

NBER WORKING PAPER SERIES

MEASURING WELFARE AND INEQUALITY WITH INCOMPLETE PRICE INFORMATION

David Atkin
Benjamin Faber
Thibault Fally
Marco Gonzalez-Navarro

Working Paper 26890
<http://www.nber.org/papers/w26890>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2020, Revised December 2020

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, Yulu Tang, and Guan Yi Lin provided excellent research assistance. This paper subsumes an earlier version titled “A New Engel on Price Index and Welfare Estimation”. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by David Atkin, Benjamin Faber, Thibault Fally, and Marco Gonzalez-Navarro. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Measuring Welfare and Inequality with Incomplete Price Information
David Atkin, Benjamin Faber, Thibault Fally, and Marco Gonzalez-Navarro
NBER Working Paper No. 26890
March 2020, Revised December 2020
JEL No. D12,E31,F63,O12

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 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 implement this approach to estimate inflation and welfare changes in rural India between 1987 and 2000, and to revisit the impacts of India's trade reforms. Our estimates reveal that lower rates of inflation for the rich erased the real income convergence documented by the existing literature that uses the subset of consumption with well-measured prices to calculate inflation.

David Atkin
MIT Department of Economics
The Morris and Sophie Chang
Building, E52-550
50 Memorial Drive
Cambridge, MA 02142
and NBER
atkin@mit.edu

Benjamin Faber
Department of Economics
University of California, Berkeley
697A Evans Hall
Berkeley, CA 94720
and NBER
benfaber@econ.berkeley.edu

Thibault Fally
Department of Agricultural and Resource Economics
University of California, Berkeley
University Hall, 7th Floor
Berkeley, CA 94720-3310
and NBER
fally@berkeley.edu

Marco Gonzalez-Navarro
Agricultural and Resource Economics
UC-Berkeley
207 Giannini Hall
Berkeley, CA 94720-3310
and University of California, Berkeley
marcog@berkeley.edu

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 (2019) and Comin et al. (forthcoming); or by allowing income-groups to have different taste parameters as in Atkin et al. (2018), Jaravel (2019), and Argente and Lee (2020).¹ 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 consumption in rural areas—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 still requires strong functional form assumptions (e.g. quality-adjusting prices or accounting for gains from variety).

In this paper, we instead propose and implement a new approach that uses rich, but widely available, expenditure survey microdata—and in particular does not require observing reliable price data for all consumption 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 implement this approach to quantify changes in household welfare for Indian districts over time, and to revisit the impacts of India’s 1991 trade reforms.

Our analysis proceeds in three steps. First, we develop the theory behind our approach. Our starting point is an environment with information about household expenditures on different goods and services as well as total (nominal) household outlays at different points in time (or potentially in different locations to measure welfare differences rather than changes).² Addi-

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

²For readability we also refer to total household outlays as income. The 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.

tionally, the researcher only observes well-measured price changes for a subset of household consumption (e.g. foodstuffs). This environment is typical of settings where researchers have access to household expenditure surveys. Short of assuming particular realizations for unobserved or poorly-measured prices, recovering changes in the full price index, and hence welfare, is challenging in this environment and clearly impossible without restrictions on preferences.

The cornerstone of the paper’s methodology is the insight that quasi-separable preferences, as defined by Gorman (1970; 1976), provide a natural (and testable) restriction that allows us to estimate income-specific welfare changes in the absence of complete price information.³ Quasi-separability requires that subsets of goods or services are separable in the expenditure function (not the utility function, hence the term “quasi”). This implies that relative outlays across goods i within some subset G of consumption are an (arbitrary) function of within- G relative prices and the level of household utility. Thus, rich and poor households facing the same prices may differ in their relative budget shares within group G . And since quasi-separable demand systems can be of any rank in the terminology of Lewbel (1991), they can accommodate arbitrarily non-linear patterns of non-homotheticity found in the data. Prices outside of G may affect total outlays on group G in a fully flexible manner, and they may also affect relative outlays within G , but this latter effect only operates through changes in household welfare (i.e. by increasing or decreasing the cost of living). This last property is key. As we lay out below, if we wish to make the fullest use of the available price data—and hence exploit changes in relative outlays within product groups where prices are well measured—then quasi-separability provides the minimal (i.e. necessary and sufficient) restriction on preferences to infer changes in welfare at any point of the income distribution.

We state our approach formally in one lemma and two propositions. To show the central role of quasi-separability, we start by holding relative prices within group G fixed. Lemma 1 proves that if, and only if, preferences are quasi-separable in group G , we can recover changes in price indices at any point of the income distribution from horizontal shifts across time periods in what we term “relative Engel curves”—i.e. relative expenditure shares within G plotted against log total outlays per capita. Intuitively, the horizontal distance between these curves at any point in the income distribution reveals the change in nominal outlays that holds relative expenditure shares within G fixed, and hence maintains the same level of utility (given that within- G relative prices do not change and prices outside G affect relative shares within G only through utility). It is then straightforward to recover changes in welfare for any household from the distance in outlays between period 0 and 1 relative expenditure shares, either traveling along period 0’s relative Engel curve (to recover the equivalent variation, EV) or period 1’s curve

³Deaton 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., forthcoming), several variants of PIGL, PIGLOG and Translog preferences (Deaton and Muellbauer, 1980), and a general class of Gorman preferences discussed in Fally (2018).

(to recover the compensating variation, CV). A similar exercise across locations rather than time reveals welfare differences across space.

Propositions 1 and 2 then relax the untenable assumption that relative prices remain fixed within a quasi-separable subset of goods G . Our main Proposition 1 applies if reliable price data are available for goods inside a G group—such as basic foodstuffs and fuels in the Indian setting or parts of supermarket retail in the US setting. In this case, we show that we can correct our welfare estimates, either to the first order or exactly under any specific demand structure within G , to remove potential bias due to confounding within- G relative price changes. Put another way, we obtain unbiased estimates of the full price index that covers *all* household consumption using only relative expenditure shares and price changes for subsets of goods for which we have reliable price data, without requiring restrictions on unobserved price changes outside of G . We argue above that in most settings it is not possible to obtain reliable price data for large swaths of the service and manufacturing sectors, in part because of difficulties valuing 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 these settings.⁴

Proposition 2 explores the special case where the researcher has no reliable price information at all. We prove that we can still recover price indices, EV and CV at each point of the income distribution by averaging multiple estimates calculated using different goods within G as long as an orthogonality condition holds across these goods: that changes in within- G relative prices are unrelated to local slopes of relative Engel curves. This result provides us with the implicit identifying assumption required in the poorest data environments where there is no reliable price information for any part of consumption.

In the second step, we form a bridge between the theoretical results and the empirical implementation. Our estimation approach follows directly from our propositions 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 horizontal shifts calculated.

A natural question in taking our approach to the data is how plausible are the assumptions behind our propositions, 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 unique and unbiased identification: i) on the invertibility of Engel

⁴Since price changes outside of G are unrestricted, we can accommodate changes in quality and variety, and hence quality- and variety-adjusted price changes, outside of G . We can also accommodate unmeasured quality or variety changes inside G by correcting the prices we use for our within- G price correction using now-standard variety and quality corrections we discuss in Section 3.2.3.

curves, ii) on implementing the within- G price corrections, iii) on aggregating up to good-level data in settings where barcode-level data are available, iv) on sample selection, and v) on preference heterogeneity across households and over time.

In the final step, we implement our methodology in two applications. First, we draw on Indian household expenditure survey microdata to quantify changes in rural welfare between 1987/88 and 1999/2000 at different points of the income distribution, and for every district in India.⁵ 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 Indian household surveys with both reliable 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 in the sample, it is reassuring that our estimates of the full price index for these households are very similar to Deaton’s estimates of what is essentially a food and fuel price index (despite coming to this conclusion in very different ways—we exploit shifts in relative Engel curves while Deaton uses observed price changes).

Looking across the income distribution, our estimates bring to light 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 a Quadratic AIDS demand system and goods with observable price data as in Almås and Kjelsrud (2017).⁶ 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 that this convergence entirely disappears once we account for the differential inflation across income groups revealed by our approach.

The most likely explanation for these findings is that higher-income Indian households disproportionately benefited from lower inflation in product 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,

⁵We 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.

⁶Almås and Kjelsrud (2017) use the same NSS expenditure data but include two categories with poorly measured prices (clothing; bedding and footwear). Additionally, as the 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 the addition of new varieties).

2015). 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, 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 (see Deaton and Kozel, 2005 for an overview).⁷ As long as the additional recall period did not change *relative* budget shares within a given food product group 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 the Great Indian Poverty Debate.

In the second application, we use our method to revisit Topalova's (2010) analysis of the local labor market impacts of India's 1991 trade reforms. Topalova's main finding is that rural poverty rates (the fraction of households below the poverty line) increased relatively more in districts where import competition rose most. While Topalova highlights effects on poverty rates, our approach uncovers adverse effects of import competition across the full income distribution, including among the highest income households.

In addition to the literatures mentioned above, our use of Engel-like relationships 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 in the economy and so is not suitable for distributional analysis. As shown in Almås et al. (2018), calculating income-specific price index changes under the existing Engel methodology that uses AIDS preferences 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.⁸

⁷Deaton (2003a) calculates poverty by adjusting food expenditure using the mapping between food and fuels expenditure (which had no recall period added) from earlier rounds. The method implicitly assumes that relative prices of food and fuels did not change. Tarozzi (2007) explores a related approach using multiple auxiliary variables.

⁸Ligon (2019) recovers the marginal utility of income ("neediness") using an approach based on constant-

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). Several of these papers calculate income-group specific price indices and are cited above. Crawford and Neary (2019) extend this approach to the product characteristic space. Finally, Redding and Weinstein (2020) show how to use CES preferences to account for changes in product quality for categories of goods where prices are observed. As we discuss further in Section 3.2.3, these recent advances complement our Proposition 1 by providing estimates of the variety and quality-adjusted prices needed to correct for within- G relative price changes when products contain multiple varieties.

We structure the paper as follows. Section 2 develops the theory. Section 3 presents our estimation approach and derives corollaries for unique and unbiased identification. Section 4 describes the data and applies our methodology in two applications. Section 5 concludes.

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 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 two main Propositions.

2.1 Data Environment

Our starting point is an environment with information on total (nominal) household outlays per-capita, y_h , for different households h coupled with their expenditure shares, $\frac{x_{hi}}{y_h}$, across the goods and services that they consume (for readability we will refer to them simply as goods)—where x_{hi} is per-capita household expenditure on good $i \in I$. In our later propositions, we will distinguish between settings where well-measured (i.e. quality and variety-adjusted) prices p_i are available for some subset of goods G but not necessarily for the remaining goods NG , and settings where no well-measured prices are available. These environments correspond to widely available expenditure survey data.

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.⁹ 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. Isomorphic results would hold across space if we replaced time periods by locations.

elasticity Frisch demand. These preferences restrict direct cross-price effects from unobserved price changes to be zero. To recover money-metric welfare measures, such as EV and CV, requires knowledge of all price changes.

⁹If household panel data are available, we can infer price index changes for individual households.

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 discussed above, existing approaches to estimating income-specific price index changes under non-homothetic preferences in principle require knowledge of (quality- and variety adjusted) price changes for all products and services entering household consumption. To have a hope of recovering welfare changes, we require that some function of the vector of observed changes in household expenditures only depends on changes in observed prices (e.g. for foodstuffs) and household utility. The fact that prices are typically missing for entire expenditure categories naturally leads us to focus on exploiting changes in expenditure shares within product groups—and thus be able to maximally leverage the price data that are available to accommodate arbitrary substitution patterns within groups.¹⁰ 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:*

$$e(p, U_h) = \tilde{e}(e_G(p_G, U_h), p_{NG}, U_h) \quad (1)$$

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

Quasi-separability is separability in the expenditure function (rather than the utility function). It 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 in 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).

To see this more clearly, 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 (see Lemma 2 in Appendix B). For example, the preferences used in Comin et al. (forthcoming) 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

¹⁰For example, if instead we want to exploit changes in expenditure shares across product groups, we require knowledge of impossible-to-estimate elasticities of substitution between groups with and without observed prices.

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 systems that allow for rich substitution effects within G (captured by function F_G) and between goods within and outside G (function K).¹¹

Second, we define what we term “relative Engel curves” as follows.

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

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. $V(p^0, y_h^0) = V(p^1, \frac{y_h^0}{P^0(y_h^0)})$).¹²

These two price indices are intimately 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 a given income level in period 0, $\log P^0(y_h^0)$, is equal to the horizontal shift in the relative Engel curve of any good $i \in G$ at that income level, such that*

¹¹For example, we do not need to impose that (Allen-Uzawa) price elasticities are constant across goods within G as in Comin et al. (forthcoming).

¹²The two price indices are closely related: $y_h^1 = y_h^0 / P^0(y_h^0)$ implies $y_h^0 = y_h^1 / P^1(y_h^1)$.

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

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

Lemma 1 states that we can infer changes in exact price indices and welfare at any given point of the initial or final income distribution by observing: i) relative expenditure shares across some subset G of goods, and ii) total outlays. Appendix C provides the proofs.

To aid intuition, Figure 1 graphically illustrates Lemma 1. 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 as non-homotheticity is the only factor driving changes in relative outlays.¹³ 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-parallel shifts in relative Engel curves. The key role of quasi-separability is to ensure that these outside- G price changes only affect within- G relative expenditures through changing utility. Thus, shifts in relative Engel curves reveal changes in the price index at any point of the income distribution.

To make these statements precise, we lay out several steps of the proof of Lemma 1. Under quasi-separability, relative expenditure on good i within group G can be written as a compensated function $H_{iG}(p_G, U)$ of utility and within- G relative prices alone:

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

using the homogeneous scalar function $e_G(p_G, U)$ introduced in equation (A.7).

¹³Here we abstract from preference (taste) heterogeneity but discuss this possibility in detail in Section 3.2.5.

How do horizontal shifts in relative Engel curves (across $\log y_h$ space) identify changes in price indices? To obtain $P^1(p^0, p^1, y_h^1)$, start with the period 1 relative budget share on the relative Engel curve in period 1:

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

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 .¹⁴ 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 . As we show in Lemma 2 of Appendix C, quasi-separability is required to express within-group expenditure shares as a function of utility and within-group relative prices (i.e. to ensure the existence of the function $H_{iG}(p_G, U)$ above).¹⁵ 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 standard Engel curves (i.e. expenditure shares plotted against \log total outlays). In Lemma 4 in Appendix B, 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 standard Engel curves do recover price indices as we show in Lemma 3). These results connect to Almås et al. (2018) who show that the standard 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.

¹⁴Note that H_{iG} is homogeneous of degree zero in prices p_G^t within group G .

¹⁵Lemma 2 draws on existing results (Blackorby et al. 1978), but provides a more direct proof.

2.3 Recovering Income-Specific Welfare Changes From Expenditure Survey Data

Our two main propositions relax the assumption of constant relative prices within subset G in the context of two different data environments.

Environments With Partial Price Information

When price information is available for goods i within a quasi-separable set G , but not for goods outside that set, our first proposition shows that information on the within- G price changes alone is sufficient to fully adjust relative Engel curves to account for any confounding response to changing prices, holding utility constant. We can then infer changes in price indices from horizontal shifts in these adjusted curves. This proposition is valuable in typical expenditure survey contexts 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.

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 the adjusted relative Engel curve of any good $i \in G$ at that income level, such that*

$$E_{iG}(p^0, \frac{y_h^1}{P^1(y_h^1)}) = E_{iG}(p^1, y_h^1) \times \frac{H_{iG}(p_G^0, U_h^1)}{H_{iG}(p_G^1, U_h^1)} \quad (2)$$

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) *EV for a given income level is revealed by the horizontal distance along period 0's relative Engel curve between new and old expenditure shares, such that*

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

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

This proposition describes how to adjust relative Engel curves to account for confounding vertical shifts due to changes in within- G relative prices.¹⁶ Specifically, we can recover $\log P^1(y_h^1)$ from the horizontal difference in relative Engel curves from the period 1 expenditure share after 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 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.

¹⁶Note 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 $E_{iG}(p^1, y_h^1) = H_{iG}(p_G^1, U_h^1)$.

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 any information on the structure of preferences or prices for goods outside G . As long as there is a group G of goods for which preferences are quasi-separable and reliable price data are available, we can uncover changes in price indices and welfare.¹⁷ As described in Section 3.1.2, we implement the price adjustment in Proposition 1 in two different ways: as a first order approximation, evaluating elasticities in the base period, and in its exact form after specifying within-group demands H_{iG} .

Environments With No Price Information

In data-poor environments where no price data are available, Proposition 2 shows how to exploit the multiplicity of price index estimates—one for each of the goods $i \in G$ —to recover unbiased estimates of price index and welfare changes, subject to an orthogonality condition on relative price changes within G . This proposition is most relevant for settings where only expenditure and not quantity data are collected, no separate price survey data are available, and where knowledge of likely price shocks allows for partial tests of the orthogonality condition.

We first write equation (2) in logs and take a first-order approximation of changes in $\log H_{iG}$ due to relative price changes, holding utility fixed.¹⁸ Subsequently inverting the relative Engel curve at the period 1 expenditure share, for any good $i \in G$ we obtain:¹⁹

$$\log(y_h^1) - \log E_{iG}^{-1}\left(p^0, \frac{x_{ih}^1}{x_{Gh}^1}\right) \approx \log(P^1) + (\beta_{ih}^0)^{-1} \log \frac{H_{iG}(p_G^0, U_h^1)}{H_{iG}(p_G^1, U_h^1)} \quad (3)$$

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 the initial set of prices p^0 and $\frac{H_{iG}(p_G^0, U_h^1)}{H_{iG}(p_G^1, U_h^1)}$ is the price-induced vertical shifts defined as in Proposition 1 above.²⁰ The first term on the right-hand side of (3) 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 .

While any individual estimate can be biased, these biases may cancel out if we average over multiple estimates of $\log P^1$ derived from different goods $i \in G$. Proposition 2 states this orthogonality condition formally:

¹⁷To 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 C): $\frac{\partial \log H_{iG}}{\partial \log p_j} = \frac{\partial \log E_{iG}}{\partial \log p_j} + E_{jG} \frac{x_G}{y} \frac{\partial \log E_{iG}}{\partial \log y}$. Estimating these terms only requires information on household total outlays, expenditures on goods within group G , and prices among these goods.

¹⁸I.e. assuming that the vertical shifts in relative Engel curves due to within- G relative price changes are proportional to those price changes.

¹⁹Symmetrically for P^0 : $\log(y_h^0) - \log E_{iG}^{-1}\left(p^1, \frac{x_{ih}^0}{x_{Gh}^0}\right) \approx \log(P^0) + (\beta_{ih}^1)^{-1} \log \frac{H_{iG}(p_G^1, U_h^0)}{H_{iG}(p_G^0, U_h^0)}$.

²⁰Using the compensated cross-price elasticities $\sigma_{ijh} = \frac{\partial \log H_{iG}}{\partial \log p_j}$ of relative expenditures, the vertical shift can be rewritten as a function of relative price changes: $\log \frac{H_{iG}(p_G^1, U_h^1)}{H_{iG}(p_G^0, U_h^1)} \approx \sum_{j \in G} \sigma_{ijh} \Delta \log p_j$ with $\sum_{j \in G} \sigma_{ijh} = 0$.

Proposition 2. *Assume that, to the first order, vertical shifts in relative Engel curves due to within- G relative price changes are orthogonal to their slopes: $\frac{1}{G} \sum_{i \in G} (\beta_{ih}^0)^{-1} \log \frac{H_{iG}(p_G^1, U_h^1)}{H_{iG}(p_G^0, U_h^1)} = 0$. 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, $\log P^1(y_h^1)$, corresponds to the average horizontal shift in the relative Engel curves of goods $i \in G$ at that income level, calculated as in Lemma 1.*
- ii) EV for a given income level is revealed by the average horizontal distance along period 0's relative Engel curves of goods $i \in G$ between new and old expenditure shares calculated as in Lemma 1.*

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

Thus, if relative prices are changing within G but are unobserved, we can still recover unbiased estimates of price indices and welfare as long as the within- G price effects are not systematically related to the inverse of the local (i.e. income-specific) slopes of relative Engel curves. To illustrate this proposition with an example, imagine G includes both luxuries and necessities (in terms of relative outlays within G). If the relative price of the luxuries rise, households would purchase relatively more necessities and we would falsely infer that they became poorer on average. However, if some luxuries within G rose in price and others fell, averaging over estimates would recover the true welfare change.

Such a condition is informative in the most challenging contexts where no reliable price data are available for any product group. Absent relative price changes systematically related to within- G income elasticities, our methodology still provides unbiased estimates. For example, if we are interested in the impacts of shocks or policies that differ by product, such as tariff changes, our price index estimates are likely to be unbiased if the tariff changes are unrelated to relative Engel slopes (a condition that is testable even if reliable price data are not available).

3 From Theory to Estimation

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 microdata. We then turn to identification and derive corollaries to our theoretical propositions describing testable conditions for unique and unbiased identification. We draw on these results in our applications to perform a number of validation exercises and robustness checks.

3.1 Estimation Approach

Suppose that we want to estimate the welfare change between two periods for specific percentile of the household income distribution. The graphical exposition of Lemma 1 in Figure

1 provides a simple estimation approach. First, we use non-parametric methods to estimate flexible relative Engel curves separately for both periods and for each location (what we will call a market). We can then recover changes in exact income-specific price indices as well as household welfare from the horizontal shift in these curves at different points of the income distribution, potentially adjusting curves or adding a correction term to account for within- G relative price changes. Repeating this procedure for multiple goods generates many estimates that can be combined to increase precision (and potentially accommodate the misclassification of goods into quasi-separable sets or good-specific taste shocks as we discuss below).

3.1.1 Environments With No Price Information

Since our price correction approaches build on it, we first describe our approach in the absence of well-measured price data, even within group G (i.e. assuming that the orthogonality condition in Proposition 2 holds).

No Price Correction Approach

The first step is to use expenditure survey microdata to estimate kernel-weighted local polynomial regressions of relative expenditure shares, $x_{ih'm}^t/x_{Gh'm}^t$, on log total outlays per capita, $\log y_{h'}^t$, for every good $i \in G$, period t and market m , where h' indexes the individual households in the expenditure surveys. This provides estimates of x_{ihm}^t/x_{Ghm}^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.²¹

With these relative Engel curves in hand, 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 for period 1 provides a point estimate of relative expenditures for households at this percentile of the initial income distribution, x_{ihm}^1/x_{Ghm}^1 . The next step is to estimate the period 0 income level $\widehat{E}_{iG}^{-1}(p^0, x_{ihm}^1/x_{Ghm}^1)$ 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 $\widehat{x_{ihm}^0/x_{Ghm}^0}$ and take the corresponding income of this income percentile h , $\widehat{\log y_h^0}$.²² As we discuss in Section 3.2.1, we restrict attention to monotonic relative Engel curves to ensure this crossing point is unique.

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

²¹We 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 applications explore alternative bandwidth choices.

²²We take the two closest percentiles and linearly interpolate between them to obtain $\widehat{\log y_h^0}$.

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 y_h^0 is the actual period 0 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, $\log P^0(p^0, p^1, y_h^0)$, we follow the same procedure but going 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 . As discussed below (and previewed in Proposition 2), we combine these estimates by taking an average across different goods i at each percentile of the income distribution.²³

3.1.2 Environments With Partial Price Information

If price data are available within group G , Proposition 1 shows how to correct the observed horizontal shifts in relative Engel curves for confounding relative price effects. For this reason, we will focus on product groups with well-measured price information in our applications. Exact estimation under Proposition 1 requires knowledge of the shape of function $H_{iG}(p_G, U)$, for which we propose two functional forms below. First, we present a simple and transparent price correction term that holds to the first-order for any preference structure within G .

First-Order Price Correction Approach With information on price changes for $i \in G$, the orthogonality condition we derived in Proposition 2 can be calculated and tested. Furthermore, as we see from averaging equation (3) over $i \in G$, this covariance term provides a first-order approximation of the bias from confounding vertical shifts in relative Engel curves due to within- G relative price changes for any preference structure $H_{iG}(p_G, U)$.

Specifically, we add the slope-to-price-change correlation term in Proposition 2 to the the price index estimate derived from horizontal shifts in relative Engel curves described in Section 3.1.1 above. For example, if we assume a constant elasticity of substitution σ_G within group G , the (market-by-percentile-specific) bias correction term for $\log P^1$ takes the simple form:

$$\frac{1}{G} \sum_{i \in G} (\beta_{ih}^0)^{-1} (1 - \sigma_G) (\Delta \log p_i - \overline{\Delta \log p_G}). \quad (4)$$

The bias and hence correction term is small if relative price changes are close to uncorrelated with slopes of relative Engel curves, if within- G elasticities are small, or if within- G price changes are similar.

²³Ultimately, 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).

Exact Price Correction Approach For exact correction terms, recall from Proposition 1 that we must adjust either the period 0 or period 1 relative Engel curves to account for within- G relative price changes and then calculate horizontal shifts using the adjusted curve. Thus, we proceed as in Section 3.1.1 above, but modifying the appropriate relative Engel curve before calculating horizontal differences.

We propose two practical specifications that only require estimating a single elasticity parameter. The first is to specify a constant (compensated) elasticity of substitution between goods within group G , with an expenditure function that satisfies:²⁴

$$e(p, U_h) = \tilde{e}\left(\sum_{j \in G} A_j(U) p_j^{1-\sigma_G}, p_{NG}, U_h\right) \quad (5)$$

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. (forthcoming) by allowing for flexible substitution patterns outside of group G . The confounding vertical shifts of relative Engel curves due to within- G relative price changes that need to be adjusted for are:

$$\frac{H_{iG}(p_G^1, U_h^1)}{H_{iG}(p_G^0, U_h^1)} = e^{(1-\sigma_G)[\Delta \log p_i - \overline{\Delta \log p_G}]} \quad (6)$$

where $\overline{\Delta \log p_G} = \log \left[\sum_{j \in G} (p_j^1/p_j^0)^{1-\sigma_G} (x_{jh}^1/x_{Gh}^1) \right]^{\frac{1}{1-\sigma_G}}$ 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.

Alternatively, recall from footnote 16 that the correction term in Proposition 1 can also be 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:

$$H_{iG}(p_G^1, U_h^1) - H_{iG}(p_G^0, U_h^1) = -\xi_G \times [\Delta \log p_i - \overline{\Delta \log p_G}]. \quad (7)$$

3.2 Identification

In this subsection, we derive a number of corollaries and results related to unique and unbiased identification when taking Propositions 1 and 2 to the data.

3.2.1 Invertibility of Relative Engel Curves

The first result derives necessary and sufficient conditions under which relative Engel curve functions are invertible, and hence our price indices are identified.

Corollary. *Under the same conditions as Propositions 1 and 2:*

²⁴The corresponding utility function can be implicitly defined as: $K \left(\sum_{j \in G} A_j(U)^{1/\sigma_G} q_j^{\sigma_G/(\sigma_G-1)}, q_{NG}, U_h \right) = 1$.

i) *The necessary condition to recover unique estimates of changes in price indices and welfare is that different levels of household utility map into unique vectors of relative budget shares within the subset of goods $i \in G$ at any given set of prices.*

ii) *A sufficient condition for i) to hold is that the relative Engel curve $E_{iG}(p, y_h)$ is monotonic for at least one good $i \in G$.*

The first condition is weaker than the second. The practical advantage of the second is that it is readily verifiable in the data, and turns out to be true empirically for all markets and time periods we consider in our applications. For the estimation approach that we outline in Section 3.1 above, we restrict attention to good-market combinations where relative Engel curves are monotonic (ensuring that estimates of shifts are unique for each good-market combination).²⁵

3.2.2 Quasi-Separability and Misclassification

Bias from Violations of Quasi-Separability Although our main propositions assume preferences are quasi-separable in group G , violations of this assumption have to be systematically related to price elasticities and slopes of relative Engel curves to induce bias in our welfare estimates. 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 the 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 then directly affect within- G relative outlays (rather than only affect relative outlays through utility, given quasi-separability).

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:*

$$\frac{1}{G} \sum_{i \in G} \log E_{iG}^{-1} \left(p^0, \frac{x_{ih}^1}{x_{Gh}^1} \right) - \log \left(\frac{y_h^1}{P^1} \right) \approx \frac{1}{G} \sum_{i \in G} (\beta_{ih}^0)^{-1} \times \sum_{k \in NG} (\Delta \log p_k - \overline{\Delta \log p_G}) \frac{\partial \log H_{iG}}{\partial \log p_k}, \quad (8)$$

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

For correctly classified goods, $\frac{\partial \log H_{iG}}{\partial \log p_k} = 0$ and there is no bias.²⁶ If good $k' \in NG$ is actually a G good, $\frac{\partial \log H_{iG}}{\partial \log p_{k'}} \neq 0$ for some is . If good $i' \in G$ is actually a NG good, $\frac{\partial \log H_{i'G}}{\partial \log p_k} \neq 0$ for some ks .

Since we are averaging across multiple i estimates, these violations of quasi-separability only generate bias if the direction and magnitude of the confounding (compensated) cross-

²⁵Specifically, as non-parametrically estimated Engel curves are often noisy at the extreme tails where there are few households across large ranges of outlays, we restrict attention to good-market combinations where relative 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).

²⁶Equation (8) 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.

price effects from unobserved NG price changes $(\sum_{k \in NG} (\Delta \log p_k - \overline{\Delta \log p_G}) \frac{\partial \log H_{iG}}{\partial \log p_k})$ 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. For our applications, this result motivates both averaging over multiple i estimates and exploring the sensitivity of our estimates across alternative classifications of goods into quasi-separable nests G .

Testing for Quasi-Separability Using Lemma 2 from Appendix C, we provide a test for quasi-separability using expenditure survey data.

Corollary. *If, and only if, preferences are quasi-separable in group G , the price elasticity of the uncompensated expenditure share $x_{iG} \equiv \frac{x_i}{x_G}$ in the price of any good $j \notin G$ equals the slope of the relative Engel curve multiplied by good j 's overall budget share:*

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

This corollary can be tested by measuring the vertical shift in relative Engel curves induced by the change in the price of good j , $\left. \frac{\partial \log x_{iG}}{\partial \log p_j} \right|_y$, if reliable price data are available for some goods outside subset G . Alternatively, we can explore the horizontal shift induced by this price change which is equal to the ratio $\left. \frac{\partial \log x_{iG}}{\partial \log p_j} \right|_y / \left. \frac{\partial \log x_{iG}}{\partial \log y} \right|_y$. Under quasi-separability, this ratio coincides with the marginal effect of a good j price-change on the price indices P^0 and P^1 . This result generates a second and more practical test for quasi-separability:

Corollary. *The elasticity of the exact price index P^t , $t \in \{0, 1\}$ with respect to the price of any good $j \notin G$ equals the overall expenditure share of good j :*

$$\frac{\partial \log P^t}{\partial \log p_j^t} = \frac{p_j^t q_{jh}^t}{y_h^t}. \quad (9)$$

This equality is simply Shephard's Lemma applied to our price indices. Since we do not use prices from non- G goods to estimate our price indices, our estimation strategy does not guarantee that this equality holds. One additional benefit of this test is that it also serves as a smell test for our general approach. Recall that we are calculating a price index covering the full consumption bundle from relative expenditures within some group G . The test asks whether our estimated price index responds to price changes for goods outside group G as it must (and will only do so fully if quasi-separability holds).

3.2.3 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, pack-sizes or flavors), either because that is the level the data are reported at

or because specific varieties are not consumed widely enough given the number of households sampled. Fortunately, our lemmas and propositions can also be applied to aggregates of varieties rather than individual varieties, even if demands for those varieties are non-homothetic within g . For instance, there may be fancy packaged sea salts and simple table salt that are consumed in different proportions by the rich and the poor.

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:*

$$E_{g,G}(p^1, y_h^1) = E_{g,G}(p^0, \frac{y_h^1}{P^1(y_h^1)}) \quad \text{and} \quad E_{g,G}(p^0, y_h^0) = E_{g,G}(p^{t1}, \frac{y_h^0}{P^0(y_h^0)}).$$

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 within each subgroup g can be aggregated across the i 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 .²⁷

Several remarks are in order. First, note that these subgroup price indices can be non-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 the price correction approaches in Section 3.1.2; or assume that the mis-measured or omitted relative price changes satisfy an orthogonality condition similar to expression (8) above. Finally, a more practical consideration is that relative Engel curves for subgroup g may be strictly monotonic while consumption of specific varieties i is zero (and thus flat) for some locations, periods, and/or ranges of income.

Taken together, these aggregation results are particularly valuable when implementing our approach to estimate price indices and welfare from highly-disaggregated data that are available for some subset of consumption G —e.g. barcode-level retail scanner data.

²⁷For example, the price aggregates derived in Redding and Weinstein (2020) could be used for $P_g(p_g, U)$, assuming within- g preferences have their CES structure.

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 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 this 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 we can calculate horizontal shifts in relative Engel curves for) can potentially eliminate bias from relative price shocks within the G group (or taste shocks as we discuss below). However, in the presence of such shocks, averaging over the subset of goods for which there is overlap 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, if there is no true overlap when estimating P^0 for poor 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 poor period 1 households appear as poor as they were in period 0.

To address such sample selection concerns, we exploit the fact that we observe whether or not a given good has missing overlap at a given income percentile and if so, whether the estimate this good would provide is censored 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 do so, 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).²⁸ In the rare cases where the median is unobserved due to most estimates being censored, we can impose 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, symmetry alone proves sufficient to solve selection issues in our Indian application.

A different type of sample selection arises if we don't observe any goods for which relative

²⁸We rank estimates, placing unobserved estimates below the lowest or above the highest 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.

Engel curves overlap for a given percentile and market. In this case, we face a market-level sample selection issue when aggregating across markets. For example, there may be missing markets among poor percentiles for P^0 and rich percentiles for P^1 if real incomes grew. 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 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 and Taste Changes

Finally, we formally consider three concerns related to taste heterogeneity and taste changes.

Omitted Variable Bias in Engel Curve Estimation

The first issue is common to all Engel curve estimates: if taste differences are correlated with income, Engel curves (both standard and relative) will be biased. For example, more educated households may both value certain goods more and have higher incomes. This may bias our price index and welfare estimates, as changes in real income over time would not affect budget shares in the way our estimated relative Engel curves predict. This bias can be addressed either 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 applications).

Heterogeneous Price Index Changes

The second issue is that if tastes for goods differ across household types within a given income percentile in a way that correlates with relative price changes across goods, then price index and welfare changes for a given income percentile will differ by household type. In this case, we show that our method yields a weighted average change: $\frac{y_h^1}{\tilde{P}^1(y_h^1)}$, where $\tilde{P}^1(y_h^1) \approx \sum_z w_z^1(y_h^1) P_z^1(y_h^1)$ with weights given by the relative Engel slopes of household type z : $w_z \equiv \frac{\sum_i (\beta_{i,z}^1 / \bar{\beta}_i^1)}{\sum_{z'} \sum_i (\beta_{i,z'}^1 / \bar{\beta}_i^1)}$ (see Appendix C.7). In this scenario, if one is interested in 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.

Changes in Tastes Over Time

The third issue arises when household tastes change over time. Such taste changes are only 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 on unobserved relative price changes within G that underlies Proposition 2 and equation (4). Denoting taste shocks—i.e. shifts in within- G budget shares conditional on

prices and income—by $\Delta \log \alpha_{ih}$ and abstracting from relative price changes, we obtain:²⁹

$$\log \left(\frac{y_h^1}{P^1} \right) \approx \frac{1}{G} \sum_{i \in G} \log \widehat{E}_{iG}^{-1} \left(p^0, \frac{x_{ih}^1}{x_{ih}^0} \right) - \frac{1}{G} \sum_{i \in G} (\beta_{ih}^0)^{-1} \Delta \log \alpha_{ih} .$$

If taste shocks across i within subset G are orthogonal to the local slope of the 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. One scenario that may violate this condition is if household types have different tastes and there are compositional changes over time (e.g. increases in education). This concern will only be problematic if different household types have different price index changes over time, a condition that we can (and do) explicitly test by separately estimating and comparing price index changes for different household types.

4 Applications

In the final section, we apply our methodology to estimate changes in rural Indian welfare over time and to re-visit the welfare impacts of India’s 1991 trade reforms.

4.1 Data

Following the Great Indian Poverty Debate and Topalova (2010), we draw on rural households in two of India’s “thick” NSS survey rounds covering 1987/88 (43rd round) and 1999/2000 (55th round).³⁰ 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 above, 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 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

²⁹To ensure shares sum to unity within G , we assume that these taste shocks sum to zero. These shocks can be defined using shifters multiplying e_G in equation (1) or as price shifters as in Redding and Weinstein (2020).

³⁰As 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.

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.³¹ As we discuss above, this subset of goods with reliable price are particularly valuable for our estimation since they allow us to implement Proposition 1 and compute exact or first order price corrections if the orthogonality condition is violated, as well as to assess potential bias from violations of quasi-separability.

Finally, it is important to 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 reports. 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 we rely on relative consumption patterns within product groupings that we show is unaffected by the additional 7-day recall question.

³¹As 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 shows that our methodology accommodates scenarios like this, where the price vector may depend on utility.

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 3.2.3 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 (Appendix Table A.2 for descriptive statistics of price changes over time across these 34 goods).³² This aggregation dramatically reduces the share of empty market-by-product cells (from 50 percent to less than 15 percent as shown in Appendix Figure A.1), and moving to the next highest level of aggregation (8 goods) provides little additional benefit.³³

We divide these 34 aggregate g goods into three broader consumption groups: raw foodstuffs (e.g. rice, leafy vegetables), other food products (e.g. milk, edible oils) and fuels (e.g. firewood, kerosene) as shown in the first column of Table 1. 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 3.2.4.³⁴ 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 price index and welfare 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).³⁵ 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 faster for the poor than for the rich. Our non-homothetic price indices allow us to determine whether this nominal income convergence translated into con-

³²We use a Stone price index to aggregate the observed price changes for the 132 items i to 34 goods g (using survey-weighted mean initial expenditure shares across the $i \in g$ to compute 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.

³³Appendix Figure A.3 reports qualitatively similar inflation estimates using these alternate levels of aggregation.

³⁴In principle, comparing estimates obtained from different G groups provides an over-identification test (i.e. 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 tastes and prices 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 some goods do not have both monotonic and overlapping relative Engel curves), these conditions are unlikely to be satisfied without pooling the estimates.

³⁵For 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.

vergence in standards of living between the rich and the poor.

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).³⁶ The left panel presents estimates applying Proposition 2, the approach for environments with no price information detailed in Section 3.1.1—i.e. assuming the orthogonality condition holds and making no price correction for within- G relative price changes. As above, we plot population-weighted averages across districts by decile. The remaining two panels of Figure 3 apply the approach described in Proposition 1 and Section 3.1.2 for environments with partial price information—i.e. implementing first-order and exact price corrections calculated using 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. For the first-order price correction in the middle panel, 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. For the exact price correction in the right panel, we use the isoelastic correction (non-homothetic CES) in equation (6) 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. Put another way, recall that our first-order price correction term also serves as a test of our orthogonality condition. Thus, the fact the estimates change little 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, we focus our remaining analysis on the no price correction approach (labeled “AFFG NPC price index”). In all cases, we would 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 (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.³⁷ 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

³⁶As 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.

³⁷As 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.

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 an approach that only exploits 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 dincome-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 well-measured price changes for the full consumption basket to implement.³⁸ 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, 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.³⁹ 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. As touched upon in the introduction, lower inflation in these specific categories is consistent with the fact that the Indian trade

³⁸See footnote 6 for a detailed description of how Almås and Kjelsrud (2017) the ratio of the Indian food and non-food CPI to navigate the lack of well-measured price data for categories beyond food and fuels.

³⁹Averaging across rounds, the richest decile spent a third of their expenditure on these categories.

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).

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. India's 1999-2000 NSS added an additional 7-day recall period for expenditures on food products 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 A.1 shows that the choice of recall period did not affect relative consumption shares within our G groups.⁴⁰ 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).

4.3.1 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 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 market across 108 sensible splits of our g goods into plausibly quasi-separable groupings G .⁴¹ 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 plausible G groupings. Reassuringly, our baseline is close to the mean for every decile of the income distribution. In addition, the 2.5th–97.5th percentile ranges are reasonably tight—suggesting that the conditions under which misclassification bias is small (equation 8) are met in our setting.

Next, we present the second and simpler of the two tests of quasi-separability in Section

⁴⁰In addition, Appendix Figure A.4 shows similar patterns of pro-rich inflation between the 1987/88 and 1994/95 survey rounds when the questionnaire was unchanged.

⁴¹As 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).

3.2.2. The test predicts that the elasticity of the price index—calculated using only the subset of goods in G —with respect to prices of an outside good j should equal the expenditure share of the outside good (equation 9). Given that we have reliable price data for only foods and fuels, we implement the test by re-estimating the log of the price index from food expenditures only (i.e. using only 2 of the 3 G s) by district and by decile, and regressing these indices on log fuel price changes interacted with fuel expenditure shares.⁴² Assuming fuel price changes across districts are independent of other unobserved price changes, we expect a coefficient equal to unity. Note that this test goes to heart of our methodology that recovers the complete price index for all goods despite only using relative consumption for a subset of goods for which we have reliable price data. In particular, it asks whether relative consumption within a particular group (food in this case) successfully captures price changes outside that group (fuels in this case, but more generally manufactures and services where prices are poorly measured).

We show the results of this test in Table 2. Panel A uses our baseline price index estimates, both P^0 and P^1 , calculated using only food groups (i.e. excluding fuels). Panel B additionally applies the exact price correction approach. Columns 2 and 4 include district fixed effects and thus exploit within-market variation across deciles—i.e. do our estimated price indices increase with fuel prices relatively more for deciles with larger expenditure shares on fuel? In support of our quasi-separability assumption, coefficients are close to unity and in no case can we reject a coefficient of one.

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 A.5-A.7 illustrate and assess these sample selection issues. The left panel of Appendix Figure A.5 presents the 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 there is most likely to be no true overlap in the presence of economic growth.⁴³

The middle panel of Appendix Figure A.5 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

⁴²Of course, this is only a partial test, as we do not know if quasi-separability holds with respect to goods for which we do not observe prices.

⁴³Figure A.6 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.

any market-decile pairs for which the median ranked good has no overlap. As shown in Appendix Figure A.7, 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 across time (see Section 3.2.5). Appendix Figure A.8 recalculates price indices using non-parametric Engel curves that condition on a standard set of linear controls for household characteristics.⁴⁴ Reassuringly, results change little, suggesting that systematic bias in estimates of cross-sectional Engel curves is unlikely to be driving our findings.

Appendix Figure A.9 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 conditional on their position in the income distribution (i.e. due to taste heterogeneity). Third, we can address concerns that the composition of household types may have changed over time, biasing 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 A.10 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 A.11 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.

⁴⁴In 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).

4.4 Revisiting the Impacts of India's 1991 Trade Reforms

In this section, 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 above) 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.⁴⁵ 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.⁴⁶ For expositional purposes, we focus on the log of our equivalent variation welfare metric.⁴⁷ 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 Figure 7 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 A.12). One potential explanation is that hard hit areas did not experience the same increases in retail-

⁴⁵This specification corresponds to column 8 in Table 3a) of Topalova (2010).

⁴⁶We 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.

⁴⁷As discussed above, we have more overlapping Engel curves and so less noise when calculating P^1 and EV.

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.

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 does not require price information for the full consumption basket. Instead, our approach estimates changes in household price indices and welfare across the income distribution from horizontal shifts in relative Engel curves. In poor data environments without any reliable price information, we can uncover theory-consistent and exact price indices as well as welfare changes as long as relative price changes within some quasi-separable product group are uncorrelated with slopes of relative Engel curves. Where reliable price data do exist for some subset of goods, we can relax the restrictions on relative price changes within the quasi-separable group, as well as validate our assumptions on preferences.

We apply this new method to measure changes in household welfare and revisit the effects of trade over India's reform period. 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. Second, going beyond poverty rates, our estimation reveals that the adverse effects of import competition in India are borne by households across the entire income distribution, not just the poor.

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

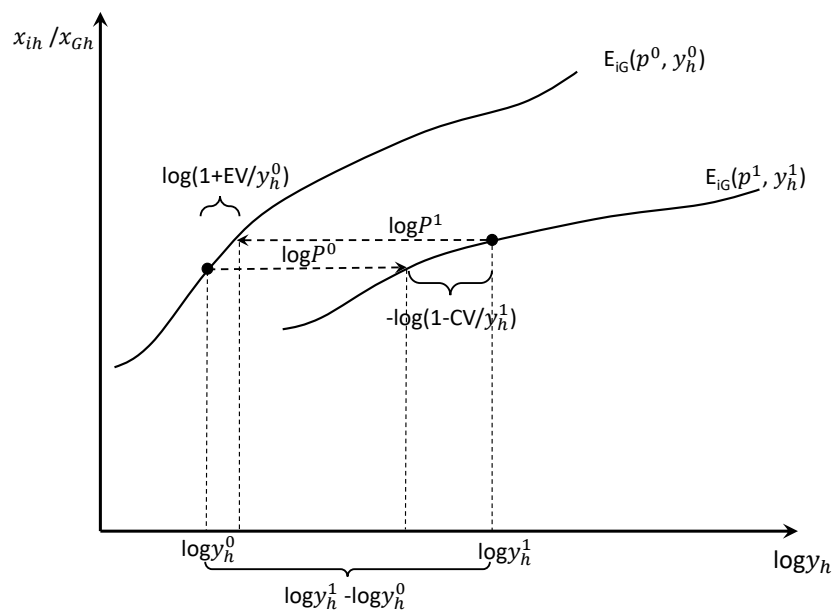
- ALMÁS, I. (2012): "International income inequality: Measuring PPP bias by estimating Engel curves for food," *American Economic Review*, 102(2), 1093–1117.
- ALMÁS, I., T. BEATTY, AND T. CROSSLEY (2018): "Lost in translation: What do Engel curves tell us about the cost of living?," *IFS working paper 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 (2020): “Measuring the Cost of Living in Mexico and the US,” *NBER Working Paper 27806*.
- 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.
- BLACKORBY, C., D. PRIMONT, AND R. R. RUSSELL (1978): *Duality, separability, and functional structure: Theory and economic applications*, vol. 2. Elsevier Science Ltd.
- BORUSYAK, K., AND X. JARAVEL (2018): “The distributional effects of trade: Theory and evidence from the United States,” *Harvard University working paper*.
- COMIN, D. A., D. LASHKARI, AND M. MESTIERI (forthcoming): “Structural change with long-run income and price effects,” *Econometrica*.
- 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, I., AND J. P. NEARY (2019): “New Characteristics and Hedonic Price Index Numbers,” *Economics Papers 2019-W02*, Economics Group, Nuffield College, University of Oxford.
- 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,” *The Quarterly Journal of Economics*, 131(3), 1113–1180.
- FALLY, T. (2018): “Integrability and generalized separability,” *CEPR discussion paper 12667*.
- FEENSTRA, R. C. (1994): “New product varieties and the measurement of international prices,” *American Economic Review*, 84(1), 155–57.
- GOLDBERG, P. K., A. K. KHANDELWAL, N. PAVCNİK, 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. (2019): "Are Poor Cities Cheap for Everyone? Non-homotheticity and the Cost of Living Across US Cities," *NBER Working Paper 26574*.
- HANOCH, G. (1975): "Production and demand models with direct or indirect implicit additivity," *Econometrica*, pp. 395–419.
- HOTTMAN, C., AND R. MONARCH (2018): "Estimating unequal gains across U.S. consumers with supplier trade data," *Federal Reserve Board working paper*.
- JARAVEL, X. (2019): "The unequal gains from product innovations: Evidence from the us retail sector," *The Quarterly Journal of Economics*, 134(2), 715–783.
- LEWBEL, A. (1991): "The rank of demand systems: Theory and nonparametric estimation," *Econometrica*, pp. 711–730.
- 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," *The Quarterly Journal of Economics*, 135(1), 503–560.
- SARHAN, A. (1955): "Estimation of the mean and standard deviation by order statistics. Part II," *The 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 $E_{iG}(p^0, y_h^0)$ and $E_{iG}(p^1, y_h^1)$, 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

93



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 described in Proposition 2 and Section 3.1.1 for environments with no price information—i.e. assuming the orthogonality condition holds and making no price correction for within- G relative price changes. Middle panel applies the first-order price correction and right panel applies the exact correction, both described in Proposition 1 and Section 3.1.2, 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



38

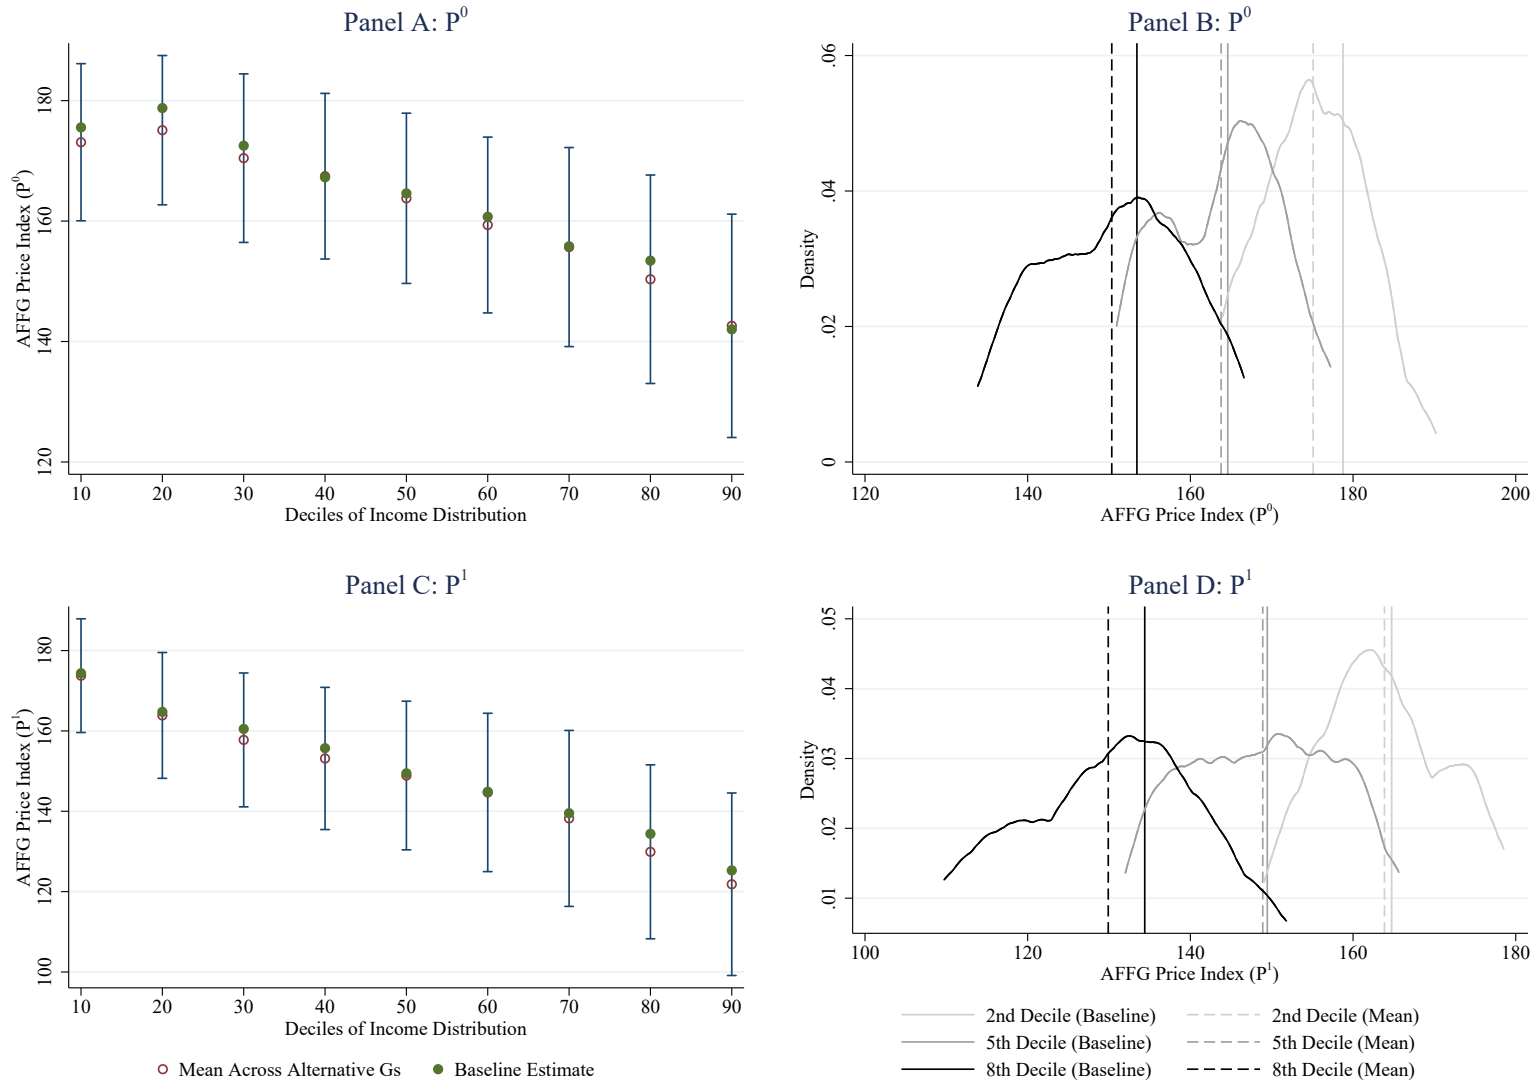
Notes: Figure shows the percentage change in the rural price index between 1987/88 and 1999/2000 for each decile of the local per-capita outlay distribution (averaged across districts using population weights). Left panel plots our AFFG NPC price index changes estimated from horizontal shifts in relative 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 AFFG 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 4.3.1 for further discussion.

Figure 7: 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 4.4 for further discussion.

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	Cereals - rice	Rice; chira; khoi, lawa; muri; other rice products.
Raw food products	Cereals	Cereals - wheat	Cereals - wheat	Wheat, atta, wheat/atta PDS; maida; suji, rawa; sewai (noodles); bread (bakery).
Raw food products	Cereals	Cereals - coarse	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	Gram (full grain/whole); gram products.
Raw food products	Gram and pulses	Pulses - besan, moong	Pulses - besan, moong	Besan; moong; soyabean; other pulse products.
Raw food products	Gram and pulses	Pulses - urd, masur	Pulses - urd, masur	Urd; masur; arhar (tur); khasari; peas (dry); gram (split); other pulses.
Raw food products	Meat, fish and eggs	Meat, fish and eggs	Meat	Goat meat, mutton; beef, buffalo meat; pork; poultry.
Raw food products	Meat, fish and eggs	Meat, fish and eggs	Fish, prawn	Fish, prawn.
Raw food products	Fruits and vegetables	Fruits and vegetables	Eggs	Eggs, egg products.
Raw food products	Fruits and vegetables	Fruits and vegetables	Vegetable - root vegetables	Potato; arum; radish; carrot; turnip; beet; sweet potato; onion; other root vegetables.
Raw food products	Fruits and vegetables	Fruits and vegetables	Vegetable - gourds	Pumpkin; gourd; bitter gourd; cucumber; parwal/patal; jhinga/torai; snake gourd; other gourds.
Raw food products	Fruits and vegetables	Fruits and vegetables	Vegetable - leafy vegetables	Cauliflower; cabbage; brinjal; lady's finger; french beans, barbati; tomato; palak/other leafy vegetables.
Raw food products	Fruits and vegetables	Fruits and vegetables	Vegetable - other vegetables	Peas (fresh); chilli (green); capsicum; plantain(green); jackfruit (green).
Raw food products	Fruits and vegetables	Premium Fruits	Premium Fruits	Apple; grapes; leechi; orange/mausani; pineapple; pears (naspatti); mango; watermelon.
Raw food products	Fruits and vegetables	Other fresh fruits	Other fresh fruits	Banana; jackfruit; singara; papaya; kharboza; berries.
Raw food products	Fruits and vegetables	Dry fruits and nuts	Dry fruits and nuts	Coconut (copra); groundnut; dates; cashewnut; walnut; raisin (kishmish, monacca, etc.).
Other food products	Dairy products and edible oils	Ghee	Ghee	Ghee; butter.
Other food products	Dairy products and edible oils	Milk	Milk	Milk (liquid).
Other food products	Dairy products and edible oils	Milk (condensed/powder)	Milk (condensed/powder)	Milk (condensed/powder); curd; baby food.
Other food products	Dairy products and edible oils	Vanaspati, margarine	Vanaspati, margarine	Vanaspati, margarine.
Other food products	Dairy products and edible oils	Edible oils	Edible oils	Ground nut oil; mustard oil; coconut oil; other edible oils.
Other food products	Sugar, salt, and spices	Sugar	Sugar	Sugar; gur; sugar candy (misri); sugar(others sources); honey.
Other food products	Sugar, salt, and spices	Salt	Salt	Salt.
Other food products	Sugar, salt, and spices	Spices	Spices	Turmeric; black pepper; dry chillies; garlic; tamarind; ginger; curry powder.
Other food products	Refreshments and intoxicants	Beverages	Beverages	Tea (leaf); coffee (cups); coconut (green).
Other food products	Refreshments and intoxicants	Processed food	Processed food	Cooked meals; pickles; sauce; jam, jelly.
Other food products	Refreshments and intoxicants	Pan	Pan	Pan (finished); supari; lime; katha.
Other food products	Refreshments and intoxicants	Tobacco	Tobacco	Bidi; cigarettes; leaf tobacco; snuff.
Other food products	Refreshments and intoxicants	Intoxicants	Intoxicants	Country liquor; beer; foreign liquor or refined liquor.
Fuels	Fuels	Coke, coal, charcoal	Coke, coal, charcoal	Coke; coal; charcoal.
Fuels	Fuels	Kerosene	Kerosene	Kerosene.
Fuels	Fuels	Firewood and chips	Firewood and chips	Firewood and chips.
Fuels	Fuels	Electricity	Electricity	Electricity.
Fuels	Fuels	Matches	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 C). Some NSS items were not consistently classified over rounds. Specifically: (Concorded) Rice uses individual items from R43 (Rice; Paddy) and R55 (Rice; Paddy). (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). (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).

Table 2: Quasi-Separability Test

<i>Panel A</i>	Dep. var: log AFFG Price Index (no price correction, calculated excluding fuels)			
	log P ⁰	log P ⁰	log P ¹	log P ¹
$\Delta\log(\text{Price Fuel}) \cdot \text{Exp. Share Fuel}$	1.136*** (0.180)	1.115*** (0.180)	0.928*** (0.196)	0.877*** (0.199)
<i>p</i> -value test $\beta=1$	0.45	0.52	0.71	0.54
Obs	2926	2926	2986	2986
R2	0.0350	0.0350	0.0316	0.0310
Decile specific $\Delta\log(\text{Price Fuel})$	No	Yes	No	Yes
Income decile FE	Yes	Yes	Yes	Yes

<i>Panel B</i>	Dep. var: log AFFG Price Index (exact price correction, calculated excluding fuels)			
	log P ⁰	log P ⁰	log P ¹	log P ¹
$\Delta\log(\text{Price Fuel}) \cdot \text{Exp. Share Fuel}$	1.011*** (0.174)	0.981*** (0.173)	0.913*** (0.189)	0.867*** (0.191)
<i>p</i> -value test $\beta=1$	0.95	0.91	0.65	0.49
Obs	2934	2934	2986	2986
R2	0.0215	0.0212	0.0262	0.0256
Decile specific $\Delta\log(\text{Price Fuel})$	No	Yes	No	Yes
Income decile FE	Yes	Yes	Yes	Yes


Notes: Table performs the quasi-separability test described in subsection 3.2.2. Dependent variable in Panel A is the log AFFG NPC price index, either $\log P^0$ or $\log P^1$, estimated using only food items (i.e. excluding fuels). In Panel B the dependent variable is the log AFFG price index (exact price correction approach) estimated using only food items correcting for relative price changes using the isoelastic correction with $\sigma_G = 0.7$. The explanatory variable in columns 1 and 3 is the log change in the price of fuels (calculated using a Paasche price index of fuel items where weights are given by mean district-level expenditure shares across items within the fuels category) multiplied by the district-by-decile expenditure share on fuels. The explanatory variable in columns 2 and 4 uses the decile-specific log change in the price of fuels (calculated using a Paasche price index of fuel items where weights are given by district-by-decile mean expenditure shares across items within the fuels category) multiplied by the district-by-decile expenditure share on fuels. The first row of the bottom panel reports the *p*-value on the test of the coefficient of interest being equal to 1, as required by the quasi-separability test. Regressions are weighted using district weights. Robust standard errors in parenthesis.

Appendix

A Additional Figures and Tables


Appendix Figures

Figure A.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 A.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 A.3: Rural Indian Cost of Living Inflation 1987/88–1999/2000: Estimates Using Alternative Levels of Good Aggregation

8




Notes: Figure shows the average percentage change in the rural AFGG 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 A.4: 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 AFGG 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 A.5: 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 4.3.1 for further discussion.

Figure A.6: 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 A.5. See Section 4.3.1 for further discussion.

Figure A.7: 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 4.3.1 for further discussion.

Figure A.8: 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 4.3.1 for further discussion.

Figure A.9: Rural Indian Cost of Living Inflation 1987/88–1999/2000: Heterogeneity by Household Type




Notes: Figure shows rural AFGG 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 A.10: 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 A.11: 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.

Figure A.12: 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 Figure 7 and Section 4.4, 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 4.4 for discussion.

Appendix Tables

Table A.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 A.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.

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.¹ 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

¹These 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.²

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.

C Theory Appendix

C.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 the section C.4 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.

²While “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.

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 $H_{iG}(p'_G, U) = H_{iG}(p_G, U)$ if relative prices remain constant: $p'_G = \lambda_G p_G$ across all goods in group G . For a consumer at initial utility u , income y and price p , notice that:

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

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

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

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

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

- i) For part i), 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:

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

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 $P^0(p^0, p^1, y_h^0)$:

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

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.

- 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:

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

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

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

where the last term refers to the within-group G expenditure share of good i in period 1.

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 :

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

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):

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

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:

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

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.

C.2 Proof of Proposition 1

As we have seen for the proof of Lemma 1 (equality A.2), we have: $H_{iG}(p_G, V(p, y)) = E_{iG}(p, y)$ 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:

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

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 .

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:

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

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

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:

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

where $\frac{\partial \log E_{iG}}{\partial \log p_j}$ and $\frac{\partial \log E_{iG}}{\partial \log y}$ are the uncompensated elasticities which can be more directly estimated. To prove this result, using $H_{iG}(p, U) = E_{iG}(p, e(p, U))$ and 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:

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

C.3 Proof of Proposition 2

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

$$\log E_{iG}(p^0, y^1 / P^1(p^0, p^1, y_h^1)) = \log E_{iG}(p^1, y_h^1) + \log \frac{H_{iG}(p_G^0, U_h^1)}{H_{iG}(p_G^1, U_h^1)} \quad (\text{A.4})$$

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:

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

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:

$$\log \left(\frac{y_h^1}{P^1} \right) \approx \frac{1}{G} \sum_{i \in G} \log E_{iG}^{-1}(p^0, \frac{x_{ih}^1}{x_{Gh}^1}) + \frac{1}{G} \sum_{i \in G} (\beta_{ih}^0)^{-1} \log \frac{H_{iG}(p_G^0, U_h^1)}{H_{iG}(p_G^1, U_h^1)} \quad (\text{A.5})$$

Hence the average of the horizontal shift $\log \left(\frac{y_h^1}{P^1} \right) - \log E_{iG}^{-1}(p^0, \frac{x_{ih}^1}{x_{Gh}^1})$ 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 .

C.4 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.

• **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:

$$\frac{\partial \log e_G}{\partial p_i} = \frac{1}{p_i} H_{iG}(p_G, U) \quad (\text{A.6})$$

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.6.

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).³ Using our constructed $e_G(p_G, U)$, notice that:

$$\frac{\partial e}{\partial p_i} = \frac{\partial e_G}{\partial p_i} \cdot x_G(p, U) \quad (\text{A.7})$$

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:

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

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.

³Lemma 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 .

First, using these two functions, let us define:

$$K(Q_G, q_{NG}, U) = \min_{e_G^*, p_{NG}^*} \left\{ \frac{Q_G e_G^* + \sum_{i \notin G} p_i^* q_i}{\tilde{e}(e_G^*, p_{NG}^*, U)} \right\} \quad (\text{A.8})$$

and:

$$F_G(q_G, U) = \min_{p_G} \left\{ \frac{\sum_{i \in G} p_i^* q_i}{e_G(p_G^*, U)} \right\} \quad (\text{A.9})$$

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:

$$K(F_G(q_G, U), q_{NG}, U) = 1 \quad (\text{A.10})$$

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.8) and (A.9) 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.10) 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.10) 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.8) and (A.9) 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.9), then $\lambda_G p_i^*$ also minimizes (A.9) 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.8), then $\lambda_N e_G^*$ and $\lambda_N p_j^*$ also minimizes (A.9) 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.8) and (A.9), 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.10). Using equations (A.10), (A.8) and (A.9), 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.8) and (A.9), it is also easy to verify that the expenditure shares implied by utility defined in A.10 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.10), (A.8) and (A.9) 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).

C.5 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 any good i at that income level, such that*

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

- 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 $\frac{x_{hi}^0}{y_h^0} = E_i(p^1, y_h^1 - CV_h)$ and $\frac{x_{hi}^1}{y_h^1} = E_i(p^0, y_h^0 + EV_h)$.*

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 $E_i(p^t, y) = p_i q_i(p, y)/y$ the Engel curve for good i as a function of income y for a given set of prices p_t , and denote $V(p^t, y_h^t)$ the indirect utility function. In 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

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

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:

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

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:

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

and

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

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 $y_h^1 - CV_h = \lambda y^0$:

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

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

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

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:

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

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 $y_h^0 + EV_h$.

C.6 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 $E_i(p^t, y_h^t)$ (expenditure share x_{ih}^t/y_h^t plotted against total outlays y_h^t) reflects the price index change for any realization of price changes across periods and any y , i.e. $E_i(p^1, y) = E_i(p^0, y/P^1(y))$. 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:

$$E_i(p^1, y_h^1) = E_i(p^0, y_h^1/P^1(y_h^1)) \quad (\text{A.11})$$

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.11, we can see that it no longer depends on prices p^1 once we condition on utility U_h^1 :

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

(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 .

C.7 Proofs for Section 5.2

First-Order Correction Terms for Relative Price Changes

For instance, as shown with Proposition 2 for P^1 (see equation A.5 above), we have:

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

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 $\sum_{j \in G} \sigma_{ijht} = 0$ 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:

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

Exact Correction Terms for Relative Price Changes

Starting from Proposition 1, we now impose specific forms of demand. 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 initial consumption shares, we can predict consumption shares for all goods i within G for any change in relative prices, 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:

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

Another simple case is a special case of the EASI demand system (Lewbel and Pendakur, 2009). With EASI, H_{iG} can be written as:

$$H_{iG}(p_G, U) = \frac{A_i(U) + B_i(p_G) + UD_i(p_G)}{\sum_{j \in G} A_j(U) + B_j(p_G) + UD_j(p_G)}$$

A specification that is linear in prices yields:

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

since $\sum_{k \in G} B_{kj}(U) = 0$ if preferences are required to be quasi-separable in group G . Price semi-elasticities are given by:

$$\xi_{iJ}(U) \equiv \frac{\partial H_{iG}}{\partial \log p_J} = \frac{B_{iJ}(U)}{\sum_k A_k(U)}$$

where $\sum_j \xi_{ij}(U) = 0$.

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

$$\begin{aligned}
H_{iG}(p'_G, U) &= \frac{A_i(U) + \sum_j B_{ij}(U) \log p'_j}{\sum_k A_k(U)} \\
&= H_{iG}(p_G, U) + \frac{\sum_j B_{ij}(U) (\log p'_j - \log p_j)}{\sum_k A_k(U)} \\
&= H_{iG}(p_G, U) + \sum_j \xi_{ij}(U) (\log p'_j - \log p_j)
\end{aligned}$$

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

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

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:

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

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

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 :

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

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

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

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:

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

Test of Quasi-Separability

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 $i \in G$ do not depend on the price of any good $j \notin G$ if we hold utility U constant:

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

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

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

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 result which holds if and only if preferences are quasi-separable:

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

Note that it is possible to provide an alternative proof using Slutsky decomposition for good i and compare to the sum of other goods $i' \in G$.

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: $E_{i,G}(p^1, y_h^1) = E_{i,G}(p^0, \frac{y_h^1}{P^1(y_h^1)})$, and $E_{i,G}(p^0, y_h^0) = E_{i,G}(p^{t1}, \frac{y_h^0}{P^0(y_h^0)})$. Taking the sum across varieties $i \in g$ of a subgroup g , we obtain:

$$E_{g,G}(p^1, y_h^1) = \sum_{i \in g} E_{i,G}(p^1, y_h^1) = \sum_{i \in g} E_{i,G}(p^0, \frac{y_h^1}{P^1(y_h^1)}) = E_{g,G}(p^0, \frac{y_h^1}{P^1(y_h^1)}) \quad (\text{A.16})$$

and:

$$E_{g,G}(p^0, y_h^0) = \sum_{i \in g} E_{i,G}(p^0, y_h^0) = \sum_{i \in g} E_{i,G}(p^{t1}, \frac{y_h^0}{P^0(y_h^0)}) = E_{g,G}(p^{t1}, \frac{y_h^0}{P^0(y_h^0)}) \quad (\text{A.17})$$

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.16) and (A.17) 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 price 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 $H_{g,G}(P_G, V(p, y)) = \sum_{i \in g} H_{i,G}(p_G, V(p, y)) = \sum_{i \in g} E_{iG}(p, y) = E_{gG}(p, y)$, we obtain, as in Proposition 1:

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

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

Finally, note that we can also reformulate the orthogonality condition in Proposition 2 across subgroups, using price indices across subgroups instead of good-level prices.

Heterogeneous Preferences

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:

$$x_{iG,h,z}^1 = E_{iG,z}(p^1, y_h^1) = E_{iG,z}(p^0, \frac{y_h^1}{P^1(y_h^1)})$$

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 :

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

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:

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

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:

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

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

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

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 $\sum_i \beta_{i,z}^1 / \bar{\beta}_i^1$.