Flexible Estimation of Demand Systems: A Copula Approach*

– Supplementary Materials –

Mateo Velásquez-Giraldo † Gustavo Canavire-Bacarreza‡ Kim P. Huynh§
David T. Jacho-Chávez¶

July 18, 2018

1 Summary

This document serves as the supplementary material for the article "Flexible Estimation of Demand Systems: A Copula Approach."

Section 2 describes our estimation methodology and the specifications that we use in detail. Section 3 conducts an empirical analysis of the data that guides the selection of the specifications that are used in our meta-analysis. Section 4 discusses an alternative measure of model fitness. Finally, Section 5 presents our replication of the results in Chang and Serletis (2014)

2 Methodology

This section describes the methodology that we use for estimating Deaton and Muellbauer's (1980) Almost Ideal Demand System (AIDS), Banks et al.'s (1997) Quadratic AIDS (QUAIDS), and Barnett's (1983) Minflex Laurent (ML) models to demand systems and their implied elasticities. The specification of the conditional means models, and the methodology for imposing regularity on their estimates is the same that Chang and Serletis (2014) used, with minor adjustments to parameter normalizations

---

*The views expressed in this paper are those of the authors. No responsibility for them should be attributed to the Bank of Canada. The opinions expressed in this paper are those of the authors and do not necessarily reflect the views of the Inter-American Development Bank, its Board of Directors, or the countries they represent.

†School of Economics and Finance, Universidad EAFIT, Carrera 49 N 7 sur 50, Medellín, Colombia. Email: mateovegi@gmail.com.

‡Inter-American Development Bank, 1200 New York NW Ave, Washington, DC 20577. E-mail: gcan@iadb.org.

§Bank of Canada, 234 Wellington Ave., Ottawa ON, K1A 0G9, Canada. E-mail: kim@huynh.tv.

¶Corresponding Author: Department of Economics, Emory University, Rich Building 306, 1602 Fishburne Dr., Atlanta, GA 30322-2240, USA. E-mail: djachocha@emory.edu.

to aid numerical procedures. We summarize the conditional mean model specifications and regularity impositions from Chang and Serletis (2014), but the reader should refer to the original article for a more complete account.

We extend the estimation procedure in Chang and Serletis (2014) to allow for a more flexible specification of the budget share residuals, whose likelihood is maximized.

2.1 Maximum likelihood estimation

Chang and Serletis (2014) use the following specification to estimate all the models for household h:

$$s_h = g(p_h, y_h, \vartheta) + u_h, \quad h = 1, \ldots, H. \tag{2.1}$$

where H is the total number of households, $s_h = (s_{h;1}, s_{h;2}, s_{h;3})'$ is the observed budget share vector for household h, $p_h = (p_{h;1}, p_{h;2}, p_{h;3})'$ is the price vector it is exposed to, and y_h is its total expenditure on the three goods. The shares' conditional means are obtained through each model $g(\cdot, \cdot, \vartheta) = (g_1(\cdot, \cdot, \vartheta), g_2(\cdot, \cdot, \vartheta), g_3(\cdot, \cdot, \vartheta))'$ (AIDS, QUAIDS, and ML) and u_h is a stochastic zero-mean vector of errors which they assume to follow a trivariate normal distribution.

Since $\sum_{i=1}^{3} s_{h;i} = 1$, and because of the multivariate normality assumption, they proceed to estimate ϑ in (2.1) after arbitrarily dropping the last equation in the system (see Barten, 1969) using nonlinear full-information maximum likelihood.

We extend this methodology to allow for non-normally distributed errors, considering the empirical distributions of budget shares (see Figure 1 in the main text) and the disconnection between the normality assumption and the distribution of the residuals that it produces (see Figure 2). To achieve this, we use copula functions to construct more flexible joint residual distributions.

Since all shares add up to one, we select only two goods to be directly modeled. As Chang and Serletis (2014), we select Good 1 (gasoline) and Good 2 (local transportation) to be directly modeled in our estimation.

We model the marginal distributions of Good 1 and Good 2's budget shares as $F_1(s_{h;1}|g_1(p_h, y_h, \vartheta), \psi_1)$ and $F_2(s_{h;2}|g_2(p_h, y_h, \vartheta), \psi_2)$, where ψ_1 and ψ_2 correspond to the marginals' specific unknown parameters. We relate these marginal distributions to the joint distribution, $F_{1;2}(\cdot, \cdot)$, through a copula function, $C(\cdot, \cdot; \theta)$, i.e.,

$$F_{1;2}(s_{h;1}, s_{h;1}) = C(F_1(s_{h;1}|g_1(p_h, y_h\vartheta), \psi_1), F_2(s_{h;2}|g_2(p_h, y_h, \vartheta), \psi_2); \theta), \tag{2.2}$$

where θ represents the specific copula's dependence parameter. As discussed in Section 3, we consider the Normal and Skewed Normal distributions for the marginals $F_1(\cdot|\cdot, \psi_1)$ and $F_2(\cdot|\cdot, \psi_2)$, and the Normal, Frank, Clayton, and Gumbel copulas for $C(\cdot, \cdot; \theta)$. To allow for different dependence structures, specially to model negative dependence through the Gumbel and Frank copulas, for each choice of copula C we evaluate $F_{1;2}$ as in (2.2) and the following rotation as well:

Figure 1: Specification alternatives.

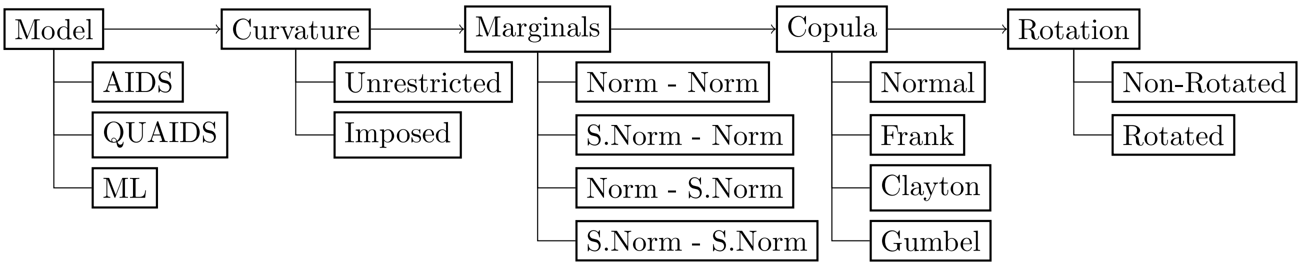

Note: (a) The options produce a total of 168 specifications for each dataset. (b) N stands for Normal and S.N for Skewed Normal.

$$\begin{aligned} F_{1;2}(s_{h;1}, s_{h;1}) &= C(F_1(s_{h;1} - g_1(p_h, y_h \vartheta)|\psi_1), 1 - F_2(s_{h;2} - g_2(p_h, y_h, \vartheta)|\psi_2); \theta) \\ &= C^*(F_1(s_{h;1}|g_1(p_h, y_h\vartheta), \psi_1), F_2(s_{h;2}|g_2(p_h, y_h, \vartheta), \psi_2); \theta) \end{aligned} \tag{2.3}$$

Copulas C^* are therefore obtained by rotating the second good in the Normal, Frank, Clayton, and Gumbel copulas. With this in mind, a total of 168[1] specifications for each type of household are estimated using the copula R package. Figure 1 presents all the specification alternatives which we consider. Note that the models considered by Chang and Serletis (2014) are obtained using Normal - Normal marginals and the Normal copula.

Given a full parameter vector $\{\vartheta^*, \psi_1^*, \psi_2^*, \theta^*\}$, we first obtain the estimated budget share means through the selected demand system model for each household h:

$$s_h^* = \{s_{h;1}^*, s_{h;2}^*\} = \{g_1(p_h, y_h, \vartheta^*), g_2(p_h, y_h, \vartheta^*)\}. \tag{2.4}$$

We then obtain $u_h^* = \{u_{h;1}^*, u_{h;2}^*\} = \{s_{h;1} - s_{h;1}^*, s_{h;2} - s_{h;2}^*\}$ for each household. The log-likelihood of each household's budget allocation is calculated through these residuals based on (2.2), assuming $g(p_h, y_h, \vartheta^*)$ is the conditional mean of budget shares. The marginals of residuals are therefore restricted to have zero means. We compute each household's residuals' log-density and add them up to obtain our total log-likelihood, which we maximize:

$$\ell(\vartheta^*, \psi_1^*, \psi_2^*, \theta^*) = \sum_{h=1}^{H} \ln c(u_{h;1}^*, u_{h;2}^* | \vartheta^*, \psi_1^*, \psi_2^*, \theta^*) \tag{2.5}$$

where $c(\cdot)$ is the density function associated with model (2.2). We construct and evaluate $c(\cdot)$ using the `copula` package in the R statistical language.

[1] Specifications that use the Normal copula are invariant to our rotation. The number of specifications is therefore $3 \times 2 \times 4 \times (3 \times 2 + 1) = 168$.

2.2 Normalizations and restrictions

We now discuss the normalizations adopted for parameter estimation and the way in which restrictions for local and global curvature imposition are implemented.

2.2.1 AIDS and QUAIDS models

The AIDS and QUAIDS models are estimated with the following parameter normalizations and restrictions:

$$\sum_{i=1}^{3} \alpha_i = 1; \quad \sum_{i=1}^{3} \beta_i = 0; \quad \sum_{i=1}^{3} \lambda_i = 0; \quad \sum_{i=1}^{3} \gamma_{i,j} = 0 \, \forall j; \quad \gamma_{i,j} = \gamma_{j,i} \, \forall i, j, \tag{2.6}$$

where $\lambda_i = 0 \, \forall i$ in the case of AIDS. This restrictions are included in our implementation by omitting α_3, β_3, λ_3, $\gamma_{2,1}$, $\gamma_{3,1}$, $\gamma_{3,2}$ and $\gamma_{3,3}$ as parameters in the optimization procedure. We optimize over the rest of parameters (8 for AIDS and 10 for QUAIDS) and obtain the omitted parameters using (2.6) in every iteration, allowing us to evaluate the objective function.

Ryan and Wales (1998) derive a simple way of restricting the models so that curvature is satisfied at the point $p^1 = p^2 = p^3 = y = 1$. The procedure consists of expressing γ as a function of an upper triangular matrix κ and the rest of the parameters, in a way which ensures that the Slutsky matrix is negative semi-definite at the aforementioned point. The transformation is specified in equations (9) and (13) in Chang and Serletis (2014). The number of free parameters in κ is the same as that of γ for the unrestricted case.

For the restricted models, we again optimize over the free parameters (κ replaces γ), obtaining an estimated γ, and estimates of all the omitted parameters at every iteration. We use the delta method with numerical derivatives to obtain standard errors for estimates of γ.

2.2.2 ML model

The ML model has three different sets of parameters: a 3×1 vector a, and two 3×3 symmetric matrices A and B. The following restrictions are imposed:

$$b_{i,i} = 0 \, \forall i; \quad a_{i,j} * b_{i,j} = 0 \, \forall i, j. \tag{2.7}$$

Sign restrictions are also necessary: in the base case, the off-diagonal elements of A and B need to be non-negative, and in order to impose global curvature, all parameters have to be non-negative. Chang and Serletis (2014) impose these restrictions by replacing restricted parameters with their squares in the model's equations. We also adopt this strategy.

It is important to decide which parameters are 0 in the $a_{i,j} * b_{i,j} = 0 \, \forall i, j$ restriction. Following Chang and Serletis (2014), we set $A_{1,2} = A_{2,1} = A_{2,3} = A_{3,2} = B_{1,3} = B_{3,1} = 0$. This, along with the symmetry restriction implies that the 9 parameters left for estimation are a_1, a_2, a_3, $A_{1,1}$, $A_{1,3}$, $A_{2,2}$, $A_{3,3}$, $B_{1,2}$, and $B_{2,3}$.

The ML share equation is homogeneous of degree 0 in its parameters, that is: all parameters can be multiplied by a constant without modifying the predicted shares. To fix the scale, Chang and Serletis (2014) impose the following normalization:

$$2\sum_{i=1}^{3} a_i + \sum_{i=1}^{3}\sum_{j=1}^{3} A_{i,j} - \sum_{i=1}^{3}\sum_{j=1}^{3} B_{i,j} = 1, \tag{2.8}$$

where restricted parameters are replaced by their squares. The authors seem to leave $B_{2,3}$ out of the estimation (as they do not report it) and recover it at every iteration from (2.8). We opt against this strategy, as for some parameter values (2.8) can produce negative implied values for $B_{2,3}$, which needs to be non-negative. This would require us to restrict the feasible 'free' parameter values, complicating our problem.

We instead drop the normalization in (2.8) and adopt the following strategy. At every loop of our optimization scheme (see Subsection 2.3 below) we rescale the parameter vector so that the sum of all values in a, A and B is 1. The parameter a_1 is fixed on the value implied by this normalization for the rest of the loop.

2.3 Optimization

Maximization of the log-likelihood function expressed in (2.5) is not trivial. This section describes the strategy and numerical methods we used in order to arrive at the estimates that we report.

We used the estimates from the replication section of this study (Tables 5, 6, and 7) as starting values for the demand systems' model parameters ϑ. We denote these particular values as ϑ_0. These models were estimated using the `mle2` function in the `bbmle` package in R. Estimation converged from arbitrary starting points for AIDS and QUAIDS models. For ML models, we used the `CRS` (controlled random search) method from the R package `nloptr` to find starting points.

Initial values for ψ_1, ψ_2 and θ are obtained maximizing the log-likelihood in (2.5) with ϑ fixed on ϑ_0 using `nloptr`'s `DIRECT-L`. Starting values for this optimization are arbitrary (but feasible). With these values, we can construct a full initial parameter vector $\{\vartheta_0, \psi_{1;0}, \psi_{2;0}, \theta_0\}$.

After constructing an initial value, we initiate an estimation loop in which we alternate optimization methods trying to improve our estimates. The loop was designed considering that using a single method often resulted in local optimums and ill-defined variance covariance matrices. Estimation stops if a well defined covariance matrix is obtained or after a maximum of ten iterations. The nth iteration consists of the following steps:

1. Maximize $\ell(\cdot)$ using `nloptr`, with $\{\vartheta_{n-1}, \psi_{1;n-1}, \psi_{2;n-1}, \theta_{n-1}\}$ as starting values. If $n = 1$ use the

Table 1: Specifications that did not converge per model and copula

Copula	**AIDS**		**AIDS**(R)		**QUAIDS**		**QUAIDS**(R)		**ML**		**ML**(R)		**Total**	
	N	Ratio	N	Ratio	N	Ratio	N	Ratio	N	Ratio	N	Ratio	N	Ratio
Normal	12	0.000	12	0	12	0.167	12	0.42	12	0.25	12	0.33	72	0.19
Frank	24	0.042	24	0	24	0.125	24	0.21	24	0.21	24	0.50	144	0.18
Clayton	24	0.000	24	0	24	0.083	24	0.21	24	0.42	24	0.38	144	0.18
Gumbel	24	0.083	24	0	24	0.083	24	0.21	24	0.29	24	0.29	144	0.16
Total	84	0.036	84	0	84	0.107	84	0.24	84	0.30	84	0.38	504	0.18

Note: (a) Columns N mark the total number of specifications that were estimated for a given copula-demand model combination. Columns "Ratio" specify the fraction of those models that did not converge. (b) (R) marks demand models with curvature impositions. (c) The Normal copula always has half of the number of specifications as others, as it is not modified by the rotation.

method `SBPLX`, else use `COBYLA`.

2. Maximize $\ell(\cdot)$ using `mle2`, which also computes the estimates' covariance matrix. The initial values are the results of step 1. The method `BFGS` is always used, but as it is based on a numerical gradient it can fail, halting with an error. In this cases, switch to the Nelder-Mead method, also supported by `mle2`. Denote the estimates obtained in this step by $\{\vartheta_n, \psi_{1;n}, \psi_{2;n}, \theta_n\}$.

3. Check if the Hessian matrix computed by `mle2` in step 2 has no missing entries, is finite and invertible. Finally check that all entries in the diagonal of the implied covariance matrix are non-negative. If these conditions hold, exit the loop taking $\{\vartheta_n, \psi_{1;n}, \psi_{2;n}, \theta_n\}$ as the final result. If the conditions do not hold and $n < 10$ go to step 1 of iteration $n + 1$.

This algorithm allows us to obtain well defined estimates of most of the specifications that we consider. Table 1 presents a descriptive analysis of the models that did not converge, disaggregated by demand system model and copula. Of all the considered specifications, 18% did not converge. No particular copula had significantly higher rates of convergence than others. The most problematic demand system to estimate was the ML, with more than 30% of estimations failing to converge.

2.4 Regularity checks

Table 2 presents the percentage of models that attained regularity out of those that converged for every demand model-copula combination. Regularity is defined as not having violations of positivity, monotonicity, or curvature. Only 45% of specifications that converged attained regularity. The restricted ML model dominates in this respect, as curvature can be imposed globally. However, for the AIDS models, more than 50% of specifications have no violations. Copulas do not seem to have a big effect, but the Frank copula has the lowest percentage of regularity violations.

Table 2: Specifications that attain each regularity condition per model and copula

Copula	**AIDS**		**AIDS**(R)		**QUAIDS**		**QUAIDS**(R)		**ML**		**ML**(R)		**Total**	
	N	Ratio	N	Ratio	N	Ratio	N	Ratio	N	Ratio	N	Ratio	N	Ratio
Regularity														
Normal	12	0.67	12	0.92	10	0	7	0.000	9	0.000	8	1	58	0.47
Frank	23	0.61	24	0.79	21	0	19	0.000	19	0.105	12	1	118	0.40
Clayton	24	0.67	24	1.00	22	0	19	0.000	14	0.143	15	1	118	0.48
Gumbel	22	0.68	24	1.00	22	0	19	0.053	17	0.000	17	1	121	0.47
Total	81	0.65	84	0.93	75	0	64	0.016	59	0.068	52	1	415	0.45
Positivity														
Normal	12	1	12	1.00	10	1.00	7	1.00	9	0.67	8	1	58	0.95
Frank	23	1	24	0.83	21	0.71	19	0.95	19	0.89	12	1	118	0.89
Clayton	24	1	24	1.00	22	0.82	19	1.00	14	0.79	15	1	118	0.94
Gumbel	22	1	24	1.00	22	0.77	19	0.89	17	0.76	17	1	121	0.91
Total	81	1	84	0.95	75	0.80	64	0.95	59	0.80	52	1	415	0.92
Monotonicity														
Normal	12	1	12	1.00	10	1.00	7	1.00	9	0.67	8	1	58	0.95
Frank	23	1	24	0.83	21	0.71	19	0.95	19	0.89	12	1	118	0.89
Clayton	24	1	24	1.00	22	0.82	19	1.00	14	0.79	15	1	118	0.94
Gumbel	22	1	24	1.00	22	0.77	19	0.89	17	0.76	17	1	121	0.91
Total	81	1	84	0.95	75	0.80	64	0.95	59	0.80	52	1	415	0.92
Curvature														
Normal	12	0.67	12	0.92	10	0	7	0.000	9	0.000	8	1	58	0.47
Frank	23	0.61	24	0.96	21	0	19	0.000	19	0.105	12	1	118	0.43
Clayton	24	0.67	24	1.00	22	0	19	0.000	14	0.143	15	1	118	0.48
Gumbel	22	0.68	24	1.00	22	0	19	0.105	17	0.000	17	1	121	0.48
Total	81	0.65	84	0.98	75	0	64	0.031	59	0.068	52	1	415	0.47

Note: (a) *Regularity* checks that the estimates comply with *Positivity*, *Monotonicity* and *Curvature*. (b) Columns N mark the total number of specifications that converged for a given copula-demand model combination. Columns "Ratio" specify the fraction of those models that attained regularity (no violations of positivity, monotonicity, or curvature). (c) (R) marks demand models with curvature impositions.

The following section presents the methodology that we use to check regularity conditions.

Positivity

We check positivity by evaluating whether for household h:

$$g_j(p_h, y_h, \widehat{\vartheta}) \geq 0 \quad \text{for} \quad j \in \{1, 2, 3\}.$$

We count the number of households for which this condition is not satisfied in every database (using the final estimated parameters) and report it in our tables.

Monotonicity

Chang and Serletis (2014) present explicit expressions for the indirect utility functions of all the models. We implement these functions and check monotonicity by evaluating whether for household h:

$$\partial h(p_h, y_h, \widehat{\vartheta})/\partial p_{h;j} < 0 \quad \text{for} \quad j \in \{1, 2, 3\},$$

where $h(\cdot)$ is the indirect utility function, p_h is a price vector and y_h is the household's total expenditure.

We check the condition for every household using R's `numDeriv` package to compute the derivative. We report the number of violations for every database.

Curvature

Ryan and Wales (1998) report that curvature conditions are equivalent to the requirement that the Slutsky matrix of the demand system be symmetric and negative semi-definite. We now provide the derivation of the expression that we use for the Slutsky matrix.

The models (AIDS, QUAIDS and ML) give us expressions for budget shares depending on prices and wealth $g_j(p, y, \vartheta)$ for $j \in \{1, 2, 3\}$. We use this expression to compute the Marshallian demand for each good $x_j(p, y, \vartheta)$:

$$x_j(p, y, \vartheta) = \frac{g_j(p, y, \vartheta) \times y}{p_j}. \tag{2.9}$$

To use the Slutsky equation, we first compute the derivatives of the Marshalian demand in (2.9) with respect to prices and wealth as follows.

$$\frac{\partial x_j}{\partial p_k}(p, y, \vartheta) = \frac{y}{p_j^2}\left[p_j * \frac{\partial g_j}{\partial p_k}(p, y, \vartheta) - g_j(p, y, \vartheta) * \delta_{i,j}\right], \tag{2.10}$$

$$\frac{\partial x_j}{\partial y}(p, y, \vartheta) = \frac{1}{p_j}\left[\frac{\partial g_j}{\partial y}(p, y, \vartheta) * y + g_j(p, y, \vartheta)\right], \tag{2.11}$$

where $\delta_{i,j} = 1$ when $i = j$ and $\delta_{i,j} = 0$ otherwise. We program analytical implementations of $\partial g_j/\partial p_k$ and $\partial g_j/\partial y$ in order to improve precision and considering that we will need them later on to compute elasticities.

We then use Slutsky's equation to compute the Slutsky matrix $S(p, y, \vartheta)$:

$$[S(p, y, \vartheta)]_{j,k} = \frac{\partial x_j}{\partial p_k}(p, y, \vartheta) + \frac{\partial x_j}{\partial y}(p, y, \vartheta) * x_k(p, y, \vartheta). \tag{2.12}$$

We check whether $S(p_h, y_h, \widehat{\vartheta})$ is symmetric and negative semi-definite for every household h and report the number of violations to this condition. The application of (2.12) may cause numerical errors that can make $S(p_h, y_h, \widehat{\vartheta})$ non-symmetric. Therefore, we consider as symmetric any matrix for which the maximum entry of $|S - S^T|$ is under 10^{-10}. For matrices which satisfy this condition, we replace $S = (S + S^T)/2$ to make them strictly symmetric before checking positive semi-definiteness.

2.5 Elasticities

We compute the income elasticity $\eta_{j,y}$ for good j and the Marshallian elasticities $\eta_{j,k}$ of good j on the price of good k using the same formulas as Chang and Serletis (2014):

$$\eta_{j,y} = 1 + \frac{y}{g_j}\frac{\partial g_j}{\partial y}, \tag{2.13}$$

$$\eta_{j,k} = \frac{p_k}{g_j}\frac{\partial g_j}{\partial p_k} - \delta_{j,k}. \tag{2.14}$$

We compute standard errors for elasticities using the delta method (with numerical derivatives) and use them to report z-test p-values.

3 Model selection

This section discusses different empirical characteristics of the marginal and joint distributions of the budget shares of Goods 1 and 2. These characteristics guide the choice of the possible specifications that we allow in our meta-analysis.

A first feature of the data that must be noted is the skewed distribution of the budget shares of local transportation (see Figure 1 in the main text). The distribution has a mean that is close to zero and is skewed to the right for all types of households. The conditional mean models do not account for this distribution. Figure 2 displays box-plots for the estimated residuals (u_h in (2.1)) for the AIDS, QUAIDS, and M.L models estimated under the trivariate normality assumption for each household type. The box plots show that the residual associated with Good 2 has a non-symmetric distribution regardless of the model that is used for conditional means.

Figure 2: Budget share residuals from estimation under normality.

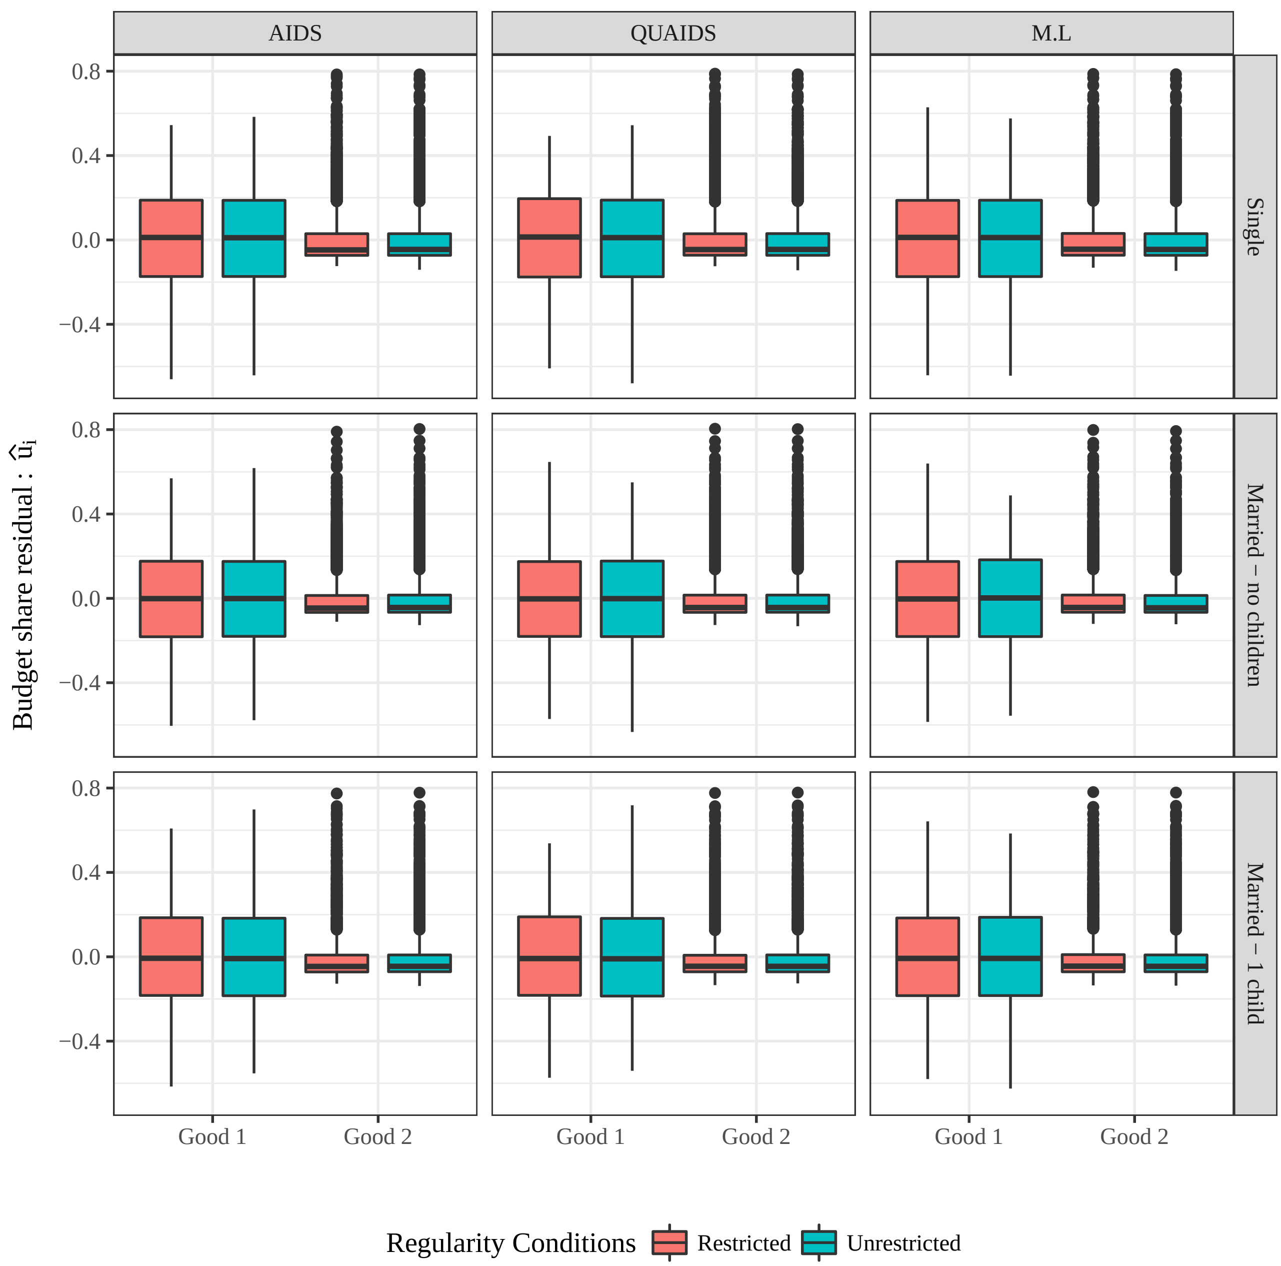

Note: (a) The graph depicts our own estimated residuals for the specifications used by Chang and Serletis (2014) assuming normality. (b) Restricted indicates the instances where local curvature is imposed for AIDS and QUAIDS models, and global curvature for Minflex Laurent models.

Figure 3: Log-likelihood of estimated marginal distributions.

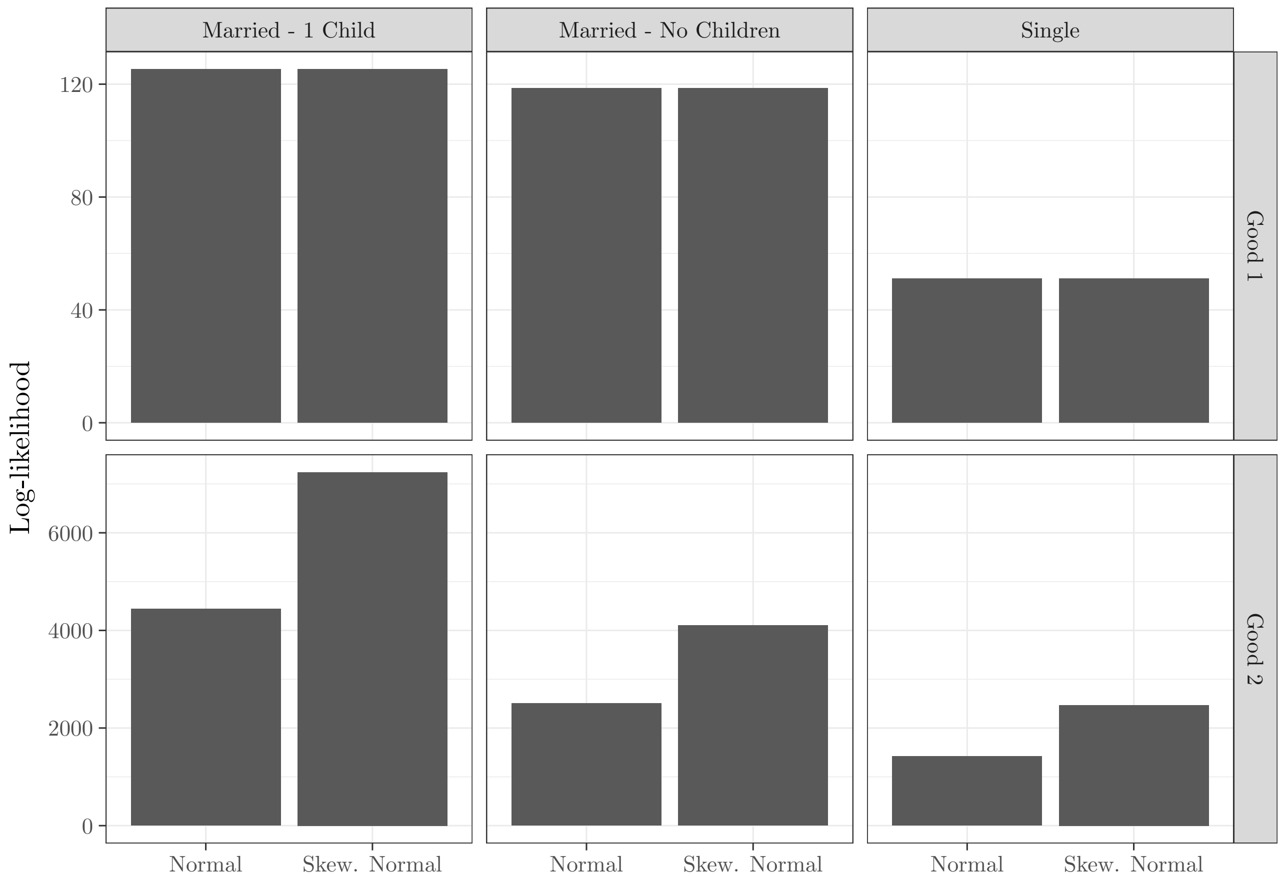

Note: (a) Each distribution was fitted to budget shares of each good using the maximum likelihood method. The reported value is the attained maximum.

To evaluate the possibility of allowing for skewness in marginal distributions of budget shares, we compare the fit of Normal and Skewed Normal marginals to the budget share data for each good and dataset. Figure 3 presents the log-likelihood of the observed budget shares under Normal and Skewed Normal marginals fitted using the maximum likelihood method. Allowing for skewness does not significantly improve the fit of the distributions to Good 1's budget shares, but it does for Good 2's. We nevertheless allow for the four possible combinations of these marginals, considering that this descriptive exercise used observed budget shares, and we are modeling unobserved residuals.

We now study the dependence between the budget shares of Goods 1 and 2. We first use the transformation estimator with log-quadratic local likelihood in Geenens et al. (2017) to illustrate the density of the bivariate relationship. Figure 4 presents contour plots this non-parametric copula density estimator, using normal marginals and with Good 2's budget share rotated. The contours show that there is a low degree of dependence and that there is no evident high asymmetry.

Table 3: Spearman's ρ and Kendall's τ for budget shares of Goods 1 and 2

Dataset	Kendall's τ	Spearman's ρ
Single member households	-0.21	-0.32
Married couples without children	-0.22	-0.32
Married couples with one child	-0.19	-0.28

Note: (a) The dependence measures were calculated on a non-parametric copula estimated for the rank transformations of each pair of budget shares for each household type.

Table 3 presents Kendal's τ and Spearman's ρ for the budget shares of Goods 1 and 2. As can be seen from the coefficients, there is negative dependence between the budget shares for all types of households. This is partially explained by the fact that all budget share must add up to one. We therefore must allow for negative dependence when constructing our copula specification.

Another important characteristic of multivariate modeling is tail dependence, which refers to increases in dependence for extreme outcomes. This characteristic is measured through the lower and upper tail dependence coefficients λ_L and λ_U, defined as

$$\begin{aligned}\lambda_L(X,Y) &= \lim_{v\to 0^+} \frac{P[F_1^{-1}(X) < v, F_2^{-1}(Y) < v]}{v} \\ \lambda_U(X,Y) &= \lim_{v\to 1^-} \frac{P[F_1^{-1}(X) > v, F_2^{-1}(Y) > v]}{1-v}\end{aligned} \tag{3.1}$$

where F_1 and F_2 are the CDFs of the random variables of interest X and Y (see Trivedi and Zimmer, 2007, Section 2.4.4). Table 4 presents estimates of the tail dependence coefficients. The CDFs were obtained as simple rank transformations, and the coefficients were estimated through the non-parametric

Figure 4: Contour plots of kernel copula density estimators of Goods 1 and 2 budget shares.

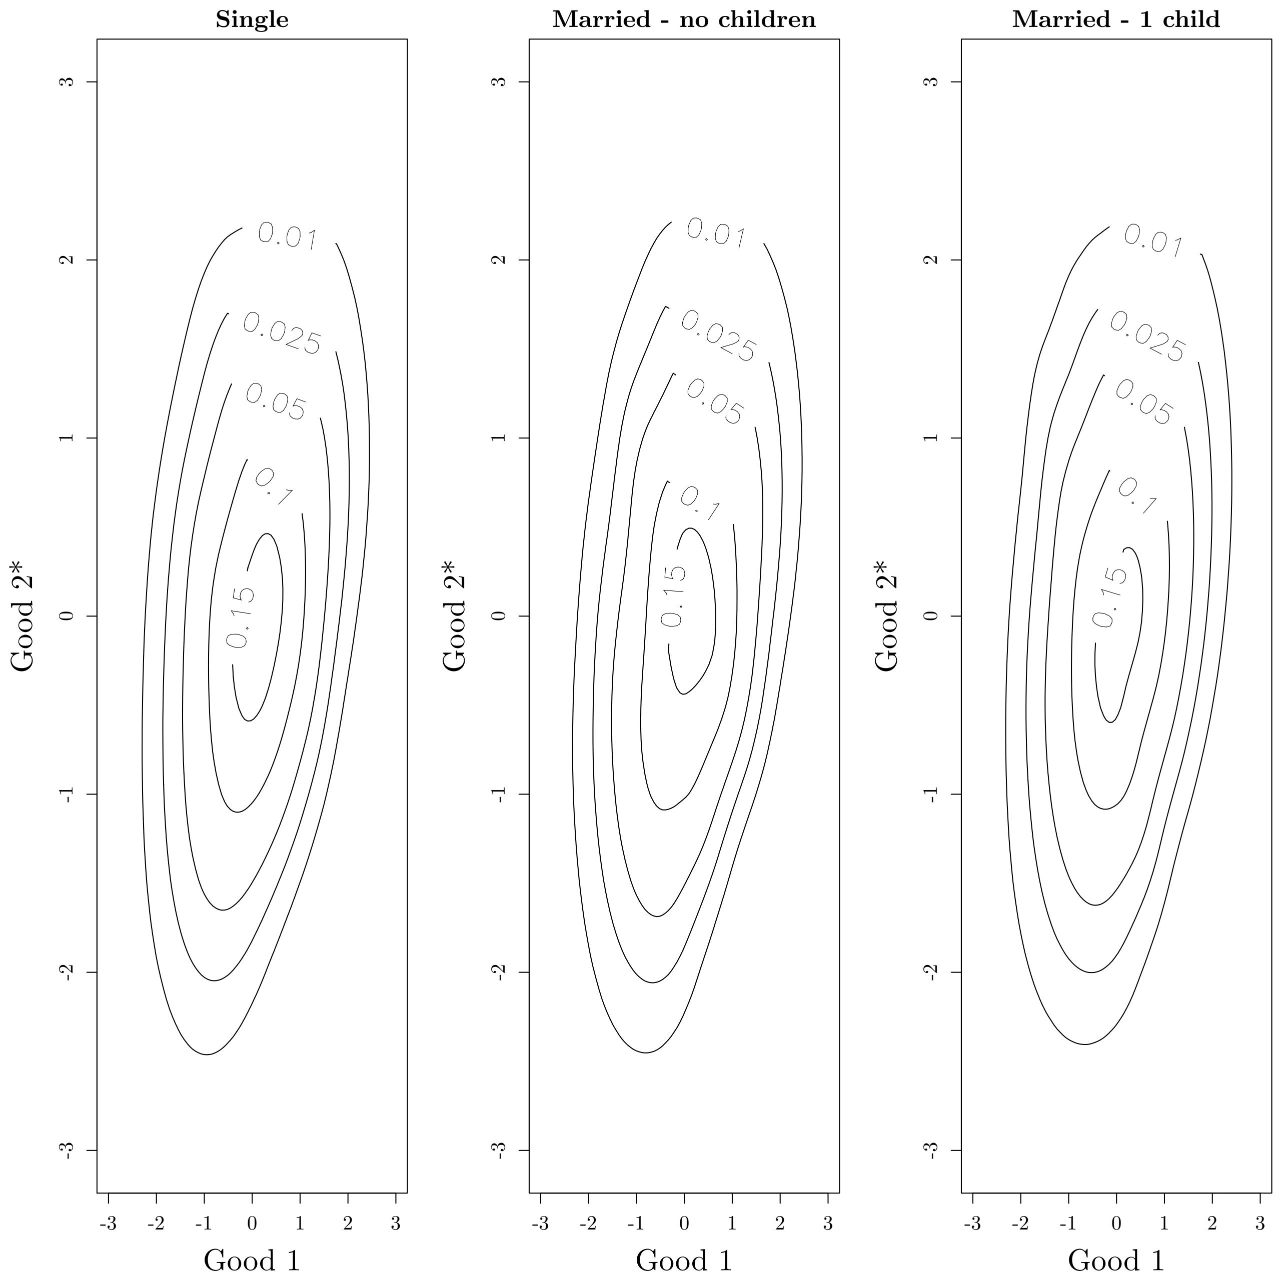

Note: (a) The symbol * denotes that Good 2's share was rotated by multiplying by -1 before applying the rank transformation. (b) Copula density estimators are obtained by applying rank transformations to the budget shares of Goods 1 and 2 and then adjusting the transformation estimator with log-quadratic local likelihood estimation proposed by Geenens et al. (2017) and implemented in the R package kdecopula. (c) Contours are represented using standard normal marginals for each good's budget shares.

Table 4: Tail dependence coefficients

Dataset	$\lambda_L(s_1, s_2)$	$\lambda_U(s_1, s_2)$	$\lambda_L(s_1, -s_2)$	$\lambda_U(s_1, -s_2)$
Single member households	0.00	0.00	0.08	0.09
Married couples without children	0.00	0.00	0.12	0.10
Married couples with one child	0.00	0.00	0.08	0.12

Note: (a) λ_L denotes the lower and λ_U the upper tail dependence coefficients. (b) Tail dependence coefficients are computed on raw budget share data for each pair of goods and household type. (c) The coefficients are computed for the budget share of a good and the negative of other as budget shares must be negatively dependent.

methods proposed by Schmid and Schmidt (2007) and implemented in the R package copula. As the dependence between good shares was shown to be negative (see Table 3), there is no tail dependence between them: if the share of Good 1 is unusually high, it would be less likely for the share of Good 2 to be high too, as they must add up to one. However, when the share of Good 2 is rotated by multiplying by -1, mild tail dependence appears in both the right and left tails.

This unconditional analysis of budget shares conveys limited information about the conditional residuals that we need to model through marginals and copulas. However, the analysis characterizes important features that at least a subset of models in our meta-analysis need to allow for: negative dependence and both left and right tail dependence. To test different specifications, we consider the Normal, Frank, Clayton, and Gumbel copulas. This set includes elliptic, archimedean, symmetric, and non-symmetric copulas that allow for no tail dependence, left tail dependence, and right tail dependence. The Clayton and Gumbel copulas do not allow for negative dependence (see Trivedi and Zimmer, 2007). To address this issue we allow for the rotation in (2.3).

4 Copula distance fitness measure

This section analyses how our tested specifications perform at capturing the bivariate relationship between Goods 1 and 2's budget shares. With this purpose, we use an alternative fitness measure that is only concerned with how close the parametric estimate of the copula in (2.2) is to a non-parametric estimate of the empirically observed data. This measure is suggested by Trivedi and Zimmer (2007, Section 4.3, p.65) for copula selection.

For every specification whose estimation converges, we have a parametric copula $C^p(\cdot,\cdot)$ (either Normal, Frank, Clayton, Gumbel, or their rotations), an estimated dependence parameter $\hat{\theta}$, parametric marginals F_i and their parameters $\hat{\psi}_i$ $(i = 1, 2)$; and estimated residuals $\hat{u}_{h;i}$ for $i = 1, 2$ and $h = 1, ..., H$. We will compare the parametric copula with a non-parametric estimator of the copula in the joint shares' CDF.

The empirical non-parametric counterpart of $C^p(\cdot, \cdot)$ is computed as in Equation 4.9 of Trivedi and Zimmer (2007, Section 4.3, p.65)

$$C^e(u^1, u^2) = \frac{1}{H} \sum_{h=1}^{H} \mathbf{1}\{\hat{u}_{h,1} \leq u^1\} \times \mathbf{1}\{\hat{u}_{h,2} \leq u^2\}. \tag{4.1}$$

where $\mathbf{1}\{\cdot\}$ represents the indicator function, which takes the value of one when its argument is true and zero otherwise. We compute the mean squared distance between C^p and C^e at all our pseudo-observations as in Equation 4.10 of Trivedi and Zimmer (2007, Section 4.3, p.66):

$$\text{Distance} = \frac{1}{H} \sum_{h=1}^{H} (C^e(\hat{u}_{h,1}, \hat{u}_{h,2}) - C^p(F_1(s_{h;1}|g_1(p_h, y_h\hat{\vartheta}), \hat{\psi}_1), F_2(s_{h;2}|g_2(p_h, y_h\hat{\vartheta}), \hat{\psi}_2); \hat{\theta}))^2. \tag{4.2}$$

To analyze if there is a discernible relationship between gasoline's estimated own-price elasticity and the accuracy of modeling of the relationship between good shares, Figures 5 and 6 reproduce Figures 2 and 3 from the main text, using the Copula Distance in (4.2) as an alternative fitness measure to the BIC.

Figure 5 does not display a clear divide in fitness as it happened with its BIC counterpart. The improvement with respect to CS's original estimates under this metric are also much smaller than with the BIC. Nevertheless, improvements can be made and the best model for every dataset again has a skewed normal marginal distribution for Good 2. No particular copula seems to do significantly better than the others under this metric.

As in the main text, Figure 6 restricts Figure 5 to those specifications that use a Skewed Normal marginal for Good 2, and marks different models for the conditional mean of shares with different colors. As in the last plot, there is not a model or copula which does clearly better than the others.

5 Chang and Serletis' (2014) Replicated tables

This section presents our own replications of each table in Chang and Serletis (2014) and compares them to the original results.

5.1 AIDS model

Table 5 displays the estimation results for the Almost Ideal Demand System (AIDS) model (Deaton and Muellbauer, 1980), both with and without local curvature imposed, for each of the three data sets. We report the number of violation for every regularity condition. We also check whether curvature is satisfied at $y = p^1 = p^2 = p^3 = 1$, to evaluate if our imposition of curvature achieves its goal.

Our parameter estimates differ from those reported by Chang and Serletis (2014). The greatest differences are found in α_0 and α. This can be due to a common issue, described by Deaton and

Figure 5: Gasoline Own-Price Elasticities vs Copula Distance.

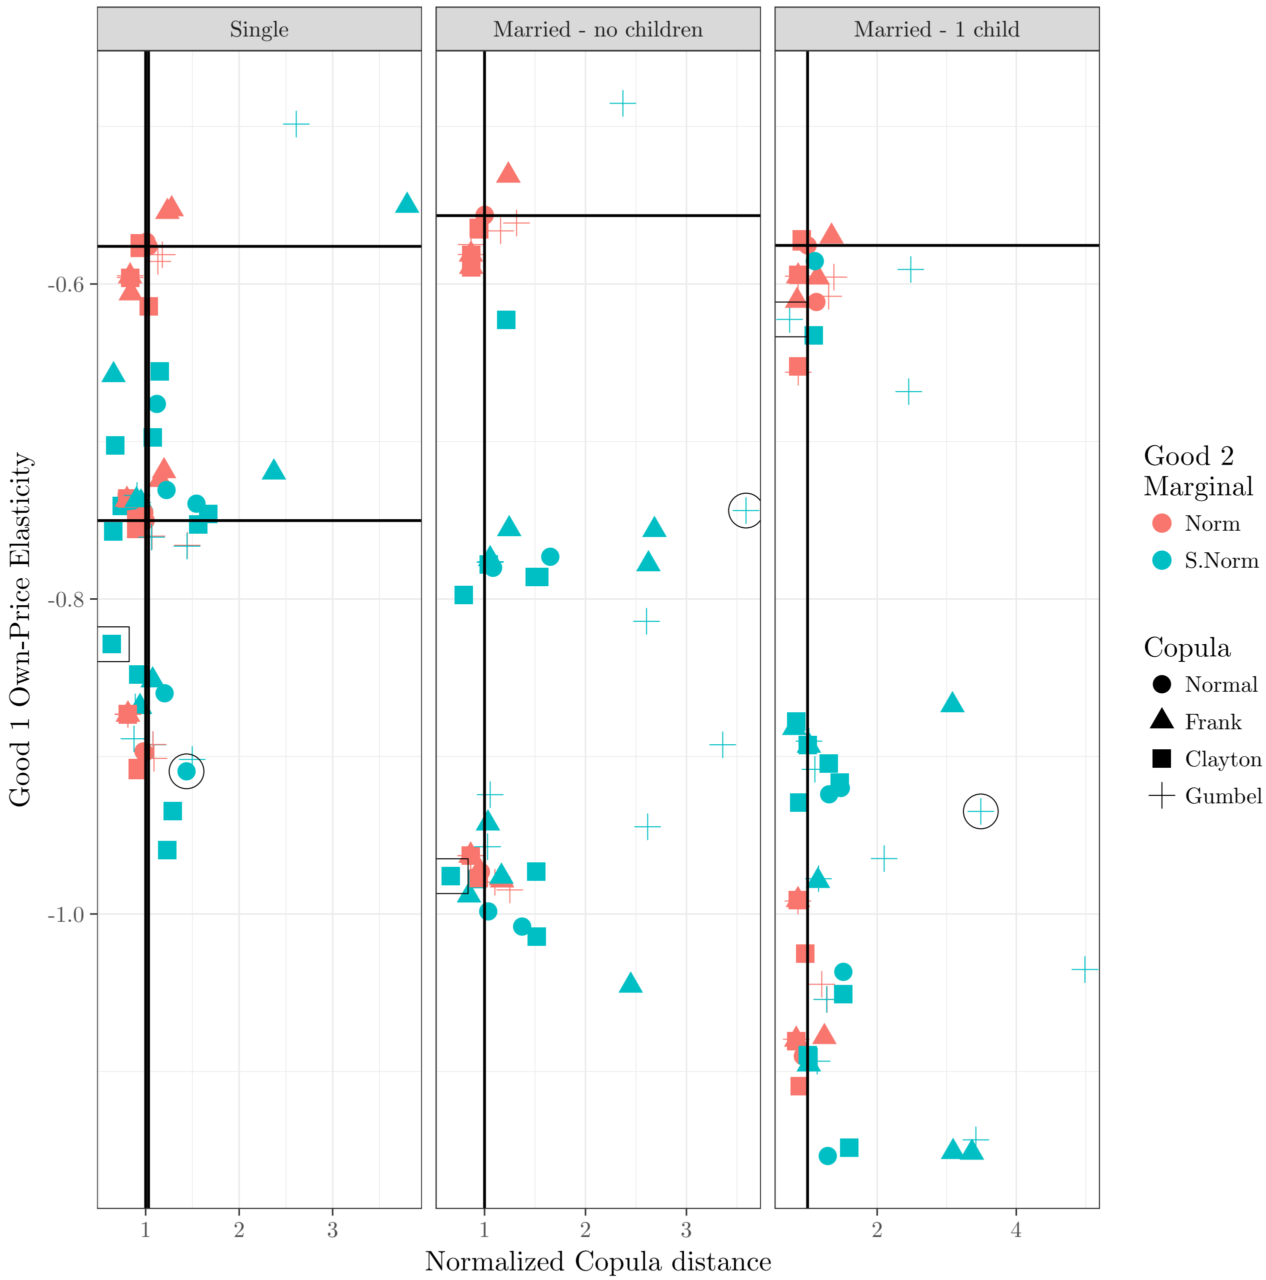

Note: (a) Elasticities are computed at sample mean prices and total expenditures. (b) Only models with no regularity violations are reported. (c) Elasticities corresponding to the models with the lowest Copula Distance are marked with a black square, and those with the lowest BIC are market with a black circle. (d) Crosses mark own replicated estimates of the models in CS (restricted AIDS and ML for single-member households, and restricted ML for the two other data sets). (e) The x-axis is normalized dividing the fitness measures by the absolute value of the best (lowest) Copula Distance of CS's specifications for each dataset. (f) Restricted models (local curvature imposed for AIDS and QUAIDS, and global curvature imposed for ML) are marked with "(R)".

Figure 6: Gasoline Own-Price Elasticities vs Copula Distance with Skewed Normal Marginals for Good 2.

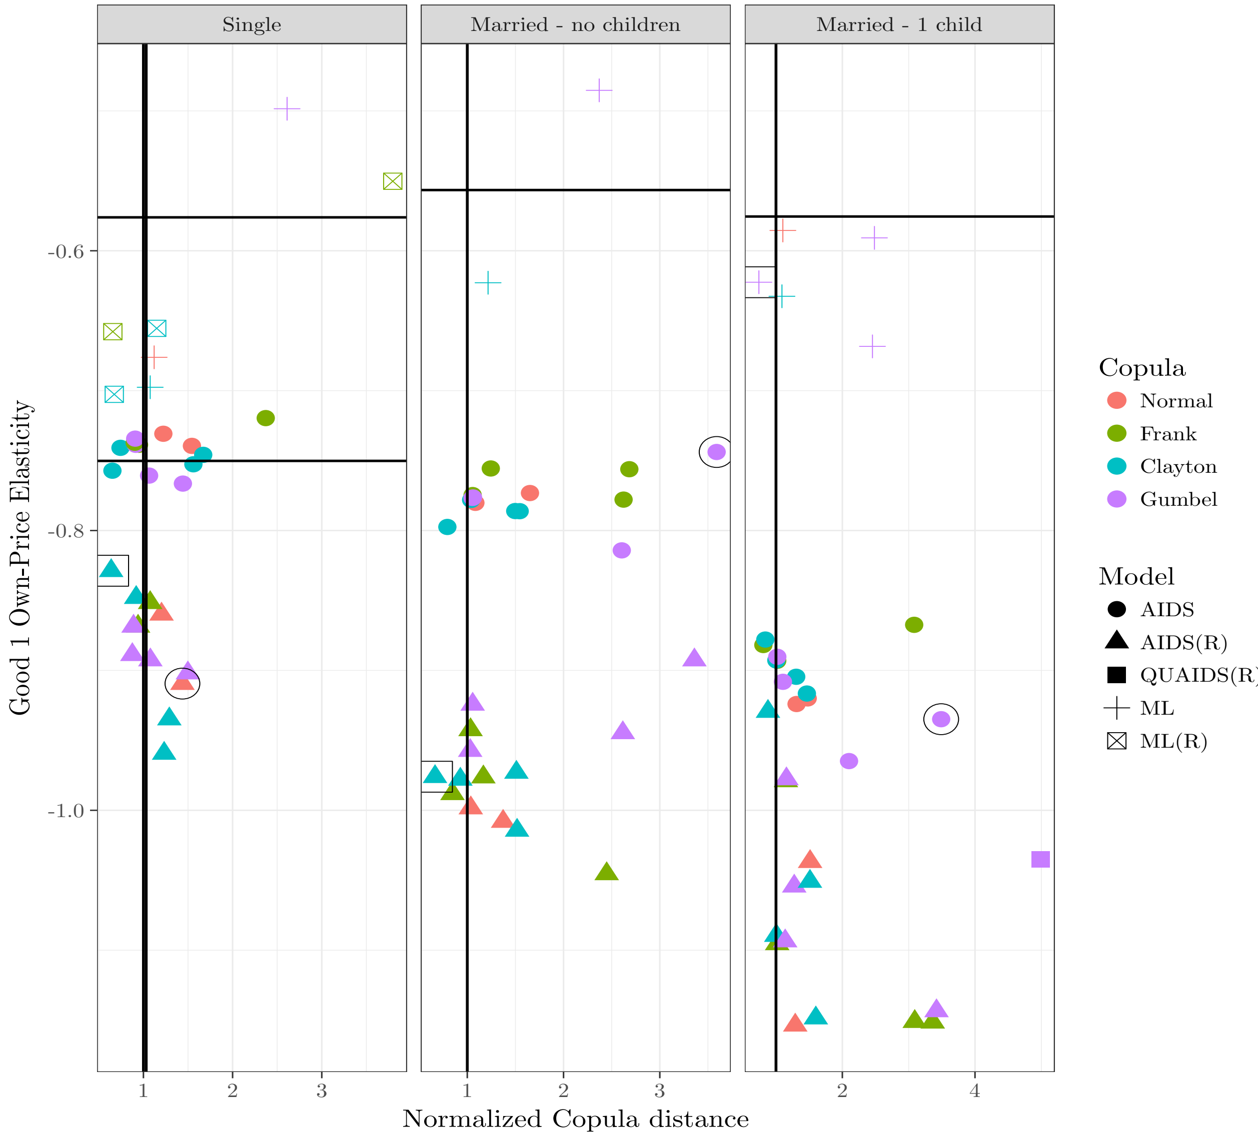

Note: (a) This figure is analogous to Figure 5 but it focuses on specifications that have a Skewed Normal distribution as the marginal of Good 2's budget share residuals. (b) Elasticities are computed at sample mean prices and total expenditures. (c) Only models with no regularity violations are reported. (d) Elasticities corresponding to the models with the lowest Copula Distance are marked with a black square, and those with the lowest BIC are market with a black circle. (e) Crosses mark own replicated estimates of the models in CS (restricted AIDS and ML for single-member households, and restricted ML for the two other data sets). (f) The x-axis is normalized dividing the fitness measures by the absolute value of the best (lowest) Copula Distance of CS's specifications for each dataset. (g) Restricted models (local curvature imposed for AIDS and QUAIDS, and global curvature imposed for ML) are marked with "(R)".

Muellbauer (1980) (Section D) in the original formulation of the AIDS model: these parameters are practically hard to identify when individual prices are closely collinear. This is applicable to our case, as the cross correlations of good prices are above 0.9 for every possible pair of goods and all datasets.

The estimates of γ sometimes have similar values but are often different. It is important to note that none of the AIDS parameter estimates reported by Chang and Serletis (2014) for any database is significant at a 5% level except β_1 and β_2. Our estimates for these parameters (β_1 and β_2) closely match theirs when local curvature is not imposed. However, imposing this restriction considerably modifies β_1 in our results, which does not happen in their estimations.

Our regularity checks for unrestricted estimations are very similar to those reported in the paper: there are no positivity or monotonicity violations, and curvature is not satisfied except for single member households, for which there are no violations. In our case, imposing local curvature eliminates violations for all databases: this is another difference, as in their estimations curvature violations are reduced to less than half, but still occur.

We also verify that curvature is not satisfied at the point of approximation ($y = p = 1$) in the unrestricted models but it is when imposed.

5.2 QUAIDS model

Table 6 presents our parameter estimates and regularity checks for the Quadratic Almost Ideal Demand System (QUAIDS) model (Banks et al., 1997). One notices that our estimates differ from those reported by Chang and Serletis (2014). They are close only in the sense that log-likelihood values are similar and that β_1 and λ_1 are always significant.
Regularity checks are also have different results: the unrestricted model generates violations of both the positivity and monotonicity conditions for single member households, but has a very low number of curvature violations. Imposing local curvature further reduces violations and makes positivity and monotonicity conditions be satisfied. For married households (with and without children) curvature conditions are violated by every observation in unrestricted estimations, and imposing local curvature does little to change this. On the other hand, the estimates from Chang and Serletis (2014) report no positivity or monotonicity violations. Their estimated curvature violations are also greatly reduced when local curvature is imposed.

5.3 ML model

We now present our results for the ML model (Barnett, 1983). Since we take a different approach and normalization than the ones used by Chang and Serletis (2014) (see Section 2), our parameter estimates are not comparable. Our estimates and regularity checks are displayed in Table 7.

None of our estimates produce positivity or monotonicity violations. However, unrestricted models have high numbers of curvature violations. These violations are corrected when global curvature is imposed. No element of B is found to be significantly different from 0 using any database, with or without global curvature imposed. This is consistent with Chang and Serletis (2014), as the only element of B that they explicitly report ($B_{1,2}$) displays the same behavior.

5.4 Elasticities

Chang and Serletis (2014) report elasticities only for model-database combinations that produce no regularity violations, which in their case are the unrestricted AIDS model for single households and the restricted ML model with all household types. For comparison, we present our estimates for the same model-database combinations in Table 8.

The unrestricted AIDS elasticities for single member households precisely match those reported by Chang and Serletis (2014). The ML estimates have differences under 5×10^{-2} for single and married without children households. Differences are greater for households with one child, especially for $\eta_{i,2}$.

References

Banks J, Blundell R, Lewbel A. 1997. Quadratic engel curves and consumer demand. *The Review of Economics and Statistics* **79**: 527–539.

Barnett WA. 1983. New indices of money supply and the flexible laurent demand system. *Journal of Business and Economic Statistics* **1**: 7–23. ISSN 07350015.

Barten AP. 1969. Maximum likelihood estimation of a complete system of demand equations. *European Economic Review* **1**: 7–73.

Chang D, Serletis A. 2014. The demand for gasoline: Evidence from household survey data. *Journal of Applied Econometrics* **29**: 291–313.

Deaton A, Muellbauer J. 1980. An almost ideal demand system. *The American Economic Review* **70**: 312–326.

Geenens G, Charpentier A, Paindaveine D. 2017. Probit transformation for nonparametric kernel estimation of the copula density. *Bernoulli* **23**: 1848–1873.
URL https://doi.org/10.3150/15-BEJ798

Ryan DL, Wales TJ. 1998. A simple method for imposing local curvature in some flexible consumer-demand systems. *Journal of Business & Economic Statistics* **16**: 331–338.

Schmid F, Schmidt R. 2007. Multivariate conditional versions of spearman's rho and related measures of tail dependence. *Journal of Multivariate Analysis* **98**: 1123 – 1140. ISSN 0047-259X.
URL http://www.sciencedirect.com/science/article/pii/S0047259X06000662

Trivedi PK, Zimmer DM. 2007. Copula modeling: An introduction for practitioners. *Foundations and Trends(R) in Econometrics* **1**: 1–111.

Table 5: AIDS model parameter estimates

	Single Members		Married Couples		Married Couples with one Child	
Parameter	Unrestricted	Local Curvature Imposed	Unrestricted	Local Curvature Imposed	Unrestricted	Local Curvature Imposed
α_0	**0.675*****	**1.235*****	**0.492*****	**1.205*****	**0.759*****	**6.706*****
	(0.003)	(0.001)	(0.004)	(0.001)	(0.003)	(0.000)
α_1	**0.929*****	**0.770*****	**1.122*****	**0.790*****	**1.226*****	**0.648*****
	(0.057)	(0.038)	(0.055)	(0.033)	(0.039)	(0.007)
α_2	**0.297*****	**0.274*****	**0.285*****	**0.252*****	**0.275*****	**0.131*****
	(0.031)	(0.028)	(0.027)	(0.025)	(0.020)	(0.005)
$\gamma_{1,1}$	0.062	0.017	0.023	−0.017	−0.051	**−0.078*****
	(0.042)	(0.036)	(0.034)	(0.030)	(0.026)	(0.023)
$\gamma_{1,2}$	−0.033	−0.038	**−0.039***	**−0.052****	**−0.054*****	**−0.065*****
	(0.022)	(0.022)	(0.016)	(0.016)	(0.013)	(0.012)
$\gamma_{2,2}$	−0.058	**−0.133***	**0.139****	−0.010	**0.144*****	−0.006
	(0.065)	(0.055)	(0.044)	(0.031)	(0.034)	(0.024)
β_1	**−0.064*****	**−0.044*****	**−0.080*****	**−0.039*****	**−0.087*****	**−0.042*****
	(0.008)	(0.006)	(0.007)	(0.005)	(0.005)	(0.003)
β_2	**−0.029*****	**−0.027*****	**−0.027*****	**−0.024*****	**−0.024*****	**−0.021*****
	(0.004)	(0.004)	(0.004)	(0.004)	(0.003)	(0.003)
Positivity violations	0	0	0	0	0	0
Monotonicity violations	0	0	0	0	0	0
Curvature violations	0	0	3326	0	6141	0
Log-likelihood	1682.069	1675.514	2951.997	2919.031	5155.302	5085.156
Num. observations	2218	2218	3326	3326	6141	6141
Curvature at $p = 1$, $y = 1$	0	1	0	1	0	1

Note: (a) This table replicates results in Table II in Chang and Serletis (2014, p. 305). (b) Numbers in parentheses are p-values. (c) Local curvature is imposed at the reference point $p^* = y = 1$ (all prices and total expenditure equal to one). (d) Stars follow the key: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. (e) Chang and Serletis (2014) do not report restricted model estimates for single member households. (f) We present an additional row checking whether curvature is satisfied at the reference point.

Table 6: QUAIDS model parameter estimates

	Single Members		Married Couples		Married Couples with one Child	
Parameter	Unrestricted	Local Curvature Imposed	Unrestricted	Local Curvature Imposed	Unrestricted	Local Curvature Imposed
α_0	**1.138*****	**1.637*****	**1.982*****	**0.114*****	**0.376*****	**1.977*****
	(0.021)	(0.014)	(0.019)	(0.011)	(0.014)	(0.008)
α_1	**−0.384*****	**0.248*****	−0.053	−0.056	**−0.205*****	**0.306*****
	(0.060)	(0.031)	(0.083)	(0.067)	(0.039)	(0.033)
α_2	0.161	**0.365****	0.097	**0.177***	**0.343***	**0.288*****
	(0.161)	(0.125)	(0.130)	(0.077)	(0.154)	(0.075)
$\gamma_{1,1}$	**−0.344*****	−0.013	**−0.204*****	**−0.193*****	**−0.416*****	**−0.144*****
	(0.056)	(0.050)	(0.059)	(0.047)	(0.030)	(0.034)
$\gamma_{1,2}$	−0.040	0.005	−0.052	−0.031	0.005	−0.019
	(0.065)	(0.030)	(0.043)	(0.022)	(0.051)	(0.019)
$\gamma_{2,2}$	−0.057	−0.065	**0.139****	**0.107***	**0.141*****	**0.128*****
	(0.065)	(0.067)	(0.045)	(0.042)	(0.036)	(0.032)
β_1	**0.356*****	**0.157*****	**0.285*****	**0.239*****	**0.305*****	**0.186*****
	(0.017)	(0.014)	(0.028)	(0.016)	(0.007)	(0.012)
β_2	0.010	−0.063	0.024	0.003	−0.041	−0.039
	(0.052)	(0.043)	(0.045)	(0.020)	(0.040)	(0.025)
λ_1	**−0.034*****	**−0.019*****	**−0.031*****	**−0.021*****	**−0.026*****	**−0.023*****
	(0.001)	(0.002)	(0.003)	(0.001)	(0.001)	(0.001)
λ_2	−0.003	0.003	−0.004	−0.002	0.001	0.001
	(0.004)	(0.004)	(0.004)	(0.001)	(0.003)	(0.002)
Positivity violations	1	0	0	0	0	0
Monotonicity violations	1	0	0	0	0	0
Curvature violations	37	2	3326	3324	6141	6141
Log-likelihood	1691.530	1689.518	2967.956	2964.580	5184.875	5181.560
Num. observations	2218	2218	3326	3326	6141	6141
Curvature at $p = 1$, $y = 1$	0	1	0	1	0	1

Note: (a) This table replicates results in Table III in Chang and Serletis (2014, p. 306). (b) Numbers in parentheses are p-values. (c) Local curvature is imposed at the reference point $p^* = y = 1$ (all prices and total expenditure equal to one). (d) Stars follow the key: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. (e) We present an additional row checking whether curvature is satisfied at the reference point.

Table 7: ML model parameter estimates

	Single Members		Married Couples		Married Couples with one Child	
Parameter	Unrestricted	Global Curvature Imposed	Unrestricted	Global Curvature Imposed	Unrestricted	Global Curvature Imposed
a_2	**−0.001*****	−0.000	**−0.001*****	−0.000	**−0.001*****	0.000
	(0.000)	(0.010)	(0.000)	(0.006)	(0.000)	(0.004)
a_3	**0.007*****	**−0.122*****	**0.004*****	**0.107*****	**0.002*****	**−0.090*****
	(0.001)	(0.006)	(0.000)	(0.001)	(0.000)	(0.001)
$A_{1,1}$	**0.374*****	**−0.797*****	**0.395*****	**−0.863*****	**0.355*****	**−0.903*****
	(0.008)	(0.038)	(0.003)	(0.002)	(0.003)	(0.001)
$A_{1,3}$	**0.521*****	**−0.294*****	**0.533*****	**0.219*****	**0.609*****	**−0.127*****
	(0.016)	(0.068)	(0.013)	(0.000)	(0.007)	(0.000)
$A_{3,3}$	**−0.092*****	**0.000****	**−0.109*****	**0.000*****	**−0.235*****	**0.000*****
	(0.021)	(0.000)	(0.013)	(0.000)	(0.007)	(0.000)
$A_{2,2}$	**0.172*****	**0.417*****	**0.147*****	**0.383*****	**0.141*****	**0.380*****
	(0.016)	(0.008)	(0.010)	(0.004)	(0.007)	(0.003)
$B_{1,2}$	0.000	−0.000	**0.000***	0.000	−0.000	−0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$B_{2,3}$	0.000	−0.000	−0.000	−0.000	−0.000	−0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Positivity violations	0	0	0	0	1	0
Monotonicity violations	0	0	0	0	1	0
Curvature violations	2215	0	3326	0	6141	0
Log-likelihood	1682.834	1677.629	2958.880	2944.850	5181.807	5131.809
Num. observations	2218	2218	3326	3326	6141	6141

Note: (a) This table replicates results in Table IV in Chang and Serletis (2014, p. 307). (b) We adopt a different parameter standardization scheme from the one used by Chang and Serletis (2014) (see Section 2), therefore parameter values are not directly comparable. (c) Numbers in parentheses are p-values. (d) Stars follow the key: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$.

Table 8: Elasticities (Only those reported in Chang and Serletis, 2014)

Good	Model		Elasticities							
			Income		Own and cross price					
(i)		Curv. Impos.	η_i		$\eta_{i,1}$		$\eta_{i,2}$		$\eta_{i,3}$	
A. Single members (2218 observations)										
(1)	AIDS	✗	0.869	(0.000)	-0.750	(0.000)	-0.029	(0.524)	-0.090	(0.263)
	M. Laurent	✓	0.852	(0.000)	-0.576	(0.000)	-0.091	(0.000)	-0.185	(0.000)
(2)	AIDS	✗	0.678	(0.000)	-0.064	(0.796)	-1.549	(0.032)	0.935	(0.208)
	M. Laurent	✓	0.777	(0.000)	-0.453	(0.000)	-0.091	(0.000)	-0.233	(0.000)
(3)	AIDS	✗	1.221	(0.000)	-0.276	(0.003)	0.151	(0.341)	-1.095	(0.000)
	M. Laurent	✓	1.221	(0.000)	-0.397	(0.000)	-0.091	(0.000)	-0.733	(0.000)
B. Married couples (3326 observations)										
(1)	M. Laurent	✓	0.826	(0.000)	-0.557	(0.000)	-0.081	(0.000)	-0.189	(0.000)
(2)	M. Laurent	✓	0.797	(0.000)	-0.501	(0.000)	-0.081	(0.000)	-0.216	(0.000)
(3)	M. Laurent	✓	1.263	(0.000)	-0.467	(0.000)	-0.081	(0.000)	-0.716	(0.000)
C. Married couples with one child (6141 observations)										
(1)	M. Laurent	✓	0.833	(0.000)	-0.576	(0.000)	-0.089	(0.000)	-0.168	(0.000)
(2)	M. Laurent	✓	0.833	(0.000)	-0.566	(0.000)	-0.089	(0.000)	-0.177	(0.000)
(3)	M. Laurent	✓	1.318	(0.000)	-0.552	(0.000)	-0.089	(0.000)	-0.677	(0.000)

Note: (a) This table replicates results in Table V in Chang and Serletis (2014, p. 309), which reports elasticities of the models they found to produce no regularity violations. (b) Numbers in parentheses are p-values. (c) 'Curv. Impos.' denotes whether curvature is imposed in the reported estimation or not.