

ESTIMATION OF AN IMPLICIT ADDITIVE INDIRECT DEMAND SYSTEM

ANTON C. YANG

Department of Economics, Yale University

PAUL V. PRECKEL

Department of Agricultural Economics, Purdue University

Quantitative economic models typically rely on CES demand or more recent non-homothetic CES systems that impose a common substitution structure across goods. We estimate a more general globally regular demand system that relaxes this restriction while preserving tractability. We combine share inversion with constrained optimization and develop inference for parameters on the boundary of the relaxed regularity domain. Using unpublished 2017 ICP expenditure and price data, we reject the common-substitution restriction and find substantial heterogeneity in substitution behavior across goods. The unrestricted system preserves the familiar ranking of agriculture, manufacturing, and services by income elasticity, but changes the low-income curvature of demand and the implied substitution structure. We further extend the estimator to a constrained panel with country-sector fixed effects and recover the same benchmark from within-country variation, which shows the approach is not specific to a single cross section. Quantitatively, we show that the empirical Engel differences matter most when shocks operate through real income and expenditure reallocation, especially in developing economies with large food shares. More broadly, we show that demand restrictions often imposed for tractability are empirical questions whose economic consequences can be measured directly from the data.

KEYWORDS: Demand system, implicit additivity, parameter estimation, constant difference of elasticity.

1. INTRODUCTION

Quantitative models in macroeconomics, international trade, and spatial economics often rely on parsimonious preference structures that remain globally regular across large changes in income and relative prices. The constant-elasticity-of-substitution (CES) family is the standard example. CES delivers closed-form demand systems and simple aggregation, which makes it attractive for quantitative analysis. But this convenience comes at a cost. CES places strong restrictions on how demand evolves with income and how substitution patterns respond to changes in income and relative prices. Those restrictions are far from innocuous when quantitative models are applied across large income ranges or over long development horizons. In such environments, the interaction between substitution patterns and non-homothetic demand, particularly the curvature of Engel responses at low income levels, shapes structural change, trade patterns, and welfare outcomes.

[Hanoch \(1975\)](#) proposes a class of implicitly additive preferences that can be viewed as a globally regular extension of CES. In its implicit additive indirect form, the constant-difference-of-elasticity (CDE) system separates substitution effects from income effects in ways that go beyond a simple relaxation of homotheticity. The system allows substitution patterns to differ across goods and can accommodate complementarity while preserving global regularity.

Anton C. Yang: anton.yang@yale.edu.

Paul V. Preckel: preckel@purdue.edu.

Anton Yang is grateful for the support provided by the Center for Global Trade Analysis at Purdue University and the Tobin Center for Economic Policy at Yale University. The authors would like to express their gratitude to the World Bank's ICP Global Office who provided the access to the unpublished data. All errors are our own.

These features make CDE attractive for applied general equilibrium analysis and for quantitative models that seek greater flexibility without giving up tractability.

Despite its flexibility, the CDE system has remained relatively underused and has rarely been estimated empirically since its introduction more than forty years ago. While the system is conceptually straightforward to economists across fields, formal estimation has remained difficult in practice. A related line of work studies more restricted non-homothetic CES (NHCES) systems within the same implicit additive tradition. [Satō \(1975\)](#) develops a non-separable non-homothetic CES class which, in the form adopted in recent quantitative work, imposes a common substitution structure across goods. Important recent contributions include [Matsuyama \(2019\)](#), who studies structural change, innovation, and trade using isoelastically non-homothetic CES preferences that satisfy implicit additivity, as well as [Comin et al. \(2021\)](#)'s NHCES framework, which introduces persistent income effects into a tractable multi-sector growth model to study structural change and the long-run reallocation of activity across agriculture, manufacturing, and services. Yet these frameworks maintain a common substitution structure across goods, and the quantitative implications of relaxing this restriction, and of moving to a more flexible but still parsimonious demand system, remain largely unexplored.

A second influential application appears in the policy world through global computable general equilibrium models such as the Global Trade Analysis Project (GTAP), which adopt the CDE system as a more flexible representation of demand and are widely used by thousands of researchers and policy institutions worldwide ([Hertel, 1997](#)). In practice, however, the preference parameters are almost never estimated directly, and quantitative counterfactuals rely heavily on calibration procedures and local approximations to handle the econometric difficulties associated with globally regular non-homothetic systems. The practical flexibility of the demand system is therefore more limited than its theoretical structure would suggest.

This paper estimates the more general globally regular system directly using unpublished 2017 World Bank ICP data covering 178 regions. Rather than imposing an aggregation in advance, we begin from disaggregated consumption data and construct broad consumption bundles guided by [Hanoch \(1975\)](#) and the empirical structure of the data. We use Engel patterns, cross-region expenditure-share variation, and within-bundle heterogeneity to identify categories that a smooth globally regular demand system can plausibly represent. The aggregation choice matters because some categories exhibit strong and systematic Engel responses, while others are shaped more by tax, regulation, culture, or institutions than by income alone.

Estimating the implicit additive system raises several econometric and computational challenges. Utility enters the share equations implicitly and is not observed, which means multiple parameter configurations can satisfy the system equally well. Because utility is defined implicitly, reduced-form estimation requires double log differencing to eliminate utility and an additional endogenous aggregation term analogous to [Anderson and Van Wincoop \(2003\)](#)'s multilateral resistance in the trade literature. We address these problems by estimating the system directly as a constrained likelihood problem in which utilities and preference parameters are solved jointly, while explicit normalizations remove the remaining degrees of freedom in the parameterization. The approach is conceptually related to inversion arguments that recover unobservables from share equations, as in [Berry \(1994\)](#), except that the inversion is embedded directly in a constrained optimization in the spirit of [Su and Judd \(2012\)](#), who develop a computational framework for estimating structural econometric models.

In estimating the unrestricted system, we find that some non-homotheticity parameters approach the lower boundary, while the constrained likelihood occasionally places substitution parameters at their lower bounds. We show that the regularity conditions on the substitution parameters can be relaxed without losing global regularity and establish conditions under which the implicit system remains globally regular on the boundary of the parameter domain, with

a unique strictly positive utility solution. These results show when a smooth globally regular demand system is supported by the data rather than imposed through a convenient aggregation.

Following [Andrews \(1999, 2000\)](#), we develop inference procedures tailored to parameters that lie on the boundary of the regularity domain. To evaluate the unrestricted and restricted systems, we implement likelihood-ratio tests that account for the fact that the unrestricted likelihood may be constrained at non-interior points along directions not being tested. The resulting statistics are thus evaluated using the mixture critical values characterized by [Chernoff \(1954\)](#), [Self and Liang \(1987\)](#), and [Andrews \(2001\)](#), rather than standard χ^2 asymptotics. We then move beyond the statistical evidence and examine whether the additional flexibility is supported by the data. We compare the income elasticities implied by the two systems to local Engel slopes estimated directly from prices, quantities, and expenditures in the ICP data after controlling for own prices and an aggregate price index for the remaining consumption bundles. All of the evidence points against the common-substitution restriction and supports substantial heterogeneity in substitution behavior across goods.

We then aggregate the system to three sectors to study broader economic implications. This also provides an external validation of the constrained-likelihood framework. Aggregating the ICP data into agriculture, manufacturing, and services, the restricted system delivers parameter estimates close to those obtained by [Comin et al. \(2021\)](#) from a postwar panel of 39 countries and preserves the familiar macroeconomic pattern in which agriculture has low income elasticity, services have high income elasticity, and manufacturing lies between them. The unrestricted system preserves this ranking but changes the substitution structure. Manufacturing and services are much more substitutable with one another than either is with agriculture, implying different counterfactual responses once relative prices change across sectors. We further extend the estimator to a panel with country-sector fixed effects using the country-year data of [Comin et al. \(2021\)](#) and recover the same benchmark from within-country variation. The constrained-likelihood framework is therefore not specific to the ICP cross section.

The substitution differences are immediate from the theoretical structure. Less obvious is how the income patterns differ, since both systems allow non-homothetic demand. The three-sector comparison in [Figure 1](#) illustrates the same qualitative pattern that appears in our main specification, with low-income curvature distributed differently across sectors, partly because the unrestricted system allows sector-specific substitution responses. The unrestricted system places more curvature in agriculture, while the restricted system places more curvature in manufacturing and services. The differences are concentrated at low expenditure levels, where non-homothetic demand matters most quantitatively.




FIGURE 1.—Three-sector income elasticities for the restricted and unrestricted implicit additive systems.

We close with a counterfactual application in a multi-region, multi-sector general equilibrium trade model. We use our estimation framework to estimate an unrestricted CDE system within GTAP using household demand data for 65 sectors and 160 regions and replace the calibrated household demand parameters with the estimates developed in this paper. We then ask how the quantitative predictions change. The empirical results predict where the counterfactuals differ. The estimated system matters most when shocks operate through real income and Engel incidence, especially in developing economies with large food shares, and less when shocks mainly work through broad relative-price changes across sectors. Relative to the calibrated system, the estimated demand system reallocates expenditure away from necessities and toward nonfood bundles when real income rises, while generating much smaller differences under multilateral trade liberalization. The broader lesson is that estimation matters most when incidence depends on how households reallocate spending between necessities, services, and imported varieties. The same logic likely applies when moving from restricted to unrestricted demand systems, where the richer system matters most along the low-income Engel curvature that shapes expenditure reallocation in developing economies. More generally, the framework developed in this paper can be combined with quantitative counterfactual methods such as [Dekle et al. \(2007, 2008\)](#), [Caliendo et al. \(2019\)](#), allowing researchers to estimate parameters from globally regular demand systems and carry them into counterfactual analysis.

2. THEORY OF THE IMPLICIT CDE FUNCTIONAL FORM

This section presents the implicit additive demand system used throughout the paper and reviews the main properties relevant for estimation, followed by a discussion of the associated identification issues.

The demand model examined in this paper is an implicit and indirect relationship that relates utility, prices and total expenditure as follows:¹

$$G\left(\frac{\mathbf{p}}{w}, u\right) = \sum_k \beta_k u^{e_k(1-\alpha_k)} \left(\frac{p_k}{w}\right)^{1-\alpha_k} \equiv 1, \quad (2.1)$$

with $\log[u^{e_k}(p_k/w)]$ replacing $u^{e_k(1-\alpha_k)}(p_k/w)^{1-\alpha_k}$ in the limiting case where α_k approaches unity and where the subscript $k \in \{1, \dots, N\}$ indexes commodities; with vector $\mathbf{p} = \{p_k\}_{k=1}^N$ denotes commodity prices; u denotes per capita utility, and w denotes the per capita total expenditure. The model parameters to be estimated, β 's, e 's, and α 's are distribution, expansion and substitution parameters, respectively ([Hanoch, 1975](#)). In addition, the levels of per capita utility for each country are estimated. While it is unusual to estimate unobservable utility in demand studies, when estimation of an underlying explicit demand system derived from a utility maximization problem (e.g., in the CES case) is undertaken, the estimation produces everything needed to calculate utility up to a strictly increasing transformation. The difference lies in the fact that with an explicit functional form, there is no need to estimate utility, whereas with an implicit functional form utility must be explicitly estimated to fully identify the model parameters. As to the data, the prices are compiled from the ICP database for 2017; w is obtained from the ICP database; and p_k/w may be interpreted as the unit-cost price or the normalized price of commodity k ; and quantities, which do not appear in equation (2.1), but are nonetheless important. The stated parametric restrictions for the demand function to be globally valid

¹[Matsuyama \(2023\)](#) provides a recent survey of non-CES aggregators and discusses Hanoch's implicit additive systems within a broader taxonomy of additive demand structures.

(monotonic and quasi-concave) are that, at all $\mathbf{p}/w \gg 0$ (i.e., unit-cost prices all strictly positive), (i) $\beta_k, e_k > 0 \forall k \in N$, and (ii) either $\alpha_k > 1$ or $0 \leq \alpha_k \leq 1 \forall k \in N$. (The weak inequalities in the second set of conditions in (ii) are justified in Section 4.3).

The model is categorized as an implicit (rather than an explicit) function because the relationship defined by equation (2.1) cannot in general be algebraically solved for utility as an explicit function of exogenous variables and parameters. The model is indirect because its indifference curves, which illustrate demand patterns, are expressed in its unit-cost prices instead of quantities. The model is closely related to other standard demand models. For example, it is easy to show that, if we set $e_k = e = 1 \forall k \in N$ and $\alpha_k = \alpha \forall k \in N$, then equation (2.1) collapses to the standard indirect CES model.

Using Roy's Identity, the derived demand correspondence is

$$q_k\left(\frac{\mathbf{p}}{w}, u\right) = \frac{w}{p_k} \frac{\theta_k}{\sum_j \theta_j} = \frac{w}{p_k} \Lambda_k, \quad (2.2)$$

where θ_k is an auxiliary variable such that

$$\theta_k = \beta_k u^{e_k(1-\alpha_k)} (1 - \alpha_k) \left(\frac{p_k}{w}\right)^{1-\alpha_k}, \quad (2.3)$$

and Λ_k is the expenditure shares of goods k as a function of θ_k , which equals

$$\Lambda_k\left(\frac{\mathbf{p}}{w}, u\right) = \frac{\theta_k}{\sum_j \theta_j}. \quad (2.4)$$

Stacking equation (2.4) across goods gives a vector-valued mapping $\mathbf{\Lambda} = \mathbf{S}(u)$ from each region's latent utility level to its expenditure shares, conditional on prices and preference parameters. This connects our setting to the share-inversion tradition in differentiated-products demand estimation (Berry, 1994, Berry and Haile, 2021, Redding and Weinstein, 2020), where the structure of the share equations is exploited to recover unobservables. The latent object in this system is the utility level, which enters the indirect system implicitly through (2.1). The inversion is implemented through constrained optimization rather than in closed form. Section 4 carries out this step by introducing u_i as latent variables and imposing (2.1)–(2.4) as constraints in a constrained-likelihood problem, with statistical fit evaluated through the share residuals defined in (4.2).

The Allen-Uzawa elasticities of substitution σ_{km} are given by

$$\sigma_{km} = \alpha_k + \alpha_m - \sum_j \Lambda_j \alpha_j - \frac{\Delta_{km} \alpha_k}{\Lambda_k}, \quad (2.5)$$

where Δ_{km} is the Kronecker delta. Since σ_{km} depends on the share-weighted average of the substitution parameters, complementarity between goods can arise when this average is sufficiently large relative to the substitution parameter and expenditure share of a particular good.

The income elasticities η_k are given by

$$\eta_k = \frac{e_k(1 - \alpha_k) + \sum_j \Lambda_j e_j \alpha_j}{\sum_j \Lambda_j e_j} + \alpha_k - \sum_j \Lambda_j \alpha_j. \quad (2.6)$$

Equation (2.6) allows goods to be inferior as well as normal, since η_k may be negative when the share-weighted average substitution term is sufficiently large. The expression also shows that income elasticities depend on the substitution structure through the α_k terms, so allowing substitution patterns to differ across goods changes the implied Engel curvature across sectors.

A key advantage of the implicit model in question, as shown in (2.5), is its greater flexibility in comparison to explicit models, as it imposes fewer restrictions on the substitution matrix. As demonstrated by Houthakker (1960) and Hanoch (1975), explicit models exhibit two types of tight linkages between substitution and income elasticities. One is that the Allen-Uzawa substitution elasticities can always be derived as functions of the income elasticities, i.e., $\sigma_{km} = \eta_k \eta_m (\sum_j \alpha_j \Lambda_j)$ for the explicitly direct models, such as the CES model, or $\sigma_{km} = \eta_k + \eta_m + (\sum_j \alpha_j \Lambda_j - 2)$ for the explicitly indirect models. The other linkage is that there is always a ratio or additive relationship between the substitution and income elasticities, i.e., $\eta_k / \eta_m = \sigma_{kj} / \sigma_{mj}$ for the direct case, or $\eta_k - \eta_m = \sigma_{kj} - \sigma_{mj}$ for the indirect case (Hanoch (1975), Yang (2021)).² In contrast, the implicit system allows substitution and income effects to vary independently across goods. This flexibility is economically important because it permits the data to determine whether cross-good demand differences arise from non-homotheticity, substitution behavior, or both, rather than imposing a fixed relationship between the two. The system thus accommodates a broader range of demand responses while preserving global regularity.

3. IDENTIFICATION ISSUES

Identification of implicitly additive preferences raises three main issues. The first concerns reduced-form estimation. Transformations that eliminate utility and the endogenous aggregation term from the demand system also prevent identification of the full parameter vector, in addition to the usual endogeneity concerns. We extend the empirical discussion in Hanoch (1975) to motivate the constrained-likelihood approach developed in Section 4.2. We then examine price endogeneity. Appendix B.8 uses external data to characterize the amount of external price variation available in the cross section. The second issue is aggregation. Because the system is nonlinear in expenditure, regional aggregates can generate Jensen-type aggregation bias. Appendix B.9 quantifies the aggregation bias in a Monte Carlo exercise calibrated to the estimated 12-bundle CDE data-generating process (DGP). The third issue arises from the unobserved normalization constant embedded in the CDE functional form, an issue that also appears in more restricted cases within the Bergson family (Bergson, 1936). Section 4.1 shows how explicit normalizations resolve this issue.

3.1. Reduced-Form Approaches

The existing demand literature suggests that reduced-form estimation of implicit indirect demand systems requires double log differencing to eliminate utility, but doing so creates identification problems. To see this, take logs of both sides of equation (2.2) with $\xi_k = p_k/w$:

²These dependencies are true for all explicitly direct or indirect additivities (see Hanoch (1975)). In the literature of the 1960s, these relationships are sometimes referred to as relationships between cross-price derivatives and income derivatives, or substitution effects and Engel derivatives (Houthakker (1960), Powell (1966)).

$$\begin{aligned}
\ln q_k &= \ln [\beta_k u^{e_k(1-\alpha_k)} (1-\alpha_k) \xi_k^{-\alpha_k}] - \ln \left[\sum_j \theta_j \right] \\
&= \ln[\beta_k(1-\alpha_k)] + e_k(1-\alpha_k) \ln u - \alpha_k \ln \xi_k - \ln \left[\sum_j \theta_j \right].
\end{aligned} \tag{3.1}$$

Eliminating the last term in (3.1) by using the logarithmic ratio:

$$\begin{aligned}
\ln \frac{q_k}{q_1} &= \ln \frac{\beta_k(1-\alpha_k)}{\beta_1(1-\alpha_1)} + [e_k(1-\alpha_k) - e_1(1-\alpha_1)] \ln u - \alpha_k \ln \xi_k + \alpha_1 \ln \xi_1 \\
&= A_k + Z_k \ln u - \alpha_k \ln \xi_k + \alpha_1 \ln \xi_1 + \epsilon_k,
\end{aligned} \tag{3.2}$$

where $A_k = \ln\{\beta_k(1-\alpha_k)/[\beta_1(1-\alpha_1)]\}$ and $Z_k = e_k(1-\alpha_k) - e_1(1-\alpha_1)$; q_1 is the chosen good for normalization; $\epsilon_k \forall k \in N-1$ is assumed to be the random error, which is independent of ξ_k , and has mean zero and constant variance.

In deriving (3.2), we eliminate the endogenous aggregation term $\ln \left[\sum_j \theta_j \right]$, which is analogous to the multilateral resistance in [Anderson and Van Wincoop \(2003\)](#). Equation (3.2) is thus easier to estimate when u is observable, such as when estimating a production function as in [Surry \(1993\)](#). In demand estimation, however, utility remains unobserved, so equation (3.2) still cannot be estimated directly. Estimation therefore requires an additional differencing step to eliminate u . One way to proceed is to choose q_2 as a second normalization good:

$$\begin{aligned}
\ln \frac{q_2}{q_1} &= \ln \frac{\beta_2(1-\alpha_2)}{\beta_1(1-\alpha_1)} + [e_2(1-\alpha_2) - e_1(1-\alpha_1)] \ln u - \alpha_2 \ln \xi_2 + \alpha_1 \ln \xi_1 \\
&= A_2 + Z_2 \ln u - \alpha_2 \ln \xi_2 + \alpha_1 \ln \xi_1 + \epsilon_2,
\end{aligned} \tag{3.3}$$

where $A_2 = \ln\{\beta_2(1-\alpha_2)/[\beta_1(1-\alpha_1)]\}$ and $Z_2 = e_2(1-\alpha_2) - e_1(1-\alpha_1)$.

Isolating (3.3) so that only u is on the left-hand side:

$$\ln u = \frac{-A_2 + \alpha_2 \ln \xi_2 - \alpha_1 \ln \xi_1 + \ln \frac{q_2}{q_1} - \epsilon_2}{Z_2}. \tag{3.4}$$

Now substituting (3.4) into (3.2) to eliminate u without loss of generality in terms of the functional form, with $R_k = Z_k/Z_2$:

$$\begin{aligned}
\ln \frac{q_k}{q_1} &= A_k + \frac{Z_k}{Z_2} (-A_2 + \alpha_2 \ln \xi_2 - \alpha_1 \ln \xi_1 + \ln \frac{q_2}{q_1} - \epsilon_2) \\
&\quad - \alpha_k \ln \xi_k + \alpha_1 \ln \xi_1 + \epsilon_k \\
&= A_k - R_k A_2 + R_k \alpha_2 \ln \xi_2 - R_k \alpha_1 \ln \xi_1 + R_k \ln \frac{q_2}{q_1} \\
&\quad - \alpha_k \ln \xi_k + \alpha_1 \ln \xi_1 - R_k \epsilon_2 + \epsilon_i.
\end{aligned} \tag{3.5}$$

By rearranging equation (3.5), we now obtain an estimation equation as follows:³

$$\ln \frac{q_k}{q_1} = S_k - \alpha_k \ln \xi_k + R_k \alpha_2 \ln \xi_2 + (1 - R_k) \alpha_1 \ln \xi_1 + R_k \ln \frac{q_2}{q_1} + \varphi_k, \quad (3.6)$$

where $\varphi_k = \epsilon_k - R_k \epsilon_2$ is the error term $\forall k \in N - 2$, satisfying standard assumptions as for ϵ_k .

Regression estimation of equation (3.6) yields the intercept S_k and the coefficients α_k , $R_k \alpha_2$, $(1 - R_k) \alpha_1$, and R_k , which recovers α_1 and α_2 . Given R_k , the relationship between e_k and e_2 is also identified. Full identification of the system, however, still requires values for β_2 and e_2 , which cannot be done without additional identities. The reduced-form approach therefore identifies at most $N - 2$ substitution parameters α_k , while leaving the expansion parameters e_k and distribution parameters β_k unidentified. Under this specification, reduced-form estimation is insufficient for full parameter identification.

In our specification, ξ_k (the unit-cost price) is likely to be correlated with φ_k , and finding appropriate instruments is difficult given the limited data in our cross-sectional setting (Berry (1994), Berry et al. (1995)). Endogeneity of this kind would lead to biased and inconsistent estimates. The method formally introduced in Section 4.2, while computationally more demanding, offers a numerical procedure based on constrained optimization that avoids embedding a separate root-finding step for u_i inside each likelihood evaluation. This enables the calculation of u satisfying the defining relationship (2.1), while achieving the identification of other model parameters (i.e., α_k , β_k and e_k). Our proof of regularity condition in Section 4.3 shows there exists a unique solution to (2.1) up to a strictly increasing transformation of u as is introduced in Section 4.1.

3.2. Price Endogeneity and External Cost Shifters

The constrained-likelihood approach does not address price endogeneity in the same way as a product-level IV demand design. The identifying assumption is that the normalized prices of broad consumption bundles are uncorrelated with the share residuals left by the implicit additive system. This is the identifying logic used in cross-country structural demand work, where broad demand systems are estimated from country-level price, expenditure, and share variation rather than from product-level cost shocks (Fieler, 2011, Caron et al., 2014, Fajgelbaum and Khandelwal, 2016, Comin et al., 2021). Broad bundle aggregation also mitigates the idiosyncratic demand shocks of the kind that motivate instruments in differentiated-products demand. Identification then comes from cross-region variation in normalized prices, total expenditure, and expenditure shares, and the restrictions imposed by implicit additivity, Roy's identity, and the normalizations on β_k and e_k . Appendix B.8 characterizes external cost shifters for several bundles of interest. It shows that supply-side variables explain a meaningful share of the residual normalized-price variation that remains after nuisance controls, but that no individual shifter set is strong enough to replace the constrained-likelihood estimates in this cross section.

³In practice, empirical demand work also faces simultaneity biases because total expenditure is jointly determined with the quantities demanded for goods k , and the fact that the arbitrary choice of normalized goods can lead to multiple parameter estimates at length N (Surry, 1993). In this case, the estimation equation can be rearranged as $\ln \Lambda_k / \Lambda_1 = S_k + (1 - \alpha_k) \ln \xi_k + R_k \alpha_2 \ln \xi_2 + [(1 - R_k) \alpha_1 - 1] \ln \xi_1 + R_k \ln q_2 / q_1 + \varphi_k$ (by letting $\Lambda_k = q_k(p_k/w) = q_k \xi_k$ so that $\ln q_k / q_1 = \ln \xi_k + \ln \xi_1$). See also equation (B.8) in Appendix B.8, where this share-based specification is used to characterize external cost variation in normalized prices.

3.3. Aggregation Bias of Regional Data

A separate concern is that the system is nonlinear in expenditure, so regional average budget shares do not necessarily equal the budget share of a household at regional mean expenditure. Estimation on regional aggregates could in principle identify parameters that differ from the household-level primitives. The source of the aggregation bias is a Jensen effect. It scales with the curvature of the Engel curves and with within-region expenditure dispersion. Appendix B.9 quantifies this effect in a Monte Carlo calibrated to the estimated 12-bundle system. Under a household-level CDE DGP, the implied income-elasticity bias is exact at zero within-region dispersion and remains small at empirically relevant levels of inequality. At a within-region expenditure coefficient of variation (CV) equal to one, corresponding under a lognormal distribution to a consumption Gini of roughly 0.45 near the upper end of within-country inequality, the mean absolute income-elasticity error is about 0.012. This mean bias is an order of magnitude smaller than the CDE–NHCES elasticity differences of 0.2 to 0.7 that drive the structural comparisons in Section 7.3. Even under a more aggressive design that steepens the low-income necessity curves, the worst replication-level maximum error is 0.150, below the lower end of that range. The evidence shows that moderate within-region heterogeneity does not mechanically distort the income elasticities we interpret.

3.4. Excess Degrees of Freedom

Since u cannot be observed, the preference defining relationship in equation (2.1) has multiple sets of parameter values that will satisfy the relationship equally well. To see this, note that since any strictly increasing transformation of utility will not alter the ordering of preferences for alternative consumption bundles, such transformations will have no impact on the quality of the estimated relationship. For example, consider the transformation $u = \rho v^\delta$ where $v > 0$, which will be strictly increasing if $\rho > 0$ and $\delta > 0$. In this case, v will be as good as u for explaining the data. Substituting the transformed u into equation (2.1):

$$\begin{aligned}
 1 \equiv G\left(\frac{\mathbf{p}}{w}, u\right) &= \sum_k \beta_k u^{e_k(1-\alpha_k)} \left(\frac{p_k}{w}\right)^{1-\alpha_k} \\
 &= \sum_k \beta_k [\rho v^\delta]^{e_k(1-\alpha_k)} \left(\frac{p_k}{w}\right)^{1-\alpha_k} \\
 &= \sum_k \beta_k \rho^{e_k(1-\alpha_k)} v^{\delta e_k(1-\alpha_k)} \left(\frac{p_k}{w}\right)^{1-\alpha_k} \\
 &= \sum_k \tilde{\beta}_k v^{\tilde{e}_k(1-\alpha_k)} \left(\frac{p_k}{w}\right)^{1-\alpha_k},
 \end{aligned} \tag{3.7}$$

where $\tilde{\beta}_k = \beta_k \rho^{e_k(1-\alpha_k)}$, and $\tilde{e}_k = \delta e_k$. Since ρ and δ were any strictly positive constants, there is a continuum of values for the $\tilde{\beta}_k$ and $\tilde{e}_k$ that exactly satisfy equation (2.1) given that this defining relationship is satisfied for β_k and e_k . Thus, these parameters are not fully identified. For this reason, we introduce normalizations for these two sets of parameters, choosing our normalizations in such a way that the estimated parameter values are “well-scaled” as we will describe in Section 4.1.

4. ECONOMETRIC PROCEDURE

We follow the maximum-likelihood approach for implicitly additive demand systems used by Cranfield et al. (2002), Preckel et al. (2010), Gouel and Guimbard (2019), and Yang et al. (2018). We estimate the implicit indirect system as a constrained maximum-likelihood problem, imposing the defining equations and regularity restrictions directly in the optimization. The concentrated log likelihood $\log \mathcal{L}_c$, following Greene (2012), is

$$\log \mathcal{L}_c = -\frac{I}{2}[J(1 + \log 2\pi) + \log |\mathbf{R}'\mathbf{R}|] \quad (\text{with } \hat{\beta}_{\text{ML}} = \text{Min}_{\beta} \frac{1}{2} \ln |\mathbf{R}'\mathbf{R}|), \quad (4.1)$$

where I and J denote the numbers of regions and consumption bundles in the sample, respectively, and $|\mathbf{R}'\mathbf{R}|$ is the determinant of the cross-equation error covariance matrix. The disturbances d_{ik} are the residuals between observed and fitted expenditure shares:

$$d_{ik} = \Lambda_{ik} - \hat{\Lambda}_{ik}, \quad (4.2)$$

where i and k index regions and goods, respectively; $\hat{\Lambda}_{ik}$ are the fitted expenditure shares, and the components of $\mathbf{R}$, r_{nm} , are constrained by:

$$\sum_n r_{nk} r_{nm} = \sum_i \frac{d_{ik} d_{im}}{I}, \quad (4.3)$$

along with $r_{km} = 0$ for all $m > k$, making $\mathbf{R} = [r_{km}]_{k,m=1,\dots,N}$ an upper triangular Cholesky factorization of the error covariance matrix. The advantage of working with this factorization of the cross-equation error covariance matrix is that evaluation of the determinate of the covariance matrix is simple—it is the square of the product of the diagonal elements of $\mathbf{R}$.

4.1. Identification and Normalization Strategy

As shown above, the implicit additive system contains excess degrees of freedom that must be removed for parameter identification. In related work, a more restricted implicit indirect non-homothetic CES demand system takes the form

$$G\left(\frac{\mathbf{p}}{w}, u\right) = \sum_k \beta_k u^{e_k(1-\alpha)} \left(\frac{p_k}{w}\right)^{1-\alpha} = 1, \quad (4.4)$$

which can be obtained by restricting $\alpha_k = \alpha$ for all k in equation (2.1). This restricted specification is the non-homothetic (non-separable) CES system developed by Satō (1975) and adopted in recent quantitative work by Comin et al. (2021). Its key restriction is that all goods share a common substitution parameter.

Comin et al. (2021) observe that the expressions for the own price and income elasticities of demand are invariant to a multiplicative scaling of the parameters equivalent to our β_k and e_k for the direct non-homothetic CES case. In the interest of parameter identification, they remove a degree of freedom for each of these sets of parameters by normalizing these parameters to unity for one good (e.g., $k = 1$). We show in Section 3.4 the same invariance to rescaling of these parameters in the indirect case. We choose a slightly different normalization scheme by setting $\sum_k \beta_k \equiv 1$ and $\sum_k e_k \equiv N$, again removing one degree of freedom for each of these sets of parameters. Our normalization of the β_k combined with the form of our implicit preference defining relationship (2.1) will tend to improve the scaling of the terms

$u^{e_k(1-\alpha_k)}(p_k/w)^{1-\alpha_k}$, which we have observed tends to improve the scaling of the parameters α_k . Similarly, the choice to normalize the sum of the exponents e_k to be equal to the number of goods (and thus each e_k equals to one on average across the goods) tends to improve the scaling of the u levels. Our motivation, as well as that of [Comin et al. \(2021\)](#), is to obtain parameter identification for our demand system. While our normalizations are mathematically equivalent, our strategy of normalizing sums of these parameters rather than individual ones avoids the possibility of making an unfortunate choice for the parameters that are set to one, which may serve to make the other parameters either very large or small numbers. Given that we solve our estimation problems numerically using general nonlinear programming software, attention to scaling can improve our likelihood of success in estimation. In addition to its role in scaling and initialization, we also estimate this restricted specification as a formal benchmark nested within the unrestricted system. Section 7 reports estimates and a fit comparison.

4.2. Constrained Optimization

The constrained-likelihood formulation is an instance of the Mathematical Programming with Equilibrium Constraints (MPEC) approach proposed by [Su and Judd \(2012\)](#). This formulation matters for two reasons. First, the preference parameters (α, β, e) and the region-specific utility levels $\{u_i\}$ are identified jointly within a single optimization problem rather than through an inner–outer loop. A nested implementation would, for each candidate parameter vector, require solving the implicit relation (2.1) for each region’s u_i before fitted shares $\hat{\Lambda}_{ik}$ can be computed and (4.1) evaluated. We treat $\{u_i\}$ as latent variables and impose (2.1) with the Roy-identity share mapping (2.3)–(2.4) as constraints. This avoids the accumulation of numerical error from repeated inner-loop fixed-point solves and scales well with the number of regions. Second, the endogenous aggregation term $\sum_j \theta_j$ that drives the main identification challenge of the reduced-form approach in Section 3.1 is handled directly through the share-equation constraint rather than substituted out. The optimizer pins down both the numerator θ_k and the denominator $\sum_j \theta_j$ consistently with the structural model. Similar concepts have been used by [Balistreri and Hillberry \(2007\)](#) to estimate [Anderson and Van Wincoop \(2003\)](#)’s “Armington model” in the international trade literature.

The constrained optimization problem is to maximize equation (4.1), subject to the constraints (4.2), (4.3) along with $r_{km} = 0 \ \forall \ m > k$, the implicit indirect additivity relationship (2.1), the normalization equations $\sum_k \beta_k = 1$ and $\sum_k e_k = N$, the auxiliary identities (2.2)–(2.4), as well as the redefined parametric restrictions of the demand system.

As we show in Section 4.3, the strictly positive lower bounds on α_k ’s are not essential, although in the direct case corner solutions may result from the utility maximization problem. Since the quasi-concavity properties from the direct case are inherited by the indirect case, the demand system is well-behaved with these relaxed lower bounds for α_k . Consequently, we impose weak inequalities for the lower bounds on the α ’s, and our parametric constraints consist of (i) $\beta_k, e_k > 0 \ \forall \ k \in N$, and (ii) $0 \leq \alpha_k \leq 1 \ \forall \ k \in N$ (refer to Section 4.3). This choice disregards the situation where $\alpha_k > 1$, which we determined to be irrelevant for our dataset because the resulting elasticity estimates lacked credibility. (Our formulation applies equally well to the case where $\alpha_k > 1 \ \forall \ k \in N$.) Following [Hanoch \(1975\)](#), we interpret $(p_k/w)^{1-\alpha_k}$ as $\ln(p_k/w)$ when $\alpha_k = 1$. However, in our empirical analysis it emerged that $\alpha_k < 1 \ \forall \ k \in N$.

4.3. Proof of Global Regularity

In reviewing the published regularity conditions in [Hanoch \(1975\)](#) we noticed what appear to be some minor discrepancies in the conditions related to the boundaries for the parameter α_k .⁴ In this section, we readdress the proof of global regularity with an eye to whether the extreme values $\alpha_k = 0$ for some k should be included.

The CDE preference relationship is said to be globally regular if the relationship in equation (2.1) is monotone in both u and p_k/w and satisfies a quasi-concavity property. To show this, it is useful to begin from the direct form of the CDE, the constant ratio elasticity of substitution (CRES). We demonstrate that under a mildly relaxed set of parametric restrictions relative to [Hanoch \(1971, 1975\)](#), we obtain global regularity for the CRES. [Hanoch \(1975\)](#) argues that regularity of the direct CRES relationship yields regularity of the indirect CDE relationship due to symmetry between $f(\mathbf{x})$ and its *indirect reciprocal* $g(\mathbf{p}/w)$ in the CDE.

Following [Hanoch \(1975\)](#), we begin by setting up the utility maximization problem for the CRES subject to the defining constraint for the CRES relationship. This relationship has near-identical form to the CDE as follows:

$$\mathcal{H}(\mathbf{x}, u) = \sum_k \beta_k u^{-e_k(1-\alpha_k)} (x_k)^{1-\alpha_k} \equiv 1, \quad (4.5)$$

with $\log(x_i/u^{-e_k})$ replacing $u^{-e_k(1-\alpha_k)}(x_k)^{1-\alpha_k}$ in the limiting case where α_k approaches unity and where $\mathbf{x}$ denotes the levels of inputs. The stated regularity conditions are: (i) $\beta_k, e_k > 0 \forall k \in N$, and (ii) either $\alpha_k > 1$ or $0 < \alpha_k \leq 1 \forall k \in N$.

THEOREM 4.1: *The latter set of global regularity conditions in equation (4.5), i.e., condition (ii): either $\alpha_k > 1$ or $0 < \alpha_k \leq 1 \forall k \in N$, can be relaxed to either $\alpha_k > 1 \forall k \in N$, or $0 \leq \alpha_k \leq 1 \forall k \in N$.*

Theorem 4.1 relaxes the strict lower-bound conditions stated in [Hanoch \(1975\)](#). To our knowledge, neither [Hanoch \(1975\)](#) nor the subsequent non-separable CES literature following [Satō \(1975\)](#) treats the weak-boundary case $0 \leq \alpha_k \leq 1$. We show that global regularity extends to this weak-boundary case.

PROOF: The case where $\alpha_k > 1$ for all k is covered in the original work by [Hanoch \(1975\)](#) (Section 2.4, p.403) and is not addressed here.⁵ Similarly, the interpretation of $\mathcal{H}(x, u)$ when some $\alpha_k = 1$ can be treated as in [Hanoch \(1975\)](#), and does not affect the qualitative properties of $\mathcal{H}(x, u)$. Hence, it is sufficient to focus on the case where $0 \leq \alpha_k < 1$ to demonstrate regularity with the relaxed lower boundary condition on α_k . Also note that the CRES is linear in x_k for the case where $\alpha_k = 0$ for all k . This case is not of central interest here, and so we assume that $\alpha_k > 0$ for at least one k .

The utility maximization problem that motivates the demand system is:

⁴In [Hanoch \(1975\)](#) (p. 403), with different subscript i , the stated parametric restrictions for $d_i (= 1 - a_i)$ in the CRES Model (i.e., equation (2.16)) is that either $0 < d_i < 1$ or $d_i \leq 0, \forall i$, and for a_i is that either $a_i > 1$ or $0 < a_i \leq 1 \forall i$; then on p. 411, the stated conditions for $b_i (= 1 - \alpha_i)$ in the CDE system (i.e., equation (3.15)) is that either $0 < b_i < 1$ or $b_i \leq 0 \forall i$, and for α_i is that either $\alpha_i \geq 1$ or $0 < \alpha_i < 1 \forall i$. Because d_i 's and b_i 's as well as a_i 's and α_i 's are used interchangeably, the same regularity conditions for d_i and b_i should imply the same restrictions for a_i and α_i which, however, are not what we find in [Hanoch \(1975\)](#).

⁵The CRES system was introduced by [Mukerji \(1963\)](#) as a generalized version of the Solow–Minhas–Arrow–Chenery production function ([Arrow et al., 1961](#)).

$$\begin{aligned}
& \underset{x_k \geq 0 \ \forall k \in N}{\text{maximize}} && u \\
& \text{subject to:} && \mathcal{H}(x, u) = \sum_k \beta_k u^{-e_k(1-\alpha_k)} (x_k)^{1-\alpha_k} = 1 \\
& && \sum_k c_k x_k \leq I,
\end{aligned} \tag{4.6}$$

for fixed I , which has the dual expenditure minimization problem:

$$\begin{aligned}
& \underset{x_k \geq 0 \ \forall k \in N}{\text{minimize}} && \sum_k c_k x_k \\
& \text{subject to:} && \sum_k \beta_k u^{-e_k(1-\alpha_k)} (x_k)^{1-\alpha_k} = 1,
\end{aligned} \tag{4.7}$$

for fixed $u > 0$. Note that each term in the sum in the constraint of the latter problem is strictly increasing in x_k (for $0 \leq \alpha_k < 1$, $\beta_k > 0$, and $e_k > 0$), decreasing in u , and that $\mathcal{H}(x, u)$ is strictly decreasing in u . The importance of these features are two-fold. First, the fact that each term is strictly increasing in x_k and that $\mathcal{H}(x, u)$ is strictly decreasing in u implies that u is strictly increasing in each x_k . This demonstrates a strict monotonic relationship between u and each x_k . Second, the fact that $\mathcal{H}(x, u)$ is strictly decreasing in u combined with the fact that each term in $\mathcal{H}(x, u)$ increases without bound as u approaches zero and approaches zero as u increases without bound, implies that $\mathcal{H}(x, u) = 1$ will have a unique solution. Thus, the implicit relationship defined by $\mathcal{H}(x, u) = 1$ defines a single value for u for any given x .

The first-order condition w.r.t. x_k is

$$c_k - \lambda \beta_k u^{-e_k(1-\alpha_k)} (1 - \alpha_k) x_k^{-\alpha_k} \geq 0. \tag{4.8}$$

For any k such that $\alpha_k > 0$ this condition can only be satisfied with $x_k > 0$ and $\lambda > 0$, implying that the equality relationship for the constraint in the expenditure minimization problem (4.7) can be replaced with “ $\geq$ ” without loss of generality.

The Hessian of the constraint in the expenditure minimization problem is diagonal:

$$\nabla_x^2 \mathcal{H}(x, u) = \text{Diag} \left\{ \beta_k u^{-e_k(1-\alpha_k)} (1 - \alpha_k) (-\alpha_k) x_k^{-\alpha_k-1} \right\}, \tag{4.9}$$

and each diagonal element is either strictly negative if $0 < \alpha_k < 1$ or zero if $\alpha_k = 0$. Thus, $\mathcal{H}(x, u)$ is a concave function and $\{x | \mathcal{H}(x, u) \geq 1\}$ is a convex set. Finally, because $\mathcal{H}(x, u)$ is strictly increasing in u , $\{x | \mathcal{H}(x, u) \geq 1 \text{ and } u \geq a\}$ is also a convex set, and so $\mathcal{H}(x, u) = 1$ defines a quasi-concave relationship between u and x . *Q.E.D.*

In Appendix A.2, we provide solutions to the expenditure minimization problem under relaxed regularity conditions. It is lengthy, so we have omitted it here. The proof leads us to our relaxed set of parametric conditions for the CRES: $\beta_k, e_k > 0 \ \forall k \in N$, and either $\alpha_k > 1 \ \forall k \in N$, or $0 \leq \alpha_k \leq 1 \ \forall k \in N$. Again, following Hanoch (1975), due to complete formal symmetry between the direct and indirect cases, global regularity of the CRES in $\mathbf{x} \geq 0$ implies global regularity of the CDE in $\mathbf{p}/w \gg 0$. Note that the corner solutions for the CRES do not similarly imply that the CDE will generate corner solutions when some of our parametric

restrictions on α_k are binding. This is because the envelope conditions used to reclaim the demand quantities involve not only derivatives with respect to the numerator in p_k/w , but also in the denominator, which appears in other terms in our defining equation.

Theorem 4.1 establishes strict regularity on the open parameter domain in which $e_k > 0$ for all k . Empirically, however, the constrained MLE may lie on the closure of this domain, with $e_k = 0$ for some goods. We thus examine the extent to which the regularity properties extend to the boundary of the parameter domain. The result below establishes existence, uniqueness, and monotonicity on the admissible normalized-price domain. Boundary estimates should be interpreted as limits of strictly regular CDE systems rather than as points in the strict regularity domain in Hanoch (1975).

As in Theorem 4.1, we focus on the case $0 \leq \alpha_k < 1$. The limiting case $\alpha_k = 1$ is handled by logarithmic substitution in Hanoch (1975).

PROPOSITION 1—Boundary well-posedness: *Let $K_0 \subset \{1, \dots, N\}$ with $|K_0| < N$. Suppose that $\beta_k > 0$ and $0 \leq \alpha_k < 1$ for all k , that $e_k = 0$ for $k \in K_0$, and that $e_k > 0$ for $k \notin K_0$. Let $\xi_k = p_k/w$ and define*

$$\mathcal{D} \equiv \{\xi \gg 0 : S(\xi) < 1\}, \quad S(\xi) \equiv \sum_{k \in K_0} \beta_k \xi_k^{1-\alpha_k}. \quad (4.10)$$

Then for every $\xi \in \mathcal{D}$, the implicit additive relationship

$$G(\xi, u) = \sum_{k=1}^N \beta_k u^{e_k(1-\alpha_k)} \xi_k^{1-\alpha_k} = 1$$

has a unique strictly positive solution $u(\xi)$. $u(\xi)$ is weakly decreasing in each normalized price ξ_k . Equivalently, the implied indirect utility is weakly decreasing in each price p_k and strictly increasing in expenditure w .

PROOF: Fix $\xi \gg 0$. Differentiating G with respect to u gives

$$G_u(\xi, u) = \frac{\partial G(\xi, u)}{\partial u} = \sum_{k=1}^N \beta_k e_k (1 - \alpha_k) u^{e_k(1-\alpha_k)-1} \xi_k^{1-\alpha_k}. \quad (4.11)$$

Terms with $k \in K_0$ are zero because $e_k = 0$. Terms with $k \notin K_0$ are strictly positive because $\beta_k > 0$, $e_k > 0$, $0 \leq \alpha_k < 1$, $\xi_k > 0$, and $u > 0$. Since $|K_0| < N$, at least one such term exists. Hence $G_u(\xi, u) > 0$ for all $u > 0$, so $G(\xi, u)$ is strictly increasing in u .

Now decompose

$$G(\xi, u) = S(\xi) + \sum_{k \notin K_0} \beta_k u^{e_k(1-\alpha_k)} \xi_k^{1-\alpha_k}.$$

For $k \notin K_0$, we have $e_k(1 - \alpha_k) > 0$. Therefore,

$$\lim_{u \downarrow 0} G(\xi, u) = S(\xi), \quad \lim_{u \rightarrow \infty} G(\xi, u) = +\infty.$$

If $\xi \in \mathcal{D}$, then $S(\xi) < 1$. Continuity of G in u thus implies the existence of some $u(\xi) > 0$ satisfying $G(\xi, u(\xi)) = 1$. Strict monotonicity of G in u implies uniqueness.

The monotonicity of the solution follows from the implicit function theorem. For each k ,

$$G_{\xi_k}(\xi, u) = \frac{\partial G(\xi, u)}{\partial \xi_k} = \beta_k(1 - \alpha_k)u^{e_k(1 - \alpha_k)}\xi_k^{-\alpha_k} \geq 0. \quad (4.12)$$

Since $G_u(\xi, u) > 0$,

$$\frac{\partial u}{\partial \xi_k} = -\frac{G_{\xi_k}(\xi, u)}{G_u(\xi, u)} \leq 0.$$

Because $\xi_k = p_k/w$, this implies

$$\frac{\partial u}{\partial p_k} = -\frac{1}{w} \frac{G_{\xi_k}(\xi, u)}{G_u(\xi, u)} \leq 0$$

and

$$\frac{\partial u}{\partial w} = \frac{1}{w^2} \frac{\sum_j p_j G_{\xi_j}(\xi, u)}{G_u(\xi, u)} > 0.$$

The last inequality is strict because $p_j > 0$ for all j and $G_{\xi_j}(\xi, u) > 0$ for all j when $0 \leq \alpha_j < 1$. We verify the condition $S(\xi_i) < 1$ in equation (4.10) for all observed regions in Section 6. *Q.E.D.*

5. DATA

The data come from the World Bank's 2017 International Comparison Program (ICP) database and include expenditures, prices, exchange rates, and population for 216 regions. The ICP reports Purchasing Power Parities (PPPs) for 109 commodity and service categories at the level of *basic headings*, the most disaggregated expenditure categories in the ICP data.⁶ We exclude expenditures by residents abroad because they do not map naturally into the demand system. After excluding regions with insufficient expenditure or price coverage, the estimation sample contains 178 regions. The data span food, housing, transportation, communication, and other major expenditure categories. Work at the basic heading level is important for two reasons. First, it allows PPP aggregation from relatively disaggregated expenditure categories. Second, it allows direct examination of within bundle Engel patterns and expenditure share variation before aggregation to broader consumption bundles.

5.1. Implicit Additivity for Analyzing Goods and Substitution

[Hanoch \(1975\)](#) argues that implicitly additive demand systems such as CDE are most appropriate for broadly defined consumption categories rather than narrowly differentiated goods. In these systems, the Allen–Uzawa partial elasticities of substitution between good k and all other goods are proportional to a single substitution parameter (i.e., $\sigma_{km} \propto \alpha_k$). As a result, substitution patterns depend on broad category characteristics rather than on whether particular goods are close substitutes ([Hanoch, 1975](#), [Barten, 1977](#)).

To illustrate the point, consider coffee and tea versus coffee and pastries, and smartphones and tablets versus smartphones and laptops. Coffee and tea are close substitutes because they

⁶In the ICP data, a basic heading is the lowest level at which participating economies report expenditure data.

satisfy similar consumption needs and are often used interchangeably. Coffee and pastries, although frequently consumed together, satisfy different needs and therefore are not close substitutes in the same sense. Likewise, smartphones and tablets have overlapping functions and uses, whereas smartphones and laptops typically serve more distinct purposes for consumers.

Because the system is not designed to capture substitution patterns among narrowly differentiated goods, we aggregate basic headings into broad consumption bundles that the model can represent more accurately. Section 5.2 describes the empirical validation procedure used to determine the baseline 12-bundle aggregation. Appendix B.2 reports the full mapping from basic headings to bundles.

5.2. Aggregation Validation

The aggregation choice determines the dimension of the parameter vector and therefore which parameters are reliably identified given the cross-region data. In non-homothetic implicit additive demand systems, pooling a vice-type good (whose cross-region variation is likely dominated by tax, regulation, or cultural characteristics) with a normal good (whose cross-region variation tracks income) can push one or more non-homotheticity parameters toward the regularity boundary and may generate instability in the unconstrained likelihood along that direction. We therefore select the aggregation through an explicit empirical validation exercise rather than ad hoc adoption of an existing classification.

We begin from a grouping of the 109 ICP basic headings based on the United Nations Classification of Individual Consumption According to Purpose (COICOP) and refine it using three empirical criteria computed from the ICP panel. The first criterion is bundle-level share variation. For each candidate bundle, we examine the cross-region mean and dispersion of expenditure shares, the number of regions in which the share falls below 1.5% (“thin” regions), and the cross-region correlation between the bundle share and log per capita expenditure (the “Engel correlation”). We prefer bundles with nontrivial expenditure shares, limited thin-region exposure, and stable Engel patterns across regions. The second criterion is within-bundle Engel heterogeneity. For each basic heading g within a candidate bundle, we compute the cross-region correlation between expenditure shares and log income, denoted ρ_g . A wide dispersion in $\{\rho_g\}$ indicates that the bundle combines categories with different Engel patterns, which is difficult to reconcile within a smooth non-homothetic demand system. The third criterion is split feasibility. For each candidate split, we verify that the proposed sub-bundles preserve adequate cross-sectional share variation. We reject splits that produce thin sub-bundles in a substantial number of regions.

Table I reports the bundle-level share statistics and Engel correlations for the 12-bundle aggregation, as well as measures of within-bundle basic-heading (BH) heterogeneity. Each bundle shows a coherent Engel pattern across regions. Staple food and other food products behave as necessities, while service-oriented categories behave as luxuries. The vice bundle (Alcohol, Tobacco and Narcotics) stands out under the second criterion as a flat-Engel category whose within-bundle BH correlations are tightly clustered near zero, consistent with cross-region variation driven by factors other than income. Figure 2 makes these patterns transparent across income groups. Food shares generally decline with income, while service-oriented shares rise. The 30 richest countries devote the largest expenditure shares to Household Goods and Services and the smallest shares to Staples, while the 30 lowest-income countries devote the smallest shares to Recreation, Culture and Education as well as Health. Some bundles, particularly Alcohol, Tobacco and Narcotics, are less monotone across low and middle income groups.

TABLE I
EMPIRICAL VALIDATION OF THE 12-BUNDLE ICP AGGREGATION

Panel A: Bundle-Level Cross-Region Share Statistics and Engel Pattern								
Bundle	Basic Headings	Mean Share (%)	Min (%)	Max (%)	P10 (%)	P90 (%)	Thin Regions	Engel Correlation
Staples	5	6.43	0.72	30.25	1.66	14.60	16	−0.78
Fruit and Vegetables	5	6.37	0.89	37.87	1.92	13.32	5	−0.65
Meat, Seafood and Dairy	11	9.70	2.06	41.72	4.21	16.37	0	−0.62
Oils, Sweets and Non-Alcoholic Beverages	8	6.44	1.04	19.65	2.81	11.14	1	−0.72
Clothing and Footwear	5	4.76	0.11	21.51	2.40	7.12	5	−0.15
Alcohol, Tobacco and Narcotics	5	3.77	0.14	15.58	1.00	7.54	29	−0.05
Household Goods and Services	13	5.15	0.84	12.57	3.05	7.03	2	+0.20
Health	7	4.03	0.00	21.65	1.34	7.23	21	+0.28
Transport and Communication	16	14.93	3.69	32.55	7.80	20.26	0	+0.39
Recreation, Culture and Education	14	7.12	0.77	15.98	2.80	11.93	6	+0.60
Restaurants, Personal and Other Services	12	13.61	2.11	31.44	5.72	22.73	0	+0.69
Housing and Utilities	8	17.68	3.79	44.95	8.67	26.60	0	+0.60
Panel B: Within-Bundle BH-Level Engel Heterogeneity								
Bundle	Basic Headings	Min Corr.	Max Corr.	Spread	Luxury	Necessity	Flat	
Staples	5	−0.67	+0.07	0.74	0	4	1	
Fruit and Vegetables	5	−0.56	−0.16	0.41	0	4	1	
Meat, Seafood and Dairy	11	−0.51	+0.20	0.71	0	8	3	
Oils, Sweets and Non-Alcoholic Beverages	8	−0.66	+0.19	0.85	0	5	3	
Clothing and Footwear	5	−0.28	+0.09	0.37	0	2	3	
Alcohol, Tobacco and Narcotics	5	−0.18	+0.15	0.32	0	0	5	
Household Goods and Services	13	−0.19	+0.33	0.52	2	0	11	
Health	7	−0.24	+0.48	0.72	3	1	3	
Transport and Communication	16	−0.34	+0.56	0.91	5	2	9	
Recreation, Culture and Education	14	−0.10	+0.67	0.77	11	0	3	
Restaurants, Personal and Other Services	12	−0.10	+0.55	0.44	9	0	3	
Housing and Utilities	8	−0.37	+0.64	1.01	4	1	3	

Notes: Panel A reports cross-region statistics for bundle expenditure shares. “Thin Regions” is the number of regions with bundle share below 1.5%. “Engel Correlation” is the cross-region Pearson correlation between bundle share and log per-capita expenditure. Bundles with correlation ≤ -0.2 ($\geq +0.2$) are classified as necessities (luxuries), and the remainder as flat or mixed. Panel B reports the distribution of basic-heading-level Engel correlations ρ_g within each bundle (cross-region Pearson correlation between the basic-heading share and log per-capita expenditure). “Spread” is the difference between the maximum and minimum ρ_g within a bundle. “Luxury,” “Necessity,” and “Flat” report the number of basic headings with $\rho_g > +0.2$, $\rho_g < -0.2$, and $|\rho_g| \leq 0.2$, respectively.

In Appendix B.6, we consider two coarser ICP aggregations to check whether the main estimates are sensitive to empirically delicate categories.⁷ The 11-bundle specification preserves the 12-bundle structure for the standard consumption categories but excludes Alcohol, Tobacco and Narcotics from the total expenditure. The 10-bundle specification further reduces the parameter dimension. It preserves the main distinction between necessities and luxuries, but combines Alcohol, Tobacco and Narcotics with Recreation, Culture, and Education into a broad discretionary-services category that displays a luxury Engel pattern.⁸

5.3. Expenditure Shares and Aggregate Expenditure

We aggregate the ICP basic-heading expenditure data into 12 commodity and service bundles across regions. The per-capita expenditure share for commodity k in region i is $\Lambda_{i,k} = \Xi_{i,k}/w_i$,

⁷These exercises help determine whether the boundary behavior discussed later reflects a particular fine aggregation or broader features of the cross-region Engel patterns. In both robustness mappings, we keep health as a separate broad good, keep household goods and household services together because the services-only split is thin in many countries, and keep transport and communication together because the communication component displays a weak Engel pattern even though the split is technically feasible.

⁸In parallel, we construct GTAP analogues of the ICP aggregations from the GTAP database (Aguiar et al., 2022) for external robustness exercises and quantitative applications. The GTAP aggregation is based on 65 production sectors rather than ICP basic headings, and its price information is constructed from GTAP purchaser/basic-price wedges normalized within bundle rather than from the ICP GEKS PPP benchmark.




FIGURE 2.—Average expenditure shares across income groups for the twelve consumption bundles in the 2017 ICP estimation sample.

where $\Xi_{i,k}$ is per-capita expenditure on bundle k and w_i is total per-capita expenditure.⁹ We take $w_i = \sum_k \Xi_{i,k}$ directly from the data, so that w_i corresponds to the observed total expenditure entering the indirect relationship $p_{i,k}/w_i$. Similarly, the observed shares $\Lambda_{i,k}$ enter directly into the maximum likelihood objective. Given prices $p_{i,k}$, total expenditure w_i , and the model parameters, we then construct the fitted expenditure shares $\hat{\Lambda}_{i,k}$ using equations (2.2)–(2.4).

5.4. Prices and Summary Statistics

We consult the World Bank ICP Global Office and use the GEKS method for constructing aggregate PPPs above the basic-heading level across regions (The World Bank, 2021). Diewert (2013) provides a comprehensive discussion of the method, which builds on the earlier contributions of Gini (1924, 1931), Éltető and Köves (1964), and Szulc (1964). The GEKS procedure gives a transitive multilateral price system and is invariant to the choice of numeraire, which is useful for a cross-region demand system. We choose the United States as the price numeraire within each bundle. The demand system uses normalized prices $p_{i,k}/w_i$, observed per-capita expenditure shares $\Lambda_{i,k}$, and total per-capita expenditure w_i . Appendix B.1 provides the mathematical details of the GEKS construction. Appendix B.3 reports summary statistics for prices, quantities, expenditures, expenditure shares, and income.

⁹Per-capita expenditure is constructed by dividing aggregate expenditure by population and converting into a common currency using the United States as the reference currency. For consistency, we also use the United States as the numeraire when constructing aggregate prices above the basic-heading level.

6. ESTIMATION: NUMERICAL OPTIMIZATION AND INFERENCE

In estimating the unrestricted CDE specification, we find that several non-homotheticity parameters $\hat{e}_k$ lie at, or are pushed toward, the lower regularity bound $e_k = 0$. This pattern is informative about both the data and the limits of the smooth implicit additive form. At the bottom of the cross-country income distribution, expenditure shares on staple food categories rise sharply as income falls. The unconstrained objective would therefore favor near-zero or slightly negative values of e_k for these categories, which suggests curvature sharper than the smooth regular CDE form can accommodate. A separate bundle, Alcohol, Tobacco and Narcotics, behaves differently. Its share is essentially uncorrelated with per-capita income across regions, and the unconstrained likelihood is unbounded along the negative e_k direction (Section 7.1.2).¹⁰

We adopt three implementation choices that structure the remainder of this section. First, we use a multistart procedure (Section 6.1.1) to verify that the reported optimum is global on the regular domain $\{e_k \geq 0\}$, complemented by a bound-relaxation diagnostic (Section 7.1.1) that probes how each $\hat{e}_k$ moves when the regularity restriction is relaxed. These diagnostics distinguish parameters that are empirically pulled to the regularity boundary from parameters that move to finite negative values under moderate relaxation and from the unbounded-likelihood case. Second, at the boundary points the Wald approximation and the standard nonparametric paired bootstrap are inconsistent (Andrews, 1999, 2000). We thus compute one-sided 90% upper endpoints from an m -out-of- n subsampling bootstrap for boundary-robust confidence intervals on parameters that may sit on the regularity boundary (Section 6.2). Third, for likelihood-ratio tests against the restricted benchmark motivated by Comin, Lashkari, and Mestieri (2021) with the common-substitution restriction $\alpha_k = \alpha$, we use a mixture- χ^2 reference distribution that accounts for boundary parameters in the unrestricted model.

6.1. Numerical Implementation

We implement the constrained maximum likelihood problem in GAMS 47.6 and solve it with the CONOPT4 nonlinear programming solver (Drud, 1985). The estimated parameters are $(\alpha_k, \beta_k, e_k)_{k=1}^N$, the latent utility levels $\{u_i\}_{i=1}^I$, the fitted shares $\{\hat{\Lambda}_{ik}\}$, the auxiliary terms $\{\theta_{ik}\}$, the disturbances $\{d_{ik}\}$, and the upper-triangular Cholesky factor R . The constraints are the implicit defining relationship (2.1), the auxiliary identities (2.3)–(2.4), the residual definition (4.2), the Cholesky relationship (4.3), the upper-triangular structure $r_{km} = 0$ for $m > k$, the parametric restrictions $\beta_k > 0$, $0 \leq \alpha_k \leq 1$ (Section 4.3), $e_k \geq 0$ (the regularity bound), and the normalizations $\sum_k \beta_k = 1$ and $\sum_k e_k = N$ (Section 4.1).¹¹ As a computational check of the admissible normalized-price condition in Proposition 1, we evaluate $\hat{S}_i = \sum_{k \in K_0} \hat{\beta}_k (p_{ik}/w_i)^{1-\hat{\alpha}_k}$ at the baseline CDE estimates for each observed region i . All 178 ICP regions satisfy the required condition, with $\max_i \hat{S}_i = 0.250 < 1$. Thus the observed normalized-price vectors $\xi_i = p_i/w_i$ lie well inside the admissible normalized-price domain $\mathcal{D}$ on which the defining CDE relationship has a unique strictly positive solution $u(\xi)$.

¹⁰As discussed, we keep Alcohol, Tobacco and Narcotics as a separate broad good because its Engel and substitution patterns differ from those of the surrounding categories. In implicitly additive systems such as Hanoch (1975), pooling categories with different demand behavior can distort both substitution and non-homotheticity estimates.

¹¹Appendix C.1 examines the sensitivity of the income and uncompensated price elasticities to the normalization choice and verifies that alternative normalizations of the implicit defining relationship leave the likelihood and implied elasticities unchanged up to numerical precision.

6.1.1. Multistart Search for the Global Constrained Optimum

The constrained likelihood is non-convex and admits multiple local maxima. To verify that the reported optimum is global on the regular domain $\{e_k \geq 0\}$, we draw $S = 100$ random starting values, with each parameter drawn independently from a uniform distribution on the indicated interval (denoted $\mathcal{U}(a, b)$):

$$\begin{aligned} \alpha_k^{(s)} &\sim \mathcal{U}(0.01, 0.30), & \beta_k^{(s)} &\sim \mathcal{U}(0.5, 2.0), \\ e_k^{(s)} &\sim \mathcal{U}(0.1, 3.0), & u_i^{(s)} &= w_i \cdot \mathcal{U}(0.8, 1.2), \end{aligned} \tag{6.1}$$

where the $\beta_k^{(s)}$ and $e_k^{(s)}$ draws are subsequently rescaled to satisfy the normalizations $\sum_k \beta_k = 1$ and $\sum_k e_k = N$. We solve the constrained problem from each starting point and store the parameter vector that achieves the maximum objective value among replications that complete with normal solver status. The auxiliary variables $\hat{\Lambda}_{ik}^{(s)}, d_{ik}^{(s)}, R^{(s)}$ are initialized to satisfy (2.3)–(2.4), (4.2), and (4.3) given the random $(\alpha^{(s)}, \beta^{(s)}, e^{(s)}, u^{(s)})$, so that every start lies in the interior of the feasible region.

We implement the multistart procedure for both the unrestricted CDE and the nested NHCES benchmark. For each specification, we report the maximum objective value found, the number of replications that attain an objective within 10^{-3} of that maximum, and the number that converge to a strictly inferior local optimum. For the NHCES specification, we also vary the starting value of the common α across low and high anchors. Under the equality restriction, the search space is one dimensional in α , so a single sweep can be misleading.¹²

6.2. Boundary-Robust Inference

The unrestricted CDE estimates place several non-homotheticity parameters at the lower regularity bound $e_k = 0$. At the boundary, both Wald-style confidence intervals and the standard nonparametric paired bootstrap are inconsistent. Andrews (1999) shows that the marginal asymptotic distribution of $\sqrt{n}(\hat{e}_k - 0)$ is the projection of a normal random variable onto the feasible cone, which reduces to a half-normal distribution for a single parameter constrained to be nonnegative. Andrews (2000) shows that the nonparametric bootstrap is inconsistent in this setting because the bootstrap analog of the boundary projection mixes draws from above and below zero.

For boundary parameters we report one-sided 90% upper endpoints using an m -out-of- n subsampling bootstrap of Bickel et al. (1997). For a sample of n regions we set $m = \text{round}(n^{0.7})$, the nearest integer to $n^{0.7}$.¹³ For the ICP sample with $n = 178$, this gives $m = 38$. Each replication (i) draws m regions without replacement, (ii) re-estimates the constrained models on the m -region subsample, and (iii) reports the resulting parameter vector $\hat{\theta}^{(b)}$ and demand elasticities.¹⁴ Let $\hat{e}_{k,n}$ denote the full-sample estimate and let $\hat{e}_{k,m}^{(b)}$ denote the corresponding subsample

¹²Appendix C.5 develops a post-estimation latitude test that evaluates how far individual parameters can move while maintaining the fitted likelihood and structural constraints.

¹³The m -out-of- n and subsampling literature uses resample sizes satisfying $m \rightarrow \infty$ and $m/n \rightarrow 0$. See Bickel et al. (1997) and Politis et al. (1999). The power rule $m \asymp n^q$ with $q \in (0, 1)$ satisfies these conditions, and we use $q = 0.7$ as a middle-of-the-range finite-sample tuning choice. Because the choice of m is a tuning issue in finite samples, we also calculate interior-parameter inferences following the estimation. See Bickel and Sakov (2008) on the choice of m in m -out-of- n bootstrap procedures.

¹⁴In the i.i.d. sample, drawing without replacement (subsampling, in the sense of Politis et al. 1999) and with replacement give asymptotically equivalent boundary-robust intervals. We use without replacement for transparency.

estimate. The one-sided 90% upper endpoint for a boundary parameter e_k is

$$U_{k,0.90} = \max \left\{ 0, \hat{e}_{k,n} + \text{quantile}_{0.90} \left[\sqrt{m/n} \left(\hat{e}_{k,m}^{(b)} - \hat{e}_{k,n} \right) \right] \right\}.$$

We report the one-sided interval $[0, U_{k,0.90}]$. As an additional check, we also apply the same m -out-of- n procedure to interior non-homotheticity parameters identified by the bound-relaxation diagnostic in Section 7.1.1. Let k_A denote any such interior category. For these parameters, we compute the one-sided m -out-of- n endpoint

$$U_{k_A,0.90}^{m/n} = \hat{e}_{k_A,n} + \text{quantile}_{0.90} \left[\sqrt{m/n} \left(\hat{e}_{k_A,m}^{(b)} - \hat{e}_{k_A,n} \right) \right],$$

and compare it with the corresponding one-sided 90% endpoint from the ordinary paired n -out-of- n bootstrap,

$$U_{k_A,0.90}^{n/n} = \text{quantile}_{0.90} \left[\hat{e}_{k_A,n}^{*(b)} \right].$$

Because e_{k_A} is strictly interior, the ordinary paired bootstrap provides the appropriate benchmark for inference on this parameter. We use the comparison as a finite sample and implementation check on the subsampling procedure. For interior parameters, the regularity conditions of Andrews (1999) hold under standard arguments, and we report two sided 90% percentile intervals from a separate n out of n paired bootstrap with $B = 2,000$ replications.

6.3. Inference for Nested Tests with Boundary Parameters

The likelihood-ratio (LR) test of the NHCES restriction $\alpha_k = \alpha$ for all k involves $p = N - 1$ equality restrictions on the substitution parameters. Under standard conditions, the reference distribution is $\chi^2(p)$. In the unrestricted CDE specification, some non-homotheticity parameters e_k sit on the boundary $e_k = 0$, which means the likelihood is constrained at a non-interior point along directions *not* being tested. Self and Liang (1987), building on Chernoff (1954), and generalized in Andrews (2001), show that the asymptotic null distribution of the LR statistic in this case is a mixture of χ^2 random variables of varying degrees of freedom, with weights determined by which subset of the boundary parameters is binding under the null. Specifically, if $|B|$ parameters lie on a single one-sided boundary in the unrestricted model and the relevant Hessian block is regular, the LR statistic under the null is asymptotically distributed as

$$LR \rightsquigarrow \sum_{j=0}^{|B|} w_j \chi^2(p-j), \quad w_j = \binom{|B|}{j} \left(\frac{1}{2} \right)^{|B|}. \quad (6.2)$$

under independence of the boundary directions. More general mixture weights for correlated boundary directions are derived in Andrews (2001).¹⁵

The mixture (6.2) has smaller 95th percentile critical values than $\chi^2(p)$ whenever $|B| > 0$, because probability mass shifts toward distributions with fewer degrees of freedom. As a result, any rejection under the naive $\chi^2(p)$ reference also holds under the mixture reference. We report both critical values in Section 7.4.

¹⁵Equation (6.2) gives the standard independence-weight approximation when the boundary parameters are asymptotically independent of one another. We treat equation (6.2) as our main reference distribution.

7. ESTIMATION RESULTS

7.1. Baseline CDE Estimates

We estimate the 12-bundle CDE system in equation (2.1) by constrained maximum likelihood, imposing the regularity restrictions discussed above and using the multistart procedure of Section 6.1.1. The specification is unrestricted across bundles in the sense that each bundle has its own (α_k, β_k, e_k) . Table II reports the point estimates with 90% confidence intervals from the corresponding bootstrap procedures.

TABLE II
ESTIMATED CDE PARAMETERS WITH BOUNDARY-ROBUST CONFIDENCE INTERVALS (12-BUNDLE ICP)

Consumption Bundle	$\hat{\alpha}_k$	$\hat{\beta}_k$	$\hat{e}_k$
Staples	0.372 [0.135, 0.486]	0.011 [0.000, 0.017]	0.000 ^B [0, 0.158] ^{m,B}
Fruit and Vegetables	0.384 [0.192, 0.526]	0.013 [0.000, 0.021]	0.000 ^B [0, 0.158] ^{m,B}
Meat, Seafood and Dairy	0.089 [0.000, 0.436]	0.013 [0.000, 0.025]	0.707 [0.325, 0.909]
Oils, Sweets and Non-Alcoholic Beverages	0.300 [0.029, 0.564]	0.010 [0.000, 0.021]	0.483 [0.000, 0.733]
Clothing and Footwear	0.505 [0.235, 0.744]	0.016 [0.000, 0.033]	0.443 [0.000, 0.799]
Alcohol, Tobacco and Narcotics	0.771 [0.574, 0.948]	0.027 [0.000, 0.119]	0.000 ^U — ^U
Household Goods and Services	0.853 [0.611, 0.980]	0.064 [0.000, 0.401]	0.266 [0.000, 0.989]
Health	0.675 [0.435, 0.875]	0.028 [0.000, 0.063]	1.083 [0.436, 1.871]
Transport and Communication	0.521 [0.284, 0.760]	0.055 [0.000, 0.116]	1.128 [0.756, 1.426]
Recreation, Culture and Education	0.947 [0.851, 1.000]	0.267 [0.004, 0.994]	3.321 [0.000, 4.560]
Restaurants, Personal and Other Services	0.928 [0.833, 0.984]	0.376 [0.002, 0.519]	3.455 [2.493, 7.436]
Housing and Utilities	0.662 [0.459, 0.870]	0.119 [0.000, 0.206]	1.115 [0.429, 1.542]
$\hat{\ell}_{CDE}$	4,127.751		

Notes: Bracketed entries are 90% confidence intervals constructed from a 2,000-replication paired percentile bootstrap for interior parameters. Superscript ^B denotes a boundary estimate whose $\hat{e}_k$ is at the bound when $\underline{e} = 0$ but moves to a moderate negative value under relaxation; superscript ^m denotes the corresponding one-sided m -out-of- n interval $[0, U_{0.90}]$. Superscript ^U denotes a case for which no finite interval is reported because the unconstrained likelihood is unbounded below.

The substitution parameters $\hat{\alpha}_k$ vary substantially across categories. The bootstrap intervals show that this heterogeneity is not driven solely by sampling noise. The non-homotheticity parameters $\hat{e}_k$ run from zero on Staples, Fruit and Vegetables, and Alcohol, Tobacco and Narcotics to above 3 for the two recreation and services categories. The resulting ranking is economically consistent with the distinction between necessities and luxuries.¹⁶

¹⁶Our multistart test confirms that the reported optimum is global on the regular domain. Across 100 random starts, 91 attain an objective within 10^{-3} of the reported maximum and the remaining 9 converge to strictly inferior

Three of the twelve $\hat{e}_k$ are estimated at the lower regularity bound. Very small non-homotheticity estimates of this kind also appear in related implicit additive demand systems. [Comin et al. \(2021\)](#) similarly find agriculture non-homotheticity parameters close to zero under manufacturing normalization in several household and cross-country specifications.¹⁷ The analogy is only partial, because their estimates are near zero rather than boundary estimates, whereas in our constrained problem $\hat{e}_k = 0$ is a regular-domain projection. A zero estimate shows that the lower regularity constraint binds, but does not reveal whether the lower-bound-relaxed likelihood has a finite maximizer slightly below zero or continues to improve as e_k is pushed downward. We investigate these cases in Section 7.1.1.

The resulting one-sided endpoints for the two boundary parameters are small (0.158).¹⁸ These bounds suggest that the regular-domain estimates at zero are not masking large positive non-homotheticity values. They are mild boundary cases in which the unconstrained objective would move only slightly below the regularity bound. As a check on the finite-sample scaling, applying the same procedure to the interior parameter (e.g., transport and communications) gives a one-sided endpoint of 1.337, close to the ordinary paired-bootstrap one-sided 90% endpoint of 1.376 and inside the upper endpoint of the ordinary two-sided 90% interval, 1.426. The subsampling routine behaves as expected for an interior parameter while delivering boundary-robust one-sided intervals for Staples and Fruit and Vegetables. The Alcohol, Tobacco and Narcotics parameter cannot be reliably reported because the bound-relaxation diagnostic (Section 7.1.1) shows its unconstrained likelihood is unbounded along the negative direction, so the regular boundary-bootstrap theory is not the appropriate inference procedure.

7.1.1. The Bound-Relaxation Diagnostic

The baseline regular-domain estimate places three non-homotheticity parameters at the lower bound, $\hat{e}_k = 0$. To investigate what these boundary estimates represent, we solve a sequence of auxiliary problems in which only the lower bound on e_k is relaxed, $\underline{e} \in \{0, -0.1, -0.25, -0.5, -1.0\}$, while all other constraints and the normalization $\sum_k e_k = N$ are held fixed. This procedure traces the direction in which the likelihood moves outside the Hanoch regular domain (Table III and Figure 3).

We partition the twelve bundles into three regimes. *Regime A* has nine bundles whose $\hat{e}_k$ is strictly interior at $\underline{e} = 0$ and stays interior under relaxation. In this regime, the standard regularity conditions of [Andrews \(1999\)](#) hold and standard percentile bootstrap inference applies. *Regime B* has two bundles, Staples and Fresh Fruit and Vegetables, whose $\hat{e}_k$ is at the bound when $\underline{e} = 0$ but moves to a moderate negative value of approximately -0.1 to -0.4 under relaxation and remains well above the lowest relaxed bound. In these mild boundary cases, the unconstrained MLE is finite and slightly negative, but the regularity restriction $e_k \geq 0$ predicts it to zero. We apply boundary inference to this regime and follow the procedure of the one-sided m -out-of- n bootstrap discussed in Section 6.2. *Regime C* has one bundle, Alcohol, Tobacco and Narcotics, whose $\hat{e}_k$ tracks the relaxed lower bound at every step over the tested interval $\underline{e} \in [-1, 0]$. We then conclude that the unrestricted MLE does not admit an economically meaningful finite stationary point. We discuss this case separately in Section 7.1.2.

The aggregate likelihood gain from relaxation is small. The strongest relaxation, $\underline{e} = -1$, raises the likelihood value only by 2.58, equivalent to an LR statistic of 5.15, whereas the

local optima. These local optima have implausible parameter signs, in particular large negative $\hat{e}_k$ values on bundles whose data clearly support positive non-homotheticity, and we exclude them from the reported solution.

¹⁷For example, their CEX final-good expenditure specifications with region and year-quarter fixed effects imply agricultural non-homotheticity parameters of approximately 0.05 and 0.03 under $\epsilon_m = 1$.

¹⁸We report m -out-of- n bootstrap results for interior parameters in Appendix B.5.

TABLE III
BOUND-RELAXATION DIAGNOSTIC FOR THE UNRESTRICTED CDE (12-BUNDLE ICP)

Bundle	Lower Bound on e_k					Regime
	$\underline{e} = 0$	-0.10	-0.25	-0.50	-1.00	
Staples	0.000	-0.100	-0.166	-0.260	-0.242	B
Fruit and Vegetables	0.000	-0.100	-0.250	-0.389	-0.408	B
Meat, Seafood and Dairy	0.707	0.689	0.660	0.622	0.561	A
Oils, Sweets and Non-Alcoholic Beverages	0.483	0.458	0.417	0.366	0.301	A
Clothing and Footwear	0.443	0.416	0.364	0.292	0.215	A
Alcohol, Tobacco and Narcotics	0.000	-0.100	-0.250	-0.500	-1.000	C
Household Goods and Services	0.266	0.289	0.239	0.104	-0.043	A
Health	1.083	1.038	0.956	0.845	0.717	A
Transport and Communication	1.128	1.098	1.050	0.986	0.896	A
Recreation, Culture and Education	3.321	3.604	4.062	4.631	6.055	A
Restaurants, Personal and Other Services	3.455	3.636	4.007	4.414	4.193	A
Housing and Utilities	1.115	1.072	0.994	0.888	0.756	A
$\hat{\ell}_{\text{CDE}}$	4127.751	4128.703	4129.453	4129.849	4130.327	
$\Delta \hat{\ell}$ vs. $\underline{e} = 0$	0.000	0.952	1.702	2.098	2.576	

Notes: Each column reports the constrained MLE of e_k from an auxiliary problem with lower bound $\underline{e}$, holding all other constraints and the normalization $\sum_k e_k = N$ fixed. Regime A denotes parameters that remain interior; regime B denotes parameters that bind at zero but stabilize at finite negative values under relaxation; regime C denotes a parameter that tracks the imposed bound. The final rows report the log likelihood and its gain relative to $\underline{e} = 0$.




FIGURE 3.—Bound-relaxation diagnostics for the 12-bundle CDE estimates. Each panel plots the constrained MLE of $\hat{e}_k$ as the lower bound on e_k is relaxed. Regime A shows interior cases, Regime B shows mild boundary cases, Staple Foods and Fruit & Vegetables, and Regime C shows the unbounded case, Alcohol, Tobacco & Narcotics.

LR statistic for the common-substitution restriction in Section 7.4 is 46.4. Maintaining global regularity costs a small fraction of the likelihood reduction the data already deliver against the restricted benchmark.

7.1.2. A Flat Engel Category and the Limits of Smooth Demand Forms

The bound-relaxation diagnostic isolates one bundle, Alcohol, Tobacco and Narcotics, for which the smooth CDE form appears misspecified relative to the data. The cross-region Engel correlation for this bundle is approximately -0.05 , neither systematically necessity like

nor systematically luxury like, and the likelihood continues to improve as the lower bound on e_{AlcTob} is relaxed over the tested range. Unlike the regime B cases, the trajectory does not stabilize at a moderate negative value. We thus report the regular domain projection $\hat{e}_{\text{AlcTob}} = 0$ without a finite boundary interval.

The substantive interpretation is that vice demand is shaped by tax, regulation, religion, and addiction in ways that do not vary smoothly with per capita income across countries. The point is not that this category has no demand response. A more reasonable interpretation is that a smooth globally regular demand system should not be expected to explain all of its cross-region variation through income alone. The smooth implicit additive functional form therefore cannot represent these dynamics for this category, and regrouping basic headings does not eliminate the issue. We provide a direct robustness test in Section 7.3.2, where we re-estimate the system on an 11 bundle aggregation that drops the Alcohol, Tobacco and Narcotics bundle entirely, and on a 10 bundle aggregation that absorbs it into a broader residual Leisure, Vices, and Other Services category. In both robustness specifications, the main results for the remaining bundles remain stable. At the same time, we show that problematic categories cannot simply be absorbed away without affecting the remaining system.¹⁹ We keep the 12 bundle specification because separating vice from non-vice consumption is the substantively correct grouping.

7.2. Restricted Benchmark: NHCES

To quantify the extent to which the data require good-specific substitution behavior, we estimate the nested benchmark that restricts the substitution parameters to be common across categories, $\alpha_k = \alpha$ for all k . Under this restriction, the defining relationship reduces to the implicit non-homothetic CES form in (4.4). We estimate this restricted benchmark on the same sample and under the same normalization and global-regularity constraints as in the CDE case. Table IV reports the corresponding restricted-benchmark estimates with two-sided percentile and one-sided m -out-of- n bootstrap intervals as appropriate.

The common substitution parameter is estimated at $\hat{\alpha} = 0.584$, with a relatively tight confidence interval. Staple foods and Fresh Fruit and Vegetables lie at or near the lower end of the non-homotheticity hierarchy, while Housing, Health, Transport and Communication, Recreation, Culture and Education, and Restaurants and Other Services show larger e_k estimates. Meanwhile, the restriction changes the shape of the hierarchy relative to the unrestricted CDE. The highest CDE values for Recreation, Culture and Education and for Restaurants and Other Services are pulled down substantially under NHCES, while several middle categories show larger positive e_k values. Thus the restricted benchmark preserves the broad necessity-luxury pattern, but it also reallocates the dispersion in non-homotheticity across categories.

The flat-Engel category provides the clearest example. In the unrestricted CDE, the constrained estimate is $\hat{e}_{\text{AlcTob}} = 0$. The bound-relaxation results show that the relaxed estimates continue to track the imposed lower bound, so the system does not yield a reliable interior estimate or confidence interval for this bundle. In the restricted NHCES, the same category yields an interior estimate, $\hat{e}_{\text{AlcTob}} = 0.840$, with a 90% interval $[0.427, 1.232]$. The Engel evidence in Table I points in the other direction. The bundle-level Engel correlation and the within-bundle basic-heading correlations are both close to zero. A large positive income-response estimate therefore has little support in the expenditure data.

¹⁹Tables B.IV and B.V show that the estimated parameters for the remaining bundles are quantitatively similar across the three specifications. Dropping the vice bundle removes the regime C result entirely, while absorbing it into a broader discretionary services residual produces an interior estimate for the merged category ($\hat{e}_{\text{resid}} = 3.84$) that is driven primarily by the recreation and restaurants components.

TABLE IV

ESTIMATED NHCES PARAMETERS WITH BOUNDARY-ROBUST CONFIDENCE INTERVALS (12-BUNDLE ICP)

Consumption Bundle	$\hat{\alpha}$ (common)	$\hat{\beta}_k$	$\hat{e}_k$
Staples	0.584 [0.517, 0.622]	0.043 [0.038, 0.047]	0.000 ^B [0, 0.060] ^{m,B}
Fruit and Vegetables	0.584 [0.517, 0.622]	0.046 [0.042, 0.052]	0.008 [0.000, 0.221]
Meat, Seafood and Dairy	0.584 [0.517, 0.622]	0.079 [0.073, 0.083]	0.594 [0.421, 0.753]
Oils, Sweets and Non-Alcoholic Beverages	0.584 [0.517, 0.622]	0.046 [0.043, 0.049]	0.457 [0.334, 0.586]
Clothing and Footwear	0.584 [0.517, 0.622]	0.044 [0.041, 0.048]	0.781 [0.521, 1.012]
Alcohol, Tobacco and Narcotics	0.584 [0.517, 0.622]	0.032 [0.027, 0.036]	0.840 [0.427, 1.232]
Household Goods and Services	0.584 [0.517, 0.622]	0.051 [0.048, 0.053]	1.263 [1.122, 1.387]
Health	0.584 [0.517, 0.622]	0.053 [0.048, 0.058]	1.492 [1.193, 1.895]
Transport and Communication	0.584 [0.517, 0.622]	0.146 [0.141, 0.152]	1.465 [1.318, 1.605]
Recreation, Culture and Education	0.584 [0.517, 0.622]	0.082 [0.078, 0.087]	1.722 [1.521, 1.895]
Restaurants, Personal and Other Services	0.584 [0.517, 0.622]	0.151 [0.144, 0.157]	1.946 [1.724, 2.146]
Housing and Utilities	0.584 [0.517, 0.622]	0.226 [0.216, 0.239]	1.432 [1.095, 1.582]
$\hat{\ell}_{\text{NHCES}}$		4,104.547	

Notes: Bracketed entries are 90% confidence intervals constructed from a 2,000-replication paired percentile bootstrap for interior parameters. Superscript ^B denotes a boundary estimate whose $\hat{e}_k$ is at the bound when $\underline{e} = 0$ but moves to a moderate negative value under relaxation; superscript ^m denotes the corresponding one-sided m -out-of- n interval $[0, U_{0.90}]$.

The source of the difference is the common substitution restriction. In the unrestricted CDE, the substitution parameter α_{AlcTob} is free and the estimator places it at a relatively high value. This allows the substitution channel to absorb much of the cross-region variation in expenditure shares and leaves little income-gradient information with which to identify a meaningful e_k . In the restricted NHCES, the equality restriction $\alpha_k = \alpha$ closes this channel. The remaining variation must pass through e_{AlcTob} , which produces a moderate positive estimate even though the empirical Engel evidence points to a nearly flat income response.²⁰

The boundary result identifies a category for which the smooth regular demand system fits the data poorly. The restricted estimate is mechanically cleaner, but less plausible in economic terms. The unrestricted system reveals this limitation because it allows substitution parameters to vary across categories. The common-substitution restriction rules out this margin of adjustment. This result motivates a closer look at the income elasticities implied by the two systems.

²⁰The CDE lets α_k vary by bundle and estimates a relatively high $\hat{\alpha}_{\text{AlcTob}} = 0.771$. NHCES sets $\alpha_k = \alpha$ and estimates $\hat{\alpha} = 0.584$. Since utility enters the share equation through $e_k(1 - \alpha_k)$, the common substitution restriction closes one channel that CDE uses to handle category-specific price and substitution variation. The restricted model then uses e_{AlcTob} to absorb the observed price and share variation.

7.3. Income Elasticities

7.3.1. Unrestricted CDE Elasticities

Figure 4 shows substantial heterogeneity across consumption bundles under the unrestricted CDE system. Staple Foods and Fruit and Vegetables are necessities throughout the expenditure distribution. Their elasticities decline from roughly $\eta \approx 0.56$ to 0.49 for staples and from about 0.58 to 0.50 for fruit and vegetables as per-capita expenditure rises. Meat, Seafood and Dairy and Oils, Sweets and Non-Alcoholic Beverages begin as luxuries at low expenditure levels and become necessities at higher expenditure levels. Clothing and Footwear shows a similar, though less central, transition across unit elasticity. Several other categories converge toward unit elasticity at high expenditure levels. Housing and Utilities, Household Goods and Services, Health, and Transport and Communication are close to one at the top of the expenditure distribution, while Recreation, Culture and Education and Restaurants, Personal and Other Services remain luxury demands throughout. Alcohol, Tobacco and Narcotics is reported under the regular-domain projection $\hat{e}_k = 0$, the regime-C boundary case discussed in Section 7.1.2. Its elasticity schedule is nearly flat, consistent with the weak cross-region Engel correlation observed in the data. The confidence intervals are noticeably asymmetric for categories whose $\hat{e}_k$ estimates lie at the lower regularity bound, namely Staples, Fruit and Vegetables, and Alcohol, Tobacco and Narcotics. This asymmetry reflects the one sided sampling distribution induced by the regularity boundary and the corresponding m out of n bootstrap correction. The intervals for interior parameters are much more symmetric.

7.3.2. Alternative Aggregations

Appendix B.6 reports two coarser ICP aggregations. An 11-bundle system drops Alcohol, Tobacco and Narcotics, and a 10-bundle system absorbs it into a broader residual services category. The common bundles are stable across these exercises. Staples and Fresh Fruit and Vegetables remain necessities, service-oriented categories remain luxuries, and the CDE-versus-NHCS likelihood comparison continues to reject the common-substitution restriction. These exercises also show that the 12-bundle aggregation preserves economically distinct Engel behavior consistent with the aggregation validation. The two coarser aggregations distort this information in different ways. Dropping the flat-Engel category changes the economic interpretation of the system and widens the low-income confidence intervals for the boundary necessity categories, while pooling it with high-elasticity service categories hides the vice pattern rather than explaining it.

7.3.3. CDE versus NHCS

We next compare the baseline 12-bundle CDE schedules to those implied by the restricted NHCS specification, in which the substitution parameters are constrained to a common α . To isolate non-homothetic demand responses from cross-region price dispersion, we evaluate income elasticities at a common reference price vector. Holding the estimated parameter vector fixed, we compute fitted utilities and expenditure shares at the reference prices and then evaluate income elasticities using equation (2.6) at each region's observed per capita expenditure level.

Figure 5 compares the unrestricted CDE elasticities to those implied by the restricted NHCS system, in which the substitution parameters are constrained to a common α .²¹ Staple Foods and Fruit and Vegetables are nearly flat under NHCS, with elasticities close to $\eta \approx 0.58$ throughout the expenditure distribution, whereas the corresponding CDE elasticities decline modestly

²¹We report the NHCS figure with its own per-panel scaling in Appendix B.4.




FIGURE 4.—Income elasticities under the unrestricted CDE specification (12 ICP bundles). Each panel plots the model-implied income elasticity against per-capita expenditure at a common reference price vector. Point estimates are shown with 90% paired-bootstrap confidence bounds. For parameters whose point estimate sits at the lower regularity bound, the bounds are constructed from the m -out-of- n subsampling bootstrap.

with expenditure. The unrestricted CDE allows additional category-specific curvature through heterogeneity in α_k , while NHCES removes this channel through the common α restriction. The NHCES elasticities are steeper at low expenditure levels for several service-oriented categories. Recreation, Culture and Education declines from approximately $\eta \approx 1.77$ to 1.09 under NHCES, compared with 1.44 to 1.17 under CDE. Restaurants, Personal and Other Services declines from approximately 1.93 to 1.16 under NHCES, compared with 1.54 to 1.19 under CDE. With α constrained to be common across categories, more of the cross-category Engel heterogeneity is absorbed by the e_k parameters.

Alcohol, Tobacco and Narcotics moves in the opposite direction. Under NHCES, the elasticity declines from approximately $\eta \approx 1.16$ to 0.83, whereas the CDE projection at $\hat{e}_k = 0$ is nearly flat, declining only from about 0.96 to 0.89. The cross-region Engel correlation for this bundle is approximately -0.05 . The flat CDE pattern therefore aligns more closely with the underlying expenditure data, whereas the NHCES result appears to reflect the common α

restriction rather than a strong income response in this category. For the remaining categories, the two systems are relatively similar at high expenditure levels. Most of the disagreement appears at low expenditure levels, particularly below approximately \$300 in per capita expenditure. These results point against the restricted NHCES specification. Section 7.4 formalizes the comparison through a likelihood ratio test of the common α restriction.




FIGURE 5.—Income elasticities under the unrestricted CDE and restricted NHCES specifications (12 ICP bundles). Each panel overlays the CDE point estimate with its 90% bootstrap confidence interval (black line with circle markers, gray band) and the NHCES point estimate with its 90% bootstrap confidence interval (red line with diamond markers, red band). Bootstrap inference follows Figure 4.

7.4. Model Fit: Likelihood Ratio with Mixture Critical Values

We evaluate the common substitution restriction using the concentrated objective values under the two nested specifications. With $\hat{\ell}_{\text{CDE}} = 4,127.751$ and $\hat{\ell}_{\text{NHCES}} = 4,104.547$, the likelihood ratio statistic is

$$LR = 2(\hat{\ell}_{\text{CDE}} - \hat{\ell}_{\text{NHCES}}) = 46.41.$$

The naive reference distribution is $\chi^2(p)$ with $p = N - 1 = 11$, which gives a 95th percentile critical value of approximately 19.68. In the unrestricted CDE, two non homotheticity parameters, Staples and Fresh Fruit and Vegetables, are classified as regime B boundary parameters. The regime C Alcohol, Tobacco and Narcotics parameter is excluded for the reasons discussed in Section 7.1.2. Under the standard independence weight approximation, the Self-Liang/Andrews mixture reference distribution is

$$\frac{1}{4}\chi_{11}^2 + \frac{1}{2}\chi_{10}^2 + \frac{1}{4}\chi_9^2,$$

with a 95th percentile critical value of approximately 18.41.

The observed statistic exceeds both reference values by a wide margin. The data therefore reject the common substitution restriction under either reference distribution.

TABLE V
LIKELIHOOD-RATIO TEST OF THE COMMON-SUBSTITUTION RESTRICTION (12-BUNDLE ICP)

	Maximized objective $\hat{\ell}$	Free substitution params
NHCES ($\alpha_k = \alpha$)	4,104.547	1
CDE (α_k free by good)	4,127.751	12
$LR = 2(\hat{\ell}_{\text{CDE}} - \hat{\ell}_{\text{NHCES}})$		46.41
Naive $\chi^2(p)$ critical value (95%), $p = N - 1 = 11$		19.68
Regime-B boundary count used in mixture, $ B = 2$		Staples, FruitVeg
Self-Liang/Andrews mixture critical value (95%)		18.41
Naive $\chi^2(11)$ p -value		2.7×10^{-6}
Mixture p -value		1.4×10^{-6}

Notes: The naive reference uses $\chi^2(11)$. The boundary-adjusted reference uses the standard Self-Liang/Andrews mixture with $p = 11$ and $|B| = 2$.

7.5. An Engel-Slope Benchmark with Price Controls

The likelihood comparison rejects the common-substitution restriction. We next ask whether that rejection shows up in the data. The two structural systems imply income elasticities, $\hat{\eta}_k^m(w)$, for $m \in \{\text{CDE}, \text{NHCES}\}$. We compare the implied elasticities to a local Engel slope estimated directly from prices, quantities, and expenditures observed in the ICP data.

For bundle k in region i , let q_{ik} denote the observed per-capita quantity index, w_i per-capita expenditure, and p_{ik} the bundle price. Define the price index for all other bundles by

$$\log P_{i,-k} = \frac{1}{N-1} \sum_{j \neq k} \log p_{ij}. \quad (7.1)$$

At each rank position ℓ in the income distribution, we estimate²²

$$\log q_{ik} = a_{k\ell} + \hat{\eta}_{k\ell}^L \log w_i + \gamma_{1k\ell} \log p_{ik} + \gamma_{2k\ell} \log P_{i,-k} + \varepsilon_{ik}, \quad i \in \mathcal{W}_\ell. \quad (7.2)$$

²²The index ℓ denotes a rank position in the expenditure distribution. The local sample includes the 20 nearby regions above and below each interior rank position, so each interior regression uses approximately 41 regions in the expenditure ranking. Thus $\hat{\eta}_{k\ell}^L$ is the local Engel slope for bundle k estimated around expenditure rank position ℓ .

The coefficient on $\log w_i$ is the local price-controlled Engel slope. We average these local slopes within income bins and compare them with the CDE and NHCES elasticities evaluated at the same expenditure levels. Economically, this asks how quantities move with expenditure after controlling for own prices and an aggregate price index for the rest of consumption.

For each system and bundle, define

$$\text{MSE}_k^m = \frac{1}{B} \sum_{b=1}^B (\hat{\eta}_{kb}^L - \hat{\eta}_k^m(\bar{w}_b))^2, \quad (7.3)$$

where $\hat{\eta}_{kb}^L$ is the average local Engel slope for bundle k within income bin b and $\bar{w}_b$ is the corresponding mean expenditure. We calculate

$$\Delta \text{MSE}_k = \text{MSE}_k^{\text{NHCES}} - \text{MSE}_k^{\text{CDE}}. \quad (7.4)$$

Thus positive values favor CDE relative to the Engel-slope benchmark, while negative values favor NHCES. Figure 6 and Table VI show that nearly all categories favor the unrestricted CDE specification, and every interval excludes zero. The largest differences appear for Alcohol, Tobacco, and Narcotics; Restaurants, Personal, and Other Services; Meat, Seafood, and Dairy; and Recreation, Culture, and Education. Six additional categories also favor CDE, although by smaller margins. Health and Clothing and Footwear favor NHCES, but the differences are economically small. The pooled comparison points in the same direction, with an overall ΔMSE of 0.016 and a 90% interval of $[0.014, 0.018]$.




FIGURE 6.—Fit to the Engel-slope benchmark with price controls. Whiskers are 90% bootstrap intervals.

This comparison provides an additional empirical margin on which to compare the two systems, beyond the Engel-pattern evidence that motivated the aggregation and beyond the boundary-robust likelihood-based inference. On this margin, the unrestricted CDE more closely matches the reduced-form Engel-slope evidence for most bundles, including the categories where the income elasticities differ most sharply from the restricted NHCES benchmark. The additional flexibility in the unrestricted system is therefore not just theoretical. It can be traced directly to the cross-region expenditure data.

TABLE VI
CATEGORY-LEVEL FIT TO THE ROLLING ENGEL-SLOPE BENCHMARK

Category	MSE (CDE)	MSE (NHCES)	Δ MSE	90% Interval	Preferred
Alcohol, Tobacco, and Narcotics	0.379	0.426	0.047	[0.037, 0.058]	CDE
Restaurants, Personal, and Other Services	0.074	0.115	0.041	[0.026, 0.056]	CDE
Meat, Seafood, and Dairy	0.037	0.077	0.040	[0.032, 0.048]	CDE
Recreation, Culture, and Education	0.072	0.101	0.029	[0.022, 0.038]	CDE
Oils, Sweets, and Non-Alcoholic Beverages	0.040	0.051	0.011	[0.007, 0.014]	CDE
Transport and Communication	0.088	0.099	0.010	[0.008, 0.012]	CDE
Household Goods and Services	0.027	0.038	0.010	[0.005, 0.016]	CDE
Fresh Fruit and Vegetables	0.127	0.136	0.009	[0.005, 0.013]	CDE
Staple Foods	0.022	0.028	0.007	[0.005, 0.009]	CDE
Housing and Utilities	0.063	0.069	0.005	[0.004, 0.007]	CDE
Health	0.171	0.166	-0.005	[-0.007, -0.003]	NHCES
Clothing and Footwear	0.125	0.111	-0.014	[-0.017, -0.012]	NHCES

Notes: Categories are ordered by descending Δ MSE. Intervals are 90% bootstrap intervals.

8. STRUCTURAL-CHANGE BENCHMARKS AND STRUCTURED PANEL ESTIMATION

8.1. A Three-Sector Structural-Change Benchmark

We next ask whether the estimation procedure delivers economically meaningful results in a simple setting where the broad patterns are already well known in the macroeconomic literature (e.g., [Boppart, 2014](#)). We aggregate the ICP data to the same three sectors used by [Comin et al. \(2021\)](#), agriculture, manufacturing, and services.²³ Our goal is not to match their estimates coefficient by coefficient. Instead, we ask whether the cross-sectional procedure developed in this paper captures the central implication in their analysis, that income effects remain quantitatively important across large income differences, with agriculture exhibiting the lowest income elasticity, services the highest, and manufacturing lying between them.

We identify which parameters are directly comparable across different specifications. In [Comin et al. \(2021\)](#), the non-homotheticity parameters are normalized relative to manufacturing, so their coefficients of interest are $\epsilon_A/\epsilon_M - 1$ and $\epsilon_S/\epsilon_M - 1$, and the coefficient on the real-consumption index equals $(1 - \sigma)(\epsilon_i/\epsilon_M - 1)$. Under our restricted NHCES specification, the equivalent terms are $e_A/e_M - 1$ and $e_S/e_M - 1$, and the corresponding coefficients on the real-consumption index are $(1 - \alpha)(e_A/e_M - 1)$ and $(1 - \alpha)(e_S/e_M - 1)$. The unrestricted CDE is different because the substitution parameter is allowed to vary across sectors, so the system does not in general map into a single coefficient on the real-consumption index. We thus report both the primitive CDE parameters and the implied income elasticities. Doing so isolates two questions that are otherwise easy to conflate, whether the estimation procedure generates the same broad Engel patterns and whether those patterns require the common-substitution structure imposed by the NHCES specification.

8.2. External Validation

Table [VII](#) presents the results of the external validation. Using a different data source and a different aggregation structure, the constrained-likelihood NHCES estimate is close to the published aggregate benchmark. The common substitution parameter is 0.516, compared with 0.50 in their employment specification with trade controls. The normalized non-homotheticity

²³[Comin et al. \(2021\)](#) use a postwar panel of 39 countries with sectoral employment, value added, sectoral prices, and consumption per capita from the Groningen 10-Sector Database and Penn World Tables.

parameters are -0.701 for agriculture and 0.250 for services, compared with -0.89 and 0.21 . The ranking is the same, and the magnitudes are sufficiently close to rule out the concern that the constrained-likelihood procedure mechanically compresses non-homotheticity.

TABLE VII
THREE-SECTOR ICP BENCHMARK AGAINST THE CLM STRUCTURAL-CHANGE ESTIMATES

Panel A: CLM Benchmark and Three-Sector Constrained-Likelihood Estimates						
Specification	Substitution Parameter(s)			$A/M - 1$	$S/M - 1$	Mean Income Elasticities (A, M, S)
CLM, World, With Trade Controls	$\sigma = 0.50$			-0.89	0.21	(0.56, 1.03, 1.14)
ICP 3-Sector NHCES, Constrained MLE	$\alpha = 0.516$			-0.701	0.250	(0.687, 1.087, 1.230)
ICP 3-Sector CDE, Constrained MLE	$\alpha_A = 0.006$	$\alpha_M = 0.940$	$\alpha_S = 0.980$	n.d.	n.d.	(0.635, 1.109, 1.194)
Panel B: Primitive CDE Estimates for the ICP 3-Sector Specification						
Sector	$\widehat{\alpha}_k$		$\widehat{\beta}_k$	$\widehat{e}_k$	Mean Income Elasticities	
Agriculture	0.006		0.007	0.514	0.635	
Manufacturing	0.940		0.188	0.000	1.109	
Services	0.980		0.805	2.486	1.194	

Notes: A , M , and S denote agriculture, manufacturing, and services. CLM refers to [Comin et al. \(2021\)](#). The CLM row reports their aggregate employment specification with trade controls. The ICP rows use the 2017 ICP cross-section aggregated to the same three sectors. Under NHCES, $A/M - 1$ and $S/M - 1$ are $\hat{e}_A/\hat{e}_M - 1$ and $\hat{e}_S/\hat{e}_M - 1$. Under CDE, these values are not reported because manufacturing yields a boundary estimate.

The unrestricted CDE preserves the economic ordering but changes the primitive representation. The point estimate is strongly non-CES in substitution, with $\hat{\alpha}_A = 0.006$, $\hat{\alpha}_M = 0.940$, and $\hat{\alpha}_S = 0.980$. The manufacturing expansion parameter lies at the lower boundary, so the manufacturing-base ratios are not defined. We therefore do not compare the unrestricted CDE to the structural-change benchmark through these ratios. The comparison instead proceeds through the implied average income elasticities. These elasticities are close to the NHCES and benchmark averages, and they preserve the same broad pattern in which agriculture remains below one, services remain above one, and manufacturing lies in the middle.

Figure 7 compares the implied income elasticities across the two systems. Near the center of the income distribution, the two systems give similar sectoral elasticities, consistent with Table VII. The similarity in averages, however, does not reflect a difference in how curvature is distributed over the income range. In agriculture, the CDE elasticity falls sharply as expenditure rises, from about 0.85 at low expenditure to about 0.42 at high expenditure. The NHCES elasticity is much flatter, falling from about 0.76 to about 0.65. In manufacturing, the CDE elasticity stays above one and moves gradually toward unit elasticity at high expenditure. The NHCES elasticity starts much higher at low expenditure and then falls more steeply.²⁴ The service sector shows a similar pattern. NHCES implies a much stronger low expenditure service response, while the two systems nearly converge at high expenditure. The same qualitative pattern appears in the 12-bundle estimates, where the unrestricted system places more low-income curvature in necessity-oriented categories and less in service-oriented categories.

The two systems do not disagree about the broad income-demand pattern across sectors. Both systems imply the same broad movement away from agriculture and toward services as income rises. The difference is where the adjustment enters. NHCES puts more of the low income curvature into manufacturing and services. CDE places more of the curvature in agriculture

²⁴ Along the sequence of high-objective solutions through $\underline{e} = -3$, the income elasticities remain stable within ± 0.05 at all expenditure levels. The optimizer switches to lower-objective local solutions only when the lower bound is relaxed further. We thus treat $\hat{e}_M = 0$ as a primitive-parameter boundary under the normalization, not as evidence that the manufacturing income response is flat.




FIGURE 7.—Three-sector income elasticities for CDE and NHCES evaluated at mean prices. Shaded areas are 90% bootstrap intervals. Manufacturing CDE uses the m-out-of-n bootstrap; all others use paired bootstrap. Rug marks show the distribution of the 178-region income data, with darker marks indicating income deciles.

and leaves non-agricultural elasticities smoother over the income distribution. This matters for counterfactuals that focus on low-income and lower-middle-income economies, since those are the observations where the two systems differ most.

The potentially larger difference appears in substitution responses. Under the unrestricted CDE, the implied Allen–Uzawa substitution elasticities are $\sigma_{MS} = 1.268$ between manufacturing and services, $\sigma_{AS} = 0.334$ between agriculture and services, and $\sigma_{AM} = 0.294$ between agriculture and manufacturing. Relative to the unit-elastic Cobb–Douglas benchmark, manufacturing and services are gross substitutes ($\sigma > 1$), while agriculture is a gross complement to both other sectors ($\sigma < 1$). The NHCES forces all three substitution margins to equal the common parameter $\alpha = 0.516$, which places every pair at the same substitution margin.

Even at points in the income distribution where the two systems imply similar income elasticities, they imply very different substitution responses to relative-price changes. Consider a Baumol-style counterfactual in which services become more expensive relative to manufacturing. NHCES forces the service response to be governed by the same substitution parameter that governs the agriculture margins. CDE allows the manufacturing-services margin to be much more price sensitive ($\sigma_{MS} > 1$), while movements out of agriculture remain governed more by income and by the low substitutability of agriculture with the other sectors ($\sigma_{Ak} < 1$). The same non-homothetic hierarchy can therefore imply a different decomposition of structural change once relative prices move.

8.3. A Constrained Panel Estimator with Country-Sector Fixed Effects

We develop a constrained panel estimator that extends the cross-sectional framework to a panel setting. We take the country-year panel used by [Comin et al. \(2021\)](#) and estimate the restricted NHCES system with country-sector fixed effects inside the same constrained-likelihood framework used throughout this paper. If the estimator recovers economic structure rather than fitting a particular cross section, its panel implementation should reproduce the authors’ benchmark when applied to the same data under comparable identifying assumptions.

Our results confirm this prediction. The constrained NHCES panel estimator reproduces the benchmark of [Comin et al. \(2021\)](#) through a different estimation procedure. The authors estimate the restricted system as a constrained log-linear panel by generalized method of moments (GMM). We estimate the same economic restrictions as a constrained quasi-likelihood problem in which latent utilities, fitted shares, and country-sector fixed effects are solved jointly. The

two approaches deliver the same sectoral ordering and similar magnitudes. The constrained-likelihood framework is therefore not specific to the ICP cross section. It can be implemented as a panel estimator that uses within-country movements in prices, expenditure, and sectoral allocation while imposing the structural demand restrictions directly.

8.3.1. Data and Panel Specification

Let c index countries, t years, and $k \in \mathcal{K} = \{A, M, S\}$ sectors (agriculture, manufacturing, and services), with manufacturing as the normalization sector. [Comin et al. \(2021\)](#)'s panel combines sectoral employment, value added, and prices from the Groningen Ten-Sector Database with aggregate consumption expenditure per capita from the Penn World Tables. We use the same country-year panel and the same three-sector aggregation. The baseline specification without trade controls contains 1,626 country-year observations across 39 countries.

For each country-year, L_{ctk} denotes employment in sector k , p_{ctk} the sectoral price, and E_{ct} per-capita consumption expenditure. The panel estimator uses manufacturing-normalized variables,

$$\tilde{p}_{ctM} = 1, \quad \tilde{p}_{ctk} = \frac{p_{ctk}}{p_{ctM}}, \quad \tilde{E}_{ct} = \frac{E_{ct}}{p_{ctM}}, \quad (8.1)$$

which match the normalization in [Comin et al. \(2021\)](#) and identify the non-homotheticity parameters relative to manufacturing. We thus impose

$$e_M = 1, \quad (8.2)$$

and report the non-homotheticity parameters as $e_A/e_M - 1$ and $e_S/e_M - 1$.

8.3.2. The CLM Benchmark Equation

The aggregate equation in [Comin et al. \(2021\)](#) follows from the restricted NHCES share system. For each non-base sector $k \in \{A, S\}$,

$$\log \frac{L_{ctk}}{L_{ctM}} = \zeta_{ck} + (1 - \sigma) \log \frac{p_{ctk}}{p_{ctM}} + (1 - \sigma) \left(\frac{\epsilon_k}{\epsilon_M} - 1 \right) \log \frac{E_{ct}}{p_{ctM}} + \left(\frac{\epsilon_k}{\epsilon_M} - 1 \right) \log \omega_{ctM} + \nu_{ctk}, \quad (8.3)$$

where σ is the common substitution elasticity, $\epsilon_k/\epsilon_M - 1$ the relative non-homotheticity parameters, ζ_{ck} a country-sector fixed effect, and ν_{ctk} a residual relative demand shock. The fixed effects absorb time-invariant relative tastes, sectoral labor-share wedges, and country-specific technology; the parameters are then identified by within-country movements in relative employment, relative prices, manufacturing-deflated expenditure, and the manufacturing share. The estimator asks whether countries that become richer, holding country-sector levels and relative prices fixed, reallocate away from agriculture and toward services in the way NHCES implies. The authors estimate (8.3) as a constrained system, with a common price coefficient across the two equations and the share and expenditure coefficients tied to $(\sigma, \epsilon_A/\epsilon_M, \epsilon_S/\epsilon_M)$.

8.3.3. Constrained Panel Estimator

We impose the same NHCES structure inside a constrained optimization problem. For each country-year, the latent utility u_{ct} satisfies the indirect NHCES relationship,

$$\sum_{j \in \mathcal{K}} \beta_j u_{ct}^{e_j(1-\alpha)} \left(\frac{\tilde{p}_{ctj}}{\tilde{E}_{ct}} \right)^{1-\alpha} = 1, \quad (8.4)$$

and the fitted share is constructed from the auxiliary term

$$\theta_{ctj} = \beta_j(1 - \alpha)u_{ct}^{e_j(1-\alpha)} \left(\frac{\tilde{p}_{ctj}}{\tilde{E}_{ct}} \right)^{1-\alpha}, \quad \hat{s}_{ctj} = \frac{\theta_{ctj}}{\sum_{\ell \in \mathcal{K}} \theta_{ct\ell}}. \quad (8.5)$$

The latent utilities, auxiliary terms, and fitted shares are not computed in an inner loop. They enter jointly as variables in a single nonlinear program, with equations (8.4) and (8.5) imposed as constraints, exactly as in the cross-sectional constrained-likelihood formulation of Section 4.2.

The panel residual lives in log relative-employment space. For $k \in \{A, S\}$,

$$r_{ctk} = \log \frac{L_{ctk}}{L_{ctM}} - \delta_k - \zeta_{ck} - \log \frac{\hat{s}_{ctk}}{\hat{s}_{ctM}}, \quad (8.6)$$

with sector intercepts δ_k and country-sector fixed effects ζ_{ck} normalized by $\sum_c \zeta_{ck} = 0$ for each k . The normalization is numerically convenient and economically innocuous. It gives the same fitted residuals as a full set of country-sector fixed effects. The objective is a Gaussian quasi-likelihood for the log relative residuals. Under homoskedastic equation errors the point estimates solve

$$\hat{\Theta} = \arg \min_{\Theta} \sum_{(c,t) \in \mathcal{P}} \sum_{k \in \{A, S\}} r_{ctk}^2, \quad (8.7)$$

subject to (8.4)–(8.6), the fixed-effects normalization, and the regularity restrictions. The parameter vector Θ collects α , $\{e_k\}$, $\{\beta_k\}$, the latent utilities, fitted shares, and fixed effects, while $\mathcal{P}$ denotes the observed country-year panel.

The shift from the cross-sectional share likelihood to the log-relative-employment objective reflects the structure of the panel data. The cross-sectional estimator observes expenditure shares directly and fits the model to them. The [Comin et al. \(2021\)](#) panel is organized around relative employment, with country-sector fixed effects mapping employment into expenditure up to stable wedges. Placing the stochastic residual in equation (8.6) makes the constrained estimator directly comparable to the their specification while preserving the structural share and utility restrictions.

8.3.4. Equivalence and Differences between the Two Estimators

The constrained estimator does not substitute equation (8.3) into the objective. It imposes the restricted NHCES system and evaluates the residual implied by the fitted shares. Taking the ratio of the fitted shares in equation (8.5),

$$\log \frac{\hat{s}_{ctk}}{\hat{s}_{ctM}} = \log \frac{\beta_k}{\beta_M} + (1 - \alpha) \log \frac{p_{ctk}}{p_{ctM}} + (1 - \alpha) \left(\frac{e_k}{e_M} - 1 \right) \log \frac{E_{ct}}{p_{ctM}} + \left(\frac{e_k}{e_M} - 1 \right) \log \hat{s}_{ctM}, \quad (8.8)$$

which reproduces the structure of equation (8.3). [Comin et al. \(2021\)](#) use the observed manufacturing share ω_{ctM} in the final term, whereas the constrained program uses the fitted share $\hat{s}_{ctM}$, which is jointly tied to prices, expenditure, utility, and the demand parameters through the imposed constraints. The intercepts and country-sector fixed effects absorb β_k/β_M and the stable employment-to-expenditure wedge. The parameter mapping is

$$\sigma = \alpha, \quad \frac{\epsilon_k}{\epsilon_M} - 1 = \frac{e_k}{e_M} - 1, \quad k \in \{A, S\}. \quad (8.9)$$

The income elasticity implied by the restricted system is

$$\eta_{ctk} = \alpha + (1 - \alpha) \frac{e_k}{\sum_{j \in \mathcal{K}} \hat{s}_{ctj} e_j}, \quad (8.10)$$

which averages to one when weighted by fitted shares. Agriculture is a necessity when e_A lies below this share-weighted average, while services are luxuries when e_S lies above it.

The economic structure follows the line of work developed by Satō (1975), Hanoch (1975), and Comin et al. (2021). The difference lies in the estimator. Comin et al. (2021) begin from the log-linear implication of the restricted system and estimate the constrained coefficient vector by GMM. We keep the implicit system in levels, introduce country-year utilities and fitted shares as latent variables, impose the share equations directly as constraints, and estimate the fixed-effects panel residual jointly in one program.

8.3.5. Panel Estimates

The restricted NHCES estimate recovers the benchmark from within-country variation. The common substitution parameter is $\hat{\sigma} = 0.568$, compared with the benchmark estimate of 0.57, and the relative non-homotheticity parameters are $\hat{e}_A/\hat{e}_M - 1 = -1.000$ and $\hat{e}_S/\hat{e}_M - 1 = 0.341$, compared with -0.98 and 0.17 (Table VIII). The agricultural ratio lies at the regularity boundary, consistent with the near-boundary agricultural non-homotheticity reported under manufacturing normalization by Comin et al. (2021). The magnitudes are not identical, but the constrained estimator recovers the same sectoral ordering and broadly similar non-homotheticity patterns.

For comparison, we evaluate the mean income elasticities in Table VIII using the observed country-year sectoral shares. This is the closest analogue to the elasticity measure reported by Comin et al. (2021). In their theory, the income elasticity is a function of the non-homotheticity parameters and the expenditure-share-weighted average of those parameters. In the aggregate panel, they estimate relative employment equations because relative employment is proportional to relative sectoral expenditure shares under their model's market-clearing condition. Evaluated at the observed shares, the constrained NHCES panel implies mean income elasticities of (0.57, 1.02, 1.18) for agriculture, manufacturing, and services, close to the no-trade-control benchmark values of (0.58, 1.06, 1.14) implied by their primitive parameters.²⁵

We then estimate the unrestricted CDE on the same fixed-effect panel. Imposing $\alpha_A = \alpha_M = \alpha_S$ replicates the restricted NHCES panel estimator, with identical fitted log-relative allocations and residuals. This confirms that the structured CDE panel nests the restricted benchmark as a special case. Releasing the equality restriction improves the fit substantially. The unrestricted panel CDE places agriculture's expansion parameter and manufacturing's substitution parameter at the regularity boundary, with $\hat{\alpha}_A = 0.560$, $\hat{\alpha}_S = 0.960$, and $\hat{e}_S \simeq 2.99$. The solution is stable across starting values, bound choices, and alternative normalizations.

The two systems are reported under different elasticity evaluations in Table VIII. NHCES elasticities are evaluated at observed country-year shares because this matches the measure reported by Comin et al. (2021). For CDE, we report elasticities under three evaluations. The observed-share row uses the same share convention as the CLM comparison and

²⁵Comin et al. (2021) report implied expenditure elasticities of (0.56, 1.03, 1.14) for agriculture, manufacturing, and services in their world specification with trade controls. They do not separately report the income elasticities for the specification without trade controls. We compute them from their reported parameter estimates, $\sigma = 0.57$, $\epsilon_A/\epsilon_M - 1 = -0.98$, and $\epsilon_S/\epsilon_M - 1 = 0.17$. Normalizing $\epsilon_M = 1$ gives $(\epsilon_A, \epsilon_M, \epsilon_S) = (0.02, 1, 1.17)$. We then apply $\eta_{ik} = \sigma + (1 - \sigma)\epsilon_k/\bar{\epsilon}_i$, where $\bar{\epsilon}_i = \sum_j \omega_{ij} \epsilon_j$ using the observed sectoral shares from our reconstructed 1,626-observation employment panel. This gives (0.58, 1.06, 1.14) after rounding.

TABLE VIII
CLM BENCHMARK AND STRUCTURED PANEL ESTIMATES

Panel A: CLM Benchmark and Restricted Panel Comparison							
Specification	Elasticity Evaluation	Substitution Parameter		$A/M - 1$	$S/M - 1$	Mean Income Elasticities (A, M, S)	
CLM, World, No Trade Controls	Observed Shares	$\sigma = 0.57$		-0.98	0.17	(0.58, 1.06, 1.14)	
Constrained NHCES Panel	Observed Shares	$\alpha = 0.57$		-1.00	0.34	(0.57, 1.02, 1.18)	
Panel B: Unrestricted CDE Panel							
Specification	Elasticity Evaluation	Substitution Parameters			$A/M - 1$	$S/M - 1$	Mean Income Elasticities (A, M, S)
Constrained CDE Panel	Observed Shares	$\alpha_A = 0.56$	$\alpha_M = 0.00^B$	$\alpha_S = 0.96$	n.c.	n.c.	(0.77, 0.79, 1.24)
Constrained CDE Panel	Fitted Shares	$\alpha_A = 0.56$	$\alpha_M = 0.00^B$	$\alpha_S = 0.96$	n.c.	n.c.	(0.64, 1.00, 1.15)
Constrained CDE Panel	Fixed-Effect Adjusted	$\alpha_A = 0.56$	$\alpha_M = 0.00^B$	$\alpha_S = 0.96$	n.c.	n.c.	(0.70, 1.05, 1.21)

Notes: The CLM row reports the aggregate employment specification without trade controls from [Comin et al. \(2021\)](#). The NHCES and CDE rows use the reconstructed 1,626-observation employment panel with country-sector fixed effects. Under NHCES, $A/M - 1$ and $S/M - 1$ denote $\hat{e}_A/\hat{e}_M - 1$ and $\hat{e}_S/\hat{e}_M - 1$, with manufacturing normalized to $\hat{e}_M = 1$. Corresponding CDE ratios are not reported because sector-specific substitution breaks comparability with the CLM coefficients. In the constrained CDE panel, β_k is fixed because country-sector fixed effects absorb constant relative taste terms. Superscript B denotes an estimate at the implemented regularity bound.

gives (0.77, 0.79, 1.24). The fitted-share row mirrors the procedure used in the ICP cross section. It solves the structural demand system at a common reference price vector, uses the fitted shares implied by that system, and then evaluates the income-elasticity formula, giving (0.64, 1.00, 1.15). The fixed-effect-adjusted row evaluates the expenditure response of the fitted panel demand at each country's estimated sectoral-allocation point. This gives (0.70, 1.05, 1.21). The choice of share evaluation matters more for CDE than for NHCES because the share vector enters the CDE elasticity formula through both the share-weighted non-homotheticity term and the share-weighted substitution terms, whereas the common- α restriction in NHCES substantially reduces this dependence. The observed-share and fitted-share measures are thus similar under NHCES but diverge under CDE. The fixed-effect-adjusted elasticities lie between the two because they evaluate the expenditure response at each country's estimated sectoral-allocation point after accounting for the country-sector fixed effects. Since the panel estimator identifies the parameters from within-country movements around those operating points, the fixed-effect-adjusted measure is the internally consistent elasticity measure for the panel specification. Appendix D.4 discusses the three elasticity measures in detail and reports the fitted-share and fixed-effect-adjusted elasticities across the income distribution.

9. COUNTERFACTUAL APPLICATIONS

Our final application bridges the estimation framework in this paper with a multi-region, multi-sector general equilibrium model of international trade. We use GTAP, whose household demand is based on the CDE system. In applied CGE work, the household demand parameters have historically been calibrated rather than directly estimated. This makes the comparison between calibrated and estimated demand parameters important for policy counterfactuals in trade, climate, energy, and environmental applications. More broadly, the framework developed here can be combined with exact-hat quantitative methods such as [Dekle et al. \(2007, 2008\)](#), [Caliendo et al. \(2019\)](#), allowing researchers to estimate demand parameters from globally regular systems and carry them directly into counterfactual analysis. A broader structural comparison of restricted and unrestricted CDE systems within a full general equilibrium framework we leave to future researchers. In this paper, we ask how the estimated demand parameters compare with the calibrated ones, and how empirical differences affect counterfactuals.

We construct the dataset from GTAP private household purchases for 65 sectors and 160 regions in the 2017 database ([Aguilar et al., 2022](#)). Following the aggregation logic in Section 5.2,

we examine which broad bundles the data can support. The GTAP sectoral structure is coarser than the ICP basic-heading data in some dimensions and finer in others, so the ICP 12-bundle split is not the preferred specification. We evaluate candidate aggregations using bundle size, share variation, thin-region exposure, Engel correlations, and cross-region price dispersion. The preferred specification is an 8-good aggregation. It pools all food-related sectors together with beverages and tobacco into one stable necessity-oriented bundle, while keeping clothing and footwear, housing and utilities, household goods and services, health, transport and communication, recreation and education, and other services separate. For each GTAP region and aggregate good, we construct household expenditure from private purchases at consumer prices by summing the sectors assigned to that category. Price variations are constructed from the consumer-over-producer price wedge aggregated within each bundle and then normalized to the United States. We then construct per-capita expenditure, expenditure shares, and quantity indices in the same form used throughout the constrained-likelihood estimation. Appendix E.1 reports the aggregation procedure, the bundle-selection criterion, and the reconciliation of the estimated parameters with the original GTAP sectoral structure.




FIGURE 8.—Benchmark GTAP income elasticities aggregated to eight consumption bundles. Panels report bundle-level income elasticities against per-capita expenditure for the 160 GTAP regions under the calibrated, point estimates, and bootstrap mean. Sector elasticities are aggregated within bundles using benchmark private-purchase weights. Lines are kernel-smoothed averages over log per-capita expenditure.

9.1. Income Elasticities and Counterfactual Incidence

We compare the income elasticities implied by the calibrated parameters with those implied by our estimated parameters. The comparison is calculated at the benchmark general equi-

librium because the model evaluates income elasticities using benchmark expenditure shares. We aggregate the benchmark sector elasticities from the calibrated parameters, point estimates, and bootstrap means to the eight consumption bundles using benchmark weights of private consumption.²⁶ The expenditure-share evidence strongly classifies Food, Beverages, and Tobacco as a necessity. In the GTAP data, the lowest-income regions devote roughly one half of total expenditure to this bundle, while the highest-income regions devote roughly one tenth. The estimated income elasticities preserve this necessity pattern (Figure 8). The elasticities implied by the calibrated parameters are also below one, but they become U-shaped and move back toward unit elasticity at both ends of the expenditure distribution. As a result, the strongest necessity pattern in the expenditure-share data appears substantially more muted under the calibrated parameters. The U-shape arises from aggregating calibrated sector elasticities using benchmark within-bundle expenditure weights. The broad Food, Beverages, and Tobacco bundle includes sectors with different calibrated elasticities, and the composition of expenditure within the bundle changes systematically with income. At low expenditure levels, the bundle is dominated by subsistence food sectors, while at higher expenditure levels processed foods, beverages, tobacco, and animal products account for a larger share of expenditure.

In the counterfactual exercises, these elasticity differences map into a clear pattern across the expenditure distribution. For Food, Beverages, and Tobacco, the estimated system behaves as a strong necessity throughout, while the calibrated system moves substantially closer to unit elasticity at low and high expenditure levels. They are closest in the middle of the distribution and diverge most at the tails. For the nonfood and service bundles, the estimated elasticities are substantially larger than the calibrated elasticities among developing regions and gradually converge as expenditure rises. Clothing and Footwear is one exception, lying below the calibrated schedule over much of the upper expenditure range under the point estimates. The implied counterfactual differences should therefore be largest among developing regions across most bundles, with an additional food and services margin among richer regions, and smallest near the middle of the expenditure distribution. These elasticity differences are the channel through which the estimated demand system reallocates common shocks. Appendix E.4 reports additional sector-level income elasticities after reconciliation to the initial sectoral structure.

9.2. Counterfactual Scenarios

We use the counterfactuals to ask whether estimated demand systems allocate shocks differently from the calibrated GTAP specification. The counterfactuals ask not only how large the gains are, but how they are distributed across goods, regions, and trading partners. We consider four counterfactual scenarios. The first raises productivity by 10 percent in every region. Because the shock is common across regions, the scenario isolates how the demand system allocates a broad increase in real income across goods. The second raises productivity by 10 percent only in low-income and lower-middle-income economies using the 2017 World Bank income classification. This places the income gain where Engel effects are strongest in the data, since developing economies devote much larger expenditure shares to food-related goods. The third cuts bilateral import tariffs on food-related goods by 50 percent. The shock changes the relative price of the strongest necessity bundle in the data and combines expenditure reallocation with substitution across sources. The fourth cuts bilateral import tariffs by 50 percent in all sectors. This is the broadest tariff shock and asks whether the estimated demand system changes expenditure allocation and trade sourcing beyond the food sector. Appendix E.5 decomposes the counterfactual differences into income-response and substitution components by separately replacing the calibrated income and substitution parameters.

²⁶We report both world averages and expenditure-tier averages across regions in Appendix E.4.

9.3. Counterfactual Results

The estimated system keeps Food, Beverages, and Tobacco as a strong necessity across the expenditure distribution, while the calibrated system moves this bundle back toward unit income elasticity at the tails. For most nonfood and service bundles, the two systems disagree mainly among developing regions, where the estimated elasticities are higher, and converge as income rises. The gap between estimated and calibrated demand should therefore be widest when a shock raises real income in economies with large food shares, and narrowest when the shock is a broad price change spread across many sectors. The four scenarios load on this margin to different degrees. The global productivity shock raises income everywhere. The productivity shock in low-income countries (LIC) and lower-middle-income countries (LMIC) concentrates that income gain where food shares are largest. The food tariff reduction moves the price of the necessity bundle directly, while the multilateral liberalization spreads price changes across the full expenditure basket in the world.




FIGURE 9.—Counterfactual responses under estimated and calibrated demand systems. Columns report the four counterfactual scenarios. Rows report share responses aggregated to the eight consumption bundles. Each cell is the percentage-point difference between the two approaches, averaged within income groups. The GTAP model is solved at the original 65-sector level. The eight-bundle aggregation is used for reporting.

Figure 9 puts all four scenarios on a common scale. Each cell reports the estimated counterfactual minus the calibrated counterfactual in percentage points. The upper row shows that the estimated system changes final expenditure composition in the places predicted by the Engel evidence. The two productivity shocks generate the largest expenditure-share movements. Under the global productivity shock, the estimated system lowers the food share in developing regions and raises shares for household goods, transport, and other services. Under the LIC/LMIC productivity shock, the same pattern becomes more concentrated where food shares are largest. The food tariff reduction operates through the same necessity margin but with the opposite sign for food expenditure shares. A real-income gain lowers the expenditure share of a necessity, while a lower food price raises the relative demand for that bundle in developing regions. The multilateral trade liberalization shows the same qualitative margins, but at much smaller magnitude. Its panels are nearly flat. Multilateral liberalization still changes relative prices across many sectors, but the resulting expenditure-share movements are small. Estimated demand matters most when shocks operate through real income and expenditure

reallocation, and substantially less when price changes are spread broadly across the expenditure basket, with the welfare decomposition in Appendix E.5 showing that productivity shocks mainly operate through the Engel margin while the food tariff counterfactual operates more through substitution.

The lower row reports the same comparison for import shares. The import-share differences are concentrated in fewer cells and are much smaller than the shifts in expenditure shares. Estimated demand changes what households buy more than the import exposure of those bundles. Even under global trade liberalization, the estimated system does not generate large import-share reallocations. Estimation matters most when the counterfactual turns on Engel incidence, food shares, or income gains in developing economies. It matters much less for broad liberalizations that move relative prices across sectors.

The lesson is that estimation matters most when incidence depends on how households reallocate spending between necessities, services, and imported varieties. The same lesson likely applies when moving from the restricted to the unrestricted system, where the richer system matters most among developing regions and along the low-income curvature in Figure 7.

10. CONCLUSION

This paper estimates a globally regular implicit additive demand system directly from cross-region expenditure and price data. The central econometric difficulty is that utility enters the indirect system implicitly and is not observed. Reduced-form transformations that eliminate utility also eliminate part of the parameter structure. We address this problem by estimating utilities, expenditure shares, and preference parameters jointly within a constrained-likelihood framework. The approach combines share inversion with constrained optimization and develops inference procedures for parameters that lie on the boundary of the regularity domain.

Our empirical results reject the common-substitution restriction. The unrestricted system preserves the broad Engel patterns in the data, with food behaving as a necessity and service-oriented categories behaving as luxuries, but it changes where the curvature of demand appears and how substitution patterns differ across goods. In the three-sector benchmark, the unrestricted system reproduces the familiar ranking of agriculture, manufacturing, and services by income elasticity, but it implies a very different substitution structure once relative prices move across sectors. Manufacturing and services are much more substitutable with one another than either is with agriculture. The same broad non-homothetic pattern can therefore imply different counterfactual responses once the common-substitution restriction is relaxed. We further extend the estimator to a panel with country-sector fixed effects using the country-year data of Comin et al. (2021), and recover the same benchmark from within-country variation. The constrained-likelihood framework is therefore not specific to the ICP cross section.

Our estimates also identify where a smooth globally regular demand system fits the data poorly. Staples and Fresh Fruit and Vegetables sit at the lower regularity boundary because their expenditure shares rise sharply at low income levels. Alcohol, Tobacco and Narcotics behaves differently. Its cross-region variation is shaped less by income than by tax, regulation, culture, and institutions, and the unrestricted likelihood becomes unbounded outside the regular domain. The unrestricted system therefore does more than fit the data flexibly. It also shows where the flexible demand structure reaches its limits.

We then take the estimation framework into a multi-region, multi-sector general equilibrium trade model by replacing the calibrated GTAP household-demand parameters with estimates from the unrestricted CDE system. Our counterfactual exercises show where the empirical Engel differences matter quantitatively. Estimated demand changes expenditure composition most strongly when shocks operate through real income and Engel incidence, especially in developing economies with large food shares, and much less when shocks mainly operate through

broad relative-price changes across sectors. The main lesson is that estimation matters most when incidence depends on how households reallocate spending between necessities, services, and imported varieties. The same logic likely applies when moving from restricted to unrestricted demand systems, where the richer system matters most along the low-income Engel curvature that shapes expenditure reallocation in developing economies.

Our paper shows that global regularity does not need to come from imposing a common substitution structure across goods. Researchers can estimate richer implicit additive systems directly and allow the data to determine where income curvature and substitution patterns differ. The same framework can also be combined with quantitative counterfactual methods such as Dekle et al. (2007, 2008), Caliendo et al. (2019), allowing estimated globally regular demand systems to enter directly into counterfactual analysis. Extending these methods to richer panel settings beyond the benchmark in this paper, stronger external identification strategies, and broader classes of implicit systems remains an important direction for future work.

REFERENCES

- AGUIAR, ANGEL, MAKSYM CHEPELIEV, ERWIN CORONG, AND DOMINIQUE VAN DER MENSBRUGGHE (2022): “The global trade analysis project (GTAP) data base: Version 11,” *Journal of Global Economic Analysis*, 7 (2). [17, 38]
- ANDERSON, JAMES E AND ERIC VAN WINCOOP (2003): “Gravity with gravitas: A solution to the border puzzle,” *American Economic Review*, 93 (1), 170–192. [2, 7, 11]
- ANDREWS, DONALD W. K. (1999): “Estimation When a Parameter Is on a Boundary,” *Econometrica*, 67 (6), 1341–1383. [3, 19, 20, 21, 23, 77]
- (2000): “Inconsistency of the Bootstrap When a Parameter Is on the Boundary of the Parameter Space,” *Econometrica*, 68 (2), 399–406. [3, 19, 20, 77]
- (2001): “Testing When a Parameter Is on the Boundary of the Maintained Hypothesis,” *Econometrica*, 69 (3), 683–734. [3, 21]
- ARROW, KENNETH J, HOLLIS B CHENERY, BAGICHA S MINHAS, AND ROBERT M SOLOW (1961): “Capital-Labour Substitution and Economic Efficiency,” *Review of Economics and Statistics*, 43 (3), 225–250. [12]
- BALISTRERI, EDWARD J AND RUSSELL H HILLBERRY (2007): “Structural estimation and the border puzzle,” *Journal of International Economics*, 72 (2), 451–463. [11]
- BALK, BERT M. (2008): *Price and Quantity Index Numbers: Models for Measuring Aggregate Change and Difference*, New York: Cambridge University Press. [48]
- BARTEN, ANTON P. (1977): “The Systems of Consumer Demand Functions Approach: A Review,” *Econometrica*, 45 (1), 23–51. [15]
- BERGSON, A.(BURK) (1936): “Real Income, Expenditure Proportionality, and Frisch’s New Methods,” *Review of Economic Studies*, 4, 33–52. [6]
- BERRY, STEVEN, JAMES LEVINSOHN, AND ARIEL PAKES (1995): “Automobile prices in market equilibrium,” *Econometrica: Journal of the Econometric Society*, 63 (4), 841–890. [8]
- BERRY, STEVEN T (1994): “Estimating discrete-choice models of product differentiation,” *The RAND Journal of Economics*, 25 (2), 242–262. [2, 5, 8]
- BERRY, STEVEN T. AND PHILIP A. HAILE (2021): “Foundations of Demand Estimation,” in *Handbook of Industrial Organization*, Elsevier, vol. 4, 1–62. [5]
- BICKEL, PETER J., FRIEDRICH GÖTZE, AND WILLEM R. VAN ZWET (1997): “Resampling Fewer Than n Observations: Gains, Losses, and Remedies for Losses,” *Statistica Sinica*, 7 (1), 1–31. [20]
- BICKEL, PETER J. AND ANAT SAKOV (2008): “On the Choice of m in the m out of n Bootstrap and Confidence Bounds for Extrema,” *Statistica Sinica*, 18 (3), 967–985. [20]
- BOPPART, TIMO (2014): “Structural change and the Kaldor facts in a growth model with relative price effects and non-Gorman preferences,” *Econometrica*, 82 (6), 2167–2196. [32]
- CALIENDO, LORENZO, MAXIMILIANO DVORKIN, AND FERNANDO PARRO (2019): “Trade and labor market dynamics: General equilibrium analysis of the china trade shock,” *Econometrica*, 87 (3), 741–835. [4, 38, 43]
- CARON, JUSTIN, THIBAUT FALLY, AND JAMES R. MARKUSEN (2014): “International trade puzzles: A solution linking production and preferences,” *The Quarterly Journal of Economics*, 129 (3), 1501–1552. [8]
- CHERNOFF, HERMAN (1954): “On the Distribution of the Likelihood Ratio,” *The Annals of Mathematical Statistics*, 25 (3), 573–578. [3, 21]

- COMIN, DIEGO, DANIAL LASHKARI, AND MARTÍ MESTIERI (2021): “Structural Change With Long-Run Income and Price Effects,” *Econometrica*, 89 (1), 311–374. [2, 3, 8, 10, 11, 19, 23, 32, 33, 34, 35, 36, 37, 38, 42, 69]
- CRANFIELD, JOHN AL, PAUL V. PRECKEL, JAMES S. EALES, AND THOMAS W. HERTEL (2002): “Estimating Consumer Demands across the Development Spectrum: Maximum Likelihood Estimates of an Implicit Direct Additivity Model,” *Journal of Development Economics*, 68 (2), 289–307. [10, 63]
- DEKLE, ROBERT, JONATHAN EATON, AND SAMUEL KORTUM (2007): “Unbalanced trade,” *American Economic Review*, 97 (2), 351–355. [4, 38, 43]
- (2008): “Global Rebalancing with Gravity: Measuring the Burden of Adjustment,” *IMF Staff Papers*, 55 (3), 511–540. [4, 38, 43]
- DI EWERT, W. EWIN (1988): “Test approaches to international comparisons.” *Measurement in Economics: Theory and Applications of Economics Indices*, Physica-Verlag HD. [48]
- (2013): “Methods of Aggregation above the Basic Heading Level within Regions,” *Measuring the Real Size of the World Economy*, Chapter 5, 121–168. [18, 48]
- DRUD, ARNE (1985): “CONOPT: A GRG code for large sparse dynamic nonlinear optimization problems,” *Mathematical programming*, 31, 153–191. [19]
- ÉLTETŐ, O. AND P. KÖVES (1964): “On a Problem of Index Number Computation Relating to International Comparison,” *Statistikai Szemle*, 42, 507–518. [18]
- FAJGELBAUM, PABLO D AND AMIT K KHANDELWAL (2016): “Measuring the unequal gains from trade,” *The Quarterly Journal of Economics*, 131 (3), 1113–1180. [8]
- FIELER, ANA CECILIA (2011): “Nonhomotheticity and bilateral trade: Evidence and a quantitative explanation,” *Econometrica*, 79 (4), 1069–1101. [8]
- FISHER, IRVING (1922): *The making of index numbers: a study of their varieties, tests, and reliability*, Boston: Houghton Mifflin Company. [47, 48]
- GINI, CORRADO (1924): “Quelques considérations au sujet de la construction des nombres indices des prix et des questions analogues,” *Metron*, 4 (1), 3–162. [18]
- (1931): “On the Circular Test of Index Numbers,” *Metron*, 9 (9), 3–24. [18]
- GOUEL, CHRISTOPHE AND HOUSSEIN GUIMBARD (2019): “Nutrition Transition and the Structure of Global Food Demand,” *Econometrica*, 101 (2), 383–403. [10]
- GREENE, WILLIAM H. (2012): “Econometric Analysis 7th Edition International Edition,” *Essex, England: Pearson Education Limited*. [10]
- HANOCH, GIORA (1971): “CRESH Production Functions,” *Econometrica: Journal of the Econometric Society*, 39, 695–712. [12]
- (1975): “Production and Demand Models with Direct or Indirect Implicit Additivity,” *Econometrica*, 43 (3), 395–419. [1, 2, 4, 6, 11, 12, 13, 14, 15, 19, 37, 58, 77, 83]
- HERTEL, THOMAS WARREN (1997): *Global trade analysis: modeling and applications*, Cambridge university press. [2]
- HOUTHAKKER, HENDRIK S (1960): “Additive Preferences,” *Econometrica: Journal of the Econometric Society*, 244–257. [6]
- MATSUYAMA, KIMINORI (2019): “Engel’s law in the global economy: Demand-induced patterns of structural change, innovation, and trade,” *Econometrica*, 87 (2), 497–528. [2]
- (2023): “Non-CES aggregators: a guided tour,” *Annual Review of Economics*, 15, 235–265. [4]
- MUKERJI, V. (1963): “Generalized SMAC Function with Constant Ratios of Elasticities of Substitution,” *Review of Economic Studies*, 30, 233–236. [12]
- POLITIS, DIMITRIS N., JOSEPH P. ROMANO, AND MICHAEL WOLF (1999): *Subsampling*, Springer Series in Statistics, New York: Springer. [20]
- POWELL, ALAN (1966): “A Complete System of Consumer Demand Equations for the Australian Economy Fitted by a Model of Additive Preferences,” *Econometrica: Journal of the Econometric Society*, 42 (2), 661–675. [6]
- PRECKEL, PAUL V., J. A. L. CRANFIELD, AND THOMAS W. HERTEL (2010): “A Modified, Implicit, Directly Additive Demand System,” *Applied Economics*, 42 (2), 143–155. [10, 63]
- REDDING, STEPHEN J AND DAVID 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. [5]
- SATŌ, KAZUO (1975): *Production functions and aggregation*, vol. 90, North-Holland Publishing Company. [2, 10, 12, 37]
- SELF, STEVEN G. AND KUNG-YEE LIANG (1987): “Asymptotic Properties of Maximum Likelihood Estimators and Likelihood Ratio Tests Under Nonstandard Conditions,” *Journal of the American Statistical Association*, 82 (398), 605–610. [3, 21]
- SU, CHE-LIN AND KENNETH L JUDD (2012): “Constrained optimization approaches to estimation of structural models,” *Econometrica*, 80 (5), 2213–2230. [2, 11]

- SURRY, YVES (1993): "The Constant Difference of Elasticities Function with Applications to the EC Animal Feed Sector," *Journal of Agricultural Economics*, 44 (1), 110–125. [7, 8]
- SZULC, B. (1964): "Methods of Aggregation above the Basic Heading Level within Regions," *Przegląd Statystyczny*, 3, 239–254. [18]
- THE WORLD BANK (2021): "Use of the GEKS Method for Price Computation," Suggestion made in official capacity by The World Bank. [18]
- WORLD BANK (2024): "Poverty and Inequality Platform and World Development Indicators," Washington, DC: World Bank, <https://pip.worldbank.org/> and <https://databank.worldbank.org/source/world-development-indicators>. [61]
- YANG, ANTON C. (2021): "Structural Estimation of a Gravity Model of Trade with the Constant-Difference-of-Elasticities Preference," *Mimeo*. [6]
- YANG, ANTON C., CHRISTOPHE GOUEL, AND THOMAS W. HERTEL (2018): "Will Income or Population be the Main Driver of Food Demand Growth to 2050?" *Mimeo*. [10]

APPENDIX A: PROOFS AND THEORETICAL DETAILS

A.1. The Proof of Monotonicity of the CRES

PROOF: For the convenience of the reader we restate the CRES function here:

$$\mathcal{H}(\mathbf{x}, u) = \sum_k \beta_k u^{-e_k(1-\alpha_k)} (x_k)^{1-\alpha_k} \equiv 1, \quad (\text{A.1})$$

where (i) $\beta_k, e_k > 0 \forall k \in N$, and (ii) either $\alpha_k > 1$ or $0 \leq \alpha_k < 1 \forall k \in N$. (Note that this argument can be patched to handle $0 \leq \alpha_k \leq 1$ by dealing with cases. the fact that, when we interpret $x_k^{1-\alpha_k} = \ln(x_k)$ when $\alpha_k = 1$, the derivative of $x_k^{1-\alpha_k}$ with respect to x_k remains positive, means that the argument below goes through.)

Taking the total differential of (A.1) to get:

$$-\sum_k \beta_k e_k (1 - \alpha_k) u^{-e_k(1-\alpha_k)-1} x_k^{1-\alpha_k} du + \sum_k \beta_k u^{-e_k(1-\alpha_k)} (1 - \alpha_k) x_k^{-\alpha_k} dx_k \equiv 0. \quad (\text{A.2})$$

Now let all $dx_k = 0$ except for $k = m$ to derive:

$$-\sum_k \beta_k e_k (1 - \alpha_k) u^{-e_k(1-\alpha_k)-1} x_k^{1-\alpha_k} du + \beta_m u^{-e_m(1-\alpha_m)} (1 - \alpha_m) x_m^{-\alpha_m} dx_m \equiv 0. \quad (\text{A.3})$$

Then we solve for the change in u for a change in x_m :

$$\frac{du}{dx_m} = \frac{\beta_m u^{-e_m(1-\alpha_m)} (1 - \alpha_m) x_m^{-\alpha_m}}{\sum_k \beta_k e_k (1 - \alpha_k) u^{-e_k(1-\alpha_k)-1} x_k^{1-\alpha_k}}. \quad (\text{A.4})$$

Notice that the numerator has the sign of $(1 - \alpha_m)$, and every term in the sum in the denominator has the sign of $(1 - \alpha_k)$. These signs are the same by the parametric restrictions, and hence:

$$\frac{du}{dx_m} > 0. \quad (\text{A.5})$$

Thus, u is strictly monotonic in each x_k .

Q.E.D.

A.2. Solutions to the Expenditure Minimization Problem under the Relaxed Regularity Conditions

The relaxation of the regularity conditions for the CRES demand system has impacts on the nature of the solutions to the expenditure minimization problem. Again we focus on the regime where $0 \leq \alpha_k < 1$, and focus on the difference between whether zero is included as the minimum value. Under the original conditions where $\alpha_k > 0$, solutions to the utility maximization problem satisfy $x_k > 0$. If $\alpha_k = 0$ is permitted for some subset of goods, the possibility that $x_k = 0$ cannot be excluded. Here we characterize optimal solutions for that case.

We define K_0 to be the set of indices for goods with $\alpha_k = 0$, and K_+ to be the subset of indices with $\alpha_k > 0$. It is convenient to consider two cases with the first being when the solution for the expenditure minimization problem has $x_k = 0$ for all $k \in K_0$, and the second being when the solution has $x_k > 0$ for one or more $k \in K_0$.

Case 1 ($x_k = 0$ for all $k \in K_0$):

Consider the first-order conditions for x_k when $k \in K_+$. These are:

$$c_k - \lambda \beta_k u^{-e_k(1-\alpha_k)} (1 - \alpha_k) x_k^{-\alpha_k} = 0. \quad (\text{A.6})$$

The case where $x_k = 0$ cannot occur because the second term on the left-hand side of the above decreases without bound as x_k approaches zero. Solving for x_k :

$$x_k = [c_k / (\lambda \beta_k u^{-e_k(1-\alpha_k)} (1 - \alpha_k))]^{-1/\alpha_k} \quad (\text{A.7})$$

for $k \in K_+$ and set $x_k^* = 0$ for $k \in K_0$. Substituting these into the defining constraint:

$$\sum_{k \in K_0+} \beta_k u^{-e_k(1-\alpha_k)} [c_k / (\lambda \beta_k u^{-e_k(1-\alpha_k)} (1 - \alpha_k))]^{-1/\alpha_k} = 1. \quad (\text{A.8})$$

This equation can be uniquely solved for λ , and the result can be substituted back into (A.7) to determine the solution for x_k . To verify that this is the correct case, it only remains to verify that the first-order conditions for x_k for $k \in K_0$ are satisfied:

$$c_k - \lambda \beta_k u^{-e_k} \geq 0, \quad (\text{A.9})$$

for all $k \in K_0$. If these conditions are satisfied, then set $x_k = 0$ for all $k \in K_0$, and the solution is defined. If one or more of these conditions are violated, it must be that $x_k > 0$ for some $k \in K_0$, and Case 2 defines the solution.

Case 2 ($x_k > 0$ for some $k \in K_0$):

Consider the first-order conditions for x_k when $k \in K_0$. These are:

$$c_k - \lambda \beta_k u^{-e_k} \geq 0, \quad (\text{A.10})$$

and

$$x_k [c_k - \lambda \beta_k u^{-e_k}] = 0. \quad (\text{A.11})$$

Since at least one $x_k > 0$ for $k \in K_0$, the value of λ calculated in Case 1 is not correct. The correct value can be calculated as:

$$\lambda = \min_{k \in K_0} \left\{ c_k [\beta_k u^{-e_k(1-\alpha_k)}]^{-1} \right\}. \quad (\text{A.12})$$

The index set K_{0+} is defined as the $k \in K_0$ that achieve the minimum in (A.12). Using this λ , the optimal values for x_k where $k \in K_{0+}$:

$$x_k = [c_k / (\lambda \beta_k u^{-e_k(1-\alpha_k)} (1 - \alpha_k))]^{-1/\alpha_k}, \quad (\text{A.13})$$

and set $x_k = 0$ for all k in K_0 but not in K_{0+} . Any values of the $x_k \geq 0$ for $k \in K_{0+}$ that satisfy the following equation will be optimal:

$$\sum_{k \in K_{0+}} \beta_k u^{-e_k} x_k = \sum_{k \in K_{0+}} \beta_k u^{-e_k(1-\alpha_k)} x_k^{1-\alpha_k} + 1. \quad (\text{A.14})$$

This completes the characterization of the solution to the expenditure problem. Relaxing the lower bound on α_k is important for parameter estimation, where we have found that this bound may be active at times.

APPENDIX B: DATA CONSTRUCTION AND ROBUSTNESS

B.1. Price Aggregation with GEKS

The GEKS method is a standard approach for constructing PPPs at higher levels of aggregation. It is designed to produce multilateral PPPs that are transitive (circular) and base-country invariant by aggregating bilateral comparisons using geometric averaging.²⁷ For our purposes, the main advantage of the GEKS method is that it can accommodate a large number of commodities and countries while producing consistent aggregate price comparisons across regions.

As the GEKS method requires the calculations of both the *Laspeyres* index and the *Paasche* index as a preliminary step, we first construct a quantity index by using the ICP price index (at the level of basic headings, g) as follows:

$$q_{i,g} = \frac{\Xi_{i,g}}{p_{i,g}}, \quad (\text{B.1})$$

where $q_{i,g}$, $\Xi_{i,g}$, and $p_{i,g}$ refer to the quantity index, expenditure, and price of commodity g in region i , respectively. By choosing basic heading PPPs for price variations within and between regions, we obtain the corresponding quantity index for calculating the [Fisher \(1922\)](#) index P_k^{Fisher} between region i relative to i' , which is defined as the multilateral geometric mean of the *Laspeyres* price index and *Paasche* price index between countries i and i' :

$$P_k^{\text{Fisher}}(p_i, p_{i'}, q_i, q_{i'}) = [P_k^{\text{Laspeyres}}(p_i, p_{i'}, q_i, q_{i'}) \cdot P_k^{\text{Paasche}}(p_i, p_{i'}, q_i, q_{i'})]^{0.5}, \quad (\text{B.2})$$

where

$$P_k^{\text{Laspeyres}}(p_i, p_{i'}, q_i, q_{i'}) \equiv \left(\sum_{g \in \mathcal{G}_k} p_{i',g} \cdot q_{i,g} \right) / \left(\sum_{g \in \mathcal{G}_k} p_{i,g} \cdot q_{i,g} \right), \quad (\text{B.3})$$

²⁷In the PPP context, transitivity or circularity means that PPP comparisons are internally consistent across countries. For example, the implied comparison between countries A and C is consistent with the comparisons between A and B and between B and C . Base-country invariance means that the multilateral PPP vector is unchanged, up to normalization, by the choice of numeraire country.

and

$$P_k^{Paasche}(p_i, p_{i'}, q_i, q_{i'}) \equiv \left(\sum_{g \in \mathcal{G}_k} p_{i',g} \cdot q_{i',g} \right) / \left(\sum_{g \in \mathcal{G}_k} p_{i,g} \cdot q_{i',g} \right), \quad (\text{B.4})$$

where we use k to denote an aggregated commodity. For each k , there exists a set $\mathcal{G}_k$ that contains the basic heading indices that compose the aggregate commodity k , and expenditures across aggregate goods $\Xi_k = \sum_{g \in \mathcal{G}_k} \Xi_g$. The index sets $\mathcal{G}_k$ are a partition of the basic headings, i.e. the union of these non-overlapping index sets is the set of all basic headings.

There are several justifications for using the Fisher index (Fisher, 1922), including its interpretation as a symmetric average of the Laspeyres and Paasche indexes, its desirable axiomatic properties in index number theory, and its foundation in Diewert's economic approach to index number theory (Diewert, 1988, Balk, 2008, Diewert, 2013). Following Diewert (2013), we compute the aggregate PPPs as follows:

$$P_{i',k} = \prod_{i=1}^I [P_k^{Fisher}(p_i, p_{i'}, q_i, q_{i'})]^{\frac{1}{I}}, \quad i' = 1, \dots, I. \quad (\text{B.5})$$

Lastly, we obtain the aggregate PPPs using equation (B.5). From these, we select a numeraire $\delta = P_{USA}$ such that the real PPPs, adjusted for the numeraire, are defined as

$$p_{i,k} \equiv \tilde{P}_{i,k} = P_{i,k} / \delta_k, \quad i = 1, \dots, I, \quad (\text{B.6})$$

where $\tilde{P}_i$ represents the amount of country i 's currency needed to purchase one unit of U.S. currency while achieving equivalent purchasing power for the relevant aggregate (Diewert, 2013). As a computational consistency check, we also compute Laspeyres and Paasche using the equivalent share-based arithmetic/harmonic-mean formulas in Diewert (2013), which reproduces exactly the same $\tilde{P}$ calculated from equation (B.6).

B.2. Commodity Aggregation Reference Table

TABLE B.I: Commodity Aggregation Based on the Confidential 2017 ICP Dataset

Cereals and Staples

Rice; Other cereals, flour and other cereal products; Bread; Other bakery products; Pasta products and couscous.

Fresh Fruit and Vegetables

Fresh or chilled fruit; Frozen, preserved or processed fruit and fruit-based products; Fresh or chilled vegetables, other than potatoes and other tuber vegetables; Fresh or chilled potatoes and other tuber vegetables; Frozen, preserved or processed vegetables and vegetable-based products.

Meat, Seafood and Dairy

Beef and veal; Pork; Lamb, mutton and goat; Poultry; Other meats and meat preparations; Fresh, chilled or frozen fish and seafood; Preserved or processed fish and seafood; Fresh milk; Preserved milk and other milk products; Cheese and curd; Eggs and egg-based products.

Oils, Sweets and Non-Alcoholic Beverages

Table B.I continued

Butter and margarine; Other edible oils and fats; Sugar; Jams, marmalades and honey; Confectionery, chocolate and ice cream; Food products n.e.c.; Coffee, tea and cocoa; Mineral waters, soft drinks, fruit and vegetable juices.

Clothing and Footwear

Clothing materials, other articles of clothing and clothing accessories; Garments; Cleaning, repair and hire of clothing; Shoes and other footwear; Repair and hire of footwear.

Alcohol, Tobacco and Narcotics

Spirits; Wine; Beer; Tobacco; Narcotics.

Household Goods and Services

Furniture and furnishings; Carpets and other floor coverings; Repair of furniture, furnishings and floor coverings; Household textiles; Major household appliances whether electric or not; Small electric household appliances; Repair of household appliances; Glassware, tableware and household utensils; Major tools and equipment; Small tools and miscellaneous accessories; Non-durable household goods; Domestic services; Household services.

Health

Pharmaceutical products; Other medical products; Therapeutic appliances and equipment; Medical services; Dental services; Paramedical services; Hospital services.

Transport and Communication

Motor cars; Motor cycles; Bicycles; Animal drawn vehicles; Fuels and lubricants for personal transport equipment; Maintenance and repair of personal transport equipment; Other services in respect of personal transport equipment; Passenger transport by railway; Passenger transport by road; Passenger transport by air; Passenger transport by sea and inland waterway; Combined passenger transport; Other purchased transport services; Postal services; Telephone and telefax equipment; Telephone and telefax services.

Recreation, Culture and Education

Audio-visual, photographic and information processing equipment; Recording media; Repair of audio-visual, photographic and information processing equipment; Major durables for outdoor and indoor recreation; Maintenance and repair of other major durables for recreation and culture; Other recreational items and equipment; Garden and pets; Veterinary and other services for pets; Recreational and sporting services; Cultural services; Games of chance; Newspapers, books and stationery; Package holidays; Education - HHC.

Restaurants, Personal and Other Services

Catering services; Accommodation services; Hairdressing salons and personal grooming establishments; Appliances, articles and products for personal care; Prostitution; Jewellery, clocks and watches; Other personal effects; Social protection - HHC; Insurance; Financial Intermediation Services Indirectly Measured (FISIM); Other financial services n.e.c.; Other services n.e.c.

Housing and Utilities

Actual rentals for housing; Imputed rentals for housing; Maintenance and repair of the dwelling; Water supply; Miscellaneous services relating to the dwelling; Electricity; Gas; Other fuels.

B.3. Summary Statistics

TABLE B.II
SUMMARY STATISTICS FOR THE DATA AGGREGATED FROM THE UNPUBLISHED 2017 ICP DATASET

	Staples	MeatFishDairy	FruitVeg	OilsSweetsBev	AlcTob	ClothFootwear	HousingUtil	HHGoodsSvcs	Health	TransComm	RecCultEdu	RestPersOther
Panel A: Prices												
Min	0.37	0.47	0.21	0.49	0.27	0.21	0.06	0.27	0.08	0.24	0.14	0.21
Max	2.48	1.79	2.13	1.95	2.10	1.55	1.61	1.78	2.00	1.55	1.84	1.60
Mean	0.92	0.83	0.78	0.98	0.80	0.68	0.46	0.71	0.44	0.77	0.55	0.62
Std	0.35	0.25	0.36	0.29	0.38	0.30	0.36	0.29	0.31	0.27	0.34	0.29
Panel B: Quantities												
Min	19.85	15.40	20.98	17.85	6.77	1.41	77.64	5.32	10.65	29.77	16.10	28.21
Max	735.54	2206.45	1059.70	1113.94	3128.15	2219.17	12712.04	1752.07	8606.54	5854.74	6344.21	12863.37
Mean	240.99	557.77	336.06	319.88	321.38	425.82	2921.92	491.47	692.50	1405.46	1016.67	1788.85
Std	120.11	364.67	215.44	231.53	385.13	382.04	2518.43	438.05	935.70	1256.21	1123.94	2066.60
Panel C: Per-Capita Expenditures												
Min	26.20	9.41	12.12	14.12	2.20	0.50	17.95	2.39	1.61	20.19	3.53	14.15
Max	731.19	2393.95	1291.06	1719.62	2917.02	1874.67	17217.01	2043.67	8606.54	6260.26	7342.11	11756.37
Mean	222.16	495.98	262.55	330.96	292.04	351.80	1854.32	432.92	410.57	1300.96	771.92	1520.08
Std	141.03	406.14	202.49	304.32	407.31	390.97	2614.51	494.00	913.44	1463.56	1136.88	2204.87
Lower CI 95%	201.45	436.35	232.82	286.27	232.24	294.39	1470.43	360.38	276.45	1086.06	604.99	1196.33
Upper CI 95%	242.87	555.62	292.28	375.64	351.85	409.20	2238.22	505.45	544.69	1515.86	938.85	1843.83
Panel D: Expenditure Shares												
Min	0.01	0.02	0.01	0.01	0.00	0.00	0.04	0.01	0.00	0.04	0.01	0.02
Max	0.30	0.42	0.38	0.20	0.16	0.22	0.45	0.13	0.22	0.33	0.16	0.31
Mean	0.06	0.10	0.06	0.06	0.04	0.05	0.18	0.05	0.04	0.15	0.07	0.14
Std	0.06	0.05	0.05	0.04	0.03	0.02	0.08	0.02	0.03	0.05	0.03	0.07
Lower CI 95%	0.06	0.09	0.06	0.06	0.03	0.04	0.17	0.05	0.04	0.14	0.07	0.13
Upper CI 95%	0.07	0.10	0.07	0.07	0.04	0.05	0.19	0.05	0.04	0.16	0.08	0.15
Panel E: Per-Capita Expenditure in U.S. Dollars												
Min												236.94
Max												54236.48
Mean												8246.25
Std												9822.29
Lower CI 95%												6804.01
Upper CI 95%												9688.49

Notes: Rounded to two decimals. Prices are normalized by the U.S. exchange rate.

B.4. Figures




FIGURE B.1.—Income elasticity schedules under the unrestricted NHCS specification (12 ICP bundles). Each panel plots the model-implied income elasticity against per-capita expenditure at a common reference price vector. Point estimates are shown with 90% paired-bootstrap confidence bounds. For parameters whose point estimate sits at the lower regularity bound, both bounds are constructed from the m-out-of-n subsampling bootstrap.




FIGURE B.2.—Income elasticities under the unrestricted CDE specification (11 ICP bundles). Each panel plots the model-implied income elasticity against per-capita expenditure at a common reference price vector. Point estimates are shown with 90% paired-bootstrap confidence bounds. For parameters whose point estimate sits at the lower regularity bound, both bounds are constructed from the m -out-of- n subsampling bootstrap.




FIGURE B.3.—Income elasticities under the unrestricted CDE specification (10 ICP bundles). Each panel plots the model-implied income elasticity against per-capita expenditure at a common reference price vector. Point estimates are shown with 90% paired-bootstrap confidence bounds. For parameters whose point estimate sits at the lower regularity bound, both bounds are constructed from the m -out-of- n subsampling bootstrap.




FIGURE B.4.—CDE income elasticities across ICP aggregations at low and high expenditure levels. Panel A reports elasticities evaluated at approximately \$300 per capita. Panel B reports elasticities evaluated at approximately \$50,000 per capita. Markers show point estimates and bars show 90% confidence intervals.

B.5. Bound-Relaxation Diagnostic for the NHCES Specification

In this section, we re-solve the NHCES system with the lower bound on e_k relaxed to $\underline{e} \in \{0, -0.1, -0.25, -0.5, -1.0\}$, holding all other constraints and the normalization $\sum_k e_k = N$ fixed. We report the analogous diagnostic for the restricted NHCES benchmark of Section 7.2 in Table B.III.

TABLE B.III
BOUND-RELAXATION DIAGNOSTIC FOR THE NHCES SPECIFICATION (12-BUNDLE ICP)

Bundle	Lower bound on e_k					Regime
	$\underline{e} = 0$	-0.10	-0.25	-0.50	-1.00	
Staples	0.000	-0.100	-0.250	-0.365	-0.365	B
Fruit and Vegetables	0.008	-0.046	-0.128	-0.193	-0.193	B
Meat, Seafood and Dairy	0.594	0.560	0.509	0.471	0.471	A
Oils, Sweets and Non-Alcoholic Beverages	0.457	0.410	0.337	0.282	0.282	A
Alcohol, Tobacco and Narcotics	0.840	0.840	0.834	0.826	0.826	A
Clothing and Footwear	0.781	0.775	0.763	0.751	0.751	A
Housing and Utilities	1.432	1.475	1.544	1.598	1.598	A
Household Goods and Services	1.263	1.279	1.303	1.322	1.322	A
Health	1.492	1.534	1.599	1.651	1.651	A
Transport and Communication	1.465	1.494	1.539	1.575	1.575	A
Recreation, Culture and Education	1.722	1.773	1.851	1.910	1.910	A
Restaurants, Personal and Other Services	1.946	2.007	2.100	2.171	2.171	A
$\hat{\ell}_{\text{NHCES}}$	4104.547	4106.270	4107.643	4107.923	4107.923	
$\Delta \hat{\ell}$ vs. $\underline{e} = 0$	0.000	1.723	3.096	3.375	3.375	

Notes: Constrained MLE under the common-substitution restriction $\alpha_k = \alpha$ as the lower regularity bound is relaxed. Two features are notable in comparison with the corresponding CDE diagnostic in Table III. The same two bundles—staples and fresh fruit and vegetables—are in regime B under both specifications, with comparable trajectories. The alcohol/tobacco/narcotics bundle, by contrast, has a finite interior NHCES estimate ($\hat{e}_{\text{AlcTob}} = 0.840$) and stays interior under relaxation; the regime-C unbounded direction observed in the unrestricted CDE arises because the CDE's good-specific substitution parameter $\hat{\alpha}_{\text{AlcTob}} = 0.771$ allows the vice bundle to absorb the cross-region variance through a nearly-Cobb–Douglas substitution channel, which then leaves e_{AlcTob} unidentified along the negative direction. The common- α restriction shuts that channel down. The NHCES estimates for $\underline{e} = -0.5$ and $\underline{e} = -1.0$ are identical to four decimal places, confirming that the unrestricted NHCES likelihood is bounded below at all bundles.

The NHCES regime classification overlaps but does not coincide with the CDE classification. Staples and fresh fruit and vegetables are in regime B under both specifications, with comparable mild-boundary trajectories. The Alcohol, Tobacco and Narcotics bundle, however, is in regime A under NHCES with a finite interior estimate, in contrast with the regime-C unbounded direction observed in the unrestricted CDE. The common-substitution restriction $\alpha_k = \alpha$ stabilizes $\hat{e}_{\text{AlcTob}}$ by shutting down the substitution-elasticity flexibility that lets the unrestricted CDE place this bundle on a nearly Cobb–Douglas substitution channel and absorb the cross-region variance.

B.6. Robustness: Alternative Aggregations

This section reports two ICP-based robustness specifications that test whether the 12-bundle results are sensitive to different aggregations. We compare the baseline 12-bundle specification with two coarser ICP aggregations: an 11-bundle that drops Alcohol, Tobacco and Narcotics,

and a 10-bundle that absorbs it into a broader Leisure, Vices and Other Services residual. In Appendix E.2, we report the 8-bundle CDE estimated on a separately constructed GTAP dataset. We use those estimates as inputs to the counterfactuals of Section 9.

B.6.1. 11-Bundle ICP Specification: Dropping the Vice Bundle

The 11-bundle specification omits Alcohol, Tobacco and Narcotics while keeping the remaining eleven bundles. Expenditure on the omitted category is excluded from total expenditure in the share equations rather than redistributed, so the remaining shares are renormalized within the 11-bundle system. This tests whether the boundary behavior is specific to the Alcohol, Tobacco and Narcotics bundle and whether it contaminates the rest of the system.

TABLE B.IV
ROBUSTNESS: 11-BUNDLE ICP SPECIFICATION

Bundle	12-bundle headline			11-bundle robustness		
	$\hat{\alpha}_k$	$\hat{\beta}_k$	$\hat{e}_k$	$\hat{\alpha}_k$	$\hat{\beta}_k$	$\hat{e}_k$
Staples	0.372	0.011	0.000 ^B	0.333	0.006	0.000 ^B
Fruit and Vegetables	0.384	0.013	0.000 ^B	0.327	0.007	0.000 ^B
Meat, Seafood and Dairy	0.089	0.013	0.707	0.193	0.010	0.552
Oils, Sweets and Non-Alcoholic Beverages	0.300	0.010	0.483	0.366	0.008	0.300
Alcohol, Tobacco and Narcotics	0.771	0.027	0.000 ^U	(dropped from sample)		
Clothing and Footwear	0.505	0.016	0.443	0.508	0.011	0.334
Housing and Utilities	0.662	0.119	1.115	0.679	0.083	0.937
Household Goods and Services	0.853	0.064	0.266	0.842	0.040	0.016
Health	0.675	0.028	1.083	0.670	0.019	0.888
Transport and Communication	0.521	0.055	1.128	0.527	0.037	0.983
Recreation, Culture and Education	0.947	0.267	3.321	0.973	0.348	3.387
Restaurants, Personal and Other Services	0.928	0.376	3.455	0.958	0.431	3.603
$\hat{\ell}$	4,127.751			3,669.705		
N bundles	12			11		
LR vs. NHCES	46.41			38.88		

Notes: Constrained MLE on the 12-bundle and 11-bundle 2017 ICP specifications. ^B marks regime B boundary estimates; ^U marks the regime C unbounded direction in the 12-bundle CDE that motivates the robustness exercise. For the 11-bundle system, the LR statistic for the CDE against the NHCES restriction exceeds the naive $\chi^2(10)$ critical value of 18.31.

The CDE estimates for the eleven bundles remain quantitatively close to the 12-bundle results. Boundary estimates at zero remain only for Staples and Fresh Fruit and Vegetables, and the bound-relaxation diagnostic continues to place both parameters in regime B with trajectories similar to those in the 12-bundle specification. The likelihood-ratio statistic continues to reject the NHCES restriction against the unrestricted CDE on this sample. The main qualitative findings are thus preserved after dropping the regime C bundle.

B.6.2. 10-Bundle ICP Specification: Absorbing the Vice Bundle

The 10-bundle specification absorbs Alcohol, Tobacco and Narcotics, Recreation, Culture and Education, and Restaurants, Personal and Other Services into a single residual bundle, while retaining the other nine bundles. This is a more aggressive aggregation than the 11-bundle specification. It dilutes the problematic category into a much broader residual and tests whether the unbounded direction disappears under aggregation.

TABLE B.V
ROBUSTNESS: 10-BUNDLE ICP SPECIFICATION*

Bundle	12-bundle headline			10-bundle robustness		
	$\hat{\alpha}_k$	$\hat{\beta}_k$	$\hat{e}_k$	$\hat{\alpha}_k$	$\hat{\beta}_k$	$\hat{e}_k$
Staples	0.372	0.011	0.000 ^B	0.350	0.005	0.000 ^B
Fruit and Vegetables	0.384	0.013	0.000 ^B	0.358	0.006	0.000 ^B
Meat, Seafood and Dairy	0.089	0.013	0.707	0.112	0.007	0.871
Oils, Sweets and Non-Alcoholic Beverages	0.300	0.010	0.483	0.303	0.005	0.578
Alcohol, Tobacco and Narcotics	0.771	0.027	0.000 ^U	(absorbed into residual)		
Clothing and Footwear	0.505	0.016	0.443	0.509	0.009	0.523
Housing and Utilities	0.662	0.119	1.115	0.667	0.063	1.384
Household Goods and Services	0.853	0.064	0.266	0.861	0.036	0.096
Health	0.675	0.028	1.083	0.608	0.013	1.326
Transport and Communication	0.521	0.055	1.128	0.555	0.031	1.378
Recreation, Culture and Education	0.947	0.267	3.321	(absorbed into residual)		
Restaurants, Personal and Other Services	0.928	0.376	3.455	(absorbed into residual)		
Residual: Leisure, Vices, Other Services	—			0.969	0.824	3.844
$\hat{\ell}$	4,127.751			3,271.995		
N bundles	12			10		

*Constrained MLE on the 12-bundle and 10-bundle 2017 ICP specifications. ^B marks regime B boundary estimates; ^U marks the regime C unbounded direction in the 12-bundle CDE.

The CDE estimates for the nine bundles whose composition is unchanged between the 12-bundle and 10-bundle specifications remain quantitatively close to the main specification. The merged residual category absorbs the vice direction; its non-homotheticity parameter is interior, and the bound-relaxation diagnostic on the 10-bundle specification places no parameter in regime C. As in the 11-bundle case, the main qualitative conclusions are unchanged.

B.7. Aggregation and Economic Interpretation

The alternative aggregations nevertheless change the interpretation of the system in economically important ways. The 11-bundle specification removes vice expenditure from total expenditure entirely, effectively redefining the object of estimation as a conditional non-vice demand system. This loss of information appears primarily at low expenditure levels, where the confidence intervals for the boundary necessity categories become wider (Figure B.5). The 10-bundle specification dilutes the flat Engel behavior of vice consumption by pooling it with strongly luxury-oriented service categories. In the 12-bundle specification, Alcohol, Tobacco and Narcotics remains nearly flat across the expenditure distribution, while Recreation, Culture and Education and Restaurants, Personal and Other Services exhibit clear luxury responses. Once combined into a single residual category, the merged bundle appears uniformly luxury-oriented. Figure B.4 shows that the common bundles remain quantitatively stable across aggregations, but that the 12-bundle specification preserves economically distinct Engel patterns revealed in Table I that become obscured under coarser aggregation.

B.8. External Price Variation in the Reduced-Form Equations

The reduced-form equations in Section 3.1 raise the standard price-endogeneity concern that normalized prices may be correlated with unobserved demand shocks. This concern motivates instruments in product-level demand estimation. The ICP data give a different tradeoff. The




FIGURE B.5.—CDE Income elasticities across ICP aggregations. Panels compare the specifications based on 12, 11, and 10 consumption bundles for key categories and the discretionary-services residual. Bands report 90% bootstrap bounds using m -out-of- n intervals for boundary categories and standard paired-bootstrap intervals otherwise.

data provide comparable prices and expenditure shares for broad consumption bundles across many regions, but only for a single reference year. Indeed, household and panel datasets have richer within-country variation and can support stronger instruments, but they typically do not provide the internationally comparable prices needed for the class of models in [Hanoch \(1975\)](#).

We use external cost shifters to characterize the price variation in the reduced-form equations. The exercise does not replace the constrained-likelihood estimator. It asks a narrower question. After the nuisance controls in the reduced-form equation are included, how much of the remaining normalized-price variation is related to observed supply-side variables?

For a bundle of interest k , a base bundle b , and a normalization bundle r , the share version of the log-differenced equation is

$$\log \frac{\Lambda_{ik}}{\Lambda_{ib}} = S_k + (1 - \alpha_k) \log \xi_{ik} + R_k \alpha_r \log \xi_{ir} + [(1 - R_k) \alpha_b - 1] \log \xi_{ib} + R_k \log \frac{q_{ir}}{q_{ib}} + \phi_{ik}, \quad (\text{B.7})$$

where $\xi_{ik} = p_{ik}/w_i$. We focus on the normalized price for bundle k , $x_{ik} = \log \xi_{ik}$ and estimate

$$x_{ik} = W'_{ik} \pi_k + Z'_{ik} \rho_k + v_{ik}, \quad (\text{B.8})$$

where W_{ik} includes log expenditure, the base and normalization price terms, the quantity-ratio term, and region controls where applicable. The vector Z_{ik} includes external cost shifters. We report the R^2 from the nuisance controls, the remaining share of normalized-price variation after partialling out those controls, the partial R^2 of the external shifters within the remaining variation, and the corresponding partial F statistic. We do not report second-stage IV estimates because the first-stage relationships are not strong enough to support a separate instrumental-variables estimator.

For food bundles, the shifters include crop-yield and crop-suitability measures, predetermined fertilizer exposure, and 2017 weather anomalies. For clothing, household goods, and transport-related bundles, the shifters include CEPII geography variables, UN Comtrade import composition, TASTE 2017 applied tariffs, and interactions between tariffs, import exposure, and landlocked status. These variables are intended to capture supply-side variation in price, although they are not perfect instruments.²⁸ Weather may affect local productivity and quality,

²⁸For food bundles, the external shifters combine crop suitability and potential-yield measures from FAO/IIASA Global Agro-Ecological Zones (GAEZ v4), predetermined fertilizer exposure from the World Bank WDI indicator

while tariffs and import composition may partly reflect policy and institutional differences. The exercise is best interpreted as descriptive evidence on the amount of external cost-side variation available in the one-year cross section.

TABLE B.VI
EXTERNAL COST SHIFTERS AND RESIDUAL NORMALIZED-PRICE VARIATION

Panel A: Trade-Cost Shifters						
Consumption Category	Cost Shifter	Obs.	$R^2(W)$	$1 - R^2(W)$	Partial $R^2(Z W)$	Partial F
Transport and Communication	CEPII Geography	167	0.975	0.025	0.085	7.44
Household Goods and Services	TASTE Tariffs	169	0.978	0.022	0.037	6.26
Clothing and Footwear	TASTE Tariff $\times$ Import Share	157	0.949	0.051	0.035	5.49
Panel B: Food-Price Shifters						
Consumption Category	Cost Shifter	Obs.	$R^2(W)$	$1 - R^2(W)$	Partial $R^2(Z W)$	Partial F
Fruit and Vegetables	Crop Yield	117	0.948	0.052	0.034	3.71
Meat, Seafood and Dairy	Crop Yield	129	0.985	0.015	0.031	3.73
Oils, Sweets and Non-Alcoholic Beverages	Crop Yield	106	0.981	0.019	0.027	2.56
Staples	Crop Yield, Fertilizer, Weather	125	0.978	0.022	0.019	0.43

Notes: The table characterizes external supply-side variation in selected normalized prices. The dependent variable in each equation is the normalized price of bundle k , $\log \xi_{ik} = \log(p_{ik}/w_i)$. The nuisance controls W include the base-bundle normalized price, the normalization-bundle normalized price, the quantity-ratio term, log per-capita expenditure, and region controls where applicable. $R^2(W)$ reports the explanatory power of the nuisance controls. $1 - R^2(W)$ is the remaining share of normalized-price variation before adding the external shifters. Partial $R^2(Z | W)$ reports the share of this remaining variation explained by the external shifters Z . No second-stage IV coefficients are reported. The exercise characterizes external cost-side variation in the one-year ICP cross section.

Table B.VI reports the strongest shifter set for each bundle of interest, where “strongest” refers to the specification with the largest partial F statistic in the normalized-price equation under the preferred base-good choice. Panel A reports trade-cost shifters. These variables have the broadest coverage and the strongest association with residual price variation in the exercise. CEPII geography variables explain about 8.5 percent of the residual normalized-price variation for Transport and Communication. TASTE tariffs and the tariff-by-import-share interaction explain about 3.5 to 3.7 percent of the residual normalized-price variation for Household Goods and Services and Clothing and Footwear. Panel B reports food-price shifters. Crop-yield variables explain about 3 percent of the residual normalized-price variation for several food bundles, while the richer crop, fertilizer, and weather specification is much weaker for Staples.

The nuisance controls already absorb most of the normalized-price variation. In the subsample reported in Table B.VI, the nuisance-control R^2 ranges from about 0.95 to 0.99, so the partial R^2 values are measured relative to a small residual component. A value of 0.03 means that the external shifters explain about 3 percent of the normalized-price variation remaining

AG.CON.FERT.ZS, and 2017 temperature and precipitation anomalies constructed from NASA POWER meteorological data. GAEZ v4 provides crop suitability and potential yields under rain-fed and irrigated conditions on a 5 arc-minute grid. The GAEZ variables are mapped to food bundles using representative crops: wheat, rice, maize, and sorghum for staples; maize, soybean, and wheat as feed-crop proxies for meat, seafood, and dairy; tomato, banana, and citrus for fruit and vegetables; and soybean, sunflower, and sugar cane for oils, sweets, and non-alcoholic beverages. Fertilizer exposure is measured as fertilizer consumption in kilograms per hectare of arable land averaged over 2010–2015. The weather anomalies are constructed relative to a 1991–2016 baseline using country-centroid observations. For clothing, household goods, and transport-related bundles, the shifters combine CEPII geography variables, UN Comtrade import composition, and TASTE for GTAP 11 applied tariffs and trade flows. The constructed variables include landlocked status, internal distance, import shares, import-weighted distance, applied tariffs, tariff interactions with import exposure, and tariff interactions with landlocked status.

after controlling for expenditure, base-bundle prices, normalization-bundle prices, quantity ratios, and region controls. To illustrate this more cleanly, Figure B.6 plots the same partial R^2 values. The bars are ordered by the amount of residual normalized-price variation explained by the external shifters. The reference line indicates 10 percent of the residual variation, a threshold that none of the selected shifter sets reaches.




FIGURE B.6.—External cost shifters and residual normalized-price variation. Bars report the partial R^2 of the external shifters in the normalized-price equation after partialling out nuisance controls. The vertical reference line indicates 10 percent of the residual normalized-price variation.

Our conclusion is that external cost shifters explain part of the remaining normalized-price variation, but not enough to replace the constrained-likelihood estimates in this cross section. We acknowledge that the constrained-likelihood approach is not an IV estimator. It estimates the globally regular system in the ideal aggregation where cross-region variations are observed and aggregate price comparisons across regions are available. Richer panel datasets with comparable broad-category prices would support a stronger price-identification strategy, but such data are not available for the ICP consumption bundles used in this paper. Even with richer panel variation, such as firm-level or granular demand panels across regions, identification in an implicit indirect system would remain challenging because utility enters implicitly and remains tied to both the region-specific substitution and expansion parameters as well as the endogenous aggregation term. Extending the framework to richer panel settings with stronger external instruments and alternative identification strategies is thus left for future work.

B.9. Monte Carlo Evidence on Aggregation Bias

This section provides the Monte Carlo evidence referenced in Section 3.3. Because the implicit additive system is nonlinear in expenditure, regional aggregates may generate Jensen-type aggregation bias that increases with Engel curvature and within-region expenditure dispersion. We ask whether this source of bias is large enough to threaten the income elasticities interpreted in the paper. We do not claim that regional aggregates can substitute for household micro data, which observe the within-region expenditure distribution directly. The purpose of

this section is to isolate aggregation bias, so that the only source of misspecification is replacing the household expenditure distribution with regional mean expenditure while holding the household-level CDE structure fixed.²⁹

The simulation is calibrated to the 12-bundle ICP specification and uses no external data. We fix a 12-good CDE DGP at the point estimates of Table II, simulate prices and mean expenditures for 178 regions to match the estimation sample, and draw within-region household expenditure from a lognormal distribution with mean equal to regional mean expenditure and coefficient of variation $c \in \{0, 0.25, 0.50, 0.75, 1.00\}$. For each c , we run 1,000 replications. Under the lognormal, $c = 1$ corresponds to a Gini coefficient of about 0.445, near the upper end of the consumption and income inequality measures reported for low and middle income economies (World Bank, 2024), so the range covers the dispersion most relevant for the developing economies where this paper's incidence results are most exposed to the aggregation concern. The baseline DGP is an interior version of the point estimates, with boundary expansion parameters lifted to a small positive value before the expansion parameters are renormalized.³⁰

For each region i and good k , let $s_k(p_i, x; \theta_0)$ denote the household CDE budget share at prices p_i and household expenditure x under the true parameters θ_0 . The simulated regional share is the expenditure-weighted average of household shares,

$$\bar{s}_{ik} = \frac{\int x s_k(p_i, x; \theta_0) dF_i(x)}{\int x dF_i(x)}, \quad (\text{B.9})$$

where F_i is the within-region expenditure distribution. The researcher observes only $(p_i, \bar{x}_i, \bar{s}_i)$ and fits the household share evaluated at regional mean expenditure, $s_k(p_i, \bar{x}_i; \theta)$, to the aggregate shares. When $c = 0$, all households in a region have the same expenditure, aggregate shares equal shares at the mean, and the estimator recovers θ_0 .

TABLE B.VII
AGGREGATION-BIAS MONTE CARLO FOR THE 12-BUNDLE ICP SPECIFICATION

CV	Rep.	Mean Absolute Error	Mean Maximum Error	Worst Maximum Error	Mean Share Gap	Max Share Gap
0.00	1000	0.0000	0.0000	0.0000	0.0000	0.0000
0.25	1000	0.0011	0.0078	0.0113	0.0004	0.0020
0.50	1000	0.0041	0.0277	0.0402	0.0016	0.0074
0.75	1000	0.0080	0.0533	0.0767	0.0031	0.0146
1.00	1000	0.0121	0.0793	0.1133	0.0048	0.0223

Notes: CV is the within-region coefficient of variation of household expenditure. Mean Absolute Error is the mean absolute error in implied income elasticities. Mean Maximum Error is the replication-level maximum absolute error, averaged over replications. Worst Maximum Error is the largest replication-level maximum error. Share gaps compare region-level aggregate budget shares with shares evaluated at regional mean expenditure.

Table B.VII reports the baseline design. The $c = 0$ row confirms that the estimator recovers the true elasticities when within-region dispersion is absent. As within-region dispersion rises,

²⁹The estimator in this analysis is nonlinear least squares on the budget shares with the same normalizations used in the simulation. It is not the constrained-likelihood estimator of Section 4. The goal is only to isolate the bias created by replacing household shares with regional aggregates.

³⁰The lift keeps the repeated estimation problem strictly interior while preserving the necessity–luxury ordering of the estimates. Because Staples and Fruit and Vegetables are the regime-B boundary goods of Section 7.1.1, and therefore the goods with the steepest low-income Engel curvature, we also report a more demanding design that gives them a smaller positive lift, steepening their Engel curves. This gives the aggregation-bias concern its strongest test.

the error in implied income elasticities rises smoothly and at a rate that is approximately second order in the dispersion, as a Jensen effect predicts. The mean absolute income-elasticity error is 0.0041 at $c = 0.50$ and 0.0121 at $c = 1.00$. The mean maximum error across bundles at $c = 1.00$ is 0.0793, while the worst replication reaches 0.1133. The second-to-last column reports the mean absolute aggregation gap, defined as the difference between the regional aggregate share and the share evaluated at regional mean expenditure. At $c = 1.00$, the gap is 0.0048, which confirms that the estimator is absorbing a genuine aggregation gap rather than recovering θ_0 because the inputs are nearly identical.




FIGURE B.7.—Aggregation bias and within-region expenditure dispersion. The figure reports mean income-elasticity errors and mean maximum errors across 1,000 replications. The baseline design is calibrated to the 12-bundle ICP point estimates. The steep-necessity design assigns a smaller positive lift to Staples and Fruit and Vegetables before normalizing expansion parameters.

Figure B.7 shows the evolution of income-elasticity errors with within-region expenditure dispersion under the baseline and steep-necessity designs. The errors rise with dispersion but remain small. The mechanism is straightforward. Aggregation bias is limited because the estimated CDE system has mild Engel curvature, so the income elasticities move slowly with expenditure. Calibrating the simulation to the estimated 12-bundle system thus produces much smaller bias than a generic DGP with sharper curvature. Figure B.8 reports good-level recovery at $c = 0.50$. Staples and Fruit and Vegetables lie close to the 45-degree line, as do the higher-elasticity service categories. This is the relevant comparison because the paper interprets implied income elasticities rather than the primitive parameters directly. Table B.VIII reports the boundary-necessity robustness exercise. The steep-necessity design preserves the 12-bundle structure but steepens the low-income necessity curves. The errors are larger than in the baseline but remain bounded. At $c = 1.00$, the mean absolute income-elasticity error is 0.0129, the mean maximum error is 0.0939, and the worst replication reaches 0.1503.

Our central conclusion is twofold. First, household micro data would be preferable because they observe the within-region distribution directly. Second, for the curvature implied by the estimated 12-bundle system, however, the aggregation bias in implied income elasticities is small relative to the differences the paper interprets. Even at $c = 1.00$, a consumption Gini of about 0.45, the baseline mean absolute bias of 0.0121 is more than an order of magnitude below the CDE–NHCES elasticity gaps of 0.2 to 0.7 that motivate the structural comparisons in Section 7.3. The largest robustness case should be read separately rather than folded into that comparison. Under the steep-necessity design the worst single-bundle error in the worst replication is 0.1503, which remains below the smallest of those interpreted gaps. The small




FIGURE B.8.—True and recovered mean income elasticities at within-region expenditure CV = 0.50. Each point is one ICP consumption bundle. The 45-degree line indicates exact recovery.

TABLE B.VIII
STEEP-NECESSITY AGGREGATION-BIAS ROBUSTNESS

CV	Rep.	Mean Absolute Error	Mean Maximum Error	Worst Maximum Error	Mean Share Gap	Max Share Gap
0.00	1000	0.0000	0.0000	0.0002	0.0000	0.0000
0.25	1000	0.0012	0.0093	0.0152	0.0004	0.0021
0.50	1000	0.0043	0.0330	0.0538	0.0016	0.0078
0.75	1000	0.0085	0.0632	0.1021	0.0031	0.0154
1.00	1000	0.0129	0.0939	0.1503	0.0048	0.0236

Notes: Same as Table B.VII.

magnitude is not evidence that aggregation is innocuous in general. It reflects the fact that the Engel curves of the estimated specification are not sharply curved at the scale of within-region dispersion considered in the simulation.

APPENDIX C: COMPUTATIONAL AND ESTIMATION DETAILS

C.1. Numerical Scaling and Economic Robustness

Previous empirical implementations of implicitly additive demand systems, such as Cranfield et al. (2002) and Preckel, Cranfield, and Hertel (2010), note that the right-hand side of the implicit additivity defining equation can be normalized to a positive constant. Given the excess degrees of freedom described in equation (3.7), the normalization $\sum_k \beta_k = 1$ is a convenient choice that also sets the right-hand side of the defining relationship equal to one. In this section, we verify that this normalization is numerical rather than economic by jointly scaling the right-hand side of the defining equation and the normalization on the distribution parameters.

Specifically, for each $\mathcal{M} \in \{1, 10, 100, 1000, 10000\}$, we estimate under

$$G\left(\frac{p_k}{w}, u_i\right) \equiv \sum_k \beta_k \equiv \mathcal{M} \in \mathbb{R}_{++} \quad \forall i, \quad (\text{C.1})$$

while maintaining the normalization $\sum_k e_k = N$ and the same regularity restrictions used in the baseline estimation. This exercise scales all β_k 's by the same factor while leaving the share equations and elasticity formulas economically unchanged. Hence, if the normalization is only removing an excess degree of freedom, the likelihood and implied elasticities should be invariant to $\mathcal{M}$ up to numerical precision.

C.2. Log-Likelihood Values

The baseline estimates in Table II correspond to $\mathcal{M} = 1$. Re-estimating the system under alternative scaling constants $\mathcal{M} \in \{1, 10, 100, 1000, 10000\}$ yields the same maximized objective value, $\hat{\ell}_{\text{CDE}} = 4,127.751$, up to numerical precision. The solver also returns normal completion status in all cases. This confirms that the normalization does not affect the fitted demand system and only rescales the distribution parameters β_k . The estimated α_k 's and e_k 's remain unchanged, while the β_k 's scale proportionally with $\mathcal{M}$.

C.3. Testing for Income Elasticity of Demand

We compute income elasticities given by

$$\eta_k = \frac{e_k(1 - \alpha_k) + \sum_m \hat{\Lambda}_m e_m \alpha_m}{\sum_m \hat{\Lambda}_m e_m} + \alpha_k - \sum_m \hat{\Lambda}_m \alpha_m, \quad (\text{C.2})$$

where η_k is the income elasticity of commodity k and $\hat{\Lambda}_k$ denotes the fitted expenditure share for that commodity.

As in the main analysis, we evaluate income elasticities using fitted utilities and fitted shares at a common average-price reference. Holding the estimated parameters fixed, we solve

$$\begin{aligned} & \text{maximize} \quad \sum_i u_i \\ & \text{subject to:} \quad \sum_k \hat{\beta}_k u_i^{\hat{e}_k(1-\hat{\alpha}_k)} \left(\frac{\bar{p}_k}{w_i} \right)^{1-\hat{\alpha}_k} - \mathcal{M} = 0 \quad \forall i, \end{aligned} \quad (\text{C.3})$$

where $\bar{p}_k = \sum_i p_{ik}/I$ is the average price for good k . We then compute fitted shares using equations (2.3)–(2.4) and substitute them into equation (C.2).

Table C.I compares the average income elasticities under $\mathcal{M} = 1$ and $\mathcal{M} = 10,000$. The results are exactly unchanged to the precision reported here. This confirms that the economic implications of the estimated demand system are invariant to this numerical scaling of the defining relationship.

C.4. Testing for Own-Price Elasticity of Demand

The Allen-Uzawa elasticities of substitution are given by

$$\sigma_{km} = \alpha_k + \alpha_m - \sum_n \Lambda_n \alpha_n - \frac{\delta_{km} \alpha_k}{\Lambda_k}, \quad (\text{C.4})$$

where δ_{km} is the Kronecker delta (equals 0 if $k \neq m$, 1 if $k = m$). The uncompensated price elasticity of substitution is

$$\epsilon_{km}^{\text{unc}} = (\sigma_{km} - \eta_k) \hat{\Lambda}_m, \quad (\text{C.5})$$

with the own-price elasticity obtained by setting $m = k$.

Table C.II reports the average uncompensated own-price elasticities under $\mathcal{M} = 1$ and $\mathcal{M} = 10,000$. As with the income elasticities, the values are unchanged to the precision reported

TABLE C.I
ROBUSTNESS OF INCOME ELASTICITIES

Consumption Bundle	$M = 1$	$M = 10,000$	Change (%)	Abs. Change
Staples	0.53	0.53	0.00	0.00
Fruit and Vegetables	0.54	0.54	0.00	0.00
Meat, Seafood and Dairy	0.80	0.80	0.00	0.00
Oils, Sweets and Non-Alcoholic Beverages	0.75	0.75	0.00	0.00
Clothing and Footwear	0.85	0.85	0.00	0.00
Alcohol, Tobacco and Narcotics	0.93	0.93	0.00	0.00
Household Goods and Services	1.05	1.05	0.00	0.00
Health	1.14	1.14	0.00	0.00
Transport and Communication	1.15	1.15	0.00	0.00
Recreation, Culture and Education	1.26	1.26	0.00	0.00
Restaurants, Personal and Other Services	1.30	1.30	0.00	0.00
Housing and Utilities	1.15	1.15	0.00	0.00

Notes: Values are averages across regions using fitted shares evaluated at average prices. Results are invariant to the scaling normalization up to numerical precision.

here. The same invariance holds for the compensated own-price elasticities. Thus, changing $\mathcal{M}$ affects the scale of the distribution parameters but not the economic behavior of the estimated demand system.

TABLE C.II
ROBUSTNESS OF UNCOMPENSATED OWN-PRICE ELASTICITIES

Consumption Bundle	$M = 1$	$M = 10,000$	Change (%)	Abs. Change
Staples	-0.40	-0.40	0.00	0.00
Fruit and Vegetables	-0.41	-0.41	0.00	0.00
Meat, Seafood and Dairy	-0.20	-0.20	0.00	0.00
Oils, Sweets and Non-Alcoholic Beverages	-0.35	-0.35	0.00	0.00
Clothing and Footwear	-0.53	-0.53	0.00	0.00
Alcohol, Tobacco and Narcotics	-0.77	-0.77	0.00	0.00
Household Goods and Services	-0.85	-0.85	0.00	0.00
Health	-0.69	-0.69	0.00	0.00
Transport and Communication	-0.62	-0.62	0.00	0.00
Recreation, Culture and Education	-0.94	-0.94	0.00	0.00
Restaurants, Personal and Other Services	-0.93	-0.93	0.00	0.00
Housing and Utilities	-0.74	-0.74	0.00	0.00

Notes: Values are averages across regions using fitted shares evaluated at average prices. Results are invariant to the scaling normalization up to numerical precision.

C.5. Latitude for Changing Parameters

We develop a structural method to determine the latitude for changing parameters. In considering the potential identification of the parameters in our preference relationship, we demonstrated in Section 4.1 that at least two normalizations are required to remove excess degrees of freedom from the system. We now ask whether additional normalizations should be imposed. To this end, we investigate whether the same log-likelihood level can be achieved under different configurations of parameter values. While this is an exceedingly difficult question to answer without assurance that the problem is convex, we approach it by starting from the estimated parameter solution.

Let $\widehat{\theta}$ denote the baseline CDE solution, where

$$\theta = (\{\alpha_k\}_{k=1}^N, \{\beta_k\}_{k=1}^N, \{e_k\}_{k=1}^N, \{u_i\}_{i=1}^I),$$

and let $\widehat{\ell}_{\text{CDE}}$ denote the corresponding maximized objective value. For each scalar component θ_j of the estimated solution, we solve two auxiliary constrained optimization problems:

$$\bar{\theta}_j = \max_{\theta} \theta_j, \quad \underline{\theta}_j = \min_{\theta} \theta_j,$$

subject to the original CDE likelihood objective (4.1), the residual and covariance restrictions (4.2)–(4.3), the implicit CDE defining relationship (2.1), the auxiliary share identities (2.3)–(2.4), the same regularity restrictions used in the baseline estimation, and the normalizations

$$\sum_k \beta_k = 1, \quad \sum_k e_k = N.$$

We additionally require the likelihood to remain within numerical tolerance of the baseline optimum:

$$\ell(\theta) \geq \widehat{\ell}_{\text{CDE}} - \tau_{\ell},$$

where τ_{ℓ} is a small numerical tolerance, while all other parameters and latent utility levels are allowed to adjust when θ_j is maximized or minimized.

For each parameter, we summarize the resulting latitude by

$$\Delta_j = \max \left\{ |\bar{\theta}_j - \widehat{\theta}_j|, |\underline{\theta}_j - \widehat{\theta}_j| \right\}.$$

A small value of Δ_j shows that the parameter cannot move appreciably without reducing the likelihood or violating the structural constraints. A large value would indicate a remaining local flat direction and would suggest that additional normalization or reparameterization may be required.

The completed auxiliary problems show that the admissible movements are numerically small. Most structural parameters move by roughly 10^{-3} , and the largest movements are on the order of 10^{-2} . These results show that substantial changes in the estimated parameters cannot be achieved while preserving the fitted likelihood and satisfying the problem constraints. We thus find no evidence of an additional local degree of freedom beyond the two normalizations imposed on $\{\beta_k\}$ and $\{e_k\}$.

C.6. Multistart Procedure

This section reports the multistart replications that support the global optimality claim in Section 6.1.1. For each specification, we draw $S = 100$ random starting points from (6.1), solve the constrained problem from each start, and report the share of replications whose objective value lies within 10^{-3} and 10^{-2} of the best solution found.

TABLE C.III
DISTRIBUTION OF OBJECTIVE VALUES ACROSS RANDOM STARTS*

Specification	Best $\hat{\ell}$ Found	Within 10^{-3}	Within 10^{-2}	Inferior Local Optimum	Range of Accepted Obj.
<i>Unrestricted CDE</i>					
12-bundle ICP (main specification)	4,127.751	91/100	91/100	9	1.681
11-bundle ICP (drop AlcTob)	3,669.705	88/100	88/100	12	0.089
10-bundle ICP (absorb residual)	3,271.995	100/100	100/100	0	0.000
<i>Restricted NHCES ($\alpha_k = \alpha$)</i>					
12-bundle ICP	4,104.547	100/100	100/100	0	0.000
11-bundle ICP	3,650.266	100/100	100/100	0	0.000
10-bundle ICP	3,257.055	100/100	100/100	0	0.000

Notes: Each row reports diagnostics from $S = 100$ multistart replications drawn from the distribution in (6.1). Columns 3 and 4 report the share of replications whose objective is within 10^{-3} and 10^{-2} of the best objective found across all 100 starts. “Inferior local optimum” counts the replications with objective gap $\geq 10^{-2}$. The last column is the spread between the highest and lowest accepted objective values (those replications that completed with normal solver status). The CDE solution for the main specification is essentially unique on the regular domain. 91 of 100 starts reach the global optimum within 10^{-3} and the remaining 9 converge to strictly inferior local maxima with implausible parameter signs (large negative $\hat{e}_k$ on bundles whose data clearly support a positive non-homotheticity).

C.7. *Boundary Inference for Interior Parameters*

TABLE C.IV

m-OUT-OF-*n* VERSUS ORDINARY BOOTSTRAP FOR INTERIOR e_k PARAMETERS

Consumption Bundle	$\hat{e}_k$	$U_{0.90}^{m/n}$	$U_{0.90}^{n/n}$	$U_{0.95}^{n/n}$
Meat, Seafood, and Dairy	0.707	0.816	0.867	0.909
Oils, Sweets, and Non-Alc. Beverages	0.483	0.620	0.681	0.733
Clothing and Footwear	0.443	0.657	0.734	0.799
Housing and Utilities	1.115	1.465	1.440	1.542
Household Goods and Services	0.266	0.716	0.871	0.989
Health	1.083	1.807	1.690	1.872
Transport and Communication	1.128	1.337	1.376	1.426
Recreation, Culture, and Education	3.321	4.030	4.095	4.560
Restaurants, Personal, and Other Services	3.455	4.610	6.234	7.437

Notes: The table reports inferences for regime-A non-homotheticity parameters whose full-sample estimates are interior to the regularity bound. $U_{0.90}^{m/n}$ is the one-sided endpoint computed as $\hat{e}_{k,n} + q_{0.90} \{ \sqrt{m/n} (\hat{e}_{k,m}^{(b)} - \hat{e}_{k,n}) \}$ using $m = 38$ and 2,000 successful subsamples. $U_{0.90}^{n/n}$ is the ordinary paired-bootstrap one-sided 90% endpoint, and $U_{0.95}^{n/n}$ is the upper endpoint of the ordinary two-sided 90% percentile interval. The diagnostic is not used for inference on interior parameters; it checks the finite-sample behavior of the m -out-of- n routine.

The main 12-bundle CDE specification on the ICP sample is effectively unique on the regular domain. Ninety-one of 100 starting values converge to the global optimum within 10^{-3} , while the remaining nine converge to strictly inferior local maxima with implausible parameter signs. The 11-bundle specification behaves similarly, with 88 of 100 starting values converging to the global optimum. The 10-bundle specification converges to the same optimum from every accepted starting value. Across all samples, the restricted NHCES specification converges to a unique optimum from every accepted start, consistent with the lower-dimensional substitution-parameter space implied by the common- α restriction.

APPENDIX D: ADDITIONAL RESULTS FOR THE PANEL ESTIMATOR

D.1. *Country-Cluster Inference*

Because identification comes from within-country variation, inference is computed at the country level. We thus use country-clustered resampling. Each draw resamples countries and keeps all years for a sampled country together, which preserves serial dependence and the within-country source of identification. For interior parameters we report country-cluster percentile intervals. For parameters at the regularity boundary we apply the same one-sided m -out-of- n procedure used in the cross section, now resampling countries: with 39 countries we set $m = \text{round}(39^{0.7}) = 13$.

D.2. *A Partially Relaxed Substitution Structure*

To understand where the unrestricted CDE departs from the restricted benchmark, we estimate a specification that allows agriculture to have its own substitution parameter while maintaining a common substitution parameter for manufacturing and services. The best solution gives $\hat{\alpha}_A = 0.545$ and $\hat{\alpha}_M = \hat{\alpha}_S = 0.768$, with fitted mean elasticities (0.553, 0.923, 1.049).

Agriculture still sits at the lower expansion bound. Allowing the agricultural substitution parameter to differ from the non-agricultural pair improves the fit relative to the restricted NHCES, but does not deliver the full flexibility of the unrestricted CDE. The remaining gain thus reflects substitution heterogeneity within the non-agricultural sectors as well.

D.3. Alternative Normalizations of the Expansion Parameters

The implicit CDE system requires a normalization on the expansion parameters because utility is latent. The main panel estimates use $e_M = 1$ to match the normalization in [Comin et al. \(2021\)](#). We re-estimate the unrestricted panel CDE under three alternatives: the cross-sectional normalization $\sum_k e_k = 3$, a share-weighted normalization $\sum_k \bar{s}_k e_k = 1$, and an exponent normalization $\sum_k e_k(1 - \alpha_k) = 3$. To make the comparison meaningful, we replace the exact boundary condition $\alpha_M = 0$ with a small positive lower bound so that the expansion-parameter ratios is well defined under all normalizations. The scale of the estimated expansion parameters changes across specifications, but the fitted elasticities do not. Agriculture remains at the lower expansion bound, manufacturing remains at the lower substitution bound, and the fitted elasticities at (0.64, 1.00, 1.15) in all four cases (Table D.I). The boundary pattern is a feature of the data, not of the manufacturing normalization.

TABLE D.I

PANEL CDE ESTIMATES UNDER ALTERNATIVE EXPANSION-PARAMETER NORMALIZATIONS

Normalization	$\hat{\alpha}_A$	$\hat{\alpha}_M$	$\hat{\alpha}_S$	Expansion Scale $\hat{e}_M$	Relative $\hat{e}_A/\hat{e}_M$	Relative $\hat{e}_S/\hat{e}_M$	Fitted-Share Elasticities (A, M, S)
$e_M = 1$	0.560	0.020 ^B	0.960	1.000	0.000 ^B	2.993	(0.644, 0.998, 1.153)
$\sum_k e_k = 3$	0.560	0.020 ^B	0.960	0.752	0.000 ^B	2.993	(0.644, 0.998, 1.153)
$\sum_k \bar{s}_k e_k = 1$	0.560	0.020 ^B	0.960	0.548	0.000 ^B	2.993	(0.644, 0.998, 1.153)
$\sum_k e_k(1 - \alpha_k) = 3$	0.560	0.020 ^B	0.960	2.728	0.000 ^B	2.993	(0.644, 0.998, 1.153)

Notes: Each row estimates the same panel CDE system under a different normalization of the expansion vector e . Because utility is latent, the scale of e is not separately identified. The invariant objects are the relative expansion parameters and the fitted elasticities. The fourth row normalizes the utility exponents $\gamma_k = e_k(1 - \alpha_k)$. Superscript B denotes an estimate at the implemented regularity bound.

D.4. Observed-Share, Fitted-Share, and Fixed-Effect-Adjusted Income Elasticities

We discuss three income-elasticity measures specialized to the panel framework. The first is the observed-share elasticity, which provides the closest comparison to [Comin et al. \(2021\)](#). The second is the fitted-share elasticity, which mirrors the approach used in the ICP cross section. The third is the fixed-effect-adjusted elasticity, which traces the expenditure response while holding the estimated country-sector effects constant.

For a generic share vector $\Lambda_{ct} = (\Lambda_{ctA}, \Lambda_{ctM}, \Lambda_{ctS})$, the CDE income elasticity is

$$\eta_{ctk}(\Lambda) = \frac{e_k(1 - \alpha_k) + \sum_{j \in \mathcal{K}} \Lambda_{ctj} e_j \alpha_j}{\sum_{j \in \mathcal{K}} \Lambda_{ctj} e_j} + \alpha_k - \sum_{j \in \mathcal{K}} \Lambda_{ctj} \alpha_j, \quad k \in \{A, M, S\}. \quad (\text{D.1})$$

The same estimated primitive vector can imply different average elasticities depending on which share vector enters (D.1). In the observed-share evaluation, we set $\Lambda_{ctk} = \omega_{ctk}$, where ω_{ctk} is the observed country-year sectoral share. In the fitted-share evaluation, we set $\Lambda_{ctk} = \hat{s}_{ctk}$, where $\hat{s}_{ctk}$ is the fitted share from the constrained panel system.

Under the NHCES restriction, $\alpha_k = \alpha$ for all k , equation (D.1) reduces to

$$\eta_{ctk}^{NHCES}(\Lambda) = \alpha + (1 - \alpha) \frac{e_k}{\sum_{j \in \mathcal{K}} \Lambda_{ctj} e_j}. \quad (\text{D.2})$$

The share vector matters through the scalar denominator. The dependence is much stronger in CDE. With sector-specific α_k 's, the share vector enters both through the denominator $\sum_j \Lambda_{ctj} e_j$ and through the share-weighted substitution terms. The observed-share and fitted-share evaluations diverge much more under CDE than under NHCES.

With manufacturing normalization, the relevant estimates are approximately

$$(e_A, e_M, e_S) = (0, 1, 2.993), \quad (\alpha_A, \alpha_M, \alpha_S) = (0.560, 0, 0.960).$$

Substituting them into (D.1) gives

$$\eta_{ctM}(\Lambda) - \eta_{ctA}(\Lambda) = \frac{1}{\sum_j \Lambda_{ctj} e_j} - 0.560. \quad (\text{D.3})$$

The agriculture–manufacturing elasticity gap is sensitive to the share-weighted expansion term in the denominator. Under observed-share evaluation, the denominator is large enough to compress the implied difference between the two sectors. Under fitted-share evaluation, the denominator changes and manufacturing moves back toward unit elasticity.

The fitted-share evaluation is the direct analogue to the income-elasticity calculations reported for the ICP cross section. In the ICP cross section, the income elasticities are evaluated at a common reference price vector and use fitted utilities and fitted shares from the structural demand system. For the panel, we perform the same calculation before applying the country-sector fixed-effect wedges. Let $\bar{p}$ denote the reference price vector and let w vary over the observed panel expenditure range. For each w , the fitted-share evaluation solves

$$\sum_{k \in \mathcal{K}} \hat{\beta}_k u(w, \bar{p})^{\hat{e}_k(1-\hat{\alpha}_k)} \left(\frac{\bar{p}_k}{w} \right)^{1-\hat{\alpha}_k} = 1, \quad (\text{D.4})$$

then forms

$$\hat{s}_k(w, \bar{p}) = \frac{\hat{\beta}_k (1 - \hat{\alpha}_k) u(w, \bar{p})^{\hat{e}_k(1-\hat{\alpha}_k)} \left(\frac{\bar{p}_k}{w} \right)^{1-\hat{\alpha}_k}}{\sum_{j \in \mathcal{K}} \hat{\beta}_j (1 - \hat{\alpha}_j) u(w, \bar{p})^{\hat{e}_j(1-\hat{\alpha}_j)} \left(\frac{\bar{p}_j}{w} \right)^{1-\hat{\alpha}_j}}. \quad (\text{D.5})$$

Substituting $\hat{s}_k(w, \bar{p})$ into (D.1) gives the fitted-share income elasticities.

D.4.1. Fixed-Effect-Adjusted Income Elasticities

The panel estimator includes country-sector fixed effects, which absorb persistent country-sector level differences in sectoral allocation, so the relevant panel schedule holds the estimated country-sector effects fixed and varies expenditure given those estimates.

Let $s_k(w, \bar{p})$ denote the structural share implied by either CDE or NHCES system at expenditure w and reference price vector $\bar{p}$. Let $\hat{\delta}_k$ denote the sector intercept and $\hat{\zeta}_{ck}$ the country-sector

fixed effect, with $\hat{\delta}_M = \hat{\zeta}_{cM} = 0$. The fixed-effect-adjusted share is

$$\tilde{s}_{ck}(w, \bar{p}) = \frac{\exp(\hat{\delta}_k + \hat{\zeta}_{ck})s_k(w, \bar{p})}{\sum_{j \in \{A, M, S\}} \exp(\hat{\delta}_j + \hat{\zeta}_{cj})s_j(w, \bar{p})}, \quad (\text{D.6})$$

and the fixed-effect-adjusted income elasticity is

$$\tilde{\eta}_{ck}(w, \bar{p}) = \eta_k(w, \bar{p}) + 1 - \sum_{j \in \{A, M, S\}} \tilde{s}_{cj}(w, \bar{p})\eta_j(w, \bar{p}), \quad (\text{D.7})$$

where $\eta_k(w, \bar{p})$ is the structural income elasticity from the demand system before fixed-effect adjustment. The plotted income elasticities average (D.7) across countries, and bootstrap bands apply the same transformation within each draw. Boundary sectors use the one-sided m -out-of- n country-subsampling interval. Interior sectors use the standard country-cluster percentile interval.

The CDE and NHCES elasticity curves lie closer together in the panel than in the ICP cross section (Figure D.1). The cross-region Engel curvature that sharply separates the two systems in the ICP data is absorbed by the country-sector fixed effects, so identification comes primarily from within-country movements in expenditure and sectoral allocation. The sectoral ranking is robust. Agriculture is below one, services are above one, and manufacturing lies in the middle.




FIGURE D.1.—Fixed-effect-adjusted three-sector income elasticities for the constrained CDE and NHCES panel estimators. Shaded areas are 90 percent country-cluster bootstrap intervals. Boundary sectors use the one-sided m -out-of- n country-subsampling interval.

Table D.II reports these three average elasticities. Panel A uses observed shares and is the object used for comparison with the CLM benchmark in Table VIII. Panel B uses fitted shares from the constrained panel program. Panel C reports the mean fixed-effect-adjusted elasticity schedule used in Figure D.1. The three measures are closely related, but differ in how the share vector and country-sector fixed effects enter the elasticity calculation.

TABLE D.II
PANEL MEAN INCOME ELASTICITIES UNDER ALTERNATIVE SHARE EVALUATIONS

System	Evaluation	Agriculture	Manufacturing	Services
Panel A: Observed-Share				
NHCES	$\eta(\omega)$	0.568	1.021	1.175
CDE	$\eta(\omega)$	0.766	0.791	1.236
Panel B: Fitted-Share				
NHCES	$\eta(\hat{s})$	0.568	0.923	1.044
CDE	$\eta(\hat{s})$	0.644	0.998	1.153
Panel C: Fixed-Effect-Adjusted				
NHCES	$\tilde{\eta}(w, \bar{p})$	0.711	1.067	1.188
CDE	$\tilde{\eta}(w, \bar{p})$	0.696	1.052	1.206

The fixed-effect-adjusted measure in Panel C is the internally consistent elasticity measure for the panel estimator. The income elasticity in equation (8.10) is share dependent, so the elasticity reported for a country depends on its position in sectoral-allocation space. The country-sector effects place each country at a persistent operating point by absorbing relative tastes, sectoral labor-share wedges, and country-specific technology. Equation (D.7) thus measures the elasticity of the fitted demand assigned to that operating point, obtained by differentiating the fixed-effect-adjusted share with respect to log expenditure. Unlike Panel B, which evaluates an intermediate measure before the fixed-effect wedges are applied, Panel C evaluates the expenditure response of the fitted panel demand directly. It also aligns with the panel identification strategy, which holds country-sector levels fixed and identifies the parameters from within-country movements in expenditure and relative allocation.

The adjustment preserves the theoretical structure of the system. The term $1 - \sum_j \tilde{s}_{cj} \eta_j$ preserves adding-up at the adjusted shares, so the fixed-effect-adjusted elasticities satisfy $\sum_k \tilde{s}_{ck} \tilde{\eta}_{ck} = 1$ in every country. The three elasticity measures therefore share the same adding-up restriction and differ only in how non-homotheticity is allocated across sectors.

APPENDIX E: COUNTERFACTUAL DETAILS

This section reports the data construction, parameter reconciliation, additional elasticity results, and welfare decompositions underlying the GTAP counterfactual exercises in Section 9.

E.1. Data Construction and Aggregation Choice

We construct the estimation dataset from private household expenditure for 65 sectors and 160 regions in the 2017 GTAP database. For sector s in region i , let $VDPP_{si}$ and $VMPP_{si}$ denote domestic and imported household purchases at consumer prices, and let $VDPB_{si}$ and $VMPB_{si}$ denote these purchases evaluated at producer prices. We define

$$X_{is}^P = VDPP_{si} + VMPP_{si}, \quad X_{is}^B = VDPB_{si} + VMPB_{si}.$$

where X_{is}^P is total household expenditure at consumer prices, while X_{is}^B is the corresponding value at producer prices. The consumer-producer price wedge is the only source of bundle-level price variation available in the GTAP household data and thus we let them enter the estimation as the bundle price measure.

Given a candidate aggregation, we assign each GTAP sector s to an aggregate consumption bundle G_k . For aggregate bundle k , household expenditure is

$$\Xi_{ik}^{GTAP} = 10^6 \sum_{s \in G_k} X_{is}^P. \quad (\text{E.1})$$

The factor 10^6 converts GTAP million-dollar values into dollar values and does not affect expenditure shares.

The corresponding bundle-level consumer-price measure is

$$P_{ik}^{raw} = \frac{\sum_{s \in G_k} X_{is}^P}{\sum_{s \in G_k} X_{is}^B}. \quad (\text{E.2})$$

The numerator is total household expenditure at consumer prices within bundle k , while the denominator is the corresponding expenditure at producer prices. We normalize these bundle prices to the United States within each category:

$$p_{ik} = \frac{P_{ik}^{raw}}{P_{US,k}^{raw}}. \quad (\text{E.3})$$

The normalized price p_{ik} equals one for the United States in every bundle. Since the values are already expressed in current U.S. dollars, no additional exchange-rate conversion is required.

The expenditure share used in the aggregation is

$$s_{ik} = \frac{\Xi_{ik}^{GTAP}}{\sum_{\ell} \Xi_{i\ell}^{GTAP}}.$$

The share s_{ik} measures the weight of bundle k in region i 's household expenditure. For each candidate aggregation, we evaluate expenditure-share variation, thin-region exposure, cross-region price dispersion, and Engel correlations. We rank candidate aggregations by calculating

the following score

$$\begin{aligned} \text{Score} = & 100\bar{T}_{1\%} + 50T_{1\%}^{\max} + 2R_{p95/p05}^{\max} + 20\sigma_p^{\max} \\ & + 50\max\{0, 0.04 - \bar{s}^{\min}\} + 6 \times \mathbf{1}\{n_{\text{sectors}}^{\min} = 1\} \\ & + 10\max\{0, 8 - N_G\} + 1.5\max\{0, N_G - 9\} - 4|\bar{\rho}|. \end{aligned}$$

The term $\bar{T}_{1\%}$ is the average fraction of regions in which a bundle share falls below 1 percent, while $T_{1\%}^{\max}$ is the largest fraction across bundles. $R_{p95/p05}^{\max}$ is the largest 95th-to-5th percentile price ratio across bundles, and $\sigma_p^{\max}$ is the largest cross-region price standard deviation. $\bar{s}^{\min}$ is the smallest mean expenditure share across bundles, $n_{\text{sectors}}^{\min}$ is the smallest number of GTAP sectors in any bundle, N_G is the number of aggregate bundles, and $|\bar{\rho}|$ is the average absolute Engel correlation. Lower values indicate aggregations with more stable prices, more balanced expenditure shares, and clearer Engel variation.³¹ Table E.I reports ranking scores.

TABLE E.I
CANDIDATE AGGREGATIONS AND SELECTION CRITERIA

Rank	Candidate Aggregation	Goods	Min Sectors	Avg Thin < 1%	Max Thin < 1%	Max P95/P05 Price	Score
1	Food aggregate with separate services split	8	2	0.010	0.056	1.421	8.072
2	Detailed food split with separate services	10	2	0.008	0.056	1.829	13.298
3	Compact 7-good aggregation	7	2	0.009	0.056	1.421	17.856
4	Compact 6-good aggregation	6	5	0.001	0.006	1.421	24.361
5	Separate beverages and tobacco with merged services	8	1	0.014	0.056	2.495	26.783
6	Separate beverages and tobacco with service split	9	1	0.015	0.056	2.495	26.903
7	User baseline aggregation	10	2	0.029	0.219	2.295	28.548
8	Oils merged with beverages and tobacco	11	2	0.028	0.219	2.295	30.035
9	Leisure and other services merged	11	1	0.065	0.381	2.495	51.707
10	ICP-style 12-good aggregation	12	1	0.061	0.381	2.495	52.889

Specifically, for readers interested in the details, the coefficient 100 on $\bar{T}_{1\%}$ places a strong penalty on bundles that frequently account for less than 1 percent of expenditure across regions, while the coefficient 50 on $T_{1\%}^{\max}$ additionally penalizes aggregations with a single especially thin bundle. The coefficient 50 on $\max\{0, 0.04 - \bar{s}^{\min}\}$ discourages bundles with extremely small average expenditure shares. These terms place the largest penalties on weak expenditure support, since bundles with little expenditure variation are difficult to estimate reliably in demand systems. The coefficients 2 and 20 on $R_{p95/p05}^{\max}$ and $\sigma_p^{\max}$ penalize unstable cross-region prices, although less heavily than thin expenditure shares. The discrete penalty $6 \times \mathbf{1}\{n_{\text{sectors}}^{\min} = 1\}$ discourages one-sector bundles, since such categories inherit sector-specific consumer-producer price wedges rather than broader bundle-level price variation. The terms $10\max\{0, 8 - N_G\}$ and $1.5\max\{0, N_G - 9\}$ impose asymmetric penalties on the number of aggregate goods. Aggregations with fewer than eight goods are penalized much more heavily than aggregations with more than nine goods because excessively coarse aggregation compresses economically meaningful Engel heterogeneity. Having somewhat more than nine goods

³¹The weights are heuristic and chosen to prioritize stable expenditure support and clear Engel variation while penalizing thin bundles, unstable prices, and excessively coarse or fine aggregations. Thin-region exposure and very small expenditure shares receive the largest penalties because bundles with weak expenditure support are difficult to estimate reliably in the data. One-sector bundles are penalized because they inherit sector-specific consumer-producer price wedges rather than broader bundle-level price variation. Additional penalties discourage aggregations that are excessively coarse or excessively fine relative to the existing sector structure in the data. The negative coefficient on $|\bar{\rho}|$ rewards clear Engel variation across regions, which is helpful for identifying non-homothetic demand patterns.

mainly raises the risk of thin bundles, unstable prices, and one-sector categories, which are already penalized separately through expenditure support, price dispersion, and one-sector penalties. The coefficient 1.5 thus acts mainly as a soft penalty against overly fine splits unless the data actually support them. Finally, the coefficient -4 on $|\bar{\rho}|$ rewards strong Engel variation across regions, which is informative for identifying non-homothetic demand patterns.

The preferred aggregation pools all food-related sectors together with beverages and tobacco into one broad food bundle while keeping the major nonfood categories separate. Relative to the ICP basic-heading system, some GTAP categories are already highly aggregated, while others are narrow and exposed to sector-specific price wedges. Finer food splits create thin bundles and unstable one-sector categories, especially for beverages and tobacco. Coarser aggregations are statistically more stable, but they merge categories with different Engel patterns and therefore obscure the economically meaningful expenditure variation that appears clearly in the 12-bundle ICP specification in the main analysis. The selected 8-good mapping is reported in Table E.II. The table maps the 65 GTAP sectors into the eight aggregate consumption bundles used in the estimation. Table E.III provides summary statistics for the 8-good aggregation. Figure E.1 shows consumption patterns across income groups.

TABLE E.II: Mapping from GTAP 65 Sectors to the 8-Good Aggregation

GTAP	Description	8-Good Aggregate
pdr	Paddy rice	Food, Beverages, and Tobacco
wht	Wheat	Food, Beverages, and Tobacco
gro	Cereal grains nec	Food, Beverages, and Tobacco
pcr	Processed rice	Food, Beverages, and Tobacco
ofd	Food products nec	Food, Beverages, and Tobacco
ctl	Cattle, sheep, goats, horses	Food, Beverages, and Tobacco
oap	Animal products nec	Food, Beverages, and Tobacco
rmk	Raw milk	Food, Beverages, and Tobacco
cmt	Bovine meat products	Food, Beverages, and Tobacco
omt	Meat products nec	Food, Beverages, and Tobacco
fsh	Fishing	Food, Beverages, and Tobacco
mil	Dairy products	Food, Beverages, and Tobacco
v_f	Vegetables, fruit, nuts	Food, Beverages, and Tobacco
ocr	Crops nec	Food, Beverages, and Tobacco
osd	Oil seeds	Food, Beverages, and Tobacco
c_b	Sugar cane and sugar beet	Food, Beverages, and Tobacco
vol	Vegetable oils and fats	Food, Beverages, and Tobacco
sgf	Sugar	Food, Beverages, and Tobacco
b_t	Beverages and tobacco products	Food, Beverages, and Tobacco
pfb	Plant-based fibers	Clothing and Footwear
wol	Wool, silk-worm cocoons	Clothing and Footwear
tex	Textiles	Clothing and Footwear
wap	Wearing apparel	Clothing and Footwear
lea	Leather products	Clothing and Footwear
coa	Coal	Housing and Utilities
gas	Gas	Housing and Utilities
ely	Electricity	Housing and Utilities
gdt	Gas manufacture and distribution	Housing and Utilities
wtr	Water	Housing and Utilities
cns	Construction	Housing and Utilities
dwe	Dwellings	Housing and Utilities

Table E.II continued

GTAP	Description	8-good aggregate
rsa	Real estate activities	Housing and Utilities
frs	Forestry	Household Goods and Services
oxt	Minerals nec	Household Goods and Services
lum	Wood products	Household Goods and Services
ppp	Paper products, publishing	Household Goods and Services
chm	Chemical products	Household Goods and Services
rpp	Rubber and plastic products	Household Goods and Services
nmm	Mineral products nec	Household Goods and Services
i_s	Ferrous metals	Household Goods and Services
nfm	Metals nec	Household Goods and Services
fmp	Metal products	Household Goods and Services
ele	Computer, electronic and optical products	Household Goods and Services
eeq	Electrical equipment	Household Goods and Services
ome	Machinery and equipment nec	Household Goods and Services
omf	Manufactures nec	Household Goods and Services
bph	Basic pharmaceutical products	Health
hht	Human health and social work	Health
oil	Oil	Transport and Communication
p_c	Petroleum, coal products	Transport and Communication
mvh	Motor vehicles and parts	Transport and Communication
otn	Transport equipment nec	Transport and Communication
otp	Transport nec	Transport and Communication
wtp	Water transport	Transport and Communication
atp	Air transport	Transport and Communication
whs	Warehousing and support activities	Transport and Communication
cmn	Communication	Transport and Communication
ros	Recreational and other services	Recreation and Education
edu	Education	Recreation and Education
trd	Trade	Other Services
afs	Accommodation and food services	Other Services
ofi	Financial services nec	Other Services
ins	Insurance	Other Services
obs	Business services nec	Other Services
osg	Public administration and defence	Other Services

TABLE E.III
8-GOOD SUMMARY STATISTICS

Good	Sectors	Mean Share	Thin < 1%	Mean Price	P95/P05 Price	Engel Correlation	High-Low Share
Food, Beverages, and Tobacco	19	0.271	0	1.115	1.421	-0.847	-0.401
Clothing and Footwear	5	0.051	0.006	1.126	1.397	-0.299	-0.031
Housing and Utilities	8	0.140	0	1.026	1.241	0.596	0.099
Household Goods and Services	14	0.082	0	1.143	1.350	0.032	-0.003
Health	2	0.037	0.056	1.025	1.133	0.382	0.031
Transport and Communication	9	0.139	0	1.029	1.359	0.359	0.047
Recreation and Education	2	0.063	0.019	0.973	1.132	0.278	0.031
Other Services	6	0.217	0	1.006	1.074	0.772	0.227




FIGURE E.1.—Average expenditure shares across income groups for the eight consumption bundles in the 2017 GTAP estimation sample.

E.2. Estimation Results for the GTAP Sample

The GTAP estimation produces a similar boundary pattern to the main ICP specification (Table E.IV). Food, Beverages, and Tobacco stabilizes at the lower regularity bound with a finite one-sided interval, while Household Goods and Services tracks the relaxed lower bound and thus does not admit a finite interval under the bound-relaxation diagnostic over the tested range $\underline{\epsilon} \in \{0, -0.0001, -0.001, -0.005, -0.01, -0.025, -0.05, -0.10, -0.25, -0.50, -1.00\}$. We nevertheless choose the aggregation validated in Appendix E.1 because the broader lesson from the ICP specification in Section 7.1.2 and Appendix B.6 is that economically distinct Engel patterns should not be removed merely to eliminate boundary behavior. In the ICP robustness exercises, dropping the problematic category or absorbing it into a broader residual changes the economic interpretation of the system and obscures more meaningful expenditure patterns. The same logic applies here. The GTAP aggregation is chosen to preserve the main necessity, manufacturing, durable, and service margins in the data, even when some categories sit near the regularity boundary stated in Hanoch (1975). Following Andrews (1999, 2000), we report one-sided m -out-of- n inference for the boundary case that stabilizes at a finite negative value under bound relaxation, using $n = 160$ regions and $m = \text{round}(n^{0.7}) = 35$.

E.2.1. Raw Engel Evidence from the Data

To verify the boundary cases in Table E.IV directly against the raw GTAP expenditure data, we compare the estimated parameters to raw Engel patterns in the data. For each bundle k , we estimate the regression $s_{ik} = a_k + b_k \log w_i + \varepsilon_{ik}$, where s_{ik} is the budget share and w_i is per-capita expenditure. We calculate the share slope $100 \times b_k$, the corresponding t -statistic and R^2 , the Spearman rank correlation, and the difference between expenditure shares in the top and bottom income quartiles. We implement a second robustness check that replaces w_i with a PPP-quantity income measure. These measures distinguish strong necessity cases from bundles whose aggregate Engel pattern is flattened by offsetting within-bundle movements. Food, Beverages, and Tobacco is the former case, with a share slope of -10.819 percentage points per log income and a Spearman correlation of -0.847 . Household Goods and Services is

the latter case, with a nearly flat share slope of 0.093, suggesting that sectors within the bundle move in opposite directions with income.

E.2.2. Verification of the Admissible Domain

To verify the admissible normalized-price condition in Proposition 1, we compute

$$\hat{S}_i = \sum_{k \in K_0} \hat{\beta}_k \left(\frac{p_{ik}}{w_i} \right)^{1-\hat{\alpha}_k}, \quad K_0 = \{\text{FoodBeveragesTobacco}, \text{HouseholdGoodsSvcs}\},$$

at the baseline GTAP estimates for each observed region i . All 160 GTAP regions satisfy the admissibility condition $\hat{S}_i < 1$. The largest value is $\max_i \hat{S}_i = 0.684$, and the smallest admissibility gap is $\min_i (1 - \hat{S}_i) = 0.316$. Thus the observed normalized-price vectors $\xi_i = p_i/w_i$ remain inside the admissible domain $\mathcal{D}$ on which the defining CDE relationship admits a unique strictly positive solution $u(\xi)$.

TABLE E.IV
ESTIMATED CDE PARAMETERS WITH BOUNDARY-ROBUST CONFIDENCE INTERVALS (8-BUNDLE GTAP)

Consumption Bundle	$\hat{\alpha}_k$	$\hat{\beta}_k$	$\hat{e}_k$
Food, Beverages, and Tobacco	0.451 [0.173, 0.548]	0.060 [0.000, 0.078]	0.000 ^B [0, 0.095] ^{m,B}
Clothing and Footwear	0.012 [0.000, 0.538]	0.008 [0.000, 0.015]	0.434 [0.094, 0.621]
Housing and Utilities	0.715 [0.435, 0.930]	0.132 [0.000, 0.296]	0.995 [0.138, 1.594]
Household Goods and Services	0.907 [0.490, 1.000]	0.198 [0.000, 0.998]	0.000 ^U — ^U
Health	0.960 [0.860, 1.000]	0.253 [0.000, 0.998]	3.674 [0.000, 6.681]
Transport and Communication	0.536 [0.238, 0.752]	0.075 [0.000, 0.115]	0.737 [0.221, 1.115]
Recreation and Education	0.231 [0.000, 0.561]	0.023 [0.000, 0.033]	0.776 [0.299, 1.148]
Other Services	0.750 [0.438, 0.938]	0.252 [0.000, 0.397]	1.384 [0.378, 3.889]
$\hat{\ell}_{\text{CDE,GTAP8}}$	2,034.010		

Notes: Bracketed entries are 90% confidence intervals. For interior parameters, intervals are constructed from a 2,000-replication paired percentile bootstrap. Superscript ^B denotes a boundary estimate whose $\hat{e}_k$ is at the lower regularity bound but stabilizes at a finite negative value under bound relaxation. Superscript ^m denotes the corresponding one-sided m -out-of- n interval $[0, U_{0.90}]$. The GTAP sample has $n = 160$ regions, and the m -out-of- n procedure uses $m = \text{round}(n^{0.7}) = 35$ with 2,000 valid subsamples. Superscript ^U denotes a case for which no finite interval is reported because the estimate tracks the relaxed lower bound. For Household Goods and Services, the bound-relaxation path follows the imposed lower bound through the completed grid, so the ordinary paired-bootstrap interval and the m -out-of- n endpoint are not reported for $\hat{e}_k$.

E.3. Sectoral Reconciliation of Estimated Demand Parameters

The demand system in GTAP is implemented at the sector-region level. Let s denote a GTAP sector, k an aggregate consumption bundle, and r a region. For each sector s in region r , we specify an income-response parameter and a substitution parameter:

$$INCPAR_{sr} = e_{sr}, \quad SUBPAR_{sr} = 1 - \alpha_{sr}.$$

Our 8-good estimation produces one pair of parameters (α_k^*, e_k^*) for each aggregate bundle k . We anchor aggregate-level non-homotheticity and substitution in the data while preserving the within-bundle relative structure of the original calibration.

Let G_k denote the set of sectors assigned to aggregate bundle k . Let V_{sr}^P denote private household expenditure at consumer prices for sector s in region r . We define the within-bundle expenditure weight as

$$\omega_{sr} = \frac{V_{sr}^P}{\sum_{m \in G_k} V_{mr}^P}.$$

The weights sum to one within each bundle and region and determine the expenditure structure used in the reconciliation.

For the income-response parameters, we preserve the relative dispersion of the original calibrated values within each bundle and region while rescaling their weighted average to match the estimated parameter:

$$e_{sr}^{new} = \begin{cases} e_k^* e_{sr}^{old} / \sum_{m \in G_k} \omega_{mr} e_{mr}^{old}, & \sum_{m \in G_k} \omega_{mr} e_{mr}^{old} > 0, \\ e_k^*, & \text{otherwise.} \end{cases}$$

This multiplicative adjustment thus preserves the within-bundle relative structure of the original calibrated parameters.

For substitution, we calculate the original substitution parameters:

$$\alpha_{sr}^{old} = 1 - SUBPAR_{sr}^{old}.$$

We then map the parameters into logit space:

$$\ell_{sr} = \log \left(\frac{\alpha_{sr}^{old}}{1 - \alpha_{sr}^{old}} \right).$$

For each bundle k and region r , we solve for a common shift parameter δ_{kr} satisfying

$$\sum_{s \in G_k} \omega_{sr} \Lambda(\ell_{sr} + \delta_{kr}) = \alpha_k^*,$$

where $\Lambda(x) = 1/(1 + \exp(-x))$ is the logistic function. The substitution parameters mapped to sectors s and regions r are

$$\alpha_{sr}^{new} = \Lambda(\ell_{sr} + \delta_{kr}), \quad SUBPAR_{sr}^{new} = 1 - \alpha_{sr}^{new}.$$

The common logit shift preserves the within-bundle ranking of the original substitution parameters because the logistic transformation $\Lambda(x)$ is strictly increasing for all $x \in \mathbb{R}$. Sectors

that were relatively more substitutable in the original calibration thus remain relatively more substitutable, preserving the qualitative heterogeneity within each bundle while adjusting the expenditure-weighted average to match the estimated aggregate parameter. Working in logit space also avoids ad hoc truncation when the original parameters lie near the $[0, 1]$ bounds.

For numerical stability and compatibility with the implementation, we impose the bounds

$$e_{sr}^{new} \geq 10^{-6}, \quad 10^{-6} \leq \alpha_{sr}^{new} \leq 1 - 10^{-6}.$$

E.4. Additional Elasticity Results

E.4.1. Bundle-Level Elasticities by Income Group

Figure E.2 summarizes the elasticity patterns by broad income groups. The left panel reports world-average elasticities across bundles, while the remaining panels report averages within the World Bank income classifications. Food, Beverages, and Tobacco is more necessity-like under the estimated system, especially in low-income and high-income regions where the calibrated elasticities move back toward unit elasticity. The difference is larger in LIC and HIC regions and smaller in the middle-income groups.




FIGURE E.2.—Calibrated and estimated GTAP income elasticities by income group. Panels report average bundle-level income elasticities under the calibrated parameters and point estimates for the full sample and for World Bank income groups. Markers show group averages, and connecting lines highlight the difference between the calibrated and estimated systems. The largest differences appear for Food, Beverages, and Tobacco in low- and high-income regions, while most nonfood and service bundles differ primarily among developing economies.

E.4.2. *Selected Sectoral Elasticities After Reconciliation*

The nonfood and service bundles show the opposite tendency. Under the estimated system, elasticities for household goods, transport, recreation, and other services are generally larger among developing regions and become closer to the calibrated values as income rises. Clothing and Footwear again behaves differently, with the estimated elasticity remaining below the calibrated elasticity over much of the distribution. The figure therefore compresses the main counterfactual prediction into a small set of averages: the estimated demand system reallocates expenditure more strongly away from necessities and toward nonfood bundles in developing economies, while the differences become smaller for broad middle-income expenditure ranges.

Figure 8 in Section 9.1 reports income elasticities at the 8-good aggregate level. The figures in this section return to the original sectoral structure after reconciliation. The selected sectors are chosen to span necessities, manufacturing, durables, and services while avoiding sectors whose import shares are dominated by purely local activity such as dwellings, construction, water, education, or health services. The 10-sector figure provides the broadest view of the sectoral patterns. It includes staple and processed food categories, beverages and tobacco, manufacturing goods, durables, and tradable services. The 8-sector figure reports a more compact set of sectors chosen for readability while preserving the main economic margins. The 6-sector figure focuses on the clearest Engel patterns in the data: staple food, beverages and tobacco, manufacturing durables, and tradable services.

In the calibrated sector-level elasticities, both Other Food (*ofd*) and Beverages and Tobacco (*b_t*) display a shallow U-shaped pattern across the expenditure distribution. Elasticities are higher among the lowest-income regions, fall through the middle of the distribution, and rise again toward the richest regions. The estimated Food, Beverages, and Tobacco bundle in the reconciled system behaves as a more stable necessity. To verify that this pattern is not a plotting artifact, we compute average calibrated elasticities using grouped expenditure ranks from aggregate private-consumption data. For both *ofd* and *b_t*, the average elasticity is approximately 0.779 in the bottom 20 percent of regions, 0.694 in the middle 20 percent, and 0.726 in the top 20 percent. The decile averages show the same pattern, with elasticities declining from the lower deciles toward the middle of the distribution and then rising again in the upper deciles. The U-shape thus arises from the originally calibrated sector elasticities.

In addition, Figures E.3–E.5 make the nonfood reallocations more concrete. Wearing apparel, pharmaceuticals, motor vehicles, electronics, communication, air transport, and recreation services absorb part of the expenditure reallocation under the estimated system as real income rises in developing regions.




FIGURE E.3.—Income elasticities for selected sectors after reconciliation to the original sectoral structure. Panels report sector-level income elasticities against per-capita expenditure for the 160 regions under the calibrated parameters, point estimates, and bootstrap means.




FIGURE E.4.—Income elasticities for selected sectors after reconciliation to the original sectoral structure.




FIGURE E.5.—Income elasticities for selected sectors after reconciliation to the original sectoral structure.

E.5. Decomposition of Welfare Incidence

The counterfactual exercises replace two sets of household-demand parameters in the GTAP CDE system. The first is the income-response parameter, INCPAR, which corresponds to the expansion parameter e_k in equation (2.1). The second is the substitution parameter, SUBPAR, which corresponds to $1 - \alpha_k$ in equation (2.1), or equivalently the parameter d_k in Hanoch (1975). The main counterfactual results thus combine changes in both income responses and substitution responses. In this section, we separate these two channels.

For each scenario, we solve the following 2×2 decomposition in the structural model:

	INCPAR Calibrated	INCPAR Estimated
SUBPAR Calibrated	A	B
SUBPAR Estimated	C	D

Case A is the original calibrated specification. Case D replaces both parameter sets with the estimates from this paper. Case B changes only the income-response parameters, while Case C changes only the substitution parameters. We decompose the total effect as

$$D - A = (B - A) + (C - A) + I,$$

where I is the interaction term which reflects endogenous general equilibrium outcome. The data, equilibrium assumptions, shocks, and solution method are held fixed across all four cases.

TABLE E.V
DECOMPOSITION OF DEMAND EFFECTS IN COUNTERFACTUALS

Scenario	Total	INCPAR	SUBPAR	GE Interaction
Global Productivity Shock	-49.408	-50.334	-18.282	19.208
LIC/LMIC Productivity Shock	0.611	1.139	-3.074	2.546
Food Tariff Reduction	0.292	-0.298	0.733	-0.142
Multilateral Trade Liberalization	-0.688	-0.438	-0.078	-0.173

Notes: Total is $D - A$. INCPAR is $B - A$. SUBPAR is $C - A$. Entries report world equivalent variation in billions of 2017 U.S. dollars. Negative values indicate smaller welfare gains under the estimated demand system relative to the calibrated system.

Table E.V reports the decomposition for world equivalent variation. Under the global productivity shock, which raises real income worldwide, the income-response parameters account for almost the entire welfare difference of -50.3 billion dollars. The substitution component is also negative, but only about one third as large and partly offset by indirect general equilibrium channels. Under the food tariff reduction, the substitution component becomes larger than the income-response component. A food tariff reduction primarily changes relative prices for the strongest necessity bundle in the data, so the divergence between the two systems operates more through substitution than through Engel responses. This also explains the relatively flat tariff panels in Figure 9. Estimation does change the substitution structure, and the food tariff counterfactual moves through that channel. The panels remain flat because the shock generates only limited real-income reallocation.

The remaining two scenarios produce small aggregate welfare differences because the income and substitution channels offset one another. The broader pattern is that productivity shocks mainly differ through the Engel side of the demand system, while the food tariff counterfactual differs more through substitution. Broad multilateral liberalization is not highly sensitive to replacing calibrated demand with estimated demand at the aggregate welfare level,

although the estimated system still places more weight on the income-response margin than on the substitution margin.