j} &= \frac{\partial(h_i/x_G)}{\partial p_j} \\
&= \frac{1}{x_G} \frac{\partial h_i}{\partial p_j} - \frac{h_i}{x_G^2} \frac{\partial x_G}{\partial p_j} \\
&= \frac{1}{x_G} \frac{\partial h_i}{\partial p_j} - \frac{h_i}{x_G^2} \left[h_j + \sum_{g \in G} p_g \frac{\partial h_g}{\partial p_j} \right] \\
&= \frac{1}{x_G} \frac{\partial h_i}{\partial p_j} - \frac{h_i}{x_G^2} \sum_{g \in G} p_g \frac{\partial h_j}{\partial p_g} - \frac{h_i h_j}{x_G^2}
\end{aligned}$$

where the last line is obtained by using the symmetry of the Slutsky matrix: $\frac{\partial h_i}{\partial p_j} = \frac{\partial h_j}{\partial p_i}$ for any i, j . Using the homogeneity of degree zero of the compensated demand w.r.t prices, we get: $\sum_{g \in G} p_g \frac{\partial h_i}{\partial p_g} = -\sum_{k \notin G} p_k \frac{\partial h_i}{\partial p_k}$ and thus:

$$\begin{aligned}
\frac{\partial(H_i/p_i)}{\partial p_j} &= \frac{1}{x_G} \frac{\partial h_i}{\partial p_j} - \frac{h_i}{x_G^2} \sum_{g \in G} p_g \frac{\partial h_j}{\partial p_g} - \frac{h_i h_j}{x_G^2} \\
&= \frac{1}{x_G} \frac{\partial h_i}{\partial p_j} + \frac{h_i}{x_G^2} \sum_{k \notin G} p_k \frac{\partial h_j}{\partial p_k} - \frac{h_i h_j}{x_G^2}
\end{aligned}$$

Given the symmetry of the Slutsky matrix, the first term is symmetric in i and j , so is the third term. Using the assumption that $\frac{h_i}{h_j}$ does not depend on the price of non- G goods for any couple of goods $i, j \in G$ and $k \notin G$, we also obtain that the second term is symmetric in i and j : $h_i \frac{\partial h_j}{\partial p_k} = h_j \frac{\partial h_i}{\partial p_k}$. Hence:

$$\frac{\partial(H_i/p_i)}{\partial p_j} = \frac{\partial(H_j/p_j)}{\partial p_i}$$

and we can apply Frobenius theorem to find such a function \tilde{P}_G satisfying equation A.5.

Note that $\sum_{i \in G} H_i(p_G, U) = 1$ for any price vector p_G and utility U , hence \tilde{P}_G is homogeneous of degree one in p_G and can take any value in $(0, +\infty)$.

The second step of the proof is to show that the expenditure function depends on price vector p_G only through the scalar function $\tilde{P}_G(p_G, U)$. To do so, we use the same idea as in Lemma 1 of Goldman and Uzawa (1964).¹ Using our constructed $\tilde{P}_G(p_G, U)$, notice that:

$$\frac{\partial e}{\partial p_i} = \frac{\partial \tilde{P}_G}{\partial p_i} \cdot x_G(p, U) \tag{A.6}$$

Since this equality is valid for any $i \in G$ and any value of \tilde{P}_G , it must be that the expenditure function e remains invariant as long as \tilde{P}_G remains constant since the Jacobian of e w.r.t p_G is null whenever the Jacobian of \tilde{P}_G is null. Hence e can be expressed as a function of \tilde{P}_G , utility U and other prices:

$$e(p, U) = \tilde{e}(\tilde{P}_G(p_G, U), p_{NG}, U)$$

This concludes the proof that i) implies quasi-separability.

• ii) implies quasi-separability. Suppose that utility satisfies:

$$K(F_G(q_G, U), q_{NG}, U) = 1$$

¹Lemma 1 of Goldman and Uzawa (1964) states that if two multivariate functions f and g are such that $\frac{\partial f}{\partial x_i} = \lambda(x) \frac{\partial g}{\partial x_i}$ it must be that $f(x) = \Lambda(g(x))$ for some function Λ over connected sets of values taken by g .

Construct \tilde{P}_G as follows:

$$\tilde{P}_G(p_G, u) = \min_{q_G} \left\{ \sum_{i \in G} p_i q_i \mid F_G(q_G, U) = 1 \right\}$$

which is homogeneous of degree 1 in p_G . Denote by \tilde{e} the function of scalars P_G, U and price vectors p_{NG} :

$$\tilde{e}(P_G, p_{NG}, U) = \min_{Q_G, q_{NG}} \left\{ Q_G P_G + \sum_{i \notin G} p_i q_i \mid K(Q_G, q_{NG}, U) = 1 \right\}$$

The expenditure function then satisfies:

$$\begin{aligned} e(p, U) &= \min_{q_G, q_{NG}} \left\{ \sum_{i \in G} p_i q_i + \sum_{i \notin G} p_i q_i \mid K(F_G(q_G, U), q_{NG}, U) = 1 \right\} \\ &= \min_{q_G, Q_G, q_{NG}} \left\{ \sum_{i \in G} p_i q_i + \sum_{i \notin G} p_i q_i \mid F_G(q_G, U) = Q_G ; K(Q_G, q_{NG}, U) = 1 \right\} \\ &= \min_{q_G, Q_G, q_{NG}} \left\{ Q_G \sum_{i \in G} p_i q_i + \sum_{i \notin G} p_i q_i \mid F_G(q_G, U) = 1 ; K(Q_G, q_{NG}, U) = 1 \right\} \\ &= \min_{Q_G, q_{NG}} \left\{ Q_G \tilde{P}_G(p_G, U) + \sum_{i \notin G} p_i q_i \mid K(Q_G, q_{NG}, U) = 1 \right\} \\ &= \tilde{e}(\tilde{P}_G(p_G, U), p_{NG}, U) \end{aligned}$$

(going from the second to third lines uses the homogeneity of F_G) which proves that ii) implies quasi-separability.

• **Quasi-separability implies ii).** Now, assume that we have in hand two functions \tilde{P}_G (homogeneous of degree 1) and \tilde{e} that satisfies usual properties of expenditure functions. From these two functions, the goal is to:

- implicitly construct utility that satisfies ii)
- verify that $\tilde{e}(\tilde{P}_G(p_G, U), p_{NG}, U)$ is the expenditure function associated with it.

First, using these two functions, let us define:

$$K(Q_G, q_{NG}, U) = \min_{P_G^*, p_{NG}^*} \left\{ \frac{Q_G P_G^* + \sum_{i \notin G} p_i^* q_i}{\tilde{e}(P_G^*, p_{NG}^*, U)} \right\} \quad (\text{A.7})$$

and:

$$F_G(q_G, U) = \min_{p_G} \left\{ \frac{\sum_{i \in G} p_i^* q_i}{\tilde{P}_G(p_G^*, U)} \right\} \quad (\text{A.8})$$

Those functions are similar to distance functions introduced by Gorman (1970). We can also check that both F_G and K are homogeneous of degree one in q_G . For instance, we have for F_G :

$$F_G(\lambda q_G, U) = \min_{p_G} \left\{ \frac{\sum_{i \in G} \lambda p_i^* q_i}{\tilde{P}_G(p_G^*, U)} \right\} = \lambda \min_{p_G} \left\{ \frac{\sum_{i \in G} p_i^* q_i}{\tilde{P}_G(p_G^*, U)} \right\} = \lambda F_G(q_G, U)$$

If \tilde{e} and \tilde{P}_G are decreasing in U , we can see that F_G and K are decreasing in U , hence the following has a unique solution:

$$K(F_G(q_G, U), q_{NG}, U) = 1 \quad (\text{A.9})$$

Let us define utility implicitly as above. These implicitly defined preferences satisfy condition ii). The next step is to show that prices p^* that minimize the right-hand side of equations (A.7) and (A.8) also coincide with actual prices p . Then the final step is to show that the expenditure function coincides with $\tilde{e}(\tilde{P}_G(p_G, U), p_{NG}, U)$.

Utility maximization subject to the budget constraint and subject to constraint (A.9) leads to the following first-order conditions in q_i :

$$\begin{aligned}\mu \frac{\partial K}{\partial Q_G} \frac{\partial F_G}{\partial q_i} &= \lambda p_i & \text{if } i \in G \\ \mu \frac{\partial K}{\partial q_j} &= \lambda p_j & \text{if } j \in G\end{aligned}$$

where p are observed prices and where μ and λ are the Lagrange multipliers associated with (A.9) and the budget constraints respectively. Using the envelop theorem, these partial derivatives are:

$$\frac{\partial K}{\partial Q_G} = \frac{P_G^*}{\tilde{e}(P_G^*, p_{NG}^*, U)} \quad ; \quad \frac{\partial K}{\partial q_j} = \frac{p_j^*}{\tilde{e}(P_G^*, p_{NG}^*, U)} \quad ; \quad \frac{\partial F_G}{\partial q_i} = \frac{p_i^*}{\tilde{P}_G(p_G^*, U)}$$

where P_G^* and p_i^* refer to counterfactual prices that minimize the right-hand side of equations (A.7) and (A.8) that define K and F_G . Note that these counterfactual prices may potentially differ from observed prices, but we will see now that relative prices are the same. Combining the FOC and envelop theorem, we obtain:

$$\begin{aligned}\mu \frac{P_G^*}{\tilde{e}(P_G^*, p_{NG}^*, U)} \frac{p_i^*}{\tilde{P}_G(p_G^*, U)} &= \lambda p_i & \text{if } i \in G \\ \mu \frac{p_j^*}{\tilde{e}(P_G^*, p_{NG}^*, U)} &= \lambda p_j & \text{if } j \in G\end{aligned}$$

But notice that if p_i^* for $i \in G$ minimizes the right-hand side of equation (A.8), then $\lambda_G p_i^*$ also minimizes (A.8) since \tilde{P}_G is homogeneous of degree one.

With $\lambda_G = \frac{\mu}{\lambda} \frac{P_G^*}{\tilde{e}(P_G^*, p_{NG}^*, U)} \frac{1}{\tilde{P}_G(p_G^*, U)}$, it implies that we can have: $p_i^* = p_i$ for $i \in G$. Also notice that if P_G^* and p_j^* for $j \notin G$ minimize the right-hand side of equation (A.7), then $\lambda_N P_G^*$ and $\lambda_N p_j^*$ also minimizes (A.8) for any $\lambda_N > 0$ since \tilde{e} is homogeneous of degree one. With $\lambda_N = \frac{\mu}{\lambda \tilde{e}(P_G^*, p_{NG}^*, U)}$, we have $\lambda_N p_j^* = p_j$. Using the FOC for goods $j \notin G$, we obtain:

$$\frac{\mu}{\lambda} = \tilde{e}(\lambda_N P_G^*, p_{NG}, U)$$

In turn, the FOC for goods $i \in G$ yields:

$$\lambda_N P_G^* = \tilde{P}_G(p_G, U)$$

So we can also replace P_G^* by \tilde{P}_G .

Now that we have proven that observed prices are also solution of the minimization of (A.7) and (A.8), it is easy to show that $\tilde{e}(\tilde{P}_G(p_G, U), p_{NG}, U)$ is equal to the expenditure function associated with utility defined in equation (A.9). Using equations (A.9), (A.7) and (A.8), and the equality between p^* and p (as well as P_G^* and P_G), we find:

$$\begin{aligned}\tilde{e}(\tilde{P}_G(p_G, U), p_{NG}, U) &= F_G(q_G, U) P_G^* + \sum_{i \notin G} p_i^* q_i \\ &= F_G(q_G, U) P_G + \sum_{i \notin G} p_i q_i \\ &= \sum_{i \in G} p_i q_i + \sum_{i \notin G} p_i q_i\end{aligned}$$

where quantities are those maximizing utility subject to the budget constraint, therefore the expenditure function coincides with $\tilde{e}(\tilde{P}_G(p_G, U), p_{NG}, U)$. Once we know that observe price minimize (A.7) and (A.8), it is also easy to verify that the expenditure shares implied by utility defined in A.9 also correspond to expenditure shares implied by the expenditure function $\tilde{e}(\tilde{P}_G(p_G, U), p_{NG}, U)$.

This shows that utility defined by (A.9), (A.7) and (A.8) leads to the same demand system as $\tilde{e}(\tilde{P}_G(p_G, U), p_{NG}, U)$, and proves that quasi-separability implies condition ii).

Proof of Corollary 4(a)

Part i) of Proposition 2 shows that preferences are quasi-separable in G if and only if relative (compensated) expenditure shares x_i/x_G for any good $i \in G$ do not depend on the price of any good $j \notin G$ if we hold utility U constant:

$$\left. \frac{\partial \log(x_i/x_G)}{\partial \log p_j} \right|_U = 0$$

Instead, holding income constant (uncompensated), we obtain:

$$\left. \frac{\partial \log(x_i/x_G)}{\partial \log p_j} \right|_y = \frac{\partial \log(x_i/x_G)}{\partial \log U} \frac{\partial \log V}{\partial \log p_j} \quad (\text{A.10})$$

where V denotes the indirect utility function. Using Roy's identity (in terms of elasticities):

$$\frac{\partial \log V}{\partial \log p_j} = -\frac{p_j q_j}{y} \frac{\partial \log V}{\partial \log y}$$

and substituting into equation A.10, we obtain:

$$\left. \frac{\partial \log(x_i/x_G)}{\partial \log p_j} \right|_y = -\frac{p_j q_j}{y} \frac{\partial \log(x_i/x_G)}{\partial \log U} \frac{\partial \log V}{\partial \log y} \quad (\text{A.11})$$

where V is the indirect utility function. In turn, note that the elasticity of relative (uncompensated) expenditure share x_i/x_G w.r.t. income, holding prices constant, is:

$$\frac{\partial \log(x_i/x_G)}{\partial \log y} = \frac{\partial \log(x_i/x_G)}{\partial \log U} \frac{\partial \log V}{\partial \log y}$$

Substituting into equation A.11, we obtain our result which holds if and only if preferences are quasi-separable:

$$\left. \frac{\partial \log(x_i/x_G)}{\partial \log p_j} \right|_y = -\frac{p_j q_j}{y} \frac{\partial \log(x_i/x_G)}{\partial \log y}$$

Note that it is possible to provide an alternative proof using Slutsky decomposition for good i and compare to the sum of other goods $i' \in G$ to obtain Corollary 4(a) for relative expenditure shares.

C Data Appendix

[Work in progress.]