

Measuring Welfare and Inequality with Incomplete Price Information*

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

July 2023

Abstract

We propose and implement a new approach that allows us to estimate income-specific changes in household welfare in contexts where well-measured prices are not available for important subsets of consumption. Using rich but widely available expenditure survey microdata, we show that we can recover income-specific equivalent and compensating variations from horizontal shifts in what we term “relative Engel curves”—as long as preferences fall within the broad quasi-separable class (Gorman 1970; 1976). Our approach is flexible enough to allow for non-parametric estimation at each point of the income distribution. We apply the methodology to estimate inflation and welfare changes in rural India between 1987 and 2000. Our estimates reveal that lower rates of inflation for the rich erased the real income convergence found in the existing literature that uses the subset of consumption with well-measured prices to calculate inflation.

Keywords: Real income inequality, non-homothetic preferences, price indices, gains from trade.

JEL Classification: F63, O12, E31, D12.

*David Atkin and Benjamin 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, and Arnaud Costinot, Swati Dhingra, Dave Donaldson, Pablo Fajgelbaum, Ana Cecilia Fielor, Jessie Handbury, Ethan Ligon, Stephen Redding and David Weinstein for valuable discussions. Amir Hafiz, Amir Hisham, Tasnia Hussain, Jay Sayre, Yulu Tang, Thyra Tuttle, and Guan Yi Lin provided excellent research assistance. This paper subsumes an earlier version titled “A New Engel on Price Index and Welfare Estimation”.

[†]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

Measuring changes in household welfare is valuable in many contexts, both to evaluate the impacts of policies and to assess changes in well-being across time and space. Furthermore, given recent political upheaval and a renewed focus on inequality, there is increased urgency to capture not just average changes but the full distribution. But while we often have reliable data on changes in nominal income, measuring changes in the denominator of real income—the cost of living—requires detailed price information that are seldom, if ever, available.

A number of recent papers use rich consumption microdata to study income-group specific welfare changes: either under explicit non-homothetic preferences such as in Fajgelbaum and Khandelwal (2016), Handbury (2021) and Comin et al. (2021); or by allowing income groups to have different taste parameters as in Atkin et al. (2018), Jaravel (2019), and Argente and Lee (2020).¹ The dramatic increase in inflation experienced by many countries in the last year has further increased interest in calculating income-group specific inflation rates (e.g., see Jaravel and Lashkari, 2022 and Baqaee et al., 2022). Common to all these approaches is the requirement that the researcher has complete (quality- and variety-adjusted) price information. Such detail is paramount for distributional analysis since we know that different income groups consume very different bundles.

While sufficiently rich data on consumption prices and quantities are available for some countries and for some components of household welfare—e.g. US retail consumption using scanner microdata covering roughly 10 percent of consumption, or developing-country expenditure surveys on foods and fuels covering more than half of rural consumption—it is not feasible to collect such detailed data for the entire consumption basket. Accurately measuring prices, quality and variety for services (e.g. housing, healthcare and education) and differentiated manufactures (e.g. electronics) is particularly difficult. And even in the richest data environments, evaluating changes in welfare from observed price data typically still requires strong functional form assumptions (e.g. quality-adjusting prices or accounting for variety gains).

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 categories—to estimate changes in exact household price indices for the full consumption basket, as well as welfare, at every point of the income distribution. We then apply this approach to quantify changes in welfare over time for Indian households at different levels of income.

Our analysis proceeds in three steps. First, we develop the theory behind our approach. In environments with incomplete price information, recovering changes in the full price in-

¹Other recent work includes Borusyak and Jaravel (2018), Hottman and Monarch (2020), and Argente et al. (Forthcoming).

dex, and hence welfare, is not possible without restrictions on preferences. The cornerstone of our methodology is the insight that quasi-separable preferences, as defined by Gorman (1970), provide a natural and testable restriction that allows us to estimate income-specific welfare changes in the absence of complete price information. Quasi-separability requires that groups of goods or services G are separable in the expenditure function (not the utility function as under direct separability): $e(p, U) = \tilde{e}(e_G(p_G, U), p_{NG}, U)$, where p_G and p_{NG} denote the vector of prices for goods within and outside G , respectively, and U is household utility.² In our context, p_G is observable but p_{NG} may not be. While this property of quasi-separable demands makes them natural candidates for inferring welfare changes with incomplete price data, as we lay out below, we are not aware of prior work making such a connection (the primary use of these demands in the literature is to rationalize constructing a sub-group price index for group G).³

The power of this restriction on preferences is that, after conditioning on the observable prices p_G , horizontal shifts across time in what we call *relative Engel curves*—projections of relative expenditures $\frac{x_i}{x_G}$ for good i as a share of group G on log total per-capita household expenditure, $\log y$ —reveal the full price index that covers *all* household consumption. At the heart of this strong result is the fact that Hicksian relative expenditures are a function of the vector of within- G relative prices and the level of household utility but not prices outside of G , $\frac{x_i}{x_G} = H_{iG}(p_G, U)$, if and only if preferences are quasi-separable. Prices outside of G may affect total expenditures on group G in a fully flexible manner, and they may also affect relative expenditures within G , but this latter effect only operates through changes in utility.⁴ Thus, as long as these demands are invertible, households at two points in time (or two households in different locations) with the same within- G relative expenditures have the same utility after conditioning on within- G prices. Comparing total nominal household expenditures across these two points in time (or space) reveals the full price index P that keeps utility fixed when the household’s cost of living changes—and hence allows us to obtain money-metric welfare measures in the absence of well-measured prices outside of group G .

To use this insight to recover changes in welfare at each point of the income distribution—even when we do not observe households with the same relative expenditure at two points in time—we turn to relative Engel curves $E_{iG}^t(y^t)$ (period t projections of $\frac{x_i}{x_G}$ on $\log y$). Assume for now that

²Deaton and Muellbauer (1980) also refer to quasi-separability as implicit separability. Specific examples in this class include the popular non-homothetic CES preferences (e.g. Gorman, 1965; Hanoch, 1975; Comin et al., 2021), several variants of PIGL, PIGLOG and Translog preferences (Deaton and Muellbauer, 1980), and a class of Gorman preferences discussed in Fally (2022).

³Blackorby and Russell (1978) show that we can construct a price sub-index for group G if, and only if, we have quasi-separability in G , where such a subindex: i) does not depend on outside- G prices, and ii) can be combined with outside- G prices to construct the overall price index. Most applications further assume homothetic separability, where within-group expenditure shares are independent of income and utility (see Blackorby et al., 1978).

⁴Quasi-separability thus places weaker restrictions on demand than the more common assumption of homothetic separability of the expenditure function in p_G , where H_{iG} would be a function of p_G but not utility U (i.e. homothetic separability implies quasi-separability but not vice versa).

within- G prices are unchanged, $p_G^0 = p_G^1$. Moving from Hicksian to Marshallian relative demand by substituting U for the indirect utility function $V(p^t, y^t)$, the price index $P^0(y^0)$ that keeps the utility of a period 0 household constant under the full vector of period 1 prices p^1 is implicitly defined by equalizing relative demands across the two periods:

$$\underbrace{H_{iG}(p_G^0, V(p^0, y^0))}_{E_{iG}^0(y^0)} = \underbrace{H_{iG}(p_G^1, V(p^1, y^0/P^0(y^0)))}_{E_{iG}^1(y^0/P^0(y^0))} \quad (1)$$

Thus, the log of this price index change, $\log P^0(y^0)$, is simply the horizontal distance (in $\log y$ space) between period 0 and period 1 relative Engel curves. And if we relax the assumption that within- G prices are unchanged, we show that we simply need to adjust the latter curve to account for these price changes within G . It is then straightforward to recover changes in welfare for any household from the horizontal distance traveled between period 0 and 1 within- G relative expenditures, either traveling along period 0's relative Engel curve (to recover the equivalent variation, EV) or period 1's curve (to recover the compensating variation, CV).⁵

A strength of this approach comes from the fact that relative Engel curves can be estimated non-parametrically since quasi-separable demand can be of any rank (see Lewbel, 1991) and so can accommodate arbitrarily non-linear patterns of non-homotheticity within G and without imposing cross-equation restrictions on goods outside of G . Thus, we can capture potentially complex patterns of inflation that favor certain parts of the income distribution.

We state our approach formally in a lemma and a proposition. Lemma 1 lays out the logic above when relative prices within group G are held fixed. Proposition 1 relaxes this assumption by using observed price changes within G to correct the welfare estimates, either to the first order or exactly under any specific demand structure within G . We argue above that in most settings it is not possible to obtain reliable price data for large swaths of the services and manufacturing sectors, in part because of difficulties capturing quality and variety. Thus, Proposition 1 provides the minimal structure on preferences (i.e. quasi-separability) that allows us to uncover the full price index and welfare in such settings.⁶

In the second step, we form a bridge between the theoretical results and the empirical implementation by creating a manual for practitioners. Our estimation approach follows directly from our theory and uses expenditure survey microdata to estimate relative Engel curves for every location, every period and every good inside a product group G . As quasi-separability places no restrictions on the shape of these curves, they can be estimated non-parametrically and

⁵The data can come from repeated cross-sections or true household panels. In the (more common) first case, our approach recovers welfare changes at each point of the income distribution. In the second case, our approach recovers welfare changes for each household.

⁶Since price changes outside of G are unrestricted, we can accommodate arbitrary changes in quality and variety outside of G . We can also accommodate quality or variety changes inside G by adjusting the prices we use for our price correction using standard methods (see Section 3.2.2).

horizontal shifts calculated (correcting for within- G price changes and taking averages across goods to guard against measurement error). A natural question in taking our approach to the data is how plausible are the assumptions behind our proposition, most notably the assumption of quasi-separability? We show that violations of quasi-separability from misclassifying which goods are and are not in the quasi-separable set G have to be systematically related to price and income elasticities to cause bias, and provide expressions for the sign and magnitude of any bias. We also present several tests for quasi-separability using the available data. Beyond quasi-separability, we derive a set of testable requirements for unbiased identification: i) on aggregating up to good-level data in settings where barcode-level data are available, ii) on sample selection, iii) on bias in the estimation of Engel curves, and iv) on preference heterogeneity across households and over time.

In the final step, we implement our methodology using Indian expenditure survey microdata to quantify changes in rural welfare between 1987/88 and 1999/2000 at different points of the income distribution for every district in India.⁷ We compare our estimates to the leading existing Indian CPI estimates that come from Deaton (2003b) who calculates standard Paasche and Laspeyres price index numbers using changes in prices of products in the household surveys with both quantity information and no evidence of multiple varieties within a given location. For poorer deciles of the income distribution, we find very similar levels of consumer price inflation. Given that the products Deaton deems to have reliable prices—foods and fuels—cover more than 80 percent of total outlays for poorer rural households, it is reassuring that our estimates of the full price index for these households are similar to Deaton’s estimates of what is essentially a food and fuel price index (despite coming to this conclusion in different ways—we exploit shifts in relative Engel curves while Deaton uses observed price changes).

Looking across the income distribution, our estimates reveal that price inflation has been far from uniform, with significantly lower inflation rates for richer households—something that is not apparent from calculating standard price indices, even when using income-group- and district-specific expenditure weights, or from estimating non-homothetic price indices using Quadratic AIDS demand and goods with observable price data.⁸ Thus, while estimates based on standard approaches designed for settings with complete price data suggest that India saw significant convergence between poor and rich households over this period, we find no convergence once we account for the income-specific inflation uncovered by our approach.

⁷We focus on rural households because that has been the focus of the existing literature (e.g. the Great Indian Poverty Debate, or Topalova (2010)); and because well-measured food and fuel prices cover most of the consumption bundle for poor rural households, allowing us to validate our estimates against standard price indices for this group.

⁸For the latter, see Almås and Kjelsrud (2017) who use the same NSS expenditure data but include two categories with poorly measured prices (clothing; bedding and footwear). Additionally, as their method requires all prices, they assume that for miscellaneous non-food—the large residual category for which prices are not available—all relative prices change by the ratio of the non-food to food CPIs produced by the Indian government (CPIs that also struggle to account for changes in quality or variety).

The most likely explanation for these findings is that higher-income Indian households disproportionately benefited from lower inflation in categories such as services and manufactures where reliable price data are simply not available. This lower inflation is consistent with substantial increases in both the quality and variety of manufacturing products, and price declines, resulting from large reductions in tariff protection (see Goldberg et al., 2010); as well as rapid growth in the share of services in both GDP and employment over this period (Mukherjee, 2015).⁹ Standard approaches to price index estimation miss these patterns as these categories are either ignored entirely (as in Deaton, 2003b) or included without any quality or variety correction (as in India’s official CPI). Since wealthy households spend disproportionately on these categories and non-homotheticities are most pronounced within them, difficulties in measuring service and manufacturing prices have the potential to substantially change the distribution of welfare changes as we find.

This analysis sheds new light on the Great Indian Poverty Debate. Because India’s 1999-2000 National Sample Survey (NSS) added an additional 7-day recall period for food products (which inflated answers to the consistently asked 30-day consumption questions and lowered poverty measures), there has been much disagreement on how much poverty changed over the reform period.¹⁰ As long as the additional recall period did not change *relative* budget shares within our groups of food products G , our approach remains unbiased. We show that this assumption holds by exploiting the fact that the 1998 ‘thin’ survey round randomly assigned households to different recall periods. Thus, our approach provides a solution to the recall issues at the center of this debate. The Appendix also presents a second application of our methodology, revisiting Topalova’s (2010) analysis of the local labor market impacts of India’s 1991 trade reforms and uncovering adverse effects of import competition across the full income distribution, including among the richest households.

In addition to the literatures mentioned above, our approach connects to a longstanding literature using traditional Engel curves and expenditure changes on income-elastic goods—typically foodstuffs—to recover unobserved changes in real income (e.g. Hamilton, 2001; Costa, 2001; Nakamura et al., 2016; Almås 2012; Young 2012). Hamilton’s (2001) initial goal was to correct biases in the US consumer price index (CPI) arising from difficulties in measuring quality-adjusted prices in consumption categories such as services and manufactures. We address a key shortcoming in this literature. Despite relying on the non-homothetic AIDS demand system to generate non-horizontal Engel curves, this approach recovers a single price index for all households and so is neither theory-consistent nor suitable for distributional analysis. As shown in

⁹Our finding may be driven in part by a surge in product innovation in these sectors that is disproportionately targeted at rich households, a mechanism Jaravel (2019) documents for the US in the 2000s.

¹⁰See Deaton and Kozel (2005) for an overview. Deaton (2003a) calculates poverty by adjusting food expenditure using the initial mapping between food and fuels expenditure (which had no recall period added), implicitly assuming that relative prices of food and fuels did not change. Tarozzi (2007) explores a related approach.

Almås et al. (2018), calculating income-specific price index changes under the existing Engel methodology re-introduces the need to observe the full vector of price changes. We propose a new approach that leverages the broad class of quasi-separable preferences to recover theory-consistent price index and welfare changes at any point of the income distribution when price information is incomplete.¹¹

Finally, a recent literature uses barcode-level microdata for price index estimation, exploiting the granularity of these data to account for changes in product variety following Feenstra (1994). We cite above those that calculate income-group specific price indices. Crawford and Neary (Forthcoming) extend this approach to product characteristic space. Redding and Weinstein (2020) show how to use CES preferences to account for changes in product quality when prices are observed. As we discuss in Section 3.2.2, these recent advances complement our theoretical proposition by providing estimates of variety and quality-adjusted prices that can be used to correct for within- G relative price changes when products contain multiple varieties.

2 Theory

In this section we develop an approach to estimating income-specific changes in price indices and welfare that does not require reliable price data covering the full consumption basket. We first describe a data environment designed to mimic widely-available household expenditure surveys. We then introduce our approach and establish the central role of quasi-separability in a simplified setting (Lemma 1), before proceeding to our main proposition.

2.1 Data Environment

Our starting point is an environment with information on total (nominal) household outlays per-capita, y_h ,¹² for different households h coupled with their per-capita expenditures x_{hi} across the goods and services $i \in I$ that they consume (for readability we will refer to them simply as goods). Well-measured prices p_i are available for some subset of goods G but not necessarily for the remaining goods NG . This data environment corresponds to expenditure survey data where either separate price surveys or unit values calculated from well-measured quantity data provide price information for some subset of goods, such as foodstuffs or fuels.

To match our empirical setting, we focus our discussion on inferring price index changes over time for households at a given percentile h of the income distribution within a particular location.¹³ Inferring changes over time requires data for two time periods. In what follows, superscripts 0 and 1 indicate time periods and p is the full vector of consumption prices. Iso-

¹¹In related work, Ligon (2019) shows how one can recover the marginal utility of expenditure from expenditure data by imposing demands that feature a constant Frisch elasticity for each good and assuming that unobserved price changes in the full consumption basket are orthogonal to these Frisch elasticities.

¹²For readability we also refer to total household outlays per-capita as income.

¹³If household panel data are available, we can infer price index changes for individual households.

morphic results would hold across space if we replaced time periods by locations.

2.2 Basic Approach and the Role of Quasi-Separability

In this environment, recovering changes in the full price index, and hence welfare, is challenging. As the following sections document—by focusing on relative expenditures within product groups where prices are well measured—quasi-separable preferences provide the minimal restrictions necessary to recover welfare changes, and allow us to do so non-parametrically within this class of preferences.

Two definitions will be central. First, we define quasi-separable demand following Gorman's original formulation (1970; 1976).

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

$$e(p, U_h) = \tilde{e}(e_G(p_G, U_h), p_{NG}, U_h) \quad (2)$$

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

Quasi-separability is separability in the expenditure function (rather than the utility function). Two features of these preferences merit discussion (see Lemma 2 in Appendix A.3 for proofs).¹⁴

First, preferences are quasi-separable if, and only if, relative expenditures on each good i within group G — $\frac{x_{ih}}{x_{Gh}}$ where x_{Gh} is total expenditure on group G —can be written as a compensated function $H_{iG}(p_G, U)$ of utility and within- G relative prices alone:

$$\frac{x_{ih}}{x_{Gh}} = H_{iG}(p_G, U_h) = \frac{\partial \log e_G(p_G, U_h)}{\partial \log p_i}. \quad (3)$$

Second, quasi-separability imposes no restrictions on substitution patterns between goods within G , or between goods outside of G , or between consumption aggregates for group G relative to NG , but limits substitution patterns between a good within G and a good within NG to operate through a common group- G aggregator (with the flexibility of that aggregator allowing the elasticity of substitution between $i \in G$ and $j \in NG$ to be pair specific). More precisely, preferences are quasi-separable if, and only if, we can define utility implicitly by: $K(F_G(q_G, U_h), q_{NG}, U_h) = 1$, where q_G and q_{NG} denote vectors of consumption of goods in G and outside G , respectively, and the function $F_G(q_G, U_h)$ is homogeneous of degree 1 in q_G .

Several examples are instructive. The preferences used in Comin et al. (2021) and Matsuyama (2015), in which utility is implicitly defined by $\sum_i^N \left(\frac{q_i}{g_i(U)}\right)^{\frac{\sigma-1}{\sigma}} = 1$, are quasi-separable in any arbitrary subset of goods. Translog (in expenditure functions), EASI and PIGLOG demand systems also satisfy quasi-separability if there are no direct cross-price effects between goods within and outside of G . Beyond these special cases, we can construct highly flexible demand

¹⁴Lemma 2 combines existing results (see e.g. Blackorby et al. 1978) and provides a more direct proof.

systems that allow for rich substitution effects within G (captured by function F_G) and between goods within and outside G (function K).

Turning to our second definition, we define what we term “relative Engel curves” as follows.

Definition *Relative Engel curves, denoted by the function $E_{iG}^t(y_h) = \frac{x_{ih}}{x_{Gh}}$, describe how relative expenditure shares within G vary with total outlays per capita in period t (i.e. given the prevailing vector of prices p^t).*

Note that since quasi-separable demand systems can have any rank in the sense of Lewbel (1991), they can accommodate arbitrarily non-linear relative Engel curves and so allow for non-parametric estimation, as we describe and implement in Sections 3 and 4 below.

Finally, we present our price index notation and define our two welfare metrics. $P^1(p^0, p^1, y_h^1)$ (or in more concise notation $P^1(y_h^1)$ or just P^1) is the exact price index change between period 0 and period 1 prices, holding utility at period 1’s level (i.e. P^1 is defined implicitly by $V(p^1, y_h^1) = V(p^0, \frac{y_h^1}{P^1(y_h^1)})$ where V is the indirect utility function). In other words, the price index $P^1(y_h^1)$ converts the household’s period 1 nominal income to the hypothetical level of income that would make them equally well off under period 0 prices. Analogously, we define $P^0(p^0, p^1, y_h^0)$ as the exact price index change between period 1 and period 0 prices holding utility at period 0’s level (i.e. $V(p^0, y_h^0) = V(p^1, \frac{y_h^0}{P^0(y_h^0)})$).¹⁵

These two price indices are closely related to equivalent and compensating variations. $EV_h = e(p^0, U_h^1) - e(p^0, U_h^0) = \frac{y_h^1}{P^1(y_h^1)} - y_h^0$ is the amount of money that would bring a household in period 0 to their period 1 utility, and $CV_h = e(p^1, U_h^1) - e(p^1, U_h^0) = y_h^1 - \frac{y_h^0}{P^0(y_h^0)}$ is the amount of money taken away from a period 1 household to bring it back to its period 0 utility.

With these definitions in hand, we turn to our first result. Lemma 1 makes no assumptions on relative price changes outside of group G , but fixes relative prices for goods within G . This assumption—that we relax below—is convenient to highlight the key role of quasi-separability in estimating welfare changes with incomplete price information for non- G goods.

Lemma 1. *Assume that relative prices within group G are unchanged (i.e. $p_i^1 = \lambda_G p_i^0$ for all $i \in G$ and for some $\lambda_G > 0$). If, and only if, preferences are quasi-separable in subset G :*

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

$$E_{iG}^0\left(\frac{y_h^1}{P^1(y_h^1)}\right) = E_{iG}^1(y_h^1) \quad \text{and} \quad E_{iG}^1\left(\frac{y_h^0}{P^0(y_h^0)}\right) = E_{iG}^0(y_h^0).$$

- ii) When the relative Engel curve for a good $i \in G$ is strictly monotonic in income y_h :*

¹⁵The two price indices mirror each other: $y_h^1 = y_h^0 / P^0(y_h^0)$ implies $y_h^0 = y_h^1 / P^1(y_h^1)$.

- a) $\log P^1(y_h^1)$ and $\log P^0(y_h^0)$ are uniquely identified by the horizontal shift in good i 's relative Engel curve, as defined by the equalities above.
- b) EV and CV for a given income level are revealed by the horizontal distance between new and old expenditure shares along period 0's or period 1's relative Engel curve for good i , respectively, such that $E_{iG}^0(y_h^0 + EV_h) = \frac{x_{ih}^1}{x_{Gh}^1}$ and $E_{iG}^1(y_h^1 - CV_h) = \frac{x_{hi}^0}{x_{hG}^0}$.

Appendix A.1 provides the proofs.

Lemma 1 i) states that the horizontal shift in relative Engel curves at any given point of the initial or final income distribution is equal to the change in the exact price index for that group. Whether we can use this result to infer changes in the price index by observing relative expenditure shares within G and total outlays, depends on whether we can invert these relationships. If a relative Engel curve is strictly monotonic, as assumed in Lemma 1 ii), observing shifts for that single good is sufficient to infer price index and welfare changes. In contrast, if a relative Engel curve is flat (independent of income), a "horizontal shift" leaves the curve unchanged and thus the shift is uninformative. In general, invertibility requires that the vector of relative expenditure shares E_{iG} across $i \in G$ is an injective function of income y_h (such that the vector of budget shares maps to a unique level of income). A sufficient condition for this invertibility requirement to hold is that at least one relative Engel curve $i \in G$ is strictly monotonic.

To aid intuition, Figure 1 graphically illustrates Lemma 1 ii). Take as an example a household at percentile h with initial per-capita outlays of y_h^0 (the bottom-left dot in the figure). Since within- G relative prices are not changing, households with the same within- G budget shares must be equally well off (recall that quasi-separability implies that relative outlays depend only on within- G prices and utility, $\frac{x_{ih}}{x_{Gh}} = H_{iG}(p_G, U_h)$).¹⁶ Thus, the horizontal distance (in $\log y_h$ space) between their initial position on the period 0 relative Engel curve and that same budget share on the period 1 relative Engel curve equals the log of the change in the price index P^0 . The CV 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 relative Engel curve to the actual within- G budget share of that household in period 1 (the upper-right dot). The same movements in reverse reveal P^1 and EV.

Since relative Engel curves are not parallel, the price index change P^0 and CV_h may vary with the household's position in the income distribution. Relatedly, P^1 and EV_h will not be identical to P^0 and CV_h if the household's utility differs in the two periods. Why are the curves not parallel? As relative prices within G are held fixed, it is changes in prices outside of group G (e.g. prices of manufactures and services) that rotate the curves apart when these goods are consumed disproportionately by richer (or poorer) households. By not placing restrictions on price changes outside of set G , income-group specific price indices can diverge leading to non-

¹⁶Here we abstract from preference (taste) heterogeneity but discuss this possibility in Section 3.2.5.

parallel shifts in relative Engel curves.

To make these statements precise, we lay out several steps of Lemma 1's proof. To obtain $P^1(p^0, p^1, y_h^1)$, start with the period 1 relative budget share on period 1's relative Engel curve:

$$\begin{aligned}
 E_{iG}^1(y_h^1) = H_{iG}(p_G^1, U_h^1) &= H_{iG}(p_G^1, V(p^1, y_h^1)) \\
 &= H_{iG}(p_G^0, V(p^1, y_h^1)) \\
 &= H_{iG}(p_G^0, V(p^0, y_h^1/P^1(p^0, p^1, y_h^1))) \\
 &= E_{iG}^0\left(y_h^1/P^1(p^0, p^1, y_h^1)\right).
 \end{aligned}$$

The first line links this unobserved compensated Hicksian demand function to observed relative Engel curves by substituting in the indirect utility function $V(p, y)$ that connects total outlays and utility. Equality between the first and second line is an implication of the homogeneous price change $p_i^1 = \lambda_G p_i^0$ within group G .¹⁷ Equality between the second and third lines follows from the definition of $P^1(p^0, p^1, y_h^1)$ above. The final line moves back to relative Engel curve functions. Thus, the difference between percentile h 's total outlays in period 1 and the total outlays of a percentile in period 0 with the same relative budget share as h had in period 1 reveals the price index change $P^1(p^0, p^1, y_h^1)$. An analogous proof applies for $P^0(p^0, p^1, y_h^0)$.

Lemma 1 shows that quasi-separability is not only sufficient but a necessary condition to recover income-specific price indices and welfare from horizontal shifts in observed within- G outlays for arbitrary price realizations outside of G . Thus, in the absence of reliable price data outside of group G , quasi-separability provides the minimal restriction on preferences such that these unknown prices do not confound shifts in relative Engel curves.

Finally, an obvious question is why do we focus on relative Engel curves, and whether alternative preferences could allow us to recover changes in the price index from shifts in traditional Engel curves (i.e. shares of total expenditure plotted against log total outlays). In Lemma 4 in Appendix A.5, we provide an impossibility result that no such approach is consistent with rational preferences while allowing for arbitrary changes in unobserved prices (if price changes are uniform, shifts in traditional Engel curves do recover price indices). These results connect to Almås et al. (2018) who show that the traditional Engel-curve methodology for recovering price indices under AIDS preferences (e.g. Hamilton, 2001) requires information on all price changes to recover income-specific price indices. Shifting attention to relative Engel curves—and quasi-separable preferences—allows us to bypass these negative results.

2.3 Recovering Income-Specific Welfare Changes From Expenditure Survey Data

Even if preferences are quasi-separable, vertical shifts in relative Engel curves due to within- G relative price changes (assumed away in Lemma 1) can confound estimates of price index

¹⁷Note that H_{iG} is homogeneous of degree zero in prices p_G^t within group G .

changes by raising or lowering relative expenditure shares conditional on utility. Our main proposition takes advantage of the fact that reliable price information may be available for some quasi-separable set of goods—but not for all consumption—to adjust relative Engel curves to account for these confounding vertical shifts and relax the assumption that relative prices within subset G are fixed.

Proposition 1. *If, and only if, preferences are quasi-separable in the subset G of goods:*

- i) *The log price index change for a given income level in period 1, $\log P^1(y_h^1)$, is equal to the horizontal shift (in $\log y_h$ space) in the price-adjusted relative Engel curve of any good $i \in G$ at that income level, such that*

$$E_{iG}^0\left(\frac{y_h^1}{P^1(y_h^1)}\right) = E_{iG}^1(y_h^1) \times \frac{H_{iG}(p_G^0, U_h^1)}{H_{iG}(p_G^1, U_h^1)} \quad (4)$$

where $\frac{H_{iG}(p_G^1, U_h^1)}{H_{iG}(p_G^0, U_h^1)}$ is the change in expenditure shares induced by the change in (relative) prices within G evaluated along the indifference curve at period 1 utility, $U_h^1 = V(p^1, y_h^1)$.

- ii) *When the price-adjusted relative Engel curve for a good $i \in G$ is strictly monotonic in income y_h :*

- a) *$\log P^1(y_h^1)$ is uniquely identified by the horizontal shift in good i 's price-adjusted relative Engel curve, as defined by the equality above.*
- b) *EV for a given income level is revealed by the horizontal distance between new and old expenditure shares along period 0's relative Engel curve for good i , such that*

$$E_{iG}^0(y_h^0 + EV_h) = \frac{x_{ih}^1}{x_{Gh}^1} \times \frac{H_{iG}(p_G^0, U_h^1)}{H_{iG}(p_G^1, U_h^1)}.$$

Switching superscripts 0 and 1 provides the log price index change $\log P^0(y_h^0)$ and CV.

Appendix A.2 provides the proofs.¹⁸

This proposition shows that we can still infer changes in $\log P^1(y_h^1)$ from horizontal shifts in relative Engel curves but after first adjusting the period 1 curve by the term $H_{iG}(p_G^0, U_h^1)/H_{iG}(p_G^1, U_h^1)$, i.e. the compensated shift in expenditure shares due to the change in within- G prices, with: $\log \frac{H_{iG}(p_G^1, U_h^1)}{H_{iG}(p_G^0, U_h^1)} = \sum_{j \in G} \int_{p_j^0}^{p_j^1} \frac{\partial \log H_{iG}}{\partial \log p_j} d \log p_j$. EV is then the additional horizontal distance traveled along the period 0 relative Engel curve to the period 0 expenditure share.

These adjustments require some knowledge of the within-group demand structure H_{iG} and within-group relative price changes. But crucially, they do not require information on the structure of preferences or prices for goods outside G . As long as there is a group G of goods for which

¹⁸Note that Proposition 1 also holds with an additive correction term, $+ [H_{iG}(p_G^0, U_h^1) - H_{iG}(p_G^1, U_h^1)]$ instead of $\times \frac{H_{iG}(p_G^0, U_h^1)}{H_{iG}(p_G^1, U_h^1)}$, since $E_{iG}^1(y_h^1) = H_{iG}(p_G^1, U_h^1)$.

preferences are quasi-separable and reliable price data are available, we can uncover changes in price indices and welfare.¹⁹

As described in Section 3.1.3 below, we implement the price adjustment in Proposition 1 in several ways: in its exact form after specifying a range of different within-group demand structures H_{iG} ; and as a first-order approximation, evaluating elasticities in the base period. The latter approach brings two benefits. First, it does not require us to take a stand on the structure of within-group demand and second, it provides a natural and transparent two-step estimation strategy—first calculating horizontal shifts in the unadjusted relative Engel curves that present themselves directly in the data (as in Lemma 1 above), and then adding a correction term formed from local elasticities and observable within- G relative price changes.

To formalize this second approach, write equation (4) in logs and take a first-order approximation of changes in $\log H_{iG}$ due to relative price changes, holding utility fixed.²⁰ Subsequently inverting the relative Engel curve at $\frac{x_{ih}^1}{x_{Gh}^1}$, for any good $i \in G$ we obtain:²¹

$$\log(y_h^1) - \log(E_{iG}^0)^{-1} \left(\frac{x_{ih}^1}{x_{Gh}^1} \right) \approx \log(P^1) + (\beta_{ih}^0)^{-1} \log \frac{H_{iG}(p_G^0, U_h^1)}{H_{iG}(p_G^1, U_h^1)} \quad (5)$$

where $\beta_{ih}^0 = \frac{\partial \log E_{iG}}{\partial \log y_h}$ denotes the slope of the relative Engel curve (i.e. the income elasticity) evaluated at income level y_h^1/P^1 and initial prices p^0 . The left-hand side of equation (5) is the horizontal shift in the price-unadjusted relative Engel curve as in Lemma 1 above. The first term on the right-hand side of (5) is the price index change we are trying to estimate. The second term is the bias due to a potential confounder: the vertical shift in relative Engel curves due to relative price changes within G . Finally, using the local compensated cross-price elasticities of relative expenditures, $\sigma_{ijh} = \frac{\partial \log H_{iG}}{\partial \log p_j}$ with $\sum_{j \in G} \sigma_{ijh} = 0$, this vertical shift, again to the first order, can be rewritten as a function of observable relative price changes: $\log \frac{H_{iG}(p_G^0, U_h^1)}{H_{iG}(p_G^1, U_h^1)} \approx \sum_{j \in G} \sigma_{ijh} (\Delta \log p_j - \overline{\Delta \log p_G})$.

3 From Theory to Estimation: An Empirical Methodology

In this section, we build on the theoretical results above to derive an empirical methodology for estimating price indices and welfare changes using household expenditure survey micro-data with price information that covers only a subset of consumption. We then turn to iden-

¹⁹To be more precise, these vertical adjustments of relative Engel curves depend on compensated changes in expenditure shares within G , holding utility constant. One can infer compensated changes in within-group expenditures from a Slutsky-type decomposition involving slopes of relative Engel curves and uncompensated price elasticities of within-group expenditure shares (see the proof of Proposition 1 in Appendix A): $\frac{\partial \log H_{iG}}{\partial \log p_j} = \frac{\partial \log(x_i/x_G)}{\partial \log p_j} + E_{jG} \frac{x_G}{y} \frac{\partial \log E_{iG}}{\partial \log y}$. Estimating these terms only requires information on household total outlays, expenditures on goods within group G , and prices of these goods.

²⁰I.e. assuming that the vertical shifts in relative Engel curves due to within- G relative price changes are proportional to those price changes.

²¹Symmetrically for P^0 : $\log(y_h^0) - \log(E_{iG}^1)^{-1} \left(\frac{x_{ih}^0}{x_{Gh}^0} \right) \approx \log(P^0) + (\beta_{ih}^1)^{-1} \log \frac{H_{iG}(p_G^1, U_h^0)}{H_{iG}(p_G^0, U_h^0)}$.

tification and derive corollaries to Proposition 1 that define testable conditions for unbiased estimation. These results naturally suggest a number of validation exercises and robustness checks that we implement in our application in Section 4. Taken together, this section serves as a manual for practitioners to apply the methodology.

3.1 Estimation Approach

Suppose that we want to estimate the welfare change between two periods for a specific percentile of the household income distribution in a particular location. First, focusing on goods within a quasi-separable group for which reliable price data are available, we use non-parametric methods to estimate flexible relative Engel curves separately for both periods. We can then recover changes in income-specific price indices and welfare from the horizontal shift in these curves at different points of the income distribution, either by adjusting relative Engel curves to account for within- G price changes or by adding a first-order correction term. In either case, combining estimates for multiple goods within G increases precision by allowing us to accommodate measurement error in the expenditure surveys. We now discuss each of these steps.

3.1.1 Estimating Horizontal Shifts in Relative Engel Curves

We first describe the procedure for estimating shifts in (price-unadjusted) relative Engel curves from the raw household expenditure survey microdata in a given location. As we detail in Section 3.1.3, these estimates are direct inputs into our first-order approach to implement price corrections and the exact approach also builds on this procedure.

The first step is to estimate kernel-weighted local polynomial regressions of relative expenditure shares, $x_{ih'}^t/x_{Gh'}^t$, on log total outlays per capita, $\log y_{h'}^t$, for every good $i \in G$ and each period t , where h' indexes the individual households in the expenditure surveys. This provides estimates of x_{ih}^t/x_{Gh}^t for any percentile h of households across the income distribution (where y_h is the predicted income for households at this percentile). We estimate these relative Engel curves at 101 points corresponding to percentiles 0 to 100 of the local income distribution.²² Following Lemma 1 and Proposition 1 above, we restrict attention to goods where relative Engel curves are monotonic (ensuring that estimates of shifts are unique for each good).²³

Abstracting from within- G relative price changes for now, consider estimating the log price index change for income percentile h in period 1, $\log P^1(p^0, p^1, y_h^1)$. The relative Engel curve

²²We first smooth the distribution of local income using a local polynomial regression of log total outlays per capita on outlays rank divided by the number of households n (with a bandwidth equal to $101/(n-1)$) to obtain $\log y_h^t$ at the 101 percentiles. To obtain relative Engel curves, we use a bandwidth equal to one quarter of the range of the income distribution in a given market. In both cases we use an Epanechnikov kernel. Our application explores alternative bandwidth choices.

²³As 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 goods where relative 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 tail portions are non-monotonic (replacing those values with a linear extrapolation from the monotonic portion of the curve). To reduce noise in our estimates at the tails of the distribution, we linearly extrapolate the top and bottom three percentiles of all curves.

for period 1 provides a point estimate of relative expenditures for households at this percentile of the initial income distribution, x_{ih}^1/x_{Gh}^1 . The next step is to estimate the period 0 income level $\widehat{E_{iG}^0}^{-1}(x_{ih}^1/x_{Gh}^1)$ associated with this relative expenditure share from the crossing point on the period 0 relative Engel curve. To do so, we find the crossing point $\widehat{x_{ih}^0/x_{Gh}^0}$ and take the corresponding income of this income percentile $h, \widehat{\log y_h^0}$.²⁴

Given these estimates, the income-percentile specific price index change $\log P^1(p^0, p^1, y_h^1)$ is equal to the difference between $\log y_h^1$ (the period 1 level of income for percentile h) and the estimate of $\widehat{\log y_h^0}$ —this is the horizontal shift labeled $\log P^1$ in Figure 1. The welfare change for income-group h , as measured by the EV, is recovered from the relationship $\log(1 + EV_h/y_h^0) = \widehat{\log y_h^0} - \log y_h^0$, where $\widehat{\log y_h^0}$ is the estimate of the period 0 income required to obtain period 1 utility and $\log y_h^0$ is the actual period 0 log income for percentile h . This expression recovers welfare changes for a hypothetical household that stays at the same point of the income distribution in both periods. If household panel data are available, we could recover welfare changes for a specific household using this methodology. To estimate the price index change holding utility at period 0's level, $\log P^0(p^0, p^1, y_h^0)$, one applies the same procedure in the opposite direction (and recovering CV from $\log(1 - CV_h/y_h^1) = \widehat{\log y_h^1} - \log y_h^1$). Each good $i \in G$ provides a separate estimate for $\log P^0, \log P^1, CV_h$ and EV_h .

3.1.2 Averaging Estimates Across Goods

Measurement error in expenditure surveys will bias estimates calculated using shifts in the relative Engel curve of any one specific good i . Averaging across multiple goods $i \in G$ at each percentile of the income distribution reduces such bias. Denote i.i.d. measurement error in percentile h expenditures by ϵ_{ih} : $x_{ih}^*(p, y_h) = x_{ih}(p, y_h)\epsilon_{ih}$ with $\epsilon_{ih} > 0$. Taking a first-order approximation as in equation (5) and averaging horizontal shifts across $i \in G$, we obtain the bias generated by such measurement error:

$$\frac{1}{G} \sum_{i \in G} \left(\log(y_h^1) - \log(\widehat{E_{iG}^0}^{-1}(\frac{x_{ih}^1}{x_{Gh}^1})) \right) \approx \log(P^1) - \frac{1}{G} \sum_{i \in G} \left((\beta_{ih}^0)^{-1} (\Delta \log \epsilon_{ih} - \Delta \frac{1}{G} \sum_{i \in G} E_{iG} \log \epsilon_{ih}) \right). \quad (6)$$

Thus, averaging horizontal shifts over a large number of goods provides unbiased estimates—i.e. the second term on the right-hand side goes to zero—because the (demeaned) i.i.d. measurement error is uncorrelated with the slopes of relative Engel curves.²⁵ The exposition above abstracts from changes in relative prices within G as we discuss price corrections next, but a similar logic applies to measurement error in prices.

²⁴We take the two closest percentiles and linearly interpolate between them to obtain $\widehat{\log y_h^0}$.

²⁵Ultimately, we will use the median as an unbiased estimate of the mean since not all goods $i \in G$ have overlapping relative Engel curves for a particular percentile (see Section 3.2.4).

3.1.3 Price Corrections

Proposition 1 shows how to correct the price index estimates above—derived solely from horizontal shifts in relative Engel curves—when relative prices are changing within group G . The first-order approach adds a price correction term composed of local elasticities and observable price changes to the average horizontal shift. The exact approach uses knowledge of the shape of function $H_{iG}(p_G, U)$ to adjust relative Engel curves before calculating horizontal shifts. We discuss the two procedures in turn.

First-Order Price Correction Equation (5) provides an estimate of $\log P^1$ as a function of the horizontal shift in good i 's relative Engel curve and a first-order correction for vertical shifts in i expenditure due to relative price changes. Substituting $\log \frac{H_{iG}(p_G^0, U_h^1)}{H_{iG}(p_G^1, U_h^1)} \approx \sum_{j \in G} \sigma_{ijh} (\Delta \log p_j - \overline{\Delta \log p_G})$, averaging estimates across multiple goods $i \in G$ as above, and rearranging, we obtain:

$$\log(P^1) \approx \frac{1}{G} \sum_{i \in G} \left(\log(y_h^1) - \log(E_{iG}^0)^{-1} \left(\frac{x_{ih}^1}{x_{ih}^0} \right) \right) - \frac{1}{G} \sum_{i \in G} \left((\beta_{ih}^0)^{-1} \left(\sum_{j \in G} \sigma_{ijh} (\Delta \log p_j - \overline{\Delta \log p_G}) \right) \right) \quad (7)$$

The left-hand side is the price index we are trying to estimate. The first term on the right is the average estimate of horizontal shifts in relative Engel curves across multiple goods $i \in G$. The final term captures the bias: the covariance between price-induced vertical shifts across $i \in G$ and the slopes of relative Engel curves at a given income level. Combining price changes for $i \in G$ as well as local income and price elasticities, β_{ih}^0 and σ_{ijh} , this covariance term corrects for such bias to the first order. If relative price changes are only weakly related to within- G income elasticities, the bias from averaging multiple estimates of price index and welfare changes will tend to cancel out and the size of the correction will be small.

To implement this correction, we estimate the slope-to-price-change correlation term above. The price changes are observed, local income elasticities, β_{ih}^0 , come from local slopes of the relative Engel curves estimated above and, in principle, the full set of local cross-price relative expenditure elasticities σ_{ijh} can be estimated using price variation within group G .

If we assume a constant elasticity of substitution σ_G within group G , the (percentile-specific) bias correction term takes the simple form:

$$\frac{1}{G} \sum_{i \in G} (\beta_{ih}^0)^{-1} \sigma_G (\Delta \log p_i - \overline{\Delta \log p_G}). \quad (8)$$

As noted above, the correction term is small if relative price changes are weakly correlated with slopes of relative Engel curves, but also if within- G elasticities are small, or if within- G price changes are similar.

Exact Price Correction To provide an exact correction, recall from Proposition 1 that we must adjust one of the period's relative Engel curves to account for within- G relative price changes

and then calculate horizontal shifts using this adjusted curve. Thus, we proceed as in Section 3.1.1 above, but modifying the appropriate relative Engel curve before calculating horizontal differences for each good $i \in G$ and then averaging.

First, we propose two practical specifications that only require estimating a single elasticity parameter. One is to specify a constant (compensated) elasticity of substitution between goods within group G , with an expenditure function that satisfies:²⁶

$$e(p, U_h) = \tilde{e} \left(\left(\sum_{j \in G} A_j(U) p_j^{1-\sigma_G} \right)^{\frac{1}{1-\sigma_G}}, p_{NG}, U_h \right) \quad (9)$$

With such preferences, relative expenditures within G are given by $H_{iG}(p_G, U) = \frac{A_i(U) p_i^{1-\sigma_G}}{\sum_{j \in G} A_j(U) p_j^{1-\sigma_G}}$. This generalizes the preferences in Hanoch (1975) and Comin et al. (2021) by allowing for flexible substitution patterns outside of group G . The required adjustment due to confounding within- G relative price changes then takes the form:

$$\log H_{iG}(p_G^1, U_h^1) - \log H_{iG}(p_G^0, U_h^1) = (1 - \sigma_G) [\Delta \log p_i - \overline{\Delta \log p_G}] \quad (10)$$

where $\overline{\Delta \log p_G} = \log[\sum_{j \in G} (p_j^1/p_j^0)^{\sigma_G-1} (x_{jh}^1/x_{Gh}^1)]^{\frac{1}{\sigma_G-1}}$ is a CES index of relative price changes. With an estimate of the elasticity of substitution σ_G between goods of group G (which can be estimated using prices and expenditures on goods $i \in G$), we have a simple-to-compute multiplicative adjustment term.

To account for richer patterns of substitution, we can increase the number of nests in this constant elasticity structure to allow own and cross-price elasticities to differ across subgroups of goods. Consider a partition of group $G = g_1 \cup g_2 \cup \dots$ and a within-group expenditure function:

$$e(p, U_h) = \tilde{e} \left(\left(\sum_g \left(\sum_{j \in g} A_j(U) p_j^{1-\sigma_g} \right)^{\frac{1-\eta_G}{1-\sigma_g}} \right)^{\frac{1}{1-\eta_G}}, p_{NG}, U_h \right). \quad (11)$$

Adjustments for within- G relative price changes are now given by:

$$\log H_{iG}(p_G^1, U_h^1) - \log H_{iG}(p_G^0, U_h^1) = (1 - \sigma_g) [\Delta \log p_i - \overline{\Delta \log p_g}] + (1 - \eta_G) [\overline{\Delta \log p_g} - \overline{\Delta \log p_G}] \quad (12)$$

where $\overline{\Delta \log p_g} = \log[\sum_{j \in g} (p_j^1/p_j^0)^{1-\sigma_g} (x_{jh}^1/x_{gh}^1)]^{\frac{1}{1-\sigma_g}}$ is the price index change for subgroup $g \subset G$ and $\overline{\Delta \log p_G} = \log[\sum_{j \in G} e^{(1-\eta_G)\overline{\Delta \log p_g}} (x_{jh}^1/x_{Gh}^1)]^{\frac{1}{1-\eta_G}}$ is the overall price index change for group G .²⁷

Alternatively, recall from footnote 18 that the correction term in Proposition 1 can also be

²⁶The corresponding utility function can be implicitly defined as: $K \left(\sum_{j \in G} A_j(U)^{1/\sigma_G} q_j^{\sigma_G/(\sigma_G-1)}, q_{NG}, U_h \right) = 1$.

²⁷This more flexible structure allows for heterogeneous own-price elasticities and a more complex cross-price substitution matrix. For each good i , the own-price elasticity for relative expenditures is an i -specific weighted average between the subgroup-specific parameter σ_g and the upper-tier elasticity η_G (as opposed to a single parameter σ_G in the one-layer case). Within subgroup g , cross-price elasticities are also good-specific and differ from cross-price elasticities with goods in subgroups outside of g . See Appendix A.6 for a description of the full substitution matrix.

written in an additive form. Specifying that semi-elasticities ξ_G within group G are constant akin to EASI demands (Lewbel and Pendakur, 2009) provides an additive adjustment expressed in levels rather than logs of expenditure that is again simple to compute:

$$H_{iG}(p_G^1, U_h^1) - H_{iG}(p_G^0, U_h^1) = -\xi_G \times [\Delta \log p_i - \overline{\Delta \log p_G}]. \quad (13)$$

For additional flexibility, these semi-elasticities ξ_i can also be good-specific (see Appendix A.6):

$$H_{iG}(p_G^1, U_h^1) - H_{iG}(p_G^0, U_h^1) = -\xi_i \times \left[\Delta \log p_i - \sum_{j \in G} \xi_j \Delta \log p_j / \left(\sum_{k \in G} \xi_k \right) \right]. \quad (14)$$

3.2 Identification and Validation

In this subsection, we derive a number of corollaries and tests related to unbiased identification when taking Proposition 1 to the data. Specifically, we derive an expression for the potential bias from violations of quasi-separability; construct several tests of the quasi-separability assumption; and discuss how to deal with other potential biases due to data aggregation, omitted variables in Engel curve estimation, sample selection issues, and taste heterogeneity.

3.2.1 Quasi-Separability and Misspecification

Bias from Violations of Quasi-Separability Although our main propositions assume preferences are quasi-separable in group G , violations of this assumption only induce bias in our welfare estimates if they are systematically related to price elasticities and slopes of relative Engel curves. Here we make this statement precise by solving for the first-order bias.

Suppose we misclassify a good i that truly belongs in G as a non- G (NG) good (i.e. we omit a good that belongs within quasi-separable group G). Alternatively, suppose we falsely classify a NG good j as part of G . In both cases, price changes outside of what we believe to be the G group can directly affect within- G relative outlays (rather than only affect relative outlays through utility as would be true if goods were correctly classified into quasi-separable groups).

Corollary. *To the first order, the bias from taking an average over estimates from all goods i that we believe to be in G (potentially including misclassified goods) is equal to:*

$$\frac{1}{G} \sum_{i \in G} \log(E_{iG}^0)^{-1} \left(\frac{x_{ih}^1}{x_{Gh}^1} \right) - \log \left(\frac{y_h^1}{P^1} \right) \approx \frac{1}{G} \sum_{i \in G} (\beta_{ih}^0)^{-1} \times \sum_{k \in NG} (\Delta \log p_k - \overline{\Delta \log p_G}) \left. \frac{\partial \log(x_i/x_G)}{\partial \log p_k} \right|_U, \quad (15)$$

where k denotes the goods we believe to be in NG .

For correctly classified goods, $\left. \frac{\partial \log(x_i/x_G)}{\partial \log p_k} \right|_U = \frac{\partial \log H_{iG}}{\partial \log p_k} = 0$ and there is no bias.²⁸ If good $k' \in NG$ is actually a G good, $\left. \frac{\partial \log(x_i/x_G)}{\partial \log p_{k'}} \right|_U \neq 0$ for some is . If good $i' \in G$ is actually a NG good, $\left. \frac{\partial \log(x_{i'}/x_G)}{\partial \log p_k} \right|_U \neq 0$ for some ks .

²⁸Equation (15) abstracts from relative price changes within G (or assumes they all equal $\Delta \log p_G$) since, as we describe above, these relative price changes can be observed and corrected for.

Averaging across multiple i estimates, these violations of quasi-separability only generate bias if the direction and magnitude of the confounding (compensated) cross-price effects from unobserved NG price changes $(\sum_{k \in NG} (\Delta \log p_k - \overline{\Delta \log p_G}) \frac{\partial \log(x_i/x_G)}{\partial \log p_k} \Big|_U)$ are systematically related to the slopes of relative Engel curves (β_{ih}^0) for the goods within G . In addition, the bias will be small if most goods are correctly classified, if price changes are similar for G and NG goods, or if compensated cross-price elasticities are small. In our application, this result motivates both averaging over multiple i estimates and exploring the sensitivity of our estimates to alternative classifications of goods into quasi-separable nests G .

Testing for Quasi-Separability with Outside Price Data We now present a direct test of quasi-separability that relies on the key property that expenditure shares within group G , $\frac{x_i}{x_G}$, can be expressed as a function $H_{iG}(U, p_G)$ of utility and relative prices within group G —i.e. within- G expenditure shares do not depend on outside prices after conditioning on these variables. This property is a necessary and sufficient condition for quasi-separability:

Corollary QS 1. *Preferences are quasi-separable in group G if, and only if, compensated expenditure shares for good i within group G do not depend on outside prices p_j for any $j \notin G$:*

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

This corollary is a direct consequence of Lemma 2 in Appendix A.3 (discussed in Section 2.2).

We can also derive a test based on uncompensated rather than compensated demand, providing an alternative characterization of quasi-separability if one cannot condition on utility. A necessary and sufficient condition for quasi-separability in G is that the uncompensated price effect of each good j outside G on the relative expenditure share of good i within G is given by:

$$\frac{\partial \log(x_i/x_G)}{\partial \log p_j} \Big|_y = -\frac{x_j}{y} \frac{\partial \log E_{iG}}{\partial \log y}. \quad (16)$$

where $\frac{\partial \log E_{iG}}{\partial \log y}$ is the slope of the relative Engel curve for good $i \in G$, and $\frac{x_j}{y}$ the overall expenditure share on good $j \notin G$. The proof (see Appendix A.6) relies on Roy's identity linking changes in utility to changes in income and prices.²⁹

Testing for Quasi-Separability with Outside Expenditures Data The tests described above require price information for goods outside group G to test that preferences are quasi-separable in group G . We argue that reliable price data may not exist for large parts of consumption. Thus, we propose a further characterization that relies only on prices for goods within group G .

Corollary QS 2. *Preferences are quasi-separable in group G if, and only if, the elasticity of expenditures for any good $j \notin G$ with the price of good $i \in G$ is proportional to the share of good i in*

²⁹Given that these local slopes (for each i in a given period, market and income level) require imprecise non-parametric estimation, we prefer to test the null of zero effects for outside- G price changes as in Corollary QS1.

group G expenditures, i.e.:

$$\frac{\partial \log x_j}{\partial \log p_i} = \frac{x_i}{x_G} \times \gamma_G$$

for any $j \notin G$ and $i \in G$, where $\gamma_G = \sum_{k \in G} \frac{\partial \log x_j}{\partial \log p_k}$ is common across all goods i in group G .

In other words, this corollary states that the effect of prices of goods within G on expenditures for outside goods are fully captured by $\overline{\Delta \log p_G} = \sum_i (x_i/x_G) \Delta \log p_i$, the change in the relative expenditure share weighted log price change across goods in G (with coefficient γ_G).³⁰ The proof (see Appendix A.6) exploits Slutsky symmetry under rational preferences, which implies that the compensated price effects are symmetric, $\left. \frac{\partial x_j}{\partial \log p_i} \right|_U = \left. \frac{\partial x_i}{\partial \log p_j} \right|_U$, for any pair of goods i and j . This symmetry property allows us to rely on outside expenditures and prices within group G (instead of within-group expenditure and outside prices). As information on expenditures outside G is generally easier to obtain than prices, the data requirements for this test are more easily met.

3.2.2 Aggregation Across Varieties of a Good

Researchers often estimate Engel curves for a broadly-defined good (indexed here by g) that itself contains many varieties (the i s in our exposition up to now, e.g. different types, preparations, brands, sizes or flavors), either because that is the level the data are collected at or because specific varieties are not consumed widely enough given the number of households sampled. Fortunately, Lemma 1 and Proposition 1 can also be applied to aggregates of varieties rather than individual varieties, even if demands for those varieties are non-homothetic within g .

Corollary. *Suppose that G in our exposition above can be partitioned into subgroups of goods: $G = g1 \cup g2 \cup g3 \dots$ (e.g. salt, milk, lentils etc.). Denote by $E_{g,G}$ the expenditure share on subgroup g within group G . Under the assumptions of Lemma 1:*

$$E_{g,G}^1(y_h^1) = E_{g,G}^0 \left(\frac{y_h^1}{P^1(y_h^1)} \right) \quad \text{and} \quad E_{g,G}^0(y_h^0) = E_{g,G}^1 \left(\frac{y_h^0}{P^0(y_h^0)} \right).$$

In other words, the key equivalence continues to hold if we treat the subgroups g as products (instead of the individual varieties i). Furthermore, under the assumption that prices across the i s within each subgroup g can be aggregated into price indices, $P_g(p_g, U)$, we can apply Proposition 1 and the price-adjustment corollaries above to correct for relative price changes, but now using subgroup price indices $P_g(p_g, U)$ instead of individual prices p_i .³¹

Several remarks are in order. First, note that these subgroup price indices can be non-

³⁰Thus, if and only if preferences are quasi-separable, the equality $\Delta \log x_j = \gamma_G \overline{\Delta \log p_G}$ must hold to the first order when prices of goods $i \in G$ change. Note also that this result holds for both uncompensated and compensated price effects (i.e. controlling for utility), since the difference between the two is proportional to the expenditure share of good i .

³¹For example, the price aggregates derived in Redding and Weinstein (2020) could be used for $P_g(p_g, U)$ if we assume that within- g preferences have their CES structure.

homothetic: relative consumption within subgroup g can vary with utility U (and thus income); the rich and poor can even consume distinct varieties. Second, aggregation can accommodate differences in shopping amenities and store-level price differences (modeled as store-specific varieties). Third, aggregation can accommodate new and disappearing varieties within subgroup g using existing methods. For example, if a popular new variety of salt appeared between periods 0 and 1, this would lower the salt price index $P_g(p_g, U)$. If g is in the NG group then no correction is necessary, with the reduction in the salt price index raising utility, altering within- G expenditure shares, and lowering the full price index $P^1(y_h^1)$. If g is in the quasi-separable group G , we would either need to: calculate the change in the salt price index (e.g. using the share of salt expenditure spent on the new variety and the within-salt elasticity of substitution as in Feenstra, 1994) and correct for it using one of our price correction approaches; or assume that the mis-measured or omitted relative price changes satisfy an orthogonality condition similar to expression (15) above. Finally, a more practical consideration that favors aggregation is that relative Engel curves for subgroup g may be strictly monotonic while consumption of specific varieties within g are zero (and thus relative Engel curves are flat) for some locations, periods, and/or ranges of income.

Taken together, these aggregation results are particularly valuable when attempting to estimate price indices and welfare from highly-disaggregated data that are only available for some subset of consumption G —most prominently barcode-level retail scanner data.

3.2.3 Bias in Engel Curve Estimation

Omitted variables can bias estimates of relative Engel curves just as they can traditional Engel curves—biased in the sense that the estimated curve does not provide a causal estimate of how consumption patterns change with exogenous changes in income. One source of such bias is if rich and poor households (along the x-axis) pay different prices for the same goods (with relative expenditures on the y-axis).³² An important example in the Indian context is the Public Distribution System (PDS) which provides poor households with subsidized staples. However, even if the curves themselves are not causally identified, our method still uncovers unbiased estimates of price index and welfare changes, as long as the price vector faced by households is a function of real income.

To see this point, recall that the horizontal shifts in relative Engel curves recover the price index from the change in nominal income required to hold utility at either its initial (P^0) or final (P^1) level. Thus, even if price vectors differ with real income, we are correctly comparing the change in the price index holding utility fixed. Returning to the PDS example, eligibility criteria

³²Related to unobserved price heterogeneity, choke prices may result in the rich and poor being exposed to different sets of price changes. Zero consumption is admissible with quasi-separability. However, relative Engel curves will be flat at low or high income levels, violating the strict monotonicity restriction we use for estimation. Thus, our methodology would discard such products from the set we average over to obtain the estimated price index change.

are indeed based on a utility metric rather than just nominal income—specifically households below the poverty line are eligible, with the poverty line based on real needs. Therefore, when moving horizontally between period 0 and period 1 relative Engel curves, PDS eligibility does not change. A household initially at a utility level below (above) the PDS cutoff will be eligible (ineligible) in both periods at the utility level used to construct P^0 . A similar logic applies to price differences emanating from variation in store or product availability, as long as store entry and stocking are functions of real income (which many models of retail would predict). Section 3.2.5 addresses the closely-related topic of taste differences correlated with income.

As the above discussion makes clear, our method does not in general require that estimates of relative Engel curves are causally identified. However, concerns remain if the relationship between real income and the price vector is not stable across the two periods (in which case horizontal shifts will not hold utility constant over time). These remaining concerns can be addressed either by controlling for the location or household characteristics at the root of the price differences when estimating relative Engel curves or by estimating curves separately for these different types of location or households (as we do in our application).

A second concern frequently discussed when estimating Engel curves is idiosyncratic (i.e. household-level) measurement error in expenditures. As total outlays are simply the sum of expenditures, there will be correlated measurement error in the dependent variable (expenditure shares) and independent variable (total outlays per capita). Estimating relative Engel curves poses a similar problem, although potentially less severe since measurement error in expenditures outside of G (which does not appear in the denominator of relative shares) will simply attenuate the coefficient on total outlays per capita.³³ In either case, a similar logic applies to that discussed above, with causal identification of relative Engel curves not a necessary requirement for unbiased price index estimates. Specifically, if the distribution of measurement error is common across survey rounds (e.g. due to similar survey designs and implementation), the size of the horizontal shift remains unaffected as with the heterogeneous price example above.

Finally, note that the above discussion relates to using a single good's relative Engel curve for estimation. Since we average horizontal shifts across goods $i \in G$, as we describe in Section 3.1.2 above, any i -specific bias in the estimation of relative outlays as a function of incomes would also have to be systematically related to the slopes of the relative Engel curves across the $i \in G$ in order to bias the welfare estimates.

3.2.4 Unobserved Welfare Changes (Sample Selection)

Not all levels of household utility in period 0 are necessarily observed in period 1 and vice versa. For example, when evaluating price index changes P^0 for poor households in period 0, there

³³A direct solution would be to instrument for total outlays per capita with outside- G total outlays per capita which would address any remaining concerns regarding correlated measurement error.

may be no equally poor households in period 1 if there is real income growth (and similarly when evaluating P^1 for rich households in period 1). This means that Engel curves may not always overlap in budget share space for all income percentiles, and gives rise to sample selection concerns, especially at the tails.

These selection issues take two forms, missing goods and missing markets. Recall from Section 3.1.2 that averaging multiple price index estimates (one for each good for which we can measure the horizontal shift in its relative Engel curve) can potentially eliminate bias from measurement error in expenditures or prices within the G group. However, in the presence of such shocks, averaging over the subset of goods for which there is overlap in relative Engel curves at a given percentile h generates potential biases since overlapping and non-overlapping goods experienced different shocks. This is particularly problematic at the tails of the distribution. For example, suppose real income grew and so there is no true overlap when estimating P^0 for the poorest households. Any overlapping goods we do observe must have experienced large vertical shocks to relative Engel curves such that the resulting price index estimate makes the poorest period 1 households appear to have real incomes similar to the poorest in period 0.

To address such sample selection concerns, we exploit the fact that we observe whether a particular good has no overlap at a particular income percentile and if so, whether the missing estimate is censored P from above or from below (which depends on the sign of the slope of the relative Engel curve). Combining this information with the assumption that the distribution of price index estimates across different goods within G is symmetric for a given income percentile allows us to consistently estimate the price index change.

To implement this correction, we order the observed (i.e. overlapping goods) and unobserved (i.e. non-overlapping goods) price index estimates and take the median (which is an unbiased estimate of the mean).³⁴ In the rare cases where the median is unobserved due to most estimates being censored, we require a stronger assumption: that the distribution of price index estimates across different goods within G is uniform for a given income percentile. That allows us to solve for the mean as long as at least two goods overlap (see Sarhan, 1955). As we discuss below, the symmetry assumption alone proves sufficient to solve selection issues in our Indian application.

A different type of sample selection arises if, for a particular market, we don't observe any goods for which relative Engel curves overlap for a given percentile. In this case, we face a market-level sample selection issue when aggregating across markets. For example, if real in-

³⁴We rank estimates, placing unobserved estimates below the lowest or above the highest observed estimate depending on whether they were censored from below (e.g. when calculating P^0 for poor households or P^1 for rich households) or above (e.g. when calculating P^0 for rich households and P^1 for poor households). For example, if a relative Engel curve for some good i is upward sloping and the period 0 relative budget share for a particular income percentile is lower than any point on the period 1 curve, there is no equivalently-poor household in period 1. This implies that the missing estimate of the price index change for this percentile must be smaller than the lowest estimate obtained from other goods in G where we do observe overlap at this income percentile.

comes grew there may be missing markets among poor percentiles for P^0 and rich percentiles for P^1 . In practice, we find that almost no markets are missing after we implement the good-level selection correction above (i.e. we observe overlap in strictly monotonic relative Engel curves for at least two goods for close to every decile-market pair in our sample). Therefore, the good-level selection correction is sufficient to solve market-level selection issues. Were it not, we could apply existing two-step Heckman selection corrections or make assumptions on the distribution of estimates across markets to recover the missing markets for a given percentile h .

3.2.5 Taste Heterogeneity

Finally, we consider four concerns related to taste differences.

Statistical Demand

While taste heterogeneity correlated with incomes or price changes poses challenges to estimation that we discuss below, even random preference heterogeneity across households requires additional assumptions when moving from the theory in Proposition 1 to estimation. The demand patterns that we estimate in the data (for a given level of household income) are “statistical” in the sense of Lewbel (2001), i.e. a conditional expectation that may not be rational and quasi-separable even if (heterogeneous) individual preferences satisfy these conditions. Extending the approach of Lewbel (2001) to quasi-separability, we provide conditions such that these statistical demands satisfy both Slutsky conditions (i.e. are “integrable”) and equation (16), which holds if and only if preferences are quasi-separable.

Corollary. *Suppose that demand $x_i(y, p, z)$ for each individual indexed by z is both rational and quasi-separable in group G of goods. Statistical demand $X_i = E[x_i|y, p, x_G]$ is integrable and quasi-separable in G if, and only if, the conditional covariance matrix L between expenditure and income effects, i.e. a matrix with elements $L_{ij} = Cov\left(x_j, \frac{\partial x_i}{\partial \log y} \mid y, p, x_G\right)$, is symmetric with zero off-diagonal blocks in $(i, j) \in G \times NG$ and $(i, j) \in NG \times G$ (sufficient and necessary conditions). If, in addition, matrix L is semi-definite negative, statistical demand is both rational and quasi-separable (sufficient conditions).*

Symmetry and semi-definite negativity of matrix L are conditions already laid out in Lewbel (2001). The additional restrictions on the $G \times NG$ and $NG \times G$ blocks ensure quasi-separability. Appendix A.6 provides proofs and further discussion.

Taste Heterogeneity Correlated with Income

We discuss omitted variable bias in the estimation of relative Engel curves due to heterogenous price vectors in Section 3.2.3 above. Similar issues arise if taste differences are correlated with income in the cross-section. For example, households with different levels of education or family composition may both value certain goods more and have different average incomes. As was

the case for price heterogeneity by income, and with the same caveats, these relationships do not necessarily bias our estimates of horizontal shifts in relative Engel curves even if they confound attempts to estimate causal relationships between consumption patterns and income.

Specifically, if these taste differences are, directly or indirectly, functions of real income—e.g. richer households may acquire more education or have more children thereby changing their tastes—traveling horizontally between relative Engel curves at a given initial (or final) level of utility holds tastes constant. And as above, any remaining bias can be addressed directly by controlling for household characteristics when estimating relative Engel curves or by estimating curves separately for different types of household (we pursue both in our application).

Taste Heterogeneity Correlated with Price Changes

A different challenge arises if tastes differ within an income percentile and those taste differences correlate with relative price changes across goods. In this scenario, the price index and welfare changes for a given income percentile will differ by household type. More precisely, Appendix A.6 shows that our method will, to the first order, yield a weighted average of price index changes: $\tilde{P}^1(y_h^1) \approx \sum_z w_z^1(y_h^1) P_z^1(y_h^1)$ with weights given by the relative Engel slopes of household type z : $w_z \equiv \sum_i (\beta_{i,z}^1 / \bar{\beta}_i^1) / \sum_{z'} \sum_i (\beta_{i,z'}^1 / \bar{\beta}_i^1)$. If, instead, one wants to obtain the welfare change for a particular household type, such as households with large family sizes, we can carry out our procedure just for those households.

Taste Changes Over Time

The final set of issues arise when household tastes change over time. Such taste changes are problematic if they are systematically related to differences in slopes of relative Engel curves across goods. To be precise, we can derive an orthogonality condition analogous to the orthogonality condition for measurement error in expenditures in equation (6) above. Denoting taste shocks—shifts in within- G budget shares conditional on prices and income—by $\Delta \log \alpha_{ih}$ and abstracting from relative price changes, we obtain the bias on $\log P^1$:³⁵

$$\frac{1}{G} \sum_{i \in G} \left(\log(y_h^1) - \log(E_{iG}^0)^{-1} \left(\frac{x_{ih}^1}{x_{Gh}^1} \right) \right) \approx \log(P^1) - \frac{1}{G} \sum_{i \in G} \left((\beta_{ih}^0)^{-1} \Delta \log \alpha_{ih} \right). \quad (17)$$

If taste shocks across i within subset G are orthogonal to the local slope of i 's relative Engel curve in period 0 (or period 1 to identify P^0), the bias averages to zero across goods.

Unfortunately, such a condition is not in general testable. To see this, note that knowledge of expenditure and price changes within G and the shape of within- G preferences—i.e. the moments and parameters that allow estimation of taste shocks under separable but homothetic preferences—are insufficient in our context where preferences are non-homothetic. In such

³⁵To ensure shares sum to unity within G , we assume that these taste shocks sum to zero. Such shocks can be defined, for example, in terms of price shifters as in Redding and Weinstein (2020).

settings, to estimate taste shocks we must also net out the changes in within- G relative expenditures due to changes in real income, which would require observing the full vector of price changes—data we argue are absent in most if not all empirical contexts.³⁶ One scenario that may violate this orthogonality condition is if household types have different tastes and there are compositional changes over time (e.g. increases in education). We can (and do) address this concern explicitly by separately estimating and comparing price index changes for different household types.

4 Application: Rural Indian Welfare 1987–2000

We now apply our methodology to estimate changes in rural Indian welfare between 1987 and 2000. This exercise serves not just as a proof of concept but is also of independent interest given the large literature—and large amount of disagreement—regarding how real incomes changed over this period of major economic reforms, particularly for India’s almost 250 million rural poor at or below the poverty line (a literature dubbed the “Great Indian Poverty Debate” by Deaton and Kozel, 2005).

4.1 Data

Following Deaton and Kozel (2005), we draw on rural households in two of India’s “thick” NSS survey rounds covering 1987/88 (43rd round) and 1999/2000 (55th round).³⁷ Each round provides us with detailed expenditure data on approximately 80,000 households residing in more than 400 Indian districts. Households are asked about their expenditures on 310 goods and services in each survey round. Examples include wheat, turmeric, washing soap and diesel. The sum of all expenditures over 30 days provides our measure of total household outlays. Given limited saving in India this will closely approximate nominal income (and even more closely permanent income). As noted previously, we use the word outlays interchangeably with income. The surveys also contain basic household characteristics, district of residence, and survey weights that we use to make the sample nationally representative.

We use these data to estimate changes in household price indices and welfare for rural Indians between 1987 and 2000. We do this for 9 income deciles (i.e. percentiles 10, 20, ..., 90) in each district. Given the need to non-parametrically estimate relative Engel curves, we restrict attention to the 249 districts where we observe at least 100 households in both survey rounds. (As we show, results are not sensitive to this restriction.)

To obtain the subset of goods with reliable price data, we mimic the approach of Deaton and

³⁶It is precisely because these data are not typically available that we require our methodology that attributes changes in within- G demands (conditional on within- G prices) to changes in the price index. Taste shocks across goods within G are thus not separately identified when allowing for non-homotheticity.

³⁷As we discuss below, focusing on rural areas also allows us to validate our estimates since well-measured food and fuel prices cover most of the consumption bundle for poor rural households.

Tarozzi (2005) who carefully analyze NSS expenditure surveys to identify the subset of goods for which prices can be measured using unit values (i.e. expenditures divided by quantities) and the resulting prices are robust to concerns about unobserved product quality or variety.

Appendix B describes their procedure in detail, as well as data cleaning procedures to remove obvious price outliers. Here, we briefly summarize their methodology to identify goods with reliable price information. First, they exclude all goods and services categories where quantity data are not recorded. Next, they further exclude the clothing and footwear categories for which quantity data exist (e.g. 2 pairs of leather boots/shoes) but where product descriptions are too broad and styles too numerous to generate reliable unit values. The remaining goods are all food and fuel products. Third, they discard any foods and fuels where the variation in prices within localities suggests that these products likely contain multiple varieties or quality levels; either because there is strong evidence of multi-modal price distributions (e.g. liquid petroleum gas), or due to the combination of high price dispersion and broad product descriptions (e.g. “other milk products”). Finally, they discard products where changes in the unit of measurement over rounds make temporal comparisons impossible.

These restrictions leave us with a sample of 132 food and fuel goods for which we have unit values and where issues related to multiple quality levels are minimized. To alleviate the remaining concern of measurement error when using unit values, we again follow Deaton and Tarozzi (2005) and use the median unit value from each district and survey round (our market and period unit, respectively) as our price measures. We echo Deaton and Tarozzi in arguing that the combination of these procedures provides reliable price data for this subset of goods.

The final column of Table 1 lists these 132 goods that cover, on average, 75 percent of household consumption in our sample.³⁸ As we emphasize throughout the paper, this subset of goods with reliable price data are crucial for our estimation since they allow us to implement Proposition 1 and compute exact or first order price corrections, as well as to test for and assess potential bias from violations of quasi-separability.

Finally, we note that in the 55th round, the surveys included a 7-day recall period for all food products (in addition to the standard 30-day recall period asked in the 43rd round). While we only use the responses to the 30-day recall questions, Deaton (2003a,2003b) and others show that households inflated their 30-day reports to be consistent with their 7-day ones. Thus, this “recall bias” raises reported total outlays (the numerator for evaluating changes in real incomes) even using the 30-day recall data and is at the center of the Great Indian Poverty Debate. In Section 4.3, we show that our approach is robust to this recall bias as relative consumption

³⁸As the survey questionnaires change slightly over time, we aggregate a small number of goods to the most disaggregate classification reported consistently across rounds. In three cases we must combine purchases made at a discount through India’s Public Distribution System (available to households below the poverty line) and those bought at regular markets. Appendix B discusses these aggregations and Section 3.2.3 explains how our methodology accommodates price vectors that vary with real income.

patterns within product groupings are unaffected by the additional 7-day recall question.

4.2 Product Aggregation and Product Groups

To reduce measurement error when estimating relative Engel curves for rarely consumed items, we aggregate these 132 food and fuel items with well-measured prices to the second-lowest level of aggregation in the NSS surveys, which yields 34 goods indexed by g (listed in the third column of Table 1). The results in Section 3.2.2 prove that such an aggregation is admissible, and that we can implement price corrections, as long as we can measure price indices $P_g(p_g, U)$ for these 34 goods. We use a Stone price index to construct such indices (Appendix Table C.2 reports descriptive statistics for these price changes).³⁹ This aggregation dramatically reduces the share of empty product-by-period-by-market cells (from 50 percent to less than 15 percent as shown in Appendix Figure C.1), and moving to the next highest level of aggregation (8 goods) provides little additional benefit.⁴⁰

We divide these 34 aggregate g goods into three broader consumption groups shown in the first column of Table 1: raw foodstuffs (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 groups each form a quasi-separable G group, with all remaining goods (e.g. processed food, manufactures and services) excluded as part of the NG group. We combine estimates from goods within all three G groups by taking medians following the discussion in Section 3.2.4.⁴¹ As we describe below, Figure 6 explores robustness across 108 perturbations of sensible G groupings, including a single G group.

4.3 Changes in Indian Price Indices and Welfare Over Time

Before describing the results of our approach and comparing them to estimates derived from existing Indian CPI statistics, we first summarize the changes in nominal income between 1987 and 2000. Figure 2 plots growth rates in total household outlays per capita for each decile of the local income distribution (using population-weighted averages of log changes across all 249 rural districts).⁴² Nominal income growth exceeded 200 percent and there is a clear and strong pattern of convergence over this 13-year period, with outlays per capita rising substantially

³⁹Specifically, we aggregate the observed log price changes for the 132 items i to 34 goods g using survey-weighted mean initial expenditure shares across the $i \in g$ as weights. We compute price changes for each i from changes in district median unit values as described in Data Appendix B. When unit values are observed in the district for one but not the other period, we replace i 's missing price change with the state-level change.

⁴⁰Appendix Figure C.4 reports qualitatively similar inflation estimates using these alternate levels of aggregation.

⁴¹In principle, comparing estimates obtained from different G groups provides an over-identification test (price index estimates from different G groups should be identical if there is no misclassification of goods into quasi-separable groups and orthogonality conditions on measurement error and taste shocks are satisfied). However, given the limited number of products in our setting (recall we have about 11 goods in each of the 3 G groups and for a given market-decile not all goods have both strictly monotonic and overlapping relative Engel curves), these conditions are unlikely to be satisfied without pooling the estimates.

⁴²For each decile, we report percentage changes for incomes, price indices and welfare calculated by exponentiating the population-weighted mean of district-level log changes between 1987 and 2000.

faster for the poor than for the rich. Our non-homothetic price indices allow us to determine whether this nominal income convergence translated into convergence in standards of living.

Figure 3 presents our price index estimates using the methodology outlined in Section 3.1 (from hereon the “AFFG Price Index” after the authors initials).⁴³ The left panel presents our estimates absent any within- G price corrections—i.e. simply utilizing the horizontal shifts in relative Engel curves for goods in our three G groups. As above, we plot population-weighted averages across districts by decile. The remaining panels of Figure 3 apply the two variants of Proposition 1 described in Section 3.1.3 that draw on the well-measured price changes we have for goods in our food and fuels G groups to account for potentially confounding within- G relative price changes. The middle panel displays the first-order price correction where we assume a common elasticity of substitution of $\sigma_G = 0.7$ based on averages from Cornelsen et al.’s (2015) systematic review of food price elasticities in low income countries that uses similar levels of aggregation to our 34 goods. The right panel plots the exact price correction using the isoelastic correction (non-homothetic CES) in equation (10) with the same elasticity assumption.

The first thing to notice is that the estimated inflation rates across deciles change very little after adjusting for relative price changes within G groups using either the first-order or exact approach. This is not simply the result of assuming a single elasticity that limits patterns of cross-price substitution. Appendix Figure C.3 presents exact corrections using the more flexible multi-nest non-homothetic CES demands in equation (11), calibrated using two different sets of price elasticities from the literature.⁴⁴ Under all three parameterizations, estimates are almost identical to the uncorrected price index for all income deciles. Recall from equation (8) that, to the first order, our estimates are unbiased if within- G price changes are uncorrelated with slopes of relative Engel curves. Thus, the fact the estimates change little with our price corrections implies that relative price changes within our three food and fuel G groups are either small or only weakly related to income elasticities in our context. To streamline the exposition given these results, we focus our remaining analysis on the no price correction approach (labeled “AFFG NPC price index”). In all cases, we draw similar conclusions using the first-order or exact price correction estimates.

Before discussing magnitudes and differences in inflation across deciles of the income distribution, it is instructive to plot our AFFG approach alongside the leading existing CPI estimates for rural India. The left panel of Figure 4 repeats our AFFG NPC price index. The middle panel plots Paasche and Laspeyres price index estimates using the methodology of Deaton

⁴³As an example of the horizontal shifts in relative Engel curves we use to obtain our price index estimates, Appendix Figure 3.1.1 plots relative Engel curves in 1987/88 and 1999/2000 for one g good, salt, as a share of the G group “other food products” for the largest districts in the North, East, South and West of India.

⁴⁴Specifically, rather than a single elasticity governing own and cross-price elasticities, we now have nine parameters governing these elasticities, as described in Section 3.1.3. We use Kumar et al.’s (2011) demand system estimates for food items in India that are calculated using NSS data, and a meta-analysis of price elasticity estimates for various commodities by Fally and Sayre (2018). The resulting parameter values are presented in Appendix Table C.3.

(2003b) that draws on observed price changes weighted by average district-level expenditure shares for the 132 food and fuels items where price data are deemed reliable.⁴⁵ Mechanically, these price indices do not vary across the income distribution. The right panel of Figure 4 relaxes this homotheticity by using district-decile specific expenditure shares when calculating Paasche and Laspeyres price indices. We obtain bootstrapped confidence intervals for all three indices by sampling with replacement 1000 times from the distribution of households within each district-survey round and plotting the 2.5 and 97.5 percent envelope of price index estimates at each decile (bootstrapping over the entire procedure in the case of the AFFG price index, including the non-parametric estimation of relative Engel curves).

Two main findings emerge. First, our AFFG approach generates broadly similar estimates of Indian consumer price inflation among low-income deciles compared to existing CPI estimates that are based on changes in observed food and fuel prices. Specifically, all three approaches predict price rises of between 160 and 180 percent for the poorest deciles. Since food and fuels represent a sizable fraction of rural household consumption for poor households in India (more than 80 percent for the poorest decile, averaging across both survey rounds), this finding is reassuring—particularly since we are comparing a standard price index that explicitly uses observed price changes to our approach that exploits very different variation coming from horizontal shifts in relative Engel curves.

Second, we estimate that cost of living inflation has been substantially higher for poor households compared to the rich, the opposite of what one would infer from the food and fuel Paasche and Laspeyres indices which are slightly pro-poor. Figure 5 combines the estimated changes in nominal incomes and price indices to obtain welfare changes (EV and CV in our approach, and real income for the standard CPI approach). The income-specific inflation rates estimated using the AFFG approach eliminate any convergence in welfare between the rich and poor over this period. In fact, if anything, welfare grew more for rich households. This finding contrasts starkly with the changes in real income calculated using food and fuel Paasche and Laspeyres indices which slightly magnify the already substantial convergence seen in nominal incomes. This result also stands in contrast to Almås and Kjelsrud (2017) who estimate non-homothetic price indices using a Quadratic AIDS demand system that does not impose quasi-separability but requires knowledge of price changes for the full consumption basket, including manufactures and services.⁴⁶ They find that inflation was pro-poor over the period 1993–2005.

Why are our price index estimates lower for richer households? The most likely explanation is that high-income households disproportionately benefited from price drops, new varieties,

⁴⁵As above, price changes are computed from changes in district median unit values for each of the 132 items. We calculate Laspeyres and Paasche price indices using survey-weighted mean expenditure shares at the district level (thus the index is democratic not plutocratic). We replace missing district-level price changes with state-level ones.

⁴⁶See footnote 8 for a description of how Almås and Kjelsrud (2017) utilize India's non-food CPI to navigate the lack of well-measured price data for categories beyond food and fuels.

and quality increases in consumption categories where price measurement is challenging. In particular, the rich spent a large and increasing share of their budget on durables such as manufactures and on services. These are exactly the categories for which unobserved quality differences make price data unreliable and so are omitted in Deaton’s CPI approach which only covers well-measured food and fuels, and are crudely captured, if at all, by the government non-food CPI in the Almås and Kjelsrud (2017) approach. Lower inflation in these specific categories is consistent with the fact that the Indian trade reforms were centered on manufacturing intermediates which substantially raised the quality and variety of Indian manufactures (Goldberg et al., 2010); and that there was a dramatic increase in share of services in GDP over the reform period (Mukherjee, 2015). Appendix D discusses this explanation further and contains four pieces of corroborating evidence: that expenditure shares were greater for the rich in these categories; that government-measured inflation in these categories was lower; that relative Engel curves are steepest in these categories; and that these categories saw the most new product entry based on the Indian Prowess microdata used by Goldberg et al. (2010) and others.

Beyond accounting for inflation in hard-to-measure categories, our methodology is also immune to the concerns that lie at the center of the Great Indian Poverty Debate. Recall that the 1999-2000 NSS added a 7-day recall period for food expenditures which inflated answers to the consistently-asked 30-day recall questions. The most influential solution, that of Deaton (2003a), adjusts food expenditure using the mapping between food and fuels expenditure (for which no additional recall period was added) from earlier rounds. That solution requires that relative price of food and fuels did not change. In contrast, our welfare estimates are robust to the additional recall period as long as it did not change *relative* consumption shares within a given food or fuel group G . This condition is testable using the thin NSS round 54 (1998) where, in order to test proposed changes to the surveys, households were randomly assigned to different recall periods. Consistent with our claim, Appendix Table C.1 shows that the choice of recall period did not affect relative consumption shares within our G groups.⁴⁷ Thus, our finding of no convergence in real incomes has the potential to inform, and revise the conclusions of, the Great Indian Poverty Debate summarized in Deaton and Kozel (2005).

4.4 Validation Results

In this subsection, we perform a number of validation exercises that follow from our corollaries in Section 3.2, as well as reporting several additional context-specific robustness checks.

Quasi-Separability and Misclassification

We first investigate bias from potential violations of quasi-separability due to misclassifying products into G groups. To this end, we re-estimate our price indices for each decile and mar-

⁴⁷In addition, Appendix Figure C.5 shows similar patterns of pro-rich inflation between the 1987/88 and 1994/95 survey rounds when the questionnaire was unchanged.

ket across 108 sensible splits of our g goods into plausibly quasi-separable groupings G .⁴⁸ Figure 6 presents the estimation results for each decile, plotting our baseline point estimate on top of the mean and the 2.5th–97.5th percentile range of point estimates from the 108 sensible G groupings. Reassuringly, our baseline is close to the mean for every decile of the income distribution. The 2.5th–97.5th percentile ranges are also reasonably tight—suggesting that the conditions under which misclassification bias is small (equation 15) are met in our setting.

Next, we present our two preferred tests of quasi-separability from Section 3.2.1, one using proxies for price changes outside of G (Corollary QS1) and one using outside- G expenditures (Corollary QS2). In the first test, for each good $i \in G$, we regress log changes of within- G relative expenditure shares on within- G log price changes and controls for changes in household welfare: our estimates of CV and EV as well as log total outlays per capita. We further include log changes of outside- G price indices and test whether they affect within- G relative expenditures (they should not if preferences are quasi-separable in group G). For outside- G price information, we make use of available (but imperfect) state-level price indices computed for rural households by the Indian Labor Bureau for two categories: 1) clothing, bedding and footwear, and 2) miscellaneous (which includes all services and durable manufactures).

This specification generates $I - 1$ regression equations for each of our three G groups (31 in total). To obtain the correct error distribution, we first randomly draw the two outside- G price changes 500 times and perform a joint test that these (fictitious) outside- G price changes have a zero coefficient in every equation. We then compute the test statistic for joint significance of changes in the actual outside- G price indices and compare it to the χ^2 test distribution from our random draws.⁴⁹ Panel A of Figure 7 overlays the value of the test statistic (the thick red line) on top of the statistic’s empirical distribution calculated above, with further details of the test provided in the table notes. Reassuringly, we cannot reject the null that preferences are quasi-separable between goods within and outside our G groups, with a p-value of 0.44.

A natural limitation of the previous test is that it requires data on prices for non- G categories such as manufactures and services, yet our methodology is motivated by the difficulty in reliably measuring prices for these sectors. Our second test (Corollary QS2) does not require such price information. Instead, we flip the previous test and ask whether changes in outside- G log expenditures (the sum of outlays spent on outside- G goods) respond to within- G log price changes once we control for relative expenditure share weighted changes in within- G log prices

⁴⁸As shown in Table 1, the 34 g products fall into three high-level groups (raw food, other food and fuel) and 8 subgroups within those. To discipline plausibly quasi-separable nests G , we impose that a g can only be bundled together with other gs in the same high-level group. Additionally, different gs within one of the 8 subgroups cannot be grouped into more than one G (as they are likely closely related). With these restrictions, we generate 105 possible ways of allocating gs into G groups based on tuples: i.e. $(2^4 - 1) \times (2^3 - 1) \times 1 = 105$. Finally, we add: only 1 G group across all 34 products, 2 G groups (food and fuel), and 8 G groups (one for each subgroup above).

⁴⁹For both tests, we cluster standard errors at the market level, use survey weights, and perform the test at the median decile of households using cross-market variation in price and expenditure changes across survey rounds.

and the same controls for changes in utility as before (they should not if preferences are quasi-separable in group G). As above, we obtain the statistic's empirical distribution by randomly drawing the 34 within- G log price changes 500 times and testing for their joint significance (conditional on the within- G price index and utility controls). Panel B of Figure 7 marks the value of the F-test statistic with a thick red line laid on top of the distribution of the test statistic. As with our first test, we cannot reject the null that preferences are quasi-separable between goods within and outside our G groups, with a p-value of 0.33.

Sample Selection Issues

As described in Section 3.2.4, our baseline estimates address sample selection issues due to non-overlapping relative Engel curves by ranking both missing and non-missing estimates and taking the median under the assumption of a uniform distribution of estimates across $g \in G$. Appendix Figures C.6–C.8 illustrate and assess these sample selection issues. The left panel of Figure C.6 presents price index estimates that do not correct for non-overlap issues and simply average over non-missing goods. As anticipated, the biggest discrepancies with our baseline (the right panel) occur for P^0 among the poorest deciles and P^1 among the richest deciles. These are the households where we would expect no true overlap in a growing economy.⁵⁰

The middle panel of Appendix Figure C.6 implements only the first step of our selection correction, applying symmetry but not uniformity. This step alone eliminates almost all the discrepancy between P^0 and P^1 due to sample selection issues and generates very similar estimates to our uniformity baseline (right panel). However, by only imposing symmetry, we lose any market-decile pairs for which the median ranked good has no overlap. As shown in Appendix Figure C.8, a substantial number of pairs are missing when only imposing symmetry (particularly for P^0 since the distribution of log total outlays per capita is right-skewed). However, we obtain estimates for essentially all market-deciles once uniformity is imposed and so market-level selection issues do not arise under our baseline specification.

Taste Heterogeneity and Taste Changes

We now investigate concerns that our estimates may be affected by taste heterogeneity across households or taste changes over time (see Section 3.2.5). Appendix Figure C.9 recalculates price indices using non-parametric relative Engel curves that condition on a set of linear controls for household characteristics.⁵¹ Reassuringly, results change little, suggesting that systematic bias in estimates of cross-sectional Engel curves is unlikely to be driving our findings.

⁵⁰Appendix Figure C.7 illustrates this fact by showing the frequency of non-overlapping estimates by decile, broken out by type of non-overlap (censored from above or from below) that we use to rank missing estimates.

⁵¹In particular, for each good and market (pooling across both periods) we estimate coefficients on the following controls: a scheduled caste dummy, a literacy of household head dummy, log of household size, and the share of children in the household. We then use relative Engel curves for each good-period-market evaluated at the controls' market-level median (i.e. holding demographic characteristics fixed across periods).

Appendix Figure C.10 corroborates this finding by presenting separate price index estimates for different types of rural households; small versus large households, high versus low education, young versus old, and literate versus illiterate (with the last three comparisons based on characteristics of the household head). Recall from Section 3.2.5 that these exercises are informative on a number of fronts. First, by estimating Engel curves separately across demographic groups, we limit potential bias in estimates of cross-sectional Engel curves. Second, we can explore to what extent different types of household experienced different inflation rates, both on average and by income decile, as a result of taste heterogeneity. Third, we can address concerns that the composition of household types may have changed over time, biasing our estimates if taste heterogeneity across types is systematically related to slopes of relative Engel curves (e.g. if average education or household size changed over time and educated or large households have different tastes). The fact that the price index estimates show very similar patterns for different household types provides reassurance that taste heterogeneity and taste changes (at least those due to compositional changes) are not driving our findings.

Additional Robustness Checks

We report several additional robustness checks. Appendix Figure C.11 presents results for alternative bandwidth choices when non-parametrically estimating relative Engel curves and for alternative strategies to deal with noise at the tails. Appendix Figure C.12 reports results without restricting attention to markets with at least 100 household observations in both survey rounds. Reassuringly, results are qualitatively similar to our baseline estimates in both cases. Appendix E assesses our methodology via a Monte Carlo simulation. We generate a fictitious second-period dataset with the same number of households and statistical error in relative Engel curves as in our actual sample but fixing inflation to be a step function that declines by income decile. Simulating our methodology over 250 error draws, we find that the truth lies within the 95th percentile envelope of estimates for all deciles, although the addition of measurement error slightly attenuates the slope of the mean estimates with respect to income.

Application to India's 1991 Trade Reforms

Appendix F uses our methodology to revisit the impact of India's 1991 trade reforms on the welfare of rural households in India. Closely following Topalova (2010)—but replacing her outcomes (district-level rural poverty rates and per capita outlays) with our welfare estimates—we find that the adverse effects of import competition on local labor markets are borne by households across the income distribution, including by rural households in the richest income deciles.

5 Conclusion

Measuring changes in household welfare and the distribution of those changes is challenging and typically requires the researcher to observe the full vector of quality- and variety-adjusted price changes—an incredibly difficult task for categories such as manufacturing and services. In this paper, we propose and implement a new approach that only requires reliable price information for some quasi-separable subset of products G . Horizontal shifts in relative Engel curves within this group—adjusted for within- G price changes—reveal changes in household price indices and welfare across the income distribution.

We apply this new method to measure changes in household welfare in rural India. We find that consumer price inflation was substantially higher for poor households than rich, essentially eliminating the convergence seen in nominal incomes. This finding is missed by standard price indices using the subset of consumption where prices are well measured.

Beyond providing a deeper understanding of India’s economic reforms, we believe our methodology is widely applicable in the many settings where expenditure survey data are available or can be easily collected. Given the increasing availability of survey microdata over time and across space, and the growing interest in distributional analysis, the usefulness of such an approach is only likely to grow.

References


- ALMÁS, I. (2012): “International income inequality: Measuring PPP bias by estimating Engel curves for food,” *American Economic Review*, 102(2), 1093–1117.
- ALMÁS, 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*.
- ALMÁS, I., AND A. KJELSRUD (2017): “Rags and riches: Relative prices, non-homothetic preferences, and inequality in India,” *World Development*, 97, 102–121.
- ARGENTE, D., AND M. LEE (2020): “Cost of living inequality during the great recession,” *Journal of the European Economic Association*.
- ARGENTE, D. O., C.-T. HSIEH, AND M. LEE (Forthcoming): “Measuring the cost of living in Mexico and the US,” *American Economic Journal: Macroeconomics*.
- 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.
- BAQAEE, D., A. BURSTEIN, AND Y. KOIKE-MORI (2022): “A new method for measuring welfare with income effects using cross-sectional data,” *NBER Working Paper 30549*.
- BLACKORBY, C., D. PRIMONT, AND R. R. RUSSELL (1978): *Duality, separability, and functional structure: Theory and economic applications*, vol. 2. Elsevier Science Ltd.
- BLACKORBY, C., AND R. R. RUSSELL (1978): “Indices and subindices of the cost of living and the standard of living,” *International Economic Review*, 19(1), 229–240.

- BORUSYAK, K., AND X. JARAVEL (2018): "The distributional effects of trade: Theory and evidence from the United States," *Harvard University Working Paper*.
- COMIN, D., D. LASHKARI, AND M. MESTIERI (2021): "Structural change with long-run income and price effects," *Econometrica*, 89(1), 311–374.
- CORNELSEN, L., R. GREEN, R. TURNER, A. D. DANGOUR, B. SHANKAR, M. MAZZOCCHI, AND R. D. SMITH (2015): "What happens to patterns of food consumption when food prices change? Evidence from a systematic review and meta-analysis of food price elasticities globally," *Health Economics*, 24(12), 1548–1559.
- 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.
- CRAWFORD, A., AND J. P. NEARY (Forthcoming): "New characteristics and hedonic price index numbers," *Review of Economics and Statistics*.
- 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 V. KOZEL (2005): "Data and dogma: The great Indian poverty debate," *World Bank Research Observer*, 20(2), 177–200.
- DEATON, A., AND J. MUELLBAUER (1980): *Economics and consumer behavior*. Cambridge University Press.
- DEATON, A., AND A. TAROZZI (2005): "Prices and poverty in India," in *The Great Indian Poverty Debate*, ed. by A. Deaton, and V. Kozel, chap. 16. MacMillan.
- FAJGELBAUM, P., AND A. KHANDELWAL (2016): "Measuring the unequal gains from trade," *Quarterly Journal of Economics*, 131(3), 1113–1180.
- FALLY, T. (2022): "Generalized separability and integrability: Consumer demand with a price aggregator," *Journal of Economic Theory*, 203, 105471.
- FALLY, T., AND J. SAYRE (2018): "Commodity trade matters," *NBER Working Paper 24965*.
- FEENSTRA, R. C. (1994): "New product varieties and the measurement of international prices," *American Economic Review*, 84(1), 155–57.
- GOLDBERG, P. K., A. K. KHANDELWAL, N. PAVCNIK, AND P. TOPALOVA (2010): "Imported intermediate inputs and domestic product growth: Evidence from India," *Quarterly Journal of Economics*, 125(4), 1727–1767.
- GORMAN, W. M. (1965): "Production functions in which the elasticities of substitution stand in fixed proportions to each other," *Review of Economic Studies*, 32(3), 217–224.
- GORMAN, W. M. (1970): "Quasi-separable preferences, costs, and technologies," *University of North Carolina Working Paper*.
- GORMAN, W. M. (1976): "Tricks with utility functions," *Essays in economic analysis: The proceedings of the Association of University Teachers of Economics, Sheffield 1975*.
- HAMILTON, B. W. (2001): "Using Engel's Law to estimate CPI bias," *American Economic Review*, 91(3), 619–630.

- HANDBURY, J. (2021): “Are poor cities cheap for everyone? Non-homotheticity and the cost of living across U.S. cities,” *Econometrica*, 89(6), 2679–2715.
- HANOCH, G. (1975): “Production and demand models with direct or indirect implicit additivity,” *Econometrica*, 43(3), 395–419.
- HOTTMAN, C., AND R. MONARCH (2020): “A matter of taste: Estimating import price inflation across U.S. income groups,” *Journal of International Economics*, 127.
- JARAVEL, X. (2019): “The unequal gains from product innovations: Evidence from the U.S. retail sector,” *Quarterly Journal of Economics*, 134(2), 715–783.
- JARAVEL, X., AND D. LASHKARI (2022): “Measuring growth in consumer welfare with income-dependent preferences,” *CEPR Discussion Paper 16866*.
- KUMAR, P., A. KUMAR, P. SHINOJ, AND S. RAJU (2011): “Estimation of demand elasticity for food commodities in India,” *Agricultural Economics Research Review*, 24(1), 1–14.
- LEWBEL, A. (1991): “The rank of demand systems: Theory and nonparametric estimation,” *Econometrica*, pp. 711–730.
- (2001): “Demand systems with and without errors,” *American Economic Review*, 91(3), 611–618.
- LEWBEL, A., AND K. PENDAKUR (2009): “Tricks with Hicks: The EASI demand system,” *American Economic Review*, 99(3), 827–863.
- LIGON, E. (2019): “Estimating household welfare from disaggregate expenditures,” *UC Berkeley Working Paper*.
- MATSUYAMA, K. (2015): “The home market effect and patterns of trade between rich and poor countries,” *Northwestern University working paper*.
- MUKHERJEE, A. (2015): “Services sector in India: Trends, issues, and the way forward,” *Eurasian Geography and Economics*, 56(6), 635–655.
- 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.
- REDDING, S. J., AND D. E. WEINSTEIN (2020): “Measuring aggregate price indices with taste shocks: Theory and evidence for CES preferences,” *Quarterly Journal of Economics*, 135(1), 503–560.
- SARHAN, A. (1955): “Estimation of the mean and standard deviation by order statistics. Part II,” *Annals of Mathematical Statistics*, 26(3), 505–511.
- TAROZZI, A. (2007): “Calculating comparable statistics from incomparable surveys, with an application to poverty in India,” *Journal of Business & Economic Statistics*, 25(3), 314–336.
- 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.
- YOUNG, A. (2012): “The African growth miracle,” *Journal of Political Economy*, 120(4), 696–739.


6 Figures and Tables

Figure 1: Illustration of Lemma 1



Notes: Figure illustrates how price indices and welfare can be recovered from horizontal shifts in relative Engel curves (i.e. expenditure on good i as a share of total expenditure on group G plotted against log total outlays per capita) when relative prices within group G are unchanged but prices outside of G are unrestricted. Period 0 and period 1 relative Engel curves for good i denoted by $E_{iG}^0(y_h^0)$ and $E_{iG}^1(y_h^1)$, respectively. See Section 2 for further discussion.

Figure 2: Rural Indian Growth in Nominal Income 1987/88–1999/2000



Notes: Figure shows the percentage change in rural total outlays per capita between 1987/88 and 1999/2000 for each decile of the local per-capita outlay distribution (averaged across districts using population weights). Bootstrapped confidence intervals are based on sampling with replacement 1000 times from the distribution of households within each district-survey round and plotting the 2.5 and 97.5 percent envelope of nominal income estimates at each decile. See Section 4.3 for further discussion.


Figure 3: Rural Indian Cost of Living Inflation 1987/88–1999/2000: AFGG Price Index with No Price Correction, First-Order Price Correction and Exact Price Correction

88



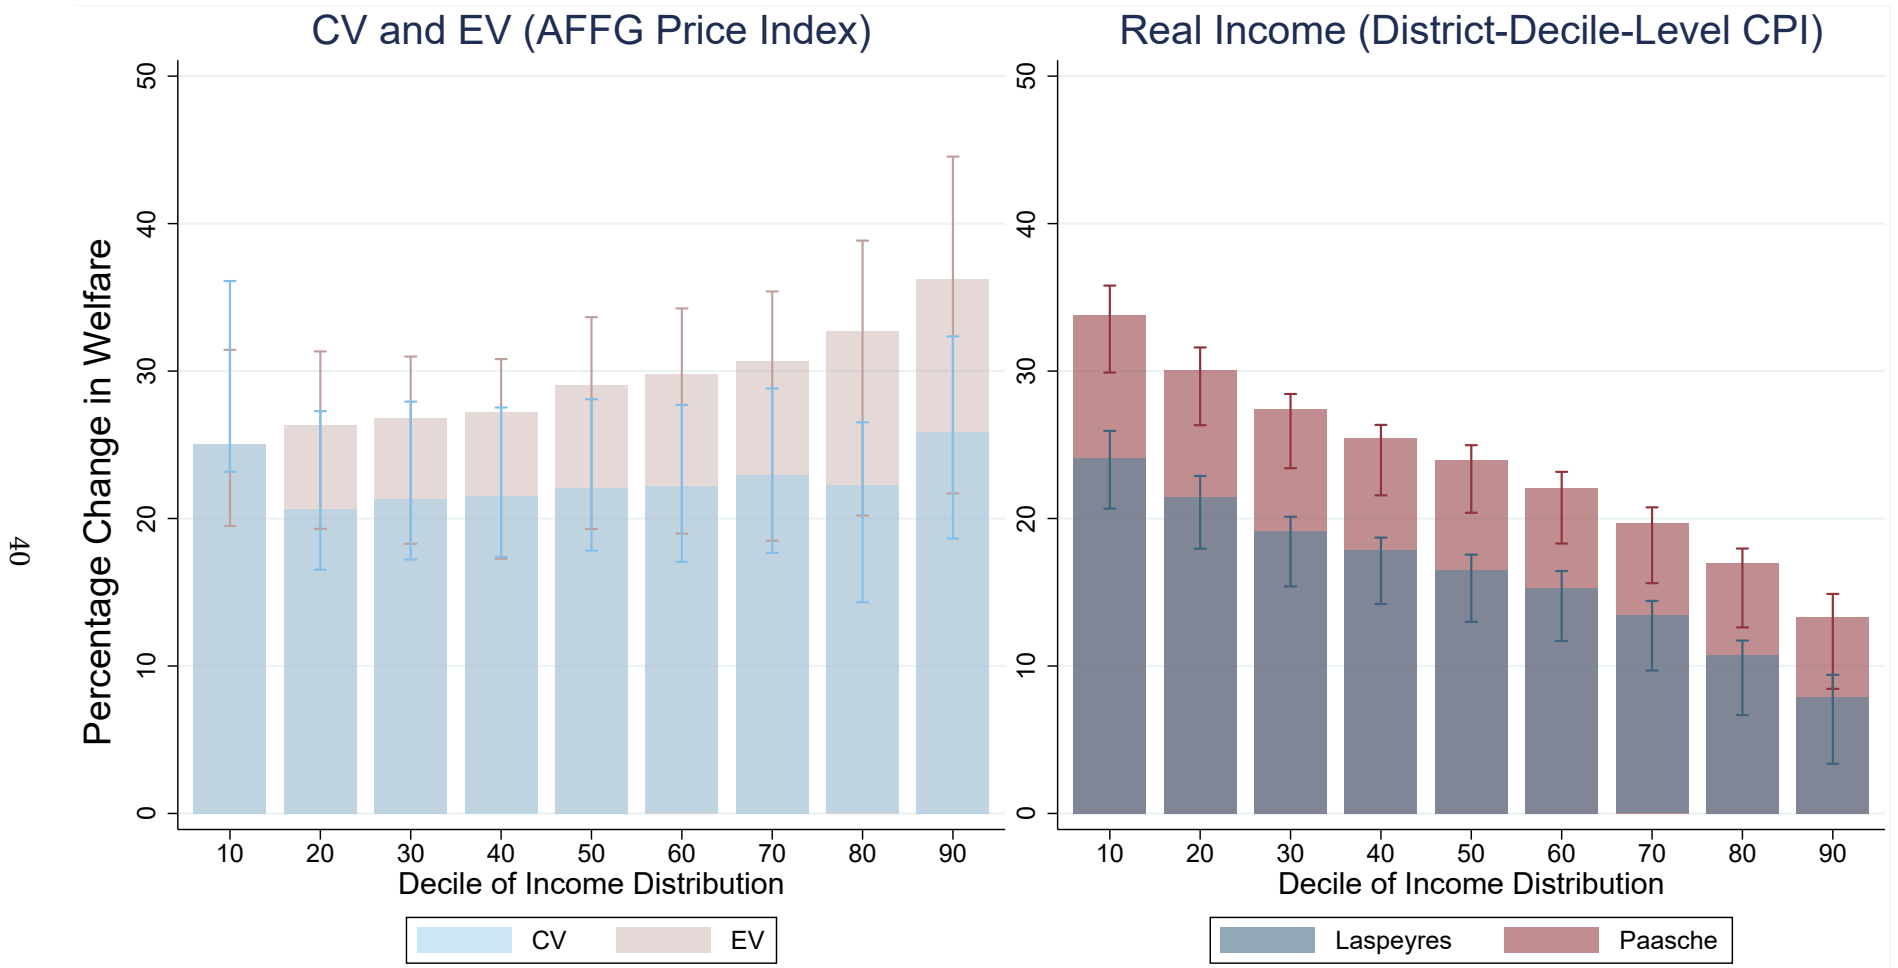
Notes: Figure shows the percentage change in the rural AFGG price index between 1987/88 and 1999/2000 for each decile of the local per-capita outlay distribution (averaged across districts using population weights). Panels show estimates both with and without corrections to account for relative price changes within G groups. Left panel reports the uncorrected price index change. Middle panel applies the first-order price correction and right panel applies the exact correction, both described in Proposition 1 and Section 3.1.3, using $\sigma_G = 0.7$. See Section 4.3 for further discussion.

Figure 4: Rural Indian Cost of Living Inflation 1987/88–1999/2000: Comparison to Existing CPI Estimates




Notes: Figure shows the percentage change in the rural price index between 1987/88 and 1999/2000 for each decile of the local per-capita outlay distribution (averaged across districts using population weights). Left panel plots our AFFG NPC price index changes estimated from horizontal shifts in relative Engel curves. Middle panel plots price index changes using Laspeyres and Paasche district-level CPIs calculated using price changes for food and fuels following Deaton (2003b). Right panel repeats the middle panel but using district-income-decile-specific budget shares to calculate the Laspeyres and Paasche indices. Bootstrapped confidence intervals are based on sampling with replacement 1000 times from the distribution of households within each district-survey round and plotting the 2.5 and 97.5 percent envelope of price index estimates at each decile. See Section 4.3 for further discussion.

Figure 5: Rural Indian Welfare Growth 1987/88–1999/2000



Notes: Figure shows the percentage change in rural welfare between 1987/88 and 1999/2000 for each decile of the local per-capita outlay distribution (averaged across districts using population weights). Left panel plots the percentage change in both equivalent and compensating variation estimated from outlay changes and horizontal shifts in relative Engel curves (the AFG NPC price index). Right panel plots the percentage change in real income calculated by deflating per-capita outlay changes by Laspeyres and Paasche price index changes (using price changes for food and fuels and district-income-decile-specific budget shares). Bootstrapped confidence intervals are based on sampling with replacement 1000 times from the distribution of households within each district-survey round and plotting the 2.5 and 97.5 percent envelope of price index estimates at each decile. See Section 4.3 for further discussion.

Figure 6: AFFG Price Index Changes Across Alternative G Groupings

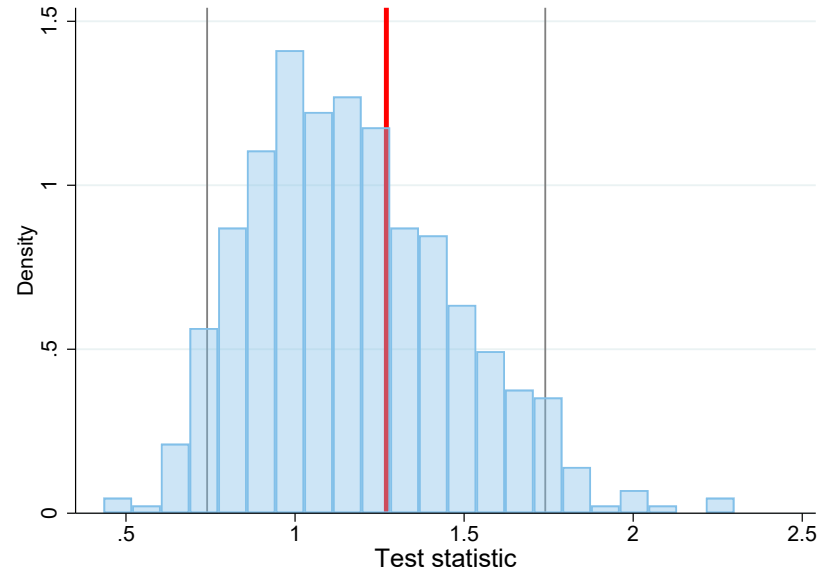



Notes: Figure reports AFFG NPC price index changes by decile of the local per-capita outlay distribution for each of 108 alternative classifications of goods into plausibly quasi-separable groupings G . Our baseline classification of three quasi-separable groups is one of the 108 classifications, and we indicate our baseline estimates in all panels. The two left panels depict for each decile the mean and the 2.5 and 97.5 percent envelope of point estimates across the 108 alternative groupings (panel A for P^0 and panel B for P^1). The two right panels depict the distribution of these estimates for the 2nd, 5th and 8th deciles of the local per-capita outlay distribution (panel B for P^0 and panel D for P^1). See Section 4.4 for further discussion.

Figure 7: Quasi-Separability Tests

Panel A: Test with Outside- G Prices

Panel B: Test with Outside- G Expenditures



Notes: Figure reports two tests of quasi-separability described in Section 3.2.1. The test in Panel A uses price changes outside of G (Corollary QS1) and the test in Panel B relies on outside- G expenditures instead (Corollary QS2). Vertical gray lines show 95 percent confidence intervals. Vertical red lines show χ^2 (Panel A) and F (Panel B) statistics obtained from sample data, respectively. Blue bars plot the empirical the distribution of QS test statistics obtained from 500 independently drawn random price datasets from a normal distribution with mean and variance identical to that of the distribution of price variation in the data. Permutation test p -values are 0.44 and 0.33 for the tests in Panel A and B, respectively. Panel A test statistic is obtained from regressing the change in log relative expenditure share of good $i \in G$ on within- G log price changes, changes in utility ($\log(1 + EV_h/y_h)$ and $\log(1 - CV_h/y_h)$, i.e. the horizontal distances illustrated in Figure 1, and log expenditures per capita) as well as two proxies for outside- G price changes from India's state-level CPI: 1) clothing, bedding and footwear, and 2) miscellaneous goods. The χ^2 statistic is obtained for the joint test that the coefficients on both outside goods prices in each of the 31 regressions are zero. Panel B test statistic is obtained from regressing the change on log expenditures on all outside- G goods on a within- G Stone price index (relative expenditure share weighted log price changes for goods in G), changes in within- G log prices and changes in household welfare as above. The F statistic is obtained from the joint test that coefficients on all $i \in G$ within- G log prices are equal to zero. In both panels we cluster standard errors at the market level, use survey weights, and perform the test at the median decile of households in each market.

Table 1: Product Groupings

3 G groups	8 G groups	34 g goods	Disaggregated NSS survey items included in the g goods
Raw food products	Cereals	Cereals - rice	Rice; chira; khoi, lawa; muri; other rice products.
Raw food products	Cereals	Cereals - wheat	Wheat, atta, wheat/atta PDS; maida; suji, rawa; sewai (noodles); bread (bakery).
Raw food products	Cereals	Cereals - coarse	Jowar, jowar products; bajra, bajra products; maize, maize products; barley, barley products; small millets, small millets products; ragi, ragi products.
Raw food products	Gram and pulses	Gram	Gram (full grain/whole); gram products.
Raw food products	Gram and pulses	Pulses - besan, moong	Besan; moong; soyabean; other pulse products.
Raw food products	Gram and pulses	Pulses - urd, masur	Urd; masur; arhar (tur); khasari; peas (dry); gram (split); other pulses.
Raw food products	Meat, fish and eggs	Meat	Goat meat, mutton; beef, buffalo meat; pork; poultry.
Raw food products	Meat, fish and eggs	Fish, prawn	Fish, prawn.
Raw food products	Meat, fish and eggs	Eggs	Eggs, egg products.
Raw food products	Fruits and vegetables	Vegetable - root vegetables	Potato; arum; radish; carrot; turnip; beet; sweet potato; onion; other root vegetables.
Raw food products	Fruits and vegetables	Vegetable - gourds	Pumpkin; gourd; bitter gourd; cucumber; parwal/patal; jhinga/torai; snake gourd; other gourds.
Raw food products	Fruits and vegetables	Vegetable - leafy vegetables	Cauliflower; cabbage; brinjal; lady's finger; french beans, barbat; tomato; palak/other leafy vegetables.
Raw food products	Fruits and vegetables	Vegetable - other vegetables	Peas (fresh); chilli (green); capsicum; plantain(green); jackfruit (green).
Raw food products	Fruits and vegetables	Premium Fruits	Apple; grapes; leechi; orange/mausani; pineapple; pears (naspati); mango; watermelon.
Raw food products	Fruits and vegetables	Other fresh fruits	Banana; jackfruit; singara; papaya; kharboosa; berries.
Raw food products	Fruits and vegetables	Dry fruits and nuts	Coconut (copra); groundnut; dates; cashewnut; walnut; raisin (kishmish, monacca, etc.).
Other food products	Dairy products and edible oils	Ghee	Ghee; butter.
Other food products	Dairy products and edible oils	Milk	Milk (liquid).
Other food products	Dairy products and edible oils	Other milk products	Milk (condensed/powder); curd; baby food.
Other food products	Dairy products and edible oils	Vanaspati, margarine	Vanaspati, margarine.
Other food products	Dairy products and edible oils	Edible oils	Ground nut oil; mustard oil; coconut oil; other edible oils.
Other food products	Sugar, salt, and spices	Sugar	Sugar; gur; sugar candy (misri); sugar(other sources); honey.
Other food products	Sugar, salt, and spices	Salt	Salt.
Other food products	Sugar, salt, and spices	Spices	Turmeric; black pepper; dry chillies; garlic; tamarind; ginger; curry powder.
Other food products	Refreshments and intoxicants	Beverages	Tea (leaf); coffee (cups); coconut (green).
Other food products	Refreshments and intoxicants	Processed food	Cooked meals; pickles; sauce; jam, jelly.
Other food products	Refreshments and intoxicants	Pan	Pan (finished); supari; lime; katha.
Other food products	Refreshments and intoxicants	Tobacco	Bidi; cigarettes; leaf tobacco; snuff.
Other food products	Refreshments and intoxicants	Intoxicants	Country liquor; beer; foreign liquor or refined liquor.
Fuels	Fuels	Coke, coal, charcoal	Coke; coal; charcoal.
Fuels	Fuels	Kerosene	Kerosene.
Fuels	Fuels	Firewood and chips	Firewood and chips.
Fuels	Fuels	Electricity	Electricity.
Fuels	Fuels	Matches	Matches.

Notes: This table details the classification of disaggregated NSS items (column 4) into various levels of aggregation: the 34 g goods used in our baseline analysis (column 3); the 8 groups that form the basis of the alternative G groupings we explore in Section 4.2 (column 2); and the 3 G groups each g good is assigned to in our baseline analysis (column 1). Different disaggregated NSS items in column 4 are separated by a semicolon. NSS items exclude those dropped by Deaton (2003b) (see Appendix A). Some NSS items were not consistently classified over rounds. Specifically: (Concorded) Rice uses individual items from R43 {Rice; Paddy} and R55 {Rice; Rice PDS}. (Concorded) Wheat uses R43 {Wheat, Atta} and R55 {Wheat, Atta PDS; Wheat, Atta other sources}. (Concorded) Jowar and Jowar products uses R43 {Jowar; Jowar products} and R55 {Jowar, Jowar products}. (Concorded) Bajra and Bajra products uses R43 {Bajra; Bajra products} and R55 {Bajra, Bajra products}. (Concorded) Maize and Maize products uses R43 {Maize; Maize products} and R55 {Maize, Maize products}. (Concorded) Barley and Barley products uses R43 {Barley; Barley products} and R55 {Barley, Barley products}. (Concorded) Small millets and Small millets products uses R43 {Small millets; Small millets products} and R55 {Small millets, Small millets products}. (Concorded) Ragi and Ragi products uses R43 {Ragi; Ragi products} and R55 {Ragi, Ragi products}. (Concorded) Beef, buffalo meat uses R43 {Beef; Buffalo meat} and R55 {Beef, Buffalo meat}. (Concorded) Goat, mutton uses R43 {Goat; Mutton} and R55 {Goat, Mutton}. (Concorded) Fish, Prawn uses R43 {Fish fresh; Fish dry} and R55 {Fish, prawn}. (Concorded) Eggs, Egg products uses R43 {Eggs; Egg products} and R55 {Eggs}. Vegetable - Gourds includes R43 {Papaya (green)} and R55 {Other gourds}. Vegetable - leafy vegetables includes R43 {Palak; Other leafy vegetables} and R55 {Palak, other leafy vegetables}. (Concorded) Vanaspati, margarine uses R43 {Vanaspati; Margarine} and R55 {Vanaspati, margarine}. (Concorded) Edible oils includes R43 {Linseed oil, Palm oil, Refined oil, Gingelly (til) oil, Rapeseed oil} and R55 {Edible oils (other)}. (Concorded) Sugar uses R43 {Sugar (crystal)} and R55 {Sugar PDS; sugar (other sources)}. (Concorded) Salt uses R43 {Sea salt; other salt} and R55 {Salt}.

Appendix

A Theory Appendix

Throughout the analysis, we assume that the expenditure function $e(p, U)$ is well behaved, i.e. it is twice continuously differentiable in prices p and utility U , strictly concave in p and increasing in U (with a positive derivative), homogeneous of degree zero in p , and globally defined for all $p \in \mathbb{R}_+^N$ and $U \in \mathbb{R}_+$.

A.1 Proof of Lemma 1

Lemma 1 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 Lemma 2 that we state and prove in Section A.3 below.

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 stated in condition i) of Lemma 2.

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

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

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

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

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

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

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

$$\frac{x_i}{x_G} = \frac{p_i h_i(p, U)}{\sum_{j \in G} p_j h_j(p, U)} = \frac{\partial \log e_G(p_G, U)}{\partial \log p_i} \equiv H_{iG}(p_G, U) \quad (\text{A.1})$$

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, we have $H_{iG}(p_G^1, U) = H_{iG}(p_G^0, U)$ if relative prices remain constant: $p_G^1 = \lambda_G p_G^0$ across all goods in group G (where λ_G is a positive scalar). For a consumer at initial utility u , income y and price p , notice that:

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

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

$$H_{iG}(p_G^t, V(p^t, y)) = E_{iG}^t(y) \quad (\text{A.2})$$

which holds for any income y (and also any price p^t and subvector p_G^t observed at time t).

Using this equality, we can now obtain each subpart of Lemma 1 on Engel curves:

- i) For the price index, define $P^1(p^0, p^1, y_h^1)$ the exact price index change at income y_h^1 for household h , implicitly defined such that $V(p^0, y^1/P^1) = V(p^1, y_h^1)$ where V is the indirect utility function. Using equality (A.2) at new and old prices, and the assumption that relative prices remain constant

within G: $p_G^1 = \lambda_G p_G^0$, we obtain:

$$\begin{aligned}
E_{iG}^0\left(y^1/P^1(p^0, p^1, y_h^1)\right) &= H_{iG}(p_G^0, V(p^0, y_h^1/P^1(p^1, p^0, y_h^1))) \\
&= H_{iG}(p_G^0, V(p^1, y^1)) \\
&= H_{iG}(p_G^1, V(p^1, y^1)) \\
&= E_{iG}^1(y_h^1)
\end{aligned}$$

where we go from the second to third line by noticing that H_{iG} is homogeneous of degree zero in prices (and $p_G^1 = \lambda_G p_G^0$ for some scalar λ_G). By switching time superscripts 1 and 0, we prove a similar equality using the other price index $P^0(p^0, p^1, y_h^0)$:

$$E_{iG}^1\left(y^0/P^0(p^0, p^1, y_h^0)\right) = E_{iG}^0(y_h^0)$$

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 0 to 1 and from 1 to 0. Note that this equality holds regardless of whether it is monotonic. If, in addition, we assume that the relative Engel curve E_{iG}^0 is strictly monotonic on the domain of interest, and thus invertible on its image, we can then uniquely characterize the price index P^1 as:

$$\log P^1(p^0, p^1, y_h^1) = \log y^1 - \log(E_{iG}^0)^{-1}(E_{iG}^1(y_h^1)) = \log y^1 - \log(E_{iG}^0)^{-1}(x_i^1/x_G^1)$$

(note that we assume that the expenditure function is well-behaved so that Engel curves are also continuous). Similarly, if the relative Engel curve E_{iG}^1 is strictly monotonic, we can uniquely characterize P^0 as:

$$\log P^0(p^0, p^1, y_h^0) = \log y^0 - \log(E_{iG}^1)^{-1}(E_{iG}^0(y_h^0)) = \log y^0 - \log(E_{iG}^1)^{-1}(x_i^0/x_G^0)$$

ii) By definition, compensating variations CV_h satisfy:

$$V(p^1, y_h^1 - CV_h) = V(p^0, y^0) = U_h^0$$

where U_h^0 denotes the utility level of household h in period 0. With the definition of CV_h and the homogeneity of function H_{iG} described above, as well as equality (A.2) for p^1 , we obtain that CV_h satisfies:

$$\begin{aligned}
E_{iG}^1(y_h^1 - CV_h) &= H_{iG}(p_G^1, V(p^1, y_h^1 - CV_h)) \\
&= H_{iG}(p_G^1, U_h^0) \\
&= H_{iG}(p_G^0, U_h^0) \\
&= x_{ih}^0/x_{hG}^0
\end{aligned}$$

where the last term refers to the within-group G expenditure share of good i in period 0. Note again that this equality holds regardless of whether the relative Engel curve is monotonic. If, in addition, we assume that the relative Engel curve E_{iG}^1 is strictly monotonic, CV_h is uniquely characterized as the horizontal movement along the E_{iG}^1 curve to reach x_{ih}^0/x_{hG}^0 :

$$CV_h = y_h^1 - (E_{iG}^1)^{-1}(x_{ih}^0/x_{hG}^0)$$

Similarly, by definition, equivalent variations EV satisfy:

$$V(p^0, y_h^0 + EV_h) = V(p^1, y_h^1) = U_h^1$$

where U_h^1 denotes to the period 1 utility level of household h .

With the definition of EV_h and the homogeneity of function H_{iG} , as well as equality (A.2) for p^1 , we

obtain that EV_h satisfies:

$$\begin{aligned}
E_{iG}^0(y_h^0 + EV_h) &= H_{iG}(p_G^0, V(p^0, y_h^0 + EV_h)) \\
&= H_{iG}(p_G^0, U_h^1) \\
&= H_{iG}(p_G^1, U_h^1) \\
&= x_{hi}^1/x_{Gh}^1
\end{aligned}$$

where the last term refers to the within-group G expenditure share of good i in period 1. If, in addition, we assume that the relative Engel curve E_{iG}^0 is strictly monotonic, EV_h is uniquely characterized as the horizontal movement along the E_{iG}^0 curve to reach x_{hi}^1/x_{Gh}^1 :

$$EV_h = (E_{iG}^0)^{-1}(x_{hi}^1/x_{Gh}^1) - y_h^0$$

Quasi-Separability as a Necessary Condition. For the shifts in Engel curves to reflect the changes in price indices, 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 2 (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^1 :

$$E_{iG}^1(y_h^1) = E_{iG}^0(y_h^1/P^1(y_h^1)) \quad (\text{A.3})$$

where $P^1(y_h^1)$ is the price index change transforming income at period 1 prices to income in 0 prices, for any change in prices that leaves within- G relative prices constant, i.e. $p_G^1 = \lambda_G p_G^0$ for some scalar λ_G . By definition of the price index, P^1 is such that $V(p^1, y_h^1) = V(p^0, y_h^1/P^1)$ where V denotes the indirect utility function. Or equivalently:

$$\frac{y_h^1}{P^1(y_h^1)} = e(V(p^1, y_h^1), p^0) = e(U_h^1, p^0)$$

using the expenditure function e , where we denote utility in period 1 by U_h^1 . Looking at the share good i in expenditures within group G , and imposing that Engel curves satisfy condition A.3, we can see that it no longer depends on prices p^1 once we condition on utility U_h^1 (aside from the subvector p_G of prices within G):

$$\frac{x_{ih}}{x_{Gh}} = E_i^1(y_h^1) = E_i^0\left(\frac{y_h^1}{P^1(y_h^1)}\right) = E_G^0(e(U_h^1, p^0))$$

Note that the expenditure share at time 1 is independent of prices p^0 in another period (as long as $p_G^1 = \lambda_G p_G^0$). Hence there exists a function H_{iG} of within- G relative prices and utility such that:

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

where p_G is the subvector of prices of p^1 and p^0 , up to a scalar factor λ_G (H_{iG} is independent of λ_G so it must be homogeneous of degree zero in p_G). This requirement corresponds to condition i) of Lemma 2. As we prove below in Lemma 2, it implies quasi-separability in G . Hence, quasi-separability in G is required if we want the shifts in relative Engel curves to reflect the changes in price indices.

A.2 Proof of Proposition 1

As we have seen for the proof of Lemma 1 (equality A.2), we have: $H_{iG}(p_G, V(p, y)) = E_{iG}(p)$ where E_{iG} is evaluated for a given set of prices p , and where $H_{iG}(p_G, U_h)$ denotes the within- G compensated expenditure share:

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

Denote utility in period 1 by $U_h^1 = V(p^1, y^1)$. We obtain:

$$\begin{aligned}
E_{iG}^0\left(y^1/P^1(p^0, p^1, y_h^1)\right) &= H_{iG}(p_G^0, V(p^0, y_h^1/P^1(p^1, p^0, y_h^1))) \\
&= H_{iG}(p_G^0, V(p^1, y^1)) \\
&= H_{iG}(p_G^1, V(p^1, y^1)) \times \frac{H_{iG}(p_G^0, U_h^1)}{H_{iG}(p_G^1, U_h^1)} \\
&= E_{iG}^1(y_h^1) \times \frac{H_{iG}(p_G^0, U_h^1)}{H_{iG}(p_G^1, U_h^1)} \\
&= E_{iG}^1(y_h^1) \times \exp \sum_{j \in G} \int_{\log p_j^1}^{\log p_j^0} \frac{\partial \log H_{iG}}{\partial \log p_j} d \log p_j
\end{aligned}$$

where each step is similar to the those of the proof of Lemma 1, aside from the new term in the last three lines, re-expressed in the last line using the derivatives $\frac{\partial \log H_{iG}}{\partial \log p_j}$ evaluated along indifference curves at utility U_h^1 . If, in addition, we assume that the relative Engel curve E_{iG}^0 is strictly monotonic, this expression uniquely characterizes P^1 , as we previously described in Lemma 1 (after adjusting for the new term $\frac{H_{iG}(p_G^0, U_h^1)}{H_{iG}(p_G^1, U_h^1)}$ multiplying expenditure shares).

For EV_h , note we have

$$V(p^1, y_h^1) = V(p^0, y_h^0 + EV_h) = V(p^0, y_h^1/P^1(p^1, p^0, y_h^1)) = U_h^1$$

Hence, using the results just above, we obtain:

$$\begin{aligned}
E_{iG}^0(y_h^1 + EV_h) &= H_{iG}(p_G^0, V(p^1, y^1)) \\
&= H_{iG}(p_G^1, V(p^1, y^1)) \times \frac{H_{iG}(p_G^0, U_h^1)}{H_{iG}(p_G^1, U_h^1)} \\
&= E_{iG}^1(y_h^1) \times \frac{H_{iG}(p_G^0, U_h^1)}{H_{iG}(p_G^1, U_h^1)} \\
&= \frac{x_{ih}^1}{x_{Gh}^1} \times \frac{H_{iG}(p_G^0, U_h^1)}{H_{iG}(p_G^1, U_h^1)}
\end{aligned}$$

If, in addition, we assume that the relative Engel curve E_{iG}^0 is strictly monotonic, this expression uniquely characterizes EV_h , as we previously described in Lemma 1 (after adjusting for the new term $\frac{H_{iG}(p_G^0, U_h^1)}{H_{iG}(p_G^1, U_h^1)}$ multiplying expenditure shares).

Symmetric arguments can be used for P^0 and CV_h by swapping the two time periods. This proves Proposition 1.

Remarks. Alternatively, note that we can also adjust relative Engel curves additively rather than multiplicatively. For instance, for price index P^1 , we obtain:

$$\begin{aligned}
E_{iG}^0\left(y^1/P^1(p^0, p^1, y_h^1)\right) &= H_{iG}(p_G^0, V(p^0, y_h^1/P^1(p^1, p^0, y_h^1))) \\
&= H_{iG}(p_G^0, V(p^1, y^1)) \\
&= H_{iG}(p_G^1, V(p^1, y^1)) + [H_{iG}(p_G^0, U_h^1) - H_{iG}(p_G^1, U_h^1)] \\
&= E_{iG}^1(y_h^1) + [H_{iG}(p_G^0, U_h^1) - H_{iG}(p_G^1, U_h^1)] \\
&= E_{iG}^1(y_h^1) + \sum_{j \in G} \int_{p_j^1}^{p_j^0} \frac{\partial H_{iG}}{\partial \log p_j} d \log p_j
\end{aligned}$$

This additive adjustment is helpful in specifications where prices enter additively (e.g. EASI). Moreover, the ratio $\frac{E_{iG}^1(y_h^1)}{H_{iG}(p_G^1, U_h^1)}$ is indeterminate in the extreme case where consumption of good i is null in period 1 while the additive adjustment can still be exploited in these cases. Is it possible for the econometrician to evaluate $\frac{\partial \log H_{iG}}{\partial \log p_j}$ without observing utility? To do so, one can use a Slutsky-type decomposition applied

to within-G expenditure shares:

$$\frac{\partial \log H_{iG}}{\partial \log p_j} = \frac{\partial \log(x_i/x_G)}{\partial \log p_j} + E_{jG} \frac{x_G}{y} \frac{\partial \log E_{iG}}{\partial \log y}$$

where $\frac{\partial \log(x_i/x_G)}{\partial \log p_j}$ and $\frac{\partial \log E_{iG}}{\partial \log y}$ are the uncompensated elasticities which can be more directly estimated. To prove this result, using $\frac{\partial \log e}{\partial \log p_j} = E_{jG} \frac{x_G}{y}$ the expenditure share of good j (Shephard's Lemma), note that we have:

$$\frac{\partial \log H_{iG}}{\partial \log p_j} = \frac{\partial \log(x_i/x_G)}{\partial \log p_j} + \frac{\partial \log e}{\partial \log p_j} \frac{\partial \log E_{iG}}{\partial \log y} = \frac{\partial \log(x_i/x_G)}{\partial \log p_j} + E_{jG} \frac{x_G}{y} \frac{\partial \log E_{iG}}{\partial \log y}$$

A.3 Lemma 2

Lemma 2 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 $K(F_G, q_{NG}, U_h)$ and $F_G(q_G, U_h)$, where $F_G(q_G, U_h)$ is homogeneous of degree 1 in q_G .

Proof of Lemma 2

Gorman (1970) and Deaton and Muellbauer (1980) have already provided a proof of the equivalence between quasi-separability and condition 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 condition 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 proof below, we drop the household subscripts and time superscripts to lighten the notation. We also assume that the expenditure function is well-behaved, and in particular twice continuously differentiable, as noted at the start of the appendix.

- **Quasi-separability implies i).** Actually we have already shown that quasi-separability implies i). In the proof of Proposition 1 above, we have shown in equation (A.1) that we have:

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

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

The most difficult part of the proof of Lemma 3 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 $e_G(p_G, U)$ such that:

$$\frac{\partial \log e_G}{\partial p_i} = \frac{1}{p_i} H_{iG}(p_G, U) \quad (\text{A.4})$$

for each i , and $e_G(p_{G0}, U) = 1$ 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 and only if the derivatives $\frac{\partial(H_i/p_i)}{\partial p_j} = \frac{\partial(H_j/p_j)}{\partial p_i}$ are symmetric. 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 $\frac{1}{x_G} \frac{\partial h_i}{\partial p_j}$ 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 : $h_i \frac{\partial h_j}{\partial p_k} = h_j \frac{\partial h_i}{\partial p_k}$ for any $k \notin G$. 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 e_G satisfying equation (A.4).

Note that $\sum_{i \in G} H_i(p_G, U) = 1$ for any price vector p_G and utility U , hence e_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 the price vector p_G only through the scalar function $e_G(p_G, U)$. To do so, we use the same idea as in Lemma 1 of Goldman and Uzawa (1964).¹ Using our constructed $e_G(p_G, U)$, notice that:

$$\frac{\partial e}{\partial p_i} = \frac{\partial e_G}{\partial p_i} \cdot x_G(p, U) \quad (\text{A.5})$$

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

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

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 .

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

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

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

Construct e_G as follows:

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

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

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

The expenditure function then satisfies:

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

(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 e_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 $\tilde{e}(e_G(p_G, U), p_{NG}, U)$ is the expenditure function associated with it.

First, using these two functions, let us define:

$$K(Q_G, q_{NG}, U) = \min_{e_G^*, p_{NG}^*} \left\{ \frac{Q_G e_G^* + \sum_{i \notin G} p_i^* q_i}{\tilde{e}(e_G^*, p_{NG}^*, U)} \right\} \quad (\text{A.6})$$

and:

$$F_G(q_G, U) = \min_{p_G} \left\{ \frac{\sum_{i \in G} p_i^* q_i}{e_G(p_G^*, U)} \right\} \quad (\text{A.7})$$

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 :

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

If \tilde{e} and e_G are decreasing in U , we can see that F_G and K are decreasing in U , hence the following has a

unique solution:

$$K(F_G(q_G, U), q_{NG}, U) = 1 \quad (\text{A.8})$$

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.6) and (A.7) also coincide with actual prices p . Then the final step is to show that the expenditure function coincides with $\tilde{e}(e_G(p_G, U), p_{NG}, U)$.

Utility maximization subject to the budget constraint and subject to constraint (A.8) leads to the following first-order conditions in q_i :

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

where p are observed prices and where μ and λ are the Lagrange multipliers associated with (A.8) and the budget constraints respectively. Using the envelop theorem, these partial derivatives are:

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

where e_G^* and p_i^* refer to counterfactual prices that minimize the right-hand side of equations (A.6) and (A.7) 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:

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

But notice that if p_i^* for $i \in G$ minimizes the right-hand side of equation (A.7), then $\lambda_G p_i^*$ also minimizes (A.7) since e_G is homogeneous of degree one. With $\lambda_G = \frac{\mu}{\lambda} \frac{e_G^*}{\tilde{e}(e_G^*, p_{NG}^*, U)} \frac{1}{e_G(p_G^*, U)}$, it implies that we can have: $p_i^* = p_i$ for $i \in G$. Also notice that if e_G^* and p_j^* for $j \notin G$ minimize the right-hand side of equation (A.6), then $\lambda_N e_G^*$ and $\lambda_N p_j^*$ also minimizes (A.7) for any $\lambda_N > 0$ since \tilde{e} is homogeneous of degree one. With $\lambda_N = \frac{\mu}{\lambda \tilde{e}(e_G^*, p_{NG}^*, U)}$, we have $\lambda_N p_j^* = p_j$. Using the FOC for goods $j \notin G$, we obtain:

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

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

$$\lambda_N e_G^* = e_G(p_G, U)$$

So we can also replace e_G^* by e_G .

Now that we have proven that observed prices are also solution of the minimization of (A.6) and (A.7), it is easy to show that $\tilde{e}(e_G(p_G, U), p_{NG}, U)$ is equal to the expenditure function associated with utility defined in equation (A.8). Using equations (A.8), (A.6) and (A.7), and the equality between p^* and p (as well as e_G^* and e_G), we find:

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

where quantities are those maximizing utility subject to the budget constraint, therefore the expenditure function coincides with $\tilde{e}(e_G(p_G, U), p_{NG}, U)$. Once we know that observe price minimize (A.6) and (A.7), it is also easy to verify that the expenditure shares implied by utility defined in A.8 also correspond to expenditure shares implied by the expenditure function $\tilde{e}(e_G(p_G, U), p_{NG}, U)$. This shows that utility defined by (A.8), (A.6) and (A.7) leads to the same demand system as $\tilde{e}(e_G(p_G, U), p_{NG}, U)$, and proves that quasi-separability implies condition ii).

A.4 Lemma 3

Before presenting the impossibility result from Lemma 4, we show here that the main idea behind Lemma 1 works for standard Engel curves when relative prices remain constant for the entire consumption basket.

Lemma 3. *Assume that prices change over time but relative prices remain unchanged, i.e. $p_i^1 = \lambda p_i^0$ for all i and some $\lambda > 0$.*

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

$$E_i^1(y_h^1) = E_i^0\left(\frac{y_h^1}{P^1(y_h^1)}\right) \quad \text{and} \quad E_i^0(y_h^0) = E_i^1\left(\frac{y_h^0}{P^0(y_h^0)}\right).$$

- ii) *EV and CV for a given income level are revealed by the horizontal distance along period 1 or period 0's Engel curves, respectively, between the new and old expenditure share, such that $\frac{x_{hi}^0}{y_h^0} = E_i^1(y_h^1 - CV_h)$ and $\frac{x_{hi}^1}{y_h^1} = E_i^0(y_h^0 + EV_h)$.*

Proof of Lemma 3

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 $E_i^t(y) = p_i q_i(p, y)/y$ 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 Lemma 3, the key property that we exploit is that q_i , E_i and V are all homogeneous of degree zero in p, y .

The first step is to show that Engel curves shift uniformly by $+\log \lambda$ if we have log total outlays (income) on the horizontal axis. By definition, we have

$$E_i^0\left(\frac{y_h^1}{\lambda}\right) = \frac{p_i^0 q_i(p^0, y_h^1/\lambda)}{(y_h^1/\lambda)} = \frac{\lambda p_i^0 q_i(p^0, y^1/\lambda)}{y^1}$$

but given that demand is homogeneous of degree zero in p, y , we have $q_i(p^0, y^1/\lambda) = q_i(\lambda p^0, y^1)$ and thus we obtain:

$$E_i^0\left(\frac{y_h^1}{\lambda}\right) = \frac{\lambda p_i^0 q_i(\lambda p^0, y^1)}{y^1} = \frac{p_i^1 q_i(p^1, y_h^1)}{y_h^1} = E_i^1(y^1)$$

Using this property, we can then show that the horizontal shift of Engel curves reflects the changes in price index and welfare:

- i) Define the price index relative to prices in period 0 implicitly as $P^1(p^0, p^1, y^1)$ such that: $V(p^1, y_h^1) = V(p^0, \frac{y_h^1}{P^1})$. With the homogeneous change in prices $p^1 = \lambda p^0$, it is immediate to verify that $P^1 = \lambda$ given that indirect utility is homogeneous of degree zero:

$$V(p^1, y_h^1) = V(\lambda p^0, y_h^1) = V(p^0, \frac{y_h^1}{\lambda})$$

Similarly, define the price index relative to prices in period 1 implicitly as $P^0(p^0, p^1, y^0)$ such that: $V(p^0, y_h^0) = V(p^1, \frac{y_h^0}{P^0})$. With the homogeneous change in prices $p^1 = \lambda p^0$, it is again immediate to verify that $P^0 = 1/\lambda$. Given that Engel curves shift by a factor λ , we obtain:

$$E_i^0\left(\frac{y_h^1}{P^1}\right) = E_i^0\left(\frac{y_h^1}{\lambda}\right) = E_i^1(y^1)$$

and

$$E_i^1\left(\frac{y_h^0}{P^0}\right) = E_i^1(\lambda y_h^0) = E_i^0(y^0)$$

hence the shift (in log) of the Engel curve from period 0 to period 1 corresponds to the price index change $\log P^1$, and the shift from period 1 to period 0 corresponds to the price index change $\log P^0$. This proves assertion i).

- ii) Compensating variations CV_h are implicitly defined as $V(p^1, y^1 - CV_h) = V(p^0, y_h^0) = U_h^0$. With the homogeneous change in prices $p^1 = \lambda p^0$, we can verify that compensating variations CV_h are such that $y_h^1 - CV_h = \lambda y_h^0$:

$$V(p^1, y^1 - CV_h) = V(p^0, y^0) = V(p^1/\lambda, y^0) = V(p^1, \lambda y_h^0)$$

Given that Engel curves shift by a factor λ , we obtain:

$$E_i^1(y_h^1 - CV_h) = E_i^1(\lambda y_h^0) = E_i^0(y_h^0)$$

hence the initial observed expenditure share $p_i^0 q_{ih}^1 / y_h^1$ of good i in period 0 corresponds to the counterfactual expenditure share of good i at new prices and total outlays $y_h^1 + EV_h$.

Equivalent variations EV_h are implicitly defined as $V(p^0, y^0 + EV_h) = V(p^1, y^1) = U_h^1$. For EV_h the proof proceeds the same way as for CV_h just by swapping periods 0 and 1 (and $1/\lambda$ instead of λ). With the homogeneous change in prices $p^1 = \lambda p^0$, we can verify that equivalent variations EV_h are such that $y_h^0 + EV_h = y^1 / \lambda$:

$$V(p^0, y^0 + EV_h) = V(p^1, y^1) = V(\lambda p^0, y^1) = V(p^0, y_h^1 / \lambda)$$

Again we obtain:

$$E_i^0(y_h^0 + EV_h) = E_i^0(y^1 / \lambda) = E_i^1(y^1)$$

hence the new observed expenditure share $p_i^1 q_{ih}^1 / y_h^1$ of good i corresponds to the counterfactual expenditure share of good i at former prices at $y_h^0 + EV_h$.

A.5 Lemma 4

Lemma 4. *Horizontal shifts in any good i 's Engel curve do not recover changes in the log price index under arbitrary changes in the price of good i relative to other goods, or groups of goods.*

Proof of Lemma 4

Suppose that for a certain good i the shift of the (standard) Engel curve $E_i^t(y_h^t)$ (expenditure share x_{ih}^t / y_h^t plotted against total outlays y_h^t for a given set of prices p^t) reflects the price index change for any realization of price changes across periods and any y , i.e. $E_i^1(y) = E_i^0(y / P^1(y))$. We know already from Lemma 3 that this is true for any preferences if we impose the price changes to be uniform across goods: $p^1 = \lambda p^0$. For it to be true for all price changes, we show that it implies:

- **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^1 , we have:

$$E_i^1(y_h^1) = E_i^0(y_h^1 / P^1(y_h^1)) \tag{A.9}$$

where $P^1(y_h^1)$ is the price index change transforming income at period 1 prices to income in 0 prices. By definition, the price index change P^1 is such that $V(p^1, y_h^1) = V(p^0, y_h^1 / P^1)$ where V denotes the indirect utility function. An equivalent characterization of the price index is:

$$\frac{y_h^1}{P^1(y_h^1)} = e(V(p^1, y_h^1), p^0) = e(U_h^1, p^0)$$

using the expenditure function e , denoting utility in period 1 by U_h^1 . Looking at the share good i in total expenditures and imposing that Engel curves satisfy condition A.9, we can see that it no longer depends on prices p^1 once we condition on utility U_h^1 :

$$\frac{x_{ih}}{y_h} = E_i^1(y_h^1) = E_i^0\left(\frac{y_h^1}{P^1(y_h^1)}\right) = E_i^0(e(U^1, p^0))$$

(note that the expenditure share at time 1 is independent of prices p^0 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 between good i and the other goods. 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 $w_i(U') > w_i(U)$ (the same argument works with $w_i(U') < w_i(U)$). We can then find $\log(p_i/p_j)$ such that:

$$\log(p_i/p_j) > \frac{\log e_0(U) - \log e_0(U')}{w_i(U') - w_i(U)}$$

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 .

A.6 Proofs for Section 3

First-order Correction Terms for Relative Price Changes

As shown in Proposition 1 (taking logs), we have for P^1 :

$$\log E_{iG}^0\left(y^1/P^1(p^0, p^1, y_h^1)\right) = \log E_{iG}^1(y_h^1) + \log \frac{H_{iG}(p_G^0, U_h^1)}{H_{iG}(p_G^1, U_h^1)} \quad (\text{A.10})$$

Note again that H_{iG} is homogeneous of degree zero in prices so a small change in relative prices will lead to only a small adjustment term $\log \frac{H_{iG}(p_G^0, U_h^1)}{H_{iG}(p_G^1, U_h^1)}$. As a first-order approximation (w.r.t relative prices), we invert the Engel curve in period 0 and obtain:

$$\begin{aligned} \frac{y_h^1}{P^1(y_h^1)} &= \log(E_{iG}^0)^{-1}\left(E_{iG}(p^1, y_h^1)\right) + (\beta_{ih}^0)^{-1} \log \frac{H_{iG}(p_G^0, U_h^1)}{H_{iG}(p_G^1, U_h^1)} \\ &= \log(E_{iG}^0)^{-1}\left(\frac{x_{ih}^1}{x_{Gh}^1}\right) + (\beta_{ih}^0)^{-1} \log \frac{H_{iG}(p_G^0, U_h^1)}{H_{iG}(p_G^1, U_h^1)} \end{aligned}$$

where $\beta_{ihm}^0 = \frac{\partial \log E_{iG}}{\partial \log y_h}$ denotes the slope of the relative Engel curve, evaluated in period 0. Taking the average across goods, we obtain:

$$\log \left(\frac{y_h^1}{P^1} \right) \approx \frac{1}{G} \sum_{i \in G} \log (E_{iG}^0)^{-1} \left(\frac{x_{ih}^1}{x_{Gh}^1} \right) + \frac{1}{G} \sum_{i \in G} (\beta_{ih}^0)^{-1} \log \frac{H_{iG}(p_G^0, U_h^1)}{H_{iG}(p_G^1, U_h^1)} \quad (\text{A.11})$$

Hence the average of the horizontal shift $\log \left(\frac{y_h^1}{P^1} \right) - \log (E_{iG}^0)^{-1} \left(\frac{x_{ih}^1}{x_{Gh}^1} \right)$ is equal to the log price index change $\log P^1$ when the adjustment term is null on average: $\frac{1}{G} \sum_{i \in G} (\beta_{ih}^0)^{-1} \log \frac{H_{iG}(p_G^0, U_h^1)}{H_{iG}(p_G^1, U_h^1)} = 0$. The same logic applies to evaluating $\log P^0$, EV_h and CV_h . As a first-order approximation w.r.t. relative price changes, note that $\log \frac{H_{iG}(p_G^0, U_h^1)}{H_{iG}(p_G^1, U_h^1)} = \sum_{j \in G} \int_{\log p_j^1}^{\log p_j^0} \frac{\partial \log H_{iG}}{\partial \log p_j} d \log p_j \approx - \sum_{j \in G} \sigma_{ijh} (\Delta \log p_j - \overline{\Delta \log p_G})$, where $\sigma_{ijh} = \frac{\partial \log H_{iG}}{\partial \log p_j}$ is the compensated price elasticity of relative consumption of i with respect to price j , $\Delta \log p_j = \log p_j^1 - \log p_j^0$ is the change in the price of good j from the base period 0, and $\overline{\Delta \log p_G}$ is the average price change within G . Note that $\sum_{j \in G} \sigma_{ijh} = 0$ due to homogeneity of degree zero of H_{iG} in all G prices so we can rewrite $\sum_{j \in G} \sigma_{ijh} \Delta \log p_j$ as $\sum_{j \in G} \sigma_{ijh} (\Delta \log p_j - \overline{\Delta \log p_G})$. Hence we obtain:

$$\log \left(\frac{y_h^1}{P^1} \right) \approx \frac{1}{G} \sum_{i \in G} \log (E_{iG}^0)^{-1} \left(\frac{x_{ih}^1}{x_{Gh}^1} \right) - \frac{1}{G} \sum_{i, j \in G} (\beta_{ih}^0)^{-1} \sigma_{ijh} (\Delta \log p_j - \overline{\Delta \log p_G})$$

Exact Correction Terms for Relative Price Changes

a) Within-group demand with implicit non-homothetic CES demand Starting from Proposition 1, we now impose specific forms of demand to account for price adjustments. First, suppose that the expenditure function takes the form:

$$e(p, U) = \tilde{e} \left(\sum_{j \in G} A_j(U) p_j^{1-\sigma_G}, p_{NG}, U \right)$$

We obtain that demand takes a constant elasticity σ_G among goods within group G (and only within G) such that:

$$H_{iG}(p_G, U) = \frac{A_i(U) p_i^{1-\sigma_G}}{\sum_{j \in G} A_j(U) p_j^{1-\sigma_G}}$$

If we have knowledge of the within- G price elasticity σ_G and consumption shares in a period of reference, we can predict consumption shares for all goods i within G for any change in relative prices (with new prices denoted p'), holding utility constant:

$$\begin{aligned} H_{iG}(p'_G, U) &= \frac{(p'_i/p_i)^{1-\sigma_G} A_i(U) p_i^{1-\sigma_G}}{\sum_{j \in G} (p'_j/p_j)^{1-\sigma_G} A_j(U) p_j^{1-\sigma_G}} \\ &= \frac{(p'_i/p_i)^{1-\sigma_G} H_{iG}(p_G, U)}{\sum_{j \in G} (p'_j/p_j)^{1-\sigma_G} H_{jG}(p_G, U)} \\ &= \frac{(p'_i/p_i)^{1-\sigma_G} (x_i/x_G)}{\sum_{j \in G} (p'_j/p_j)^{1-\sigma_G} (x_j/x_G)} \end{aligned}$$

For instance, to obtain P^1 , applying the same logic as with Proposition 1 along with such a price adjustment yields:

$$\begin{aligned}
E_{iG}^0(y^1/P^1(p^0, p^1, y_h^1)) &= H_{iG}(p_G^0, V(p^0, y_h^1/P^1(p^1, p^0, y_h^1))) \\
&= H_{iG}(p_G^0, V(p^1, y^1)) \\
&= \frac{(p_i^0/p_i^1)^{1-\sigma_G} H_{iG}(p_G^1, V(p^1, y^1))}{\sum_{j \in G} (p_j^0/p_j^1)^{1-\sigma_G} H_{jG}(p_G^1, V(p^1, y^1))} \\
&= \frac{(p_i^0/p_i^1)^{1-\sigma_G} E_{iG}^1(y_h^1)}{\sum_{j \in G} (p_j^0/p_j^1)^{1-\sigma_G} E_{jG}^1(y_h^1)} \\
&= \frac{(p_i^1/p_i^0)^{\sigma_G-1} (x_i^1/x_G^1)}{\sum_{j \in G} (p_j^1/p_j^0)^{\sigma_G-1} (x_j^1/x_G^1)}
\end{aligned}$$

We obtain equation (10) in the main text after defining a group G price index $(\sigma_G - 1)\overline{\Delta \log p_G} = \log[\sum_{j \in G} (p_j^1/p_j^0)^{\sigma_G-1} (x_j^1/x_G^1)]$.

b) Multi-tiered implicit non-homothetic CES demand In order to account for richer substitution patterns, an alternative is to create a partition of G in several sub-groups g , and have heterogeneous price elasticities across subgroups g :

$$e_G(p_G, U) = \left(\sum_g \left(\sum_{i \in g} \alpha_i(U) p_i^{1-\sigma_g} \right)^{\frac{1-\eta_G}{1-\sigma_g}} \right)^{\frac{1}{1-\eta_G}}$$

with an overall expenditure function $e(p, U) = \tilde{e}(e_G(p_G, U), p_{NG}, U)$ that remains unspecified, conditional on e_G .

Denote by H_{gG} the share of subgroup g within G as a function of p_G and U . It is given by:

$$H_{gG}(p_G, U) = \frac{\left(\sum_{i \in g} \alpha_i(U) p_i^{1-\sigma_g} \right)^{\frac{1-\eta_G}{1-\sigma_g}}}{\sum_{g'} \left(\sum_{j \in g'} \alpha_j(U) p_j^{1-\sigma_{g'}} \right)^{\frac{1-\eta_G}{1-\sigma_{g'}}}}$$

Compensated demand for a good i within subgroup g is then equal to:

$$H_{iG}(p_G, U) = \frac{\alpha_i(U) p_i^{1-\sigma_g}}{\sum_{j \in g} \alpha_j(U) p_j^{1-\sigma_g}} H_{gG}(p_G, U)$$

With this specification, the own-price elasticity for good i is:

$$\epsilon_{iG} \equiv \frac{\partial \log H_{iG}}{\partial \log p_{iG}} = (1 - \sigma_g) + (\sigma_g - \eta_G)(H_{iG}/H_{gG}) + (\eta_G - 1)H_{iG}$$

If goods $i \neq j$ are in the same subgroup g , the cross price elasticity is:

$$\frac{\partial \log H_{iG}}{\partial \log p_{jG}} = (\sigma_g - \eta_G)(H_{jG}/H_{gG}) + (\eta_G - 1)H_{jG}$$

If goods i is in subgroup g , and j in subgroup $g' \neq g$, the cross price elasticity is:

$$\frac{\partial \log H_{iG}}{\partial \log p_{jG}} = (\eta_G - 1)H_{jG}$$

To express relative expenditure shares depending on observables and relative price changes, we use

the following equalities for reallocation across subgroups g :

$$\begin{aligned}
\frac{H_{gG}(p'_G, U)}{H_{gG}(p_G, U)} &= \frac{\left(\sum_{i \in g} \alpha_i(U) (p'_i)^{1-\sigma_g}\right)^{\frac{1-\eta_G}{1-\sigma_g}} \sum_{g'} \left(\sum_{j \in g'} \alpha_j(U) p_j^{1-\sigma_{g'}}\right)^{\frac{1-\eta_G}{1-\sigma_{g'}}}}{\left(\sum_{i \in g} \alpha_i(U) p_i^{1-\sigma_g}\right)^{\frac{1-\eta_G}{1-\sigma_g}} \sum_{g'} \left(\sum_{j \in g'} \alpha_j(U) (p'_j)^{1-\sigma_{g'}}\right)^{\frac{1-\eta_G}{1-\sigma_{g'}}}} \\
&= \left[\frac{\sum_{i \in g} \alpha_i(U) p_i^{1-\sigma_g} (p'_i/p_i)^{1-\sigma_g}}{\sum_{i \in g} \alpha_i(U) p_i^{1-\sigma_g}} \right]^{\frac{1-\eta_G}{1-\sigma_g}} / \left[\frac{\sum_{g'} \left(\sum_{j \in g'} \alpha_j(U) p_j^{1-\sigma_{g'}} (p'_j/p_j)^{1-\sigma_{g'}}\right)^{\frac{1-\eta_G}{1-\sigma_{g'}}}}{\sum_{g'} \left(\sum_{j \in g'} \alpha_j(U) p_j^{1-\sigma_{g'}}\right)^{\frac{1-\eta_G}{1-\sigma_{g'}}}} \right] \\
&= \left[\sum_{i' \in g} (H_{i'G}/H_{gG}) (p'_{i'}/p_{i'})^{1-\sigma_g} \right]^{\frac{1-\eta_G}{1-\sigma_g}} / \left[\sum_{g'} H_{g'G} \left[\sum_{j \in g'} (H_{jG}/H_{g'G}) (p'_j/p_j)^{1-\sigma_{g'}} \right]^{\frac{1-\eta_G}{1-\sigma_{g'}}} \right] \\
&= e^{(1-\eta_G)\overline{\Delta \log p_g}} / \left[\sum_{g'} x_{g'h} e^{(1-\eta_G)\overline{\Delta \log p_g}} \right]
\end{aligned}$$

where:

$$\overline{\Delta \log p_g} = \log \left[\sum_{j \in g} (p'_j/p_j)^{1-\sigma_g} (x_{jh}^1/x_{gh}^1) \right]^{\frac{1}{1-\sigma_g}}$$

is the price index change for subgroup $g \subset G$, and for reallocation within subgroups:

$$\begin{aligned}
H_{iG}(p'_G, U) &= \frac{\alpha_i(U) (p'_i)^{1-\sigma_g}}{\sum_{j \in g} \alpha_j(U) (p'_j)^{1-\sigma_g}} H_{gG}(p', U) \\
&= \frac{H_{iG} (p'_i/p_i)^{1-\sigma_g}}{\sum_{j \in g} H_{jG} (p'_j/p_j)^{1-\sigma_g}} H_{gG}(p', U) \\
&= \frac{(H_{iG}/H_{gG}) (p'_i/p_i)^{1-\sigma_g}}{\sum_{i' \in g} (H_{i'G}/H_{gG}) (p'_{i'}/p_{i'})^{1-\sigma_g}} H_{gG}(p', U) \\
&= (x_{ih}/x_{ig}) e^{(1-\sigma_g)(\Delta \log p_i - \overline{\Delta \log p_g})} H_{gG}(p', U)
\end{aligned}$$

Equation (11) in the text is obtained by combining these two equations.

c) EASI within-group demand Another simple case is presented below is inspired from the EASI demand system (Lewbel and Pendakur, 2009). We can specify the within-group expenditure function as:

$$\log e_G(p_G, U) = \sum_{k \in G} A_k(U) \log p_k + \sum_{k, j \in G} B_{kj}(U) \log p_j \log p_k$$

with an overall expenditure function $e(p, U) = \tilde{e}(e_G(p_G, U), p_{NG}, U)$ that remains unspecified, conditional on e_G . Homogeneity in prices requires $\sum_{k \in G} A_k(U) = 1$ and $\sum_{k \in G} B_{kj}(U) = 0$ for each j . Integrability requires that $B_{ij}(U) = B_{ji}(U)$ is symmetric. We obtain that expenditure shares within G are:

$$H_{iG}(p_G, U) = A_i(U) + \sum_{j \in G} B_{ij}(U) \log p_j$$

Price semi-elasticities are given by:

$$\frac{\partial H_{iG}}{\partial \log p_j} = B_{ij}(U)$$

Conditional on initial expenditure shares and price semi-elasticities, we can again back out the change in expenditure shares for a given utility level:

$$H_{iG}(p'_G, U) = A_i(U) + \sum_j B_{ij}(U) \log p'_j$$

$$= H_{iG}(p_G, U) + \sum_j B_{ij}(U)(\log p'_j - \log p_j)$$

To obtain P^1 , applying Proposition 1 (now with an additive adjustment term) yields:

$$\begin{aligned} E_{iG}(p^0, y^1/P^1(p^0, p^1, y_h^1)) &= E_{iG}(p^1, y_h^1) + [H_{iG}(p_G^0, U_h^1) - H_{iG}(p_G^1, U_h^1)] \\ &= E_{iG}(p^1, y_h^1) + \sum_j B_{ij}(\log p_j^0 - \log p_j^1) \end{aligned}$$

As usual in the literature (e.g. Fajgelbaum and Khandelwal 2016), we could further specify that cross price elasticities are the same and equal to ξ_G/N_G , which leads to:

$$E_{iG}^0(y^1/P^1(p^0, p^1, y_h^1)) = E_{iG}^1(y_h^1) - \xi_G \times (\Delta \log p_i - \overline{\Delta \log p_G}) \quad (\text{A.12})$$

where $\overline{\Delta \log p_G}$ refers to the average log price change within group G .

An alternative restriction allowing for heterogeneous own-price elasticities can be:

$$B_{ij}(U) = \xi_i(U)\delta_{ij} - \frac{\xi_i(U)\xi_j(U)}{\sum_k \xi_k(U)}$$

where $\delta_{ij} = 1$ if $i = j$ and zero otherwise. This allows for non-uniform diagonal terms, yet symmetric cross-price elasticity terms generated by these diagonal terms.

Adjustments are then given by:

$$H_{iG}(p_G^1, U_h^1) - H_{iG}(p_G^0, U_h^1) = -\xi_i \left[\Delta \log p_i - \sum_j \xi_j / \left(\sum_k \xi_k \right) \Delta \log p_j \right]$$

Deviations from Quasi-Separability and Misclassification

Suppose we misclassify some goods i that would truly belong in G as a non- G good (i.e. we omit a good that belongs within the quasi-separable group G). Alternatively, suppose we falsely classify some non- G goods j as part of group G . In both cases, price changes outside of what we believe to be the G group can then directly affect within- G relative outlays (conditional on utility). In this context, we denote by $H_{iG}(p, U) = p_i h_i(p, U) / \sum_{j'} p_{j'} h_{j'}(p, U)$ the expenditure share in j within G (in terms of Hicksian demand), which now depends on the full vector of prices rather than just prices within G , but is still homogenous of degree zero in prices.

As a first-order approximation leads to the following equality, now taking sums for log price changes across all goods k :

$$\log (E_{iG}^0)^{-1} \left(\frac{x_{ih}^1}{x_{Gh}^1} \right) \approx \log \left(\frac{y_h^1}{P^1} \right) + (\beta_{ih}^0)^{-1} \sum_k (\Delta \log p_k - \overline{\Delta \log p_G}) \frac{\partial \log H_{iG}}{\partial \log p_k} \quad (\text{A.13})$$

Taking an average across goods $i \in G$, we obtain:

$$\frac{1}{G} \sum_{i \in G} \log (E_{iG}^0)^{-1} \left(\frac{x_{ih}^1}{x_{Gh}^1} \right) - \log \left(\frac{y_h^1}{P^1} \right) \approx \frac{1}{G} \sum_{i \in G} (\beta_{ih}^0)^{-1} \times \sum_{k \in NGUG} (\Delta \log p_k - \overline{\Delta \log p_G}) \frac{\partial \log H_{iG}}{\partial \log p_k}$$

The first source of bias that we already discussed is captured by the sum across goods $k \in G$ within the group: $\frac{1}{G} \sum_{i \in G} (\beta_{ih}^0)^{-1} \times \sum_{k \in G} (\Delta \log p_k - \overline{\Delta \log p_G}) \frac{\partial \log H_{iG}}{\partial \log p_k}$ (and equals zero if there is no relative price change within group G). The remaining bias is then coming from goods $k \in NG$ (classified as outside group G) as we describe in the main text:

$$\frac{1}{G} \sum_{i \in G} (\beta_{ih}^0)^{-1} \times \sum_{k \in NG} (\Delta \log p_k - \overline{\Delta \log p_G}) \frac{\partial \log H_{iG}}{\partial \log p_k}.$$

Tests of Quasi-Separability with Outside Price Data

Part i) of Lemma 2 shows that preferences are quasi-separable in G if and only if relative (compensated) expenditure shares x_i/x_G for any good $i \in G$ can be written as a function $H_{iG}(p_G, U)$ of within-group prices and utility. It is immediate to see that it is equivalent to state that 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 as stated in Corollary QS1:

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

If instead we hold income constant (uncompensated), we obtain:

$$\left. \frac{\partial \log(x_i/x_G)}{\partial \log p_j} \right|_y = \frac{\partial \log(x_i/x_G)}{\partial \log U} \frac{\partial \log V}{\partial \log p_j} \quad (\text{A.14})$$

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.14), we obtain:

$$\left. \frac{\partial \log(x_i/x_G)}{\partial \log p_j} \right|_y = -\frac{p_j q_j}{y} \frac{\partial \log(x_i/x_G)}{\partial \log U} \frac{\partial \log V}{\partial \log y} \quad (\text{A.15})$$

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.15), we obtain our condition applied to uncompensated demand, i.e. that equation (16) 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 $i' \in G$.

Tests of Quasi-Separability with Outside Expenditures Data

Corollary QS2 uses Slutsky symmetry to flip the role of goods inside group G with goods outside good G . Assuming rationality, Slutsky symmetry imposes that Slutsky substitution terms are symmetrical for any pair of good $i \neq j$. In terms of expenditures, this yields $S_{ij} = S_{ji}$ with:

$$S_{ij} = x_i \left. \frac{\partial \log x_i}{\partial \log p_j} \right|_y + \frac{x_i x_j}{y} \frac{\partial \log x_i}{\partial \log y}$$

Consider $i \in G, j \notin G$ (hence $\delta_{i=j} = 0$).

First, assume that the (necessary and sufficient) condition (16) from the previous corollary is satisfied, i.e. $\left. \frac{\partial \log(x_i/x_G)}{\partial \log p_j} \right|_y = -\frac{p_j q_j}{y} \frac{\partial \log(x_i/x_G)}{\partial \log y}$. We obtain:

$$S_{ij} = x_i \left[\left. \frac{\partial \log x_G}{\partial \log p_j} \right|_y - \frac{x_j}{y} \frac{\partial \log(x_i/x_G)}{\partial \log y} \right] + \frac{x_i x_j}{y} \frac{\partial \log x_i}{\partial \log y}$$

which simplifies into:

$$S_{ij} = x_i \left[\left. \frac{\partial \log x_G}{\partial \log p_j} \right|_y + \frac{x_j}{y} \frac{\partial \log x_G}{\partial \log y} \right]$$

On the other hand, the reverse Slutsky substitution term S_{ji} (flipping i and j) is given by:

$$S_{ji} = x_j \frac{\partial \log x_j}{\partial \log p_i} + \frac{x_i x_j}{y} \frac{\partial \log x_j}{\partial \log y}$$

Symmetry $S_{ij} = S_{ji}$ implies:

$$x_j \frac{\partial \log x_j}{\partial \log p_i} + \frac{x_i x_j}{y} \frac{\partial \log x_j}{\partial \log y} = x_i \left[\frac{\partial \log x_G}{\partial \log p_j} \Big|_y + \frac{x_j}{y} \frac{\partial \log x_G}{\partial \log y} \right] \quad (\text{A.16})$$

We can see from this expression that $x_j \frac{\partial \log x_j}{\partial \log p_i}$ is equal to a term that is proportional to x_i if we look across goods $i \in G$, from which we can conclude that:

$$\frac{\partial \log x_j}{\partial \log p_i} = (x_i/x_G) \sum_{i' \in G} \frac{\partial \log x_j}{\partial \log p_{i'}}$$

This condition is obtained for Marshallian (uncompensated) demand. Since the difference with compensated demand is proportional to the expenditure share of good i (second term of S_{ji}), the same proportionality result is obtained for compensated demand. Next, let us show that, under rational-

ity, condition from Corollary QS2 implies (sufficient) condition (16) for quasi-separability in group G . Take any good $i \in G$ and $j \notin G$. Starting from Slutsky symmetry condition (dividing by x_i) and using condition (16) between the first and second line below, we obtain:

$$\begin{aligned} \frac{\partial \log x_i}{\partial \log p_j} + \frac{x_j}{y} \frac{\partial \log x_i}{\partial \log y} &= \frac{1}{x_i} \left[\frac{\partial x_j}{\partial \log p_i} + \frac{x_i}{y} \frac{\partial x_j}{\partial \log y} \right] \\ &= \frac{1}{x_i} \left[\frac{x_i}{x_G} \sum_{k \in G} \frac{\partial x_j}{\partial \log p_k} + \frac{x_i}{y} \frac{\partial x_j}{\partial \log y} \right] \\ &= \frac{1}{x_G} \left[\sum_{k \in G} \frac{\partial x_j}{\partial \log p_k} + \frac{x_G}{y} \frac{\partial x_j}{\partial \log y} \right] \\ &= \frac{1}{x_G} \sum_{k \in G} \left[\frac{\partial x_j}{\partial \log p_k} + \frac{x_k}{y} \frac{\partial x_j}{\partial \log y} \right] \end{aligned}$$

Using again Slutsky symmetry $S_{jk} = S_{kj}$ for the term between brackets, $S_{kj} = \frac{\partial x_j}{\partial \log p_k} \Big|_y + \frac{x_k}{y} \frac{\partial x_j}{\partial \log y}$, we obtain:

$$\begin{aligned} \frac{\partial \log x_i}{\partial \log p_j} + \frac{x_j}{y} \frac{\partial \log x_i}{\partial \log y} &= \frac{1}{x_G} \sum_{k \in G} \left[\frac{\partial x_k}{\partial \log p_j} + \frac{x_j}{y} \frac{\partial x_k}{\partial \log y} \right] \\ &= \frac{1}{x_G} \left[\frac{\partial x_G}{\partial \log p_j} + \frac{x_j}{y} \frac{\partial x_G}{\partial \log y} \right] \\ &= \frac{\partial \log x_G}{\partial \log p_j} + \frac{x_j}{y} \frac{\partial \log x_G}{\partial \log y} \end{aligned}$$

which is equivalent to our sufficient condition (16) for quasi-separability in group G :

$$\frac{\partial \log(x_i/x_G)}{\partial \log p_j} = -\frac{x_j}{y} \frac{\partial \log(x_i/x_G)}{\partial \log y}$$

Aggregation across Varieties of a Good

Suppose that group G of goods can be further partitioned into subgroups of goods (varieties): $G = g1 \cup g2 \cup g3 \dots$. Denote by $E_{g,G}$ the expenditure share on subgroup g within group G . Under the assumptions of Lemma 1, we have for each variety: $E_{i,G}^1(y_h^1) = E_{i,G}^0(\frac{y_h^1}{P^1(y_h^1)})$, and $E_{i,G}^0(y_h^0) = E_{i,G}^1(\frac{y_h^0}{P^0(y_h^0)})$. Taking the

sum across varieties $i \in g$ of a subgroup g , we obtain:

$$E_{g,G}^1(y_h^1) = \sum_{i \in g} E_{i,G}^1(y_h^1) = \sum_{i \in g} E_{i,G}^0\left(\frac{y_h^1}{P^1(y_h^1)}\right) = E_{g,G}^0\left(\frac{y_h^1}{P^1(y_h^1)}\right) \quad (\text{A.17})$$

and:

$$E_{g,G}^0(y_h^0) = \sum_{i \in g} E_{i,G}^0(y_h^0) = \sum_{i \in g} E_{i,G}^1\left(\frac{y_h^0}{P^0(y_h^0)}\right) = E_{g,G}^1\left(\frac{y_h^0}{P^0(y_h^0)}\right) \quad (\text{A.18})$$

This proves the corollary.

Next, suppose that there exists a price index $P_g(p_g, U)$ summarizing prices for subgroup g , i.e. such that the expenditure function can be written: $e(p, U) = \tilde{e}(e_G(P_{g1}(p_{g1}, U), P_{g2}(p_{g2}, U), P_{g3}(p_{g3}, U), \dots))$. In this case, we can again relax the assumption of Lemma 1: equations (A.17) and (A.18) above hold if we assume that relative price indices remain constant, i.e. $P_g^1(p_g, U) = \lambda_G P_g^0(p_g, U)$ instead of assuming that the relative prices of all varieties remain constant within group G . We can use these price indices in Proposition 1 instead of the prices of individual varieties.

To see this, first note that we can express within-G compensated expenditure shares across subgroups g as a function of prices indices $P_g(p_g, U)$ instead of the full vector of within-G prices p_G :

$$H_{g,G}(P_{g1}, P_{g2}, \dots, U_h) = \frac{\sum_{i \in g} x_{hi}}{x_{hG}} = \frac{\partial \log e_G(P_{g1}, P_{g2}, \dots, U)}{\partial \log P_g}$$

(see the proof of Lemmas 1 and 2, equation A.1, for the derivation of compensated expenditure shares).

Taking the sum across varieties $i \in g$, and using $H_{g,G}(P_G, V(p, y)) = \sum_{i \in g} H_{i,G}(p_G, V(p, y)) = \sum_{i \in g} E_{iG}(y) = E_{gG}(y)$, we obtain, as in Proposition 1:

$$\begin{aligned} E_{gG}^0\left(y^1/P^1(p^0, p^1, y_h^1)\right) &= H_{g,G}(P_G^0, V(p^0, y_h^1/P^1(p^1, p^0, y_h^1))) \\ &= H_{g,G}(P_G^0, V(p^1, y^1)) \\ &= H_{g,G}(P_G^1, V(p^1, y^1)) \times \frac{H_{iG}(P_G^0, U_h^1)}{H_{iG}(P_G^1, U_h^1)} \\ &= E_{gG}^1(y_h^1) \times \exp \sum_{g' \subset G} \int_{\log P_{g'}^1}^{\log P_{g'}^0} \frac{\partial \log H_{g,G}}{\partial \log P_{g'}} d \log P_{g'} \end{aligned}$$

where we use subgroup price indices P_g instead of individual prices p_G . By homogeneity of degree zero in subgroup price indices P_g , we obtain $H_{g,G}(P_G^0, V(p^1, y^1)) = H_{g,G}(P_G^1, V(p^1, y^1))$ if $P_g^1(p_g, U) = \lambda_G P_g^0(p_g, U)$, and thus $E_{gG}^0\left(y^1/P^1(p^0, p^1, y_h^1)\right) = E_{gG}^1(y_h^1)$.

Finally, note that we can also reformulate the orthogonality condition in the first-order approximation across subgroups, using price indices across subgroups instead of good-level prices.

Statistical Demand with Heterogeneous Preferences

Lewbel (2001) asks the following question: assuming that each individual demand is rational, when is “statistical demand” (i.e. conditional expectation based on observables) also rational? Here, more appropriate to our setting, we instead assume that each individual demand is rational and quasi-separable in group G , and examine conditions under which statistical demand is also rational and quasi-separable.

Assuming that heterogeneity in demand comes from unobserved household characteristics, consider $X_i = E[x_i | y, p, x_G]$ the fitted expenditures conditional on prices p , total outlays y and total expenditures x_G on group G of goods (with NG the set of goods outside G).² Rationality requires that the Slutsky substitution matrix is i) symmetric and ii) semi-definite negative. A demand system satisfying the symmetry condition i) is said to be “integrable”. Suppose that each consumer is rational with preferences that are quasi-separable in group G of goods.

²As in Lewbel (2001), both prices and income must be included in the set of conditional observables in order to define a complete Marshallian demand system. Here we also add total expenditures on group G , $x_G = \sum_{i \in G} x_i$, to the set of controls so that we can focus our attention on heterogeneity in within-group expenditures. We can also add a subset of household characteristics in the set of conditional variables.

First, under which conditions do we have QS for aggregate demand X ? Applying the equivalence in equation (16), we need:

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

for each $i \in G$ and $j \notin G$. Since

$$\left. \frac{\partial(X_i/x_G)}{\partial \log p_j} \right|_y = \frac{1}{x_G} \frac{\partial}{\partial \log p_j} E[x_i | y, p, x_G] = E \left[\left. \frac{\partial(x_i/x_G)}{\partial \log p_j} \right| y, p, x_G \right] = E \left[-\frac{x_j}{y} \left. \frac{\partial(x_i/x_G)}{\partial \log y} \right| y, p, x_G \right]$$

(given that each individual demand is quasi-separable), the condition for quasi-separability for statistical demand is equivalent to:

$$\begin{aligned} \left. \frac{\partial(X_i/x_G)}{\partial \log p_j} \right|_y = -\frac{X_j}{y} \frac{\partial(X_i/x_G)}{\partial \log y} &\iff E \left[-\frac{x_j}{y} \left. \frac{\partial(x_i/x_G)}{\partial \log y} \right| y, p, x_G \right] = -\frac{X_j}{y} \frac{\partial(X_i/x_G)}{\partial \log y} \\ &\iff E \left[x_j \left. \frac{\partial x_i}{\partial \log y} \right| y, p, x_G \right] - E[x_j | y, p, x_G] E \left[\left. \frac{\partial x_i}{\partial \log y} \right| y, p, x_G \right] = 0 \end{aligned}$$

Hence, expressed in terms of covariance as in Lewbel (2001), we need the partial covariance matrix between x and $\frac{\partial x}{\partial \log y}$ (conditional on x_G, y and p):

$$L_{ij} = Cov \left(x_j, \left. \frac{\partial x_i}{\partial \log y} \right| y, p, x_G \right) \quad (\text{A.19})$$

to be block diagonal for G vs. non- G goods, i.e. with zero coefficients aside from the diagonal block for G goods and the diagonal block for non- G goods.

Then, applying results from Lewbel (2001), for integrability it must be also that each block in G and non- G goods is symmetric in i and j (necessary and sufficient condition). Furthermore, a sufficient condition for rationality is that this matrix is semi-definite positive. Hence, to have both quasi-separability and rationality for statistical demand, a sufficient condition is that each diagonal block of the covariance matrix A.19 (in G and non- G goods) be a symmetric semi-definite positive matrix, while off-diagonal blocks are null (note again that it is a covariance conditional on observing total expenditures on group G of goods).

Implication of Heterogeneous Preferences for Price Indices

Here we examine the role of heterogeneity in preferences across demographic groups. Denote each group by an index z .

As a first simple case, assume that each group experience the same price index change for a given level of income (yet still heterogeneous across the income distribution). With a common change in price indices, the horizontal shift is the same across groups:

$$\frac{x_{i,h,z}^1}{x_{G,h,z}^1} = E_{iG,z}^1(y_h^1) = E_{iG,z}^0 \left(\frac{y_h^1}{P^1(y_h^1)} \right)$$

It is then easy to see that the average relative Engel curve across groups also shifts by $P^1(y_h^1)$, conditional on income y_h^1 :

$$\bar{E}_{iG}^1(y_h^1) = \bar{E}_{iG}^0 \left(\frac{y_h^1}{P^1(y_h^1)} \right)$$

Hence, the average Engel curve across demographic groups we can still help us identify the price index change.

Now, suppose that $P_z^1(y_h^1)/P_{ref}^1(y_h^1) = 1 + \varepsilon_z^1(y_h^1)$. As a first-order approximation in ε , we obtain:

$$E_{iG,z}^1(y_h^1) = E_{iG,z}^0 \left(\frac{y_h^1}{P^1(y_h^1)} \right) \approx E_{iG,z}^0 \left(\frac{y_h^1}{P_{ref}^1(y_h^1)} \right) - \beta_{i,z}^1 \varepsilon_z^1$$

where $\beta_{i,z}^1(y_h^1)$ is the slope of the relative Engel curve for good i from period 1 for group z evaluated at

income $y_h^1/P_{ref}^1(y_h^1)$ in log. Taking averages across groups, we obtain:

$$\bar{E}_{iG}^1(y_h^1) \approx \bar{E}_{iG}^0\left(\frac{y_h^1}{P_{ref}^1(y_h^1)}\right) - \frac{1}{Z} \sum_z \beta_{i,z}^1 \varepsilon_z^1$$

If we use the average Engel curve, our estimated price index \tilde{P}^1 is then such that:

$$\bar{E}_{iG}^0\left(\frac{y_h^1}{\tilde{P}^1(y_h^1)}\right) \approx \bar{E}_{iG}^0\left(\frac{y_h^1}{P_{ref}^1(y_h^1)}\right) - \frac{1}{Z} \sum_z \beta_{i,z}^1 \varepsilon_z^1$$

Inverting using the average relative Engel curve, this yields:

$$\log \tilde{P}^1(y_h^1) \approx \log P_{ref}^1(y_h^1) + \frac{1}{Z} \sum_z \beta_{i,z}^1 \varepsilon_z^1 / \bar{\beta}_i^1$$

where $\bar{\beta}_i^1$ denotes the average of the derivatives: $\bar{\beta}_i^1 = \frac{1}{Z} \sum_z \beta_{i,z}^1$ (and its inverse is equal to the derivative of the inverse of the average log Engel curve). If the price index is estimated by taking an average across goods, we obtain:

$$\log \tilde{P}^1(y_h^1) \approx \log P_{ref}^1(y_h^1) + \frac{1}{Z} \sum_z \beta_{i,z}^1 \varepsilon_z^1 / \bar{\beta}_i^1$$

This shows that, if preferences are heterogeneous within a given level of income, we can interpret our naive estimator as a weighted estimator of heterogeneous price index changes, with weights proportional to $\sum_i \beta_{i,z}^1 / \bar{\beta}_i^1$.

B Data Appendix

This appendix details the various steps we took both to clean the raw survey data as well as to select goods for which we have reliable price data. We use rounds 43 (1987–88) and 55 (1999–2000) of the National Sample Surveys (NSS) produced by the Indian National Sample Survey Office.³ We extract expenditures and quantities (where available) on all goods and services alongside household identifiers, geographic indicators, survey weights, and basic household characteristics such as household size, age, education level, and literacy of the household head. As several districts split between survey rounds, we concord districts in round 55 back to the 43rd round districts.

Turning to the expenditures data, we first concord items whose descriptions changed slightly over survey rounds. The full concordance between rounds is presented in the table notes for Table 1. Mostly, concordance consists of aggregating two goods that were asked for separately in one NSS round and jointly in the other. We aggregate these into a single item to be consistent over rounds. For example, “Jowar” and “Jowar products” were two separate items in NSS 43, but then became “Jowar and Jowar products” in NSS 55, and again we aggregate these into a single item to be consistent over time.

Three concordances are related to purchases from India’s Public Distribution System (PDS): PDS and non-PDS purchases of rice, sugar and wheat were reported separately in NSS 55 but not in NSS 43, and so we pool the two types of purchase into one concorded item. As not all households are eligible to purchase goods at subsidized prices through the PDS, the assumption implicit in our methodology that all households in a location face the same price vector is violated. However, note that our methodology can accommodate price vectors that are functions of utility since horizontal shifts in relative Engel curves recover the change in nominal income required to hold utility at either its initial (P^0) or final (P^1) level. Fortunately, the eligibility criterion of the PDS program is essentially based on utility—specifically households below the poverty line are eligible, with the poverty line based on real needs not nominal incomes. Thus, when moving horizontally between period 0 and period 1 relative Engel curves to infer price index changes, PDS eligibility does not change. For example, a household initially at a utility level below (above) the PDS cutoff will be eligible (ineligible) in both periods at the utility level used to construct P^0 (and similarly for constructing P^1 but basing eligibility on the household’s final level of utility).

To calculate our measure of expenditures per capita consistently over rounds, we drop the “taxes and cesses” item that is asked in NSS55 but not in NSS43, and there is no obvious item within which it is subsumed in the latter. We also drop expenditures on three items for which we observe very few purchases (fewer than 20 purchases per round across all of India). These are jewels and pearls; other machines for household work; and other therapeutic appliances and equipment. For items with an expenditure period of 365 days (i.e. durables), we obtain the equivalent monthly measure by dividing by 365 and multiplying by 30. We then sum up monthly expenditures on all NSS items and divide by household size to obtain our measure of total outlays per capita. The NSS also provides a mean per capita expenditure variable that is not necessarily equal to the sum of monthly item expenditures we calculate. For this reason, we drop observations for which the NSS-provided per capita expenditure differs substantially from our expenditures per capita measure (a discrepancy of more than 1 SD of per capita expenditures) resulting in a reduction of about 1 percent of the sample in either round.

We obtain price data from unit values, i.e. dividing expenditures by quantities where both are reported. The following paragraphs detail how we obtain our sample of 132 goods with reliable price data.

We closely follow Deaton and Tarozzi (2005) by eliminating items for which unit values are

³These are available for download at <http://www.icsrdataservice.in/datarepository/index.php/catalog/7> and <http://www.icsrdataservice.in/datarepository/index.php/catalog/12>, respectively.

unlikely to be reliable measures of prices. Their methodology explores variation in unit prices within localities to identify products with multi-modal price distributions, suggestive of either multiple measurement units, multiple quality levels, or some combination of the above.

We implement their product exclusions by first dropping all good and service categories where quantity data are not recorded. We then further exclude the clothing and footwear categories for which quantity data exist (e.g. 2 pairs of leather boots/shoes) but where product descriptions are too broad and styles too numerous to generate reliable unit values. The remaining goods are all food and fuel products.

In the next step we drop goods listed in Deaton and Tarozzi (2005) Table A2 (other fresh fruits, other beverages; biscuits and confectionery; salted refreshments; prepared sweets; other processed food; other drugs and intoxicants; dung cake; gobar gas; other fuel and light) that lack quantity data, or have quantity data although the enumerator instructions do not request it.⁴

Next, we drop goods listed in Deaton and Tarozzi (2005) Tables A3 and A4 (Other wheat products; ice cream; other milk products; other nuts; other dry fruits; ice; fruit juice and shakes; other ingredients for pan; liquid petroleum gas; candles; cereal substitutes; other spices; other meat, birds and fish; coconut; tea (cups); coffee powder; cold beverages; cake and pastries; pan leaf; hookah tobacco; and toddy). These are items where the variation in prices within localities suggests that these products likely contain multiple varieties or quality levels; either because there is strong evidence of multi-modal price distributions (e.g. liquid petroleum gas), or due to the combination of high price dispersion and broad product descriptions (e.g. “other milk products”).


Next, we discard items listed in Deaton and Tarozzi (2005) Table A5 where changes in the unit of measurement over rounds make temporal comparisons impossible. Either items unit of measurement changed from kilos to units and vice-versa between rounds (lemon; guava), or units appear to have changed between rounds (coal gas; cheroot; zarda, kimam and surti; other tobacco products; ganja). This leaves 132 food and fuel items.

Unit values are calculated for each household by taking the ratio of expenditures to quantities. With household-level unit prices, we implement Deaton and Tarozzi’s automatic test for unit price outliers in each round, which consists of dropping unit price observations for which the absolute value of the difference between the log unit price and the mean log unit price for the particular NSS item is larger than two standard deviations of the log price. Once unit prices have been purged of outliers, we take the median price for every NSS item in a district and round as our price for the item in the district. The use of medians is recommended by Deaton and Tarozzi (2005) due to its robustness against outliers. In our final sample of 132 food and fuel items, the average household bought 26 items in round 43 and 31 items in round 55.

⁴While “Egg products” are dropped by Deaton and Tarozzi (2005), in NSS 55 the survey changed slightly and this item was merged with the larger category “Eggs”, so we decided to keep them as a single concorded item. Table 1 table notes reports the concordance we used.


C Additional Figures and Tables

Figure C.1: Sparseness Across Alternative Product Aggregations




Notes: Figure plots histogram of share of households with any observed consumption by product-period-market cell across three alternative levels of product aggregation. See Section 4.2 for further discussion.

Figure C.2: Shifts in Relative Engel Curves for Salt Over Time




Notes: Figures plot relative Engel curves for salt over time (1987/1988 NSS 43rd Round to 1999/2000 NSS 55th round) for the largest markets in the four broad regions of India (in terms of numbers of households surveyed). A market is defined as the rural area of an Indian district. Fitted relationships are based on local polynomial regressions using an Epanechnikov kernel with a bandwidth equal to one quarter of the range of the income distribution in a given market. See Section 3.1.1 for further discussion.

Figure C.3: Rural Indian Cost of Living Inflation 1987/88–1999/2000: AFGG Price Index with Alternative Exact Price Corrections




Notes: Figure shows the percentage change in the rural AFGG price index between 1987/88 and 1999/2000 for each decile of the local per-capita outlay distribution (averaged across districts using population weights). Left panel reports the uncorrected price index change. Remaining three panels apply the exact correction, described in Proposition 1 and Section 3.1.3, using elasticities shown in Appendix Table C.3. See Section 4.3 for further discussion.

Figure C.4: Rural Indian Cost of Living Inflation 1987/88–1999/2000: Estimates Using Alternative Levels of Good Aggregation



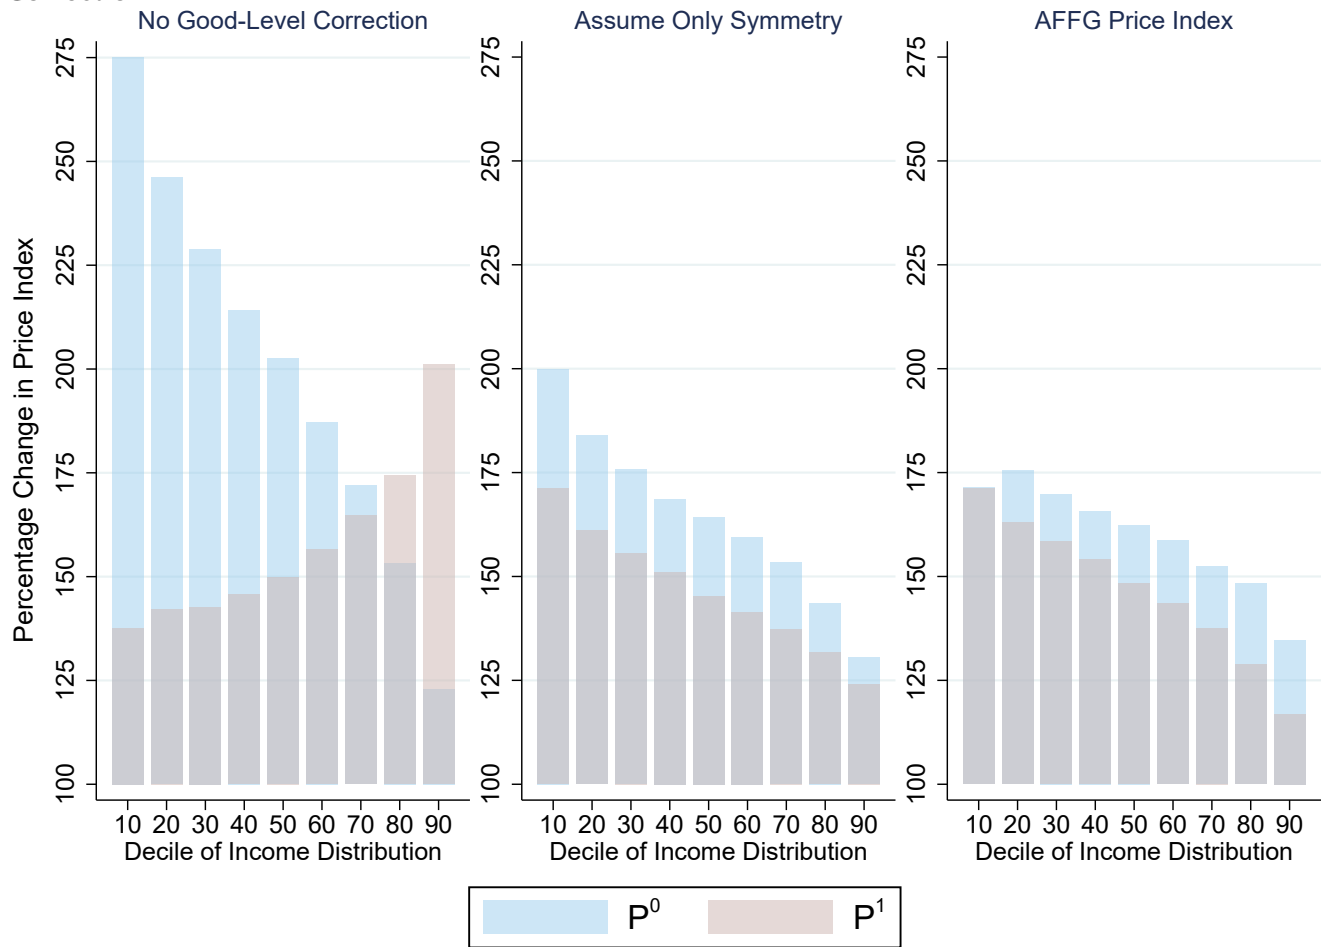
Notes: Figure shows the average percentage change in the rural AFFG NPC price index between 1987/88 and 1999/2000 for each decile of the local per-capita outlay distribution (averaged across districts using population weights). Estimates are based on horizontal shifts in relative Engel curves. The three panels use different levels of aggregation of goods in the Indian expenditure microdata. The left panel depicts our baseline estimation approach which aggregates the 132 products to 34 products (the second-lowest level of aggregation in the NSS surveys). The middle panel uses the disaggregated 132 products, while the right panel further aggregates to 8 products (the third-lowest level of aggregation in the NSS surveys). See Section 4.2 for further discussion.

Figure C.5: Recall Bias: Rural Indian Cost of Living Inflation 1987/88–1994/95




Notes: Figure shows the percentage change in the rural price index between 1987/88 and 1994/1995 for each decile of the local per-capita outlay distribution (averaged across districts using population weights). Left panel plots our AFFG NPC price index changes estimated from horizontal shifts in relative Engel curves. Middle panel plots price index changes using Laspeyres and Paasche district-level CPIs calculated using price changes of food and fuels following Deaton (2003b). Right panel repeats the middle panel but using district-income-decile-specific budget shares to calculate the Laspeyres and Paasche indices. See Section 4.3 for further discussion.

Figure C.6: Good-Level Selection Corrections (1): Price Index Changes With and Without Bias Correction



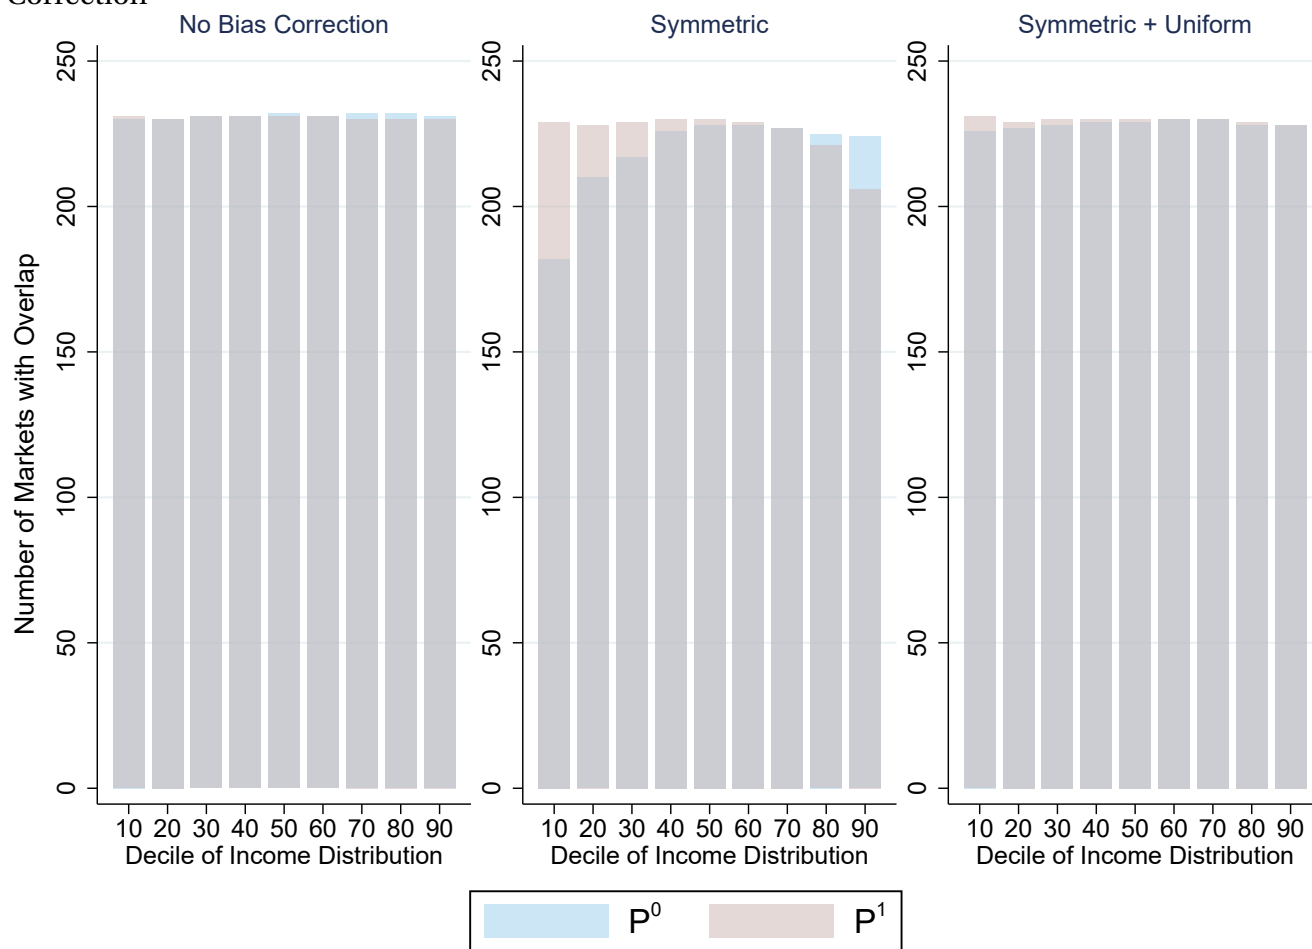
Notes: Figure shows the percentage change in the rural AFFG NPC price index between 1987/88 and 1999/2000 for each decile of the local per-capita outlay distribution (averaged across districts using population weights), both with and without correcting estimates for selection bias described in Section 3.2.4. Left panel plots estimates that are simple averages of all overlapping Engel curves for a particular market. Middle panel accounts for bias from non-overlapping Engel curves by assuming distribution of price index estimates within a market is symmetric, ordering both overlapping and non-overlapping estimates, and taking the median when observed. The Right panel, our baseline AFFG approach, further assumes the distribution is uniform to calculate medians when not observed. See Section 4.4 for further discussion.

Figure C.7: Good-Level Selection Corrections (2): Reasons for Non-Overlap




Notes: Figure shows the frequency of non-overlapping estimates by decile, broken out by type of non-overlap (censored from above or from below). This information is used to rank missing (non-overlapping) estimates and calculate the medians required for the good-level selection correction applied in both the middle and right panel of Appendix Figure C.6. See Section 4.4 for further discussion.

Figure C.8: Good-Level Selection Corrections (3): Number of Markets With and Without Bias Correction




Notes: Figure shows the number of missing market-decile pairs after applying the good-level selection correction just using symmetry (middle panel) and symmetry plus uniformity (our baseline, right panel). For comparison, left panel shows the number of market-decile pairs where we have at least one good with overlapping monotonic relative Engel curves at that decile of the income distribution and so can obtain an estimate of the price index without any bias correction. See Section 4.4 for further discussion.

Figure C.9: Rural Indian Cost of Living Inflation 1987/88–1999/2000: Including Household Controls



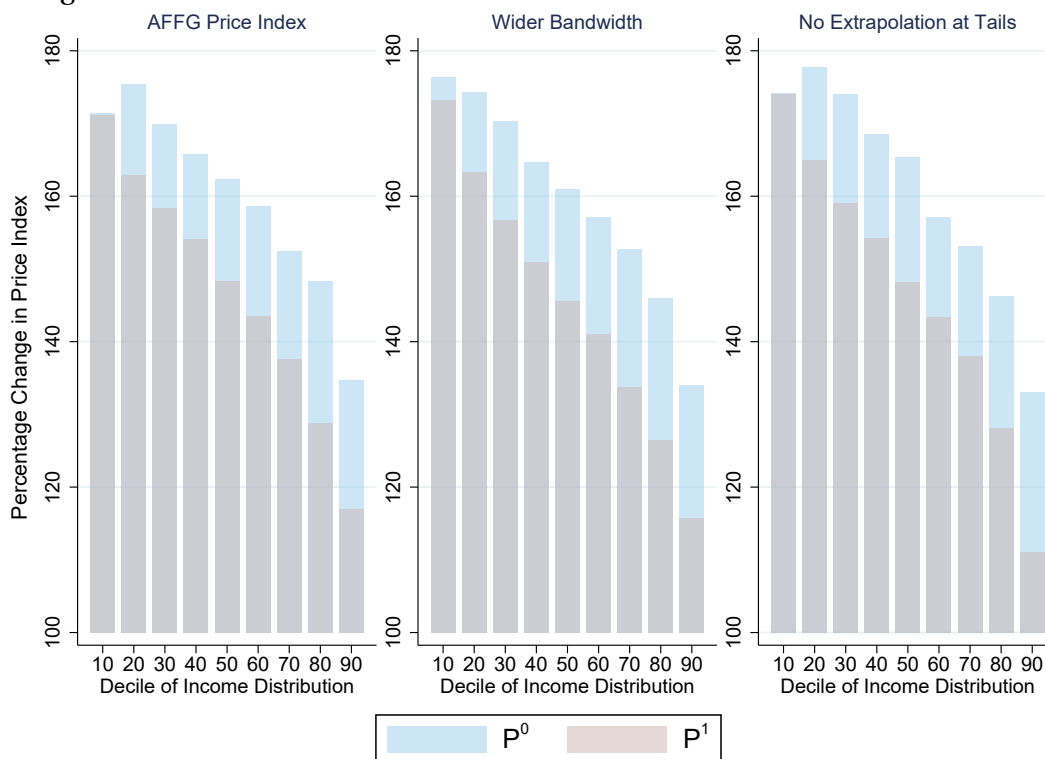
Notes: Left panel shows our baseline AFFG NPC price index estimates. Right panel shows estimates after conditioning on household controls when estimating relative Engel curves. Specifically, for each good and market, we non-parametrically regress relative budget shares against log total outlays per capita and include linear controls for household characteristics (a scheduled caste dummy, a literacy of household head dummy, log of household size, and the share of children in the household). Coefficients from these market-good specific linear controls are used to evaluate relative budget shares at the market median value (constant over time) for each characteristic. We then use these characteristic-adjusted budget shares to obtain the price index changes shown in the right panel. See Section 4.4 for further discussion.

Figure C.10: Rural Indian Cost of Living Inflation 1987/88–1999/2000: Heterogeneity by Household Type



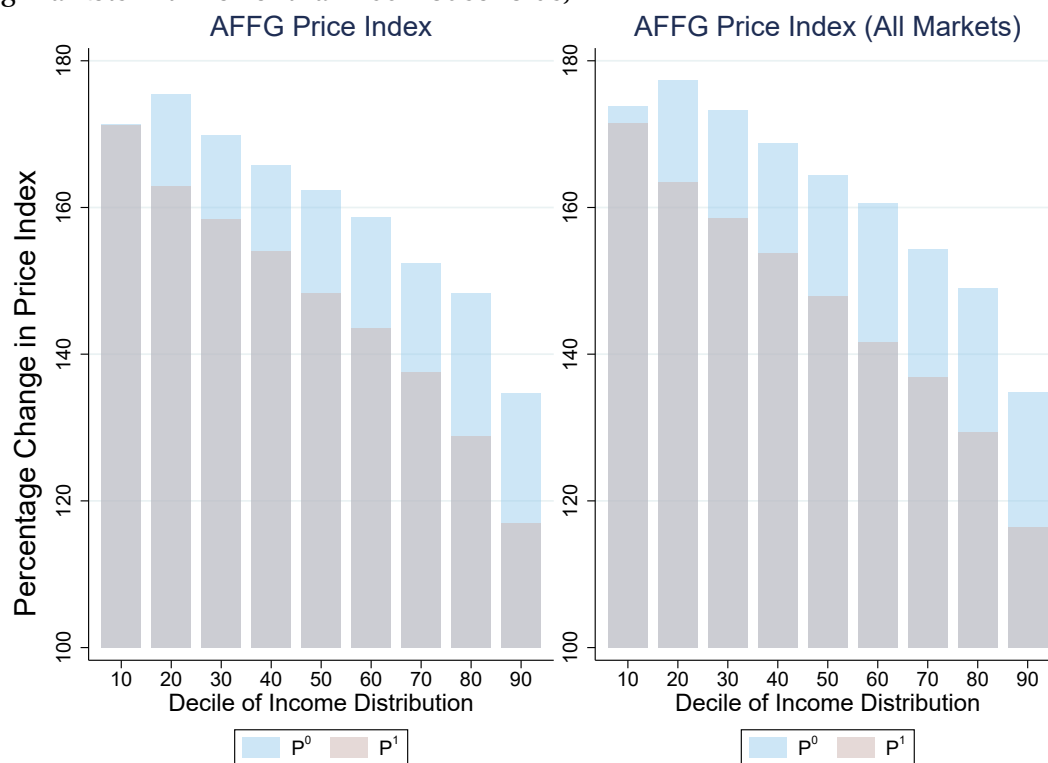
Notes: Figure shows rural AFFG NPC P^1 price index changes by decile of the local per-capita outlay distribution (averaged across districts using population weights), partitioning households within each market along four dimensions: (top-left) above and below median household size in the district, (top-right) above and below median household head education level in the district, (bottom-left) above and below median household head age in the district, (bottom-right) by literate/illiterate status of the household head.

Figure C.11: Rural Indian Cost of Living Inflation 1987/88–1999/2000: Alternative Estimates of Relative Engel Curves



Notes: Figure shows AFGG NPC price index changes using alternate methods of estimating relative Engel curves. Left panel reproduces our baseline approach. Recall that the baseline approach uses an Epanechnikov kernel for non-parametrically estimating Engel curves equal to one quarter of the range of the income distribution. Additionally, we restrict attention to good-market combinations where Engel curves in both periods are monotonic between percentiles 1 and 99 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. Middle panel extends the bandwidth of the Epanechnikov kernel used to 30 percent of the range. Right panel does not replace extreme non-monotonic values with linear extrapolations.

Figure C.12: Rural Indian Cost of Living Inflation 1987/88–1999/2000: Using All Markets (Including Markets with Fewer than 100 Households)



Notes: Figure shows the percentage change in the rural AFG NPC price index between 1987/88 and 1999/2000 for each decile of the local per-capita outlay distribution (averaged across districts using population weights). Left panel plots our baseline price index changes that exclude small markets (those with fewer than 100 households surveyed in each survey round). Right panel plots price index changes including all markets.

Table C.1: Changes in Recall Periods and Within-Group Budget Shares

7-day recall interaction	Dependent variable: Relative budget shares (34 <i>g</i> goods)			Dependent variable: Relative budget shares (136 disaggregated NSS goods)
	Coefficient	Standard error	<i>t</i> -stat	
7-day recall X Cereals - coarse	0.00281	0.00292	0.96	
7-day recall X Cereals - rice	0.00213	0.00139	1.54	
7-day recall X Cereals - wheat	0.00103	0.00149	0.69	
7-day recall X Coke, coal, charcoal	0.01053	0.00511	2.06	
7-day recall X Dry fruits and nuts	-0.00014	0.00066	-0.2	
7-day recall X Eggs	-0.00065	0.00054	-1.21	
7-day recall X Electricity	0.00029	0.00316	0.09	
7-day recall X Firewood and chips	0.00177	0.00212	0.84	
7-day recall X Fish, prawn	0.00042	0.00117	0.36	
7-day recall X Ghee	0.00146	0.00279	0.52	
7-day recall X Gram	0.00140	0.00070	1.99	
7-day recall X Intoxicants	-0.00194	0.00400	-0.48	
7-day recall X Kerosene	-0.00210	0.00256	-0.82	
7-day recall X Matches	0.00009	0.00076	0.12	
7-day recall X Meat	0.00060	0.00124	0.48	
7-day recall X Milk	0.00038	0.00179	0.21	
7-day recall X Other Fresh fruits	0.00037	0.00099	0.38	
7-day recall X Other milk products	-0.00081	0.00262	-0.31	
7-day recall X Pan	-0.00122	0.00120	-1.02	
7-day recall X Premium Fruits	0.00012	0.00065	0.18	
7-day recall X Pulses - Besan, Moong	0.00003	0.00059	0.05	
7-day recall X Pulses - Urd, Masur	-0.00012	0.00061	-0.2	
7-day recall X Tobacco	0.00289	0.00122	2.36	
7-day recall X Vanaspati, margarine	0.00164	0.00167	0.98	
7-day recall X Vegetable - gourds	0.00033	0.00050	0.67	
7-day recall X Vegetable - leafy vegetables	0.00053	0.00056	0.96	
7-day recall X Vegetable - other vegetables	0.00005	0.00036	0.15	
7-day recall X Vegetable - root vegetables	-0.00022	0.00063	-0.35	
7-day recall X beverages	-0.00055	0.00069	-0.8	
7-day recall X edible oils	0.00093	0.00150	0.62	
7-day recall X processed food	0.00183	0.00367	0.5	
7-day recall X salt	0.00060	0.00038	1.61	
7-day recall X spices	-0.00055	0.00090	-0.61	
7-day recall X sugar	0.00021	0.00084	0.25	
District X <i>g</i> good Fixed Effects	Yes			No
District X disaggregated item Fixed Effects	No			Yes
Household weights	Yes			Yes
F-stat schedule*goods=0	1.04			1.02
<i>p</i> -value schedule*goods=0	0.401			0.422
N	263663			384344

Notes: For questions regarding quantities and expenditures on food, pan, tobacco and intoxicants, the thin NSS round 54 (January-June 1998) randomized households between a 30-day and a 7-day recall period. Table tests whether reported relative budget shares (expenditure on good *i* divided by expenditures on all goods in good *i*'s *G* group) change with the recall period used. Columns 1–3 report coefficient estimates, standard errors and *t*-statistics from regression of relative budget shares on a dummy for whether the household was surveyed with a 7 day-recall period interacted with each of the 34 *i* products (after including district-product fixed effects). A significant coefficient on the interaction indicates that the recall period affected relative consumption reports for that good. The bottom of the table reports the test of joint significance for all interactions. Column 4 repeats the exercise but for the 132 disaggregated goods rather than the 34 aggregated goods we use in our baseline. Given the large number of estimates, in this case we simply report the *F*-statistic and *p*-value for joint significance at the bottom of the table.

Table C.2: Price changes for 34 g goods

34 g goods	Mean Percentage Change	SD of Percentage Change
Cereals - rice	178	13
Cereals - wheat	218	14
Cereals - coarse	205	28
Gram	243	20
Pulses - besan, moong	239	12
Pulses - urd, masur	212	11
Meat	218	11
Fish, prawn	176	29
Eggs	118	18
Vegetable - root vegetables	134	16
Vegetable - gourds	184	23
Vegetable - leafy vegetables	183	19
Vegetable - other vegetables	193	29
Premium fruits	191	29
Other fresh fruits	160	25
Dry fruits and nuts	109	35
Ghee	155	19
Milk	174	13
Other milk products	153	53
Vanaspati, margarine	59	13
Edible oils	62	13
Sugar	127	14
Salt	354	45
Spices	196	18
Beverages	203	33
Processed food	197	41
Pan	252	34
Tobacco	222	20
Intoxicants	120	87
Coke, coal, charcoal	100	50
Kerosene	26	13
Firewood and chips	118	33
Electricity	147	34
Matches	117	15

Notes: We use a Stone price index to aggregate the observed price changes of the 132 products i in the NSS to 34 goods g (using survey-weighted mean initial expenditure shares across the $i \in g$ to compute weights). Price changes for each of the 132 food and fuel items are computed from changes in district median unit values as described in Data Appendix B. When unit values are observed in the district for one but not the other period, we replace i 's missing price change with the state-level change. The first column of the table reports district-weighted means of percent changes in prices for each of the 34 goods g , along with the standard deviation of the percent change in the second column.

Table C.3: Elasticity Estimates for Two-Tier Demands

Group	Subgroup	Kumar et al. (2011)	Fally Sayre (2018)	Region	References
G1	Cereals	0.031	-0.309	World, India, EU, Argentina, US	[3], [4], [5], [6], [7]
G1	Pulses	-0.635	-0.307	World, India	[3], [4], [5], [6],
G1	F&V	-0.917	-0.515	World, US, UK	[3],
G2	Oils	-0.377	-0.190	World	[3],
G2	Sugar	-0.010	-0.236	World, India, Australia	[1], [3], [4], [5],
G2	Milk	-1.035	NA		
G1, G2	Other	-1.259	NA		
G3	Fuel	NA	-0.550	China	[2]

Notes: Table presents the elasticity estimates used for the exact price correction used in the two-tier demand structure discussed in Section 4.3 and shown in Figure C.3. The first set of estimates come from Kumar et al. (2011), the second from Fally and Sayre (2018) which reviews multiple studies (elasticities shown in in the fourth column are averages across references listed in the last column and below). For “Milk” and “Other” (not covered in Fally and Sayre 2018), we use the elasticity from Kumar et al. (2011).

- [1] Adams, F. G., and J. R. Behrman (1976): *Econometric models of world agricultural commodity markets: cocoa, coffee, tea, wool, cotton, sugar, wheat, rice*. Ballinger Publishing Company.
- [2] Burke, P. J., and H. Liao (2015): “Is the price elasticity of demand for coal in China increasing?,” *China Economic Review*, 36, 309 – 322.
- [3] Food and Agricultural Policy Research Institute (2017): *Elasticities Database*. Iowa State University.
- [4] Kumar, P., A. Kumar, P. Shinoj, and S. Raju (2011): “Estimation of demand elasticity for food commodities in India,” *Agricultural Economics Research Review*, 24(1), 1–14.
- [5] Kumar, P., P. Shinoj, S. S. Raju, A. Kumar, K. M. Rich, and S. Msangi (2010): “Factor demand, output supply elasticities and supply projections for major crops of India.,” *Agricultural Economics Research Review*, 23(1).
- [6] Roberts, M. J., and W. Schlenker (2013): “Identifying supply and demand elasticities of agricultural commodities: Implications for the US ethanol mandate,” *The American Economic Review*, 103(6), 2265–2295.
- [7] Russo, C., R. D. Green, and R. E. Howitt (2008): “Estimation of supply and demand elasticities of California commodities,” *Department of Agricultural & Resource Economics, UC Davis*.
- [8] Tiffin, R., K. Balcombe, M. Salois, and A. Kehlbacher (2011): *Estimating food and drink elasticities*. University of Reading.

D Corroborating Evidence for Pro-Rich Inflation Between 1987–2000

As discussed in Section 4.3 of the main text, the most likely explanation is that high-income households disproportionately benefited from price drops, new varieties, and quality increases in manufactures and services where price measurement is most challenging. The rich spent a large (and increasing) share of their budget on durables such as manufactures and on services. (In 1987–1988, the richest decile spent 35.8 percent of their expenditure on non food excluding fuels, durables and services on average, and the poor spent 23.5 percent. In 1999–2000, these shares rose to 42.4 and 28.3 percent, respectively.) If the rich and poor purchased different sets of products within these sectors and price changes, quality improvements and product introductions were focused on the products the rich purchased, inflation rates could differ substantially across the income distribution. We provide three pieces of suggestive evidence for this mechanism.

First, we examine category-level rural CPI data (All India Consumer Price Index for Agricultural Labourers, CPI-AL). The Indian government breaks out rural inflation into five categories (food; pan, supari, tobacco & intoxicants; fuel & light; clothing, bedding & footwear; and miscellaneous). The last two components of CPI that capture manufactures and services suffer from the same missing price issues that our methodology aims to address but, taken at face value, inflation was more moderate in these categories disproportionately consumed by the rich (202 percent for clothing and miscellaneous versus 211 percent for food, intoxicants, and fuels).

While these differences are relatively small, what matters for distributional effects is not only the difference in inflation for the category as a whole, but differential changes in inflation for goods purchased by the poor and the rich. We now report a novel finding that relative Engel curves are much steeper within durable manufactures and services than within foods, intoxicants, fuels and non durable manufactures. To do so we use our household survey data to estimate Engel slopes from regressing shares of the 301 different products within these six large product groupings (“sectors”) on log per capita household expenditures for all rural Indian households (separately by survey round). To summarize these 301 linear relative Engel curves, we take the mean of the absolute value of the slope within each sector. The first three sectors—comprising food, intoxicants and fuels for which we have good price data—have average absolute slopes of 0.011, 0.042 and 0.029 respectively in the 1987–1988 survey round. In contrast, slopes are much steeper for durables and services with average absolute slopes of 0.081 and 0.084, respectively (the remaining sector, non-durable manufactures comprised of clothing, footwear and personal care items, has an average slope of 0.016). We draw similar conclusions from the 1999–2000 round. Appendix Table D.1 reports these results. Thus, given the much more dramatic differences in what products the rich and poor buy within durable manufacturing and services, there is also much greater scope for new and improved products in these sectors as well as differential price changes in these sectors to lead to diverging inflation rates.


We argue that new product entry in manufactures and services spurred by the major economic reforms enacted in the 1990s is likely to be an important channel, with these new products targeted disproportionately at rich consumers (as documented for the US by Jaravel, 2019). While very much imperfect, we turn to the Prowess data to shine additional light on product entry in this period. These data have been used in other work to document product entry in response to trade reforms (e.g., Goldberg, Khandelwal, Pavcnik and Toplaova 2010) and we follow their methodology of defining a firm-product combination as a unique product. In the first year of the database, 1989, there are 9571 firm-product pairs, rising to 36,026 in 2000. Comparing the growth in the number of firm-products across sectors, there was dramatically more entry in the service sector (with a 12.8-fold increase in products compared to a 2.3–3.5 fold increase in other

Table D.1: Differences in Average Slopes of Relative Engel Curves Across Sectors

Sector	Average Engel Slope	
	1987/1988	1999/2000
Food	0.011	0.011
Intoxicants	0.042	0.048
Fuel	0.029	0.038
Non-durables	0.016	0.011
Durables	0.081	0.088
Services	0.084	0.055

Notes: Table presents simple averages of absolute slopes or relative Engel curves by Sector for 301 products in the 1987/1988 and 1999/2000 NSS survey rounds. Linear relative Engel curves calculated by regressing expenditure shares within a sector on log per capita household expenditures for the whole of rural India (using household survey weights).

Figure D.1: Product Growth By Sector Using Prowess Data



Notes: Figure presents the growth in products, defined as firm-product pairs, by sector using the ProwessDX dataset. For comparison purposes, sectors indexed to 100 in 1989, the first year of the sample.

sectors). Appendix Figure D.1 presents these results. This is, of course, very much suggestive since the number of firms included in Prowess increased over time and the incentives to produce annual reports that Prowess draws on may have changed over time.

Taken together, these three pieces of evidence suggest that price drops, new varieties, and quality increases in manufactures and services may have generated substantially lower inflation for the rich over this period.


E A Monte Carlo Simulation of our Empirical Methodology

In order to explore whether our empirical methodology uncovers income-specific price indices in the presence of realistic statistical error, we perform the following simulation. First, we start with the estimated relative Engel curves in the initial cross-section of our sample households (1987–88 round). Using the estimated shapes of those curves we construct fictitious curves for the later (1999–2000) round under the assumption that inflation is a step function with the highest inflation for the poorest decile (equal to 170 percent, close to what we estimate in our application) falling in increments of 5 percentage points for richer deciles so that the top decile has inflation equal to 130 percent—and that relative prices within foods and fuels are unchanged. We then create a household-level dataset containing the same counts and total expenditure distributions of the households that we observe in the data of the 1999–2000 round, but assigning them the expenditure shares from the relevant point along these fictitious relative Engel curves.

To add realistic noise to the process, we add an error term to each household’s product-specific expenditure share. We draw these errors from a normal distribution using the standard error at the relevant percentile of the estimated relative Engel curves by product and market from the 1987–88 surveys. We then rescale all expenditure shares for each household to ensure shares sum to one within the three G groups.

With this fictitious 1999–2000 household-level expenditure share dataset, we then estimate relative Engel curves and run our full procedure (smoothing the curves, checking for non-monotonicity, calculating horizontal shifts, taking medians over all admissible products etc.). Appendix Figure E.1 reports the resulting estimates from this procedure when run 250 times, drawing new error terms each time. Absent adding additional error to the fictitious household expenditures, we recover exactly the true inflation rates (the solid bars). The truth lies within the 95th percent envelope of estimates across the 250 simulations for all deciles, although at their mean or medians we obtain a slightly shallower slope with respect to income given the attenuation due to the addition of measurement error. This simulation suggests that the pro-rich inflation bias that we uncover may be an underestimate.

Figure E.1: Monte Carlo Simulation Estimates



Notes: Figure plots AFG NPC P^1 Price Index estimates using simulated data where inflation is equal to 170 percent for the poorest decile, falling in increments of 5 percent for richer deciles with the top decile facing 130 percent inflation. Solid bars report the estimates absent adding additional noise to the fictitious 55th round household expenditures. Top 97.5th percentile and bottom 2.5th percentile value of the 250 Monte Carlo simulations for each decile reported as confidence intervals alongside median and mean.

F Revisiting the Impacts of India's 1991 Trade Reforms

In this appendix, we revisit the impact of India's 1991 trade reforms on the welfare of rural households in India. We closely follow Topalova's (2010) analysis that pioneered the (now widespread) use of a shift-share instrument to identify the impacts of trade shocks.

Like her, we explore changes in rural districts across the 1987/88 and 1999/2000 NSS rounds. We focus on her most stringent specification that regresses poverty head count ratios (the dependent variable, using the Deaton, 2003a, recall bias correction discussed in Section 4.1) on district-level exposure to import tariff cuts (the independent variable). Exposure is measured as the weighted average tariff cut, with weights proportional to the initial-period sectoral employment shares in the district. She also includes district fixed effects, time fixed effects, and several time-varying district controls.⁵ To instrument the potentially-endogenous shift-share tariff exposure measure she uses both the same shift share measure but calculated only using traded industries (to deal with omitted variables correlated with initial shares of employment in traded sectors across districts) as well as a variant using the initial average level of import tariffs rather than the change (as all tariffs were brought to similar levels post reform, initially higher tariffs fell more for predetermined reasons).

We revisit this regression but replace the outcome (district-level rural poverty rates) with our welfare estimates.⁶ For expositional purposes, we focus on the log of our equivalent variation welfare metric.⁷ Again, we focus on the no price correction approach although results are insensitive to this choice. Importantly, our method allows us to calculate impacts at each decile of the local income distribution. The right panel of Appendix Figure F.1 plots the decile-specific coefficients on the tariff exposure variable (i.e. the difference in welfare growth for more exposed regions compared to less exposed). For comparison, the left panel plots estimates for the same specification but replacing the dependent variable with log total outlays per capita.

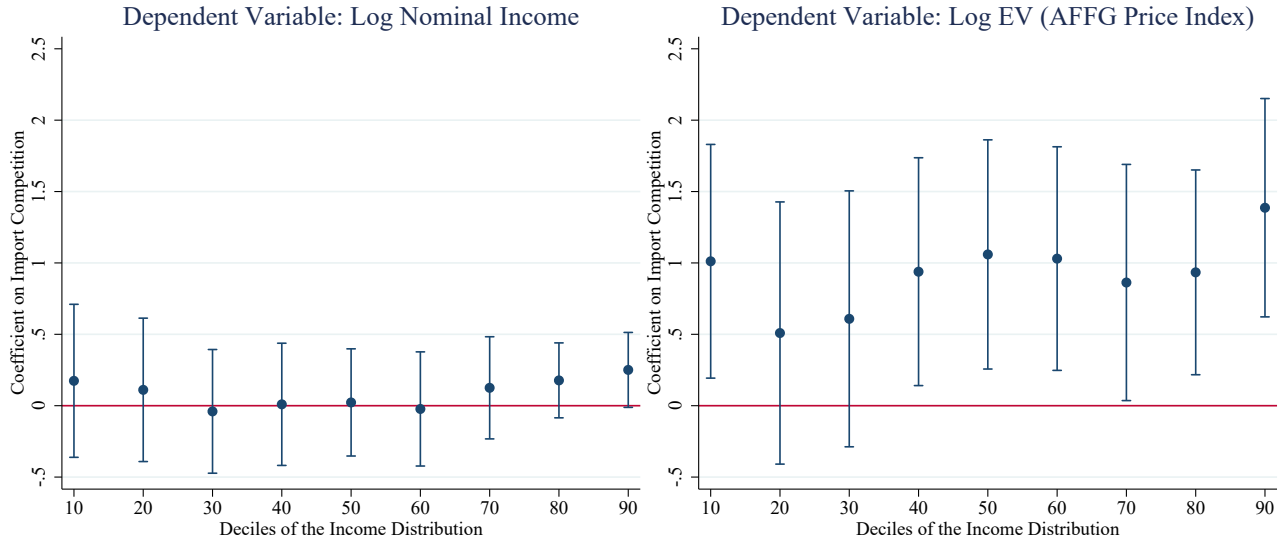
Two main findings emerge. First, while existing work has focused on the effect on poverty rates, our estimation reveals that the adverse effects of import competition on local labor markets are borne by households across the income distribution, including by rural households in the richest income deciles. Second, we find that the adverse effects on nominal outcomes are amplified when taking into account the effects on household price indices. Import competition leads to relatively higher local price inflation, particularly for richer households. This somewhat surprising finding is not simply an artifact of our approach, as it also appears when calculating a simple Laspeyres index using the raw price data for food and fuels (see Appendix Figure F.2). One potential explanation is that hard hit areas did not experience the same increases in retail-sector competition or productivity as faster-growing areas. An alternative explanation, and one beyond the scope of this paper, is that the shift-share exclusion restriction is violated.

⁵This specification corresponds to column 8 in Table 3a) of Topalova (2010).

⁶We obtain similar results using Topalova's other specifications although, as in Topalova, results are less significant. Note that Topalova does not restrict attention to markets with more than 100 survey households. Restricting Topalova's sample in this way makes her effect sizes larger.

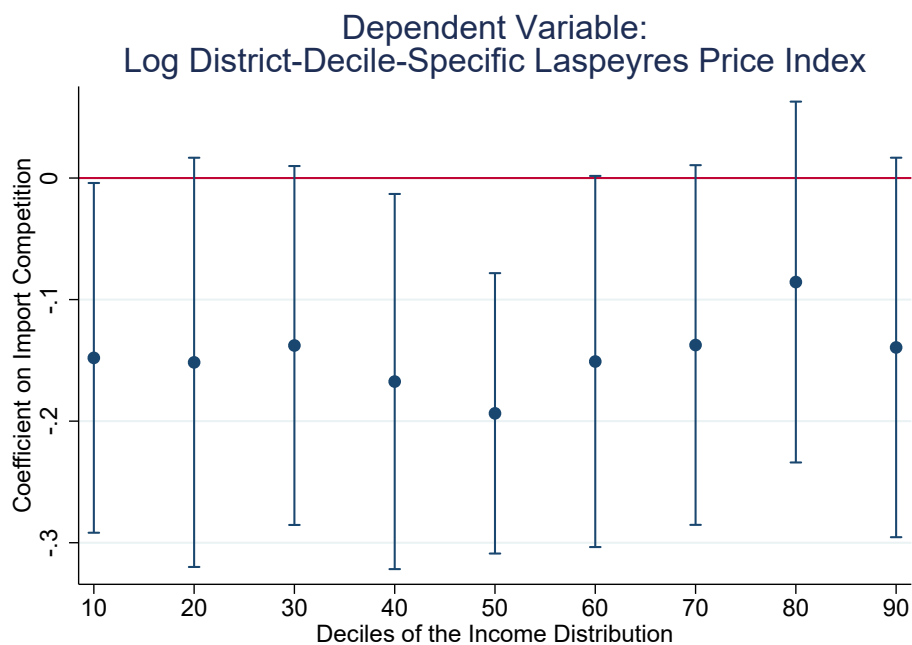
⁷As discussed in the main text, we have more overlapping Engel curves and so less noise when calculating P^1 and EV.

Figure E.1: Effect of Import Competition on Rural Nominal Income and Welfare



Notes: The left panel depicts IV point estimates of the effect of import competition on log total outlays per capita, estimated separately for each decile of the local per-capita outlay distribution. The IV regression specification follows column 8 in Table 3a) of Topalova (2010). Specifically, exposure to import competition is measured by the weighted average tariff cut, with weights proportional to the initial sectoral employment shares in the district. There are two instruments: first the same shift-share measure but calculated only using tradable industries, second this tradable shift-share but using the initial average level of import tariffs rather than the change. Regressions also include district fixed effects, time fixed effects, and additional time-varying district controls. The right panel depicts estimates from identical specifications with log welfare (measured by the log of equivalent variation using the AFG NPC price index) as the dependent variable. 95 percent confidence intervals based on standard errors clustered at the state-by-survey-round level (as in Topalova). See Section F for further discussion.

Figure F2: Effect of Import Competition on Laspeyres Price Index (Only Using Reliable Price Data)



Notes: Figure depicts IV point estimates of the effect of import competition on the log of the district-decile-specific Laspeyres price index, estimated separately for each decile of the local per-capita outlay distribution. The regression specification is identical to that described in Appendix Figure F1 and Section F, but with the log of the Laspeyres price index changes for food and fuels as the dependent variable (instead of log total outlays per capita or welfare). Laspeyres price indices calculated using district-by-decile-specific budget shares. Positive point estimates indicate negative effects of import tariff cuts. See Section F for discussion.