

INTEGRABILITY OF DEMAND ACCOUNTING FOR
UNOBSERVABLE HETEROGENEITY: A TEST ON PANEL
DATA

Mette Christensen

Integrability of Demand Accounting for Unobservable Heterogeneity: A Test on Panel Data*

Mette Christensen[†]

Abstract

In recent years it has become apparent that we must take unobservable heterogeneity into account when conducting empirical consumer demand analysis. This paper is concerned with integrability (that is, whether demand is consistent with utility maximization) of the conditional mean demand (that is, the estimated demand) when allowing for unobservable heterogeneity. Integrability is important because it is necessary in order for the demand system estimates to be used for welfare analysis. Conditions for conditional mean demand to be integrable in the presence of unobservable heterogeneity are developed in the literature. There is, however, little empirical evidence suggesting whether these conditions for integrability are likely to be met in the data or not. In this paper we exploit the fact that the integrability conditions have testable implications for panel data and use a unique long panel data set to test them. Because of the sizeable longitudinal length of the panel, we are able to identify a very flexible specification of unobservable heterogeneity: We model individual demands as an Almost Ideal Demand system and allow for unobservable heterogeneity by allowing all intercept and slope parameters of the demand system to be individual-specific. We test the conditions for integrability of the conditional mean demand of this demand system. We do not reject them. This means that the conditional mean demand generated by a population of consumers with different preferences described by different Almost Ideal Demand systems is consistent with utility maximization. Given that integrability is not rejected, we conclude by comparing the estimated demand system elasticities and welfare effects from a model with no heterogeneity (which is the model that would usually be estimated from cross sectional data) to those obtained from our heterogeneous model. We find that the homogeneous model severely overestimates income elasticities for luxury goods and that the welfare effects from the heterogeneous model exhibit a large amount of heterogeneity, but deviate with only a few percentage points from the homogeneous model at the mean.

JEL: C33, D12, D60

Keywords: Demand, Panel data, Unobservable heterogeneity, Preferences, Integrability

*I am grateful to Walter Beckert, Richard Blundell, Martin Browning, Mette Ejrnaes, Stefan Hoderlein, Arthur Lewbel, Birgitte Sloth, Frank Windmeijer, Allan Würtz and seminar participants at University of Copenhagen, Mannheim University, The 11th International Conference on Panel Data 2004, The Econometric Society European Meeting 2004, Essex University, IFS, Cemmap, Queen Mary College, Birkbeck College, Aarhus University and University of Manchester for many helpful suggestions and discussions. Financial support from the European Community's Human Potential Programme under contract HPRN-CT-2002-00235 [AGE] is gratefully acknowledged. All errors are mine.

[†]Department of Economics, University of Manchester, and The Institute for Fiscal Studies, London. E-mail: mette.christensen@manchester.ac.uk

1 Introduction

Demand system estimation provides estimates of price and income elasticities, as well as estimates of the effects of demographic variables on demands. These elasticities and effects are important inputs into many policy analyses; for example the analysis of the effects of income or commodity taxes on market demands and the implications these effects have for welfare. Indeed, one of the main motivations for estimating demand systems is to facilitate welfare analysis. For this it is necessary that demands are consistent with consumer theory; that is, it is necessary that demands are integrable¹.

In this paper we exploit a unique data set to test conditions that are necessary for the conditional mean demand function across consumers to be integrable, assuming that individual consumers all separately maximise utility. The conditional mean demand is of great importance in empirical demand analysis because this is the demand that is estimated: When we estimate a demand system, what we estimate is always the average demand, conditional on observables. To see this, consider the usual way of modelling and estimating demand systems. The usual way of modelling demand systems is by an additive model: The demand of each individual is modelled as the sum of a systematic component, i.e. some function, and an additive error term. The systematic component is functionally dependent on observables, like prices, incomes and observable demographics, and is common for all individuals: Different individuals have different values of observables, but the systematic component is the same function for all individuals. The error term is functionally independent of observables and is specific to each individual. The systematic component is then estimated from data, typically employing the assumption that the additive error term has conditional mean zero. This, together with the conditional mean zero condition on error terms, implicitly defines the systematic component to be the average budget share function. This means that what is in fact being estimated is the conditional mean budget share function. In order to use the estimated coefficients from a demand system for welfare analysis we thus need the conditional mean demand to be integrable.

There is a vast literature on demand system estimation and this literature provides an abundance of different ways of specifying the conditional mean demand to be estimated, both parametrically and nonparametrically. Widely used examples of parametric specifications are the Almost Ideal Demand system (introduced by Deaton and Muellbauer (1980)) and the Quadratic Almost Ideal Demand system (introduced by Banks, Blundell and Lewbel (1997)). Nonparametric specifications include Härdle and Jerison (1988), Lewbel (1991), Blundell, Duncan and Pendakur (1998) and Blundell, Chen and Kristensen (2007). The data used for demand system estimation is nowadays typically cross sectional household expenditure survey data, that is, data at the household level, where each household is observed only once. A couple of decades ago demand systems were estimated using aggregate data; thus Deaton and Muellbauer (1980) used aggregate data. Regardless of whether data is aggregate level data or household level data, and regardless of whether the demand system is parametrically or nonparametrically estimated, what is always estimated is the conditional mean demand.

When modelling demand systems in the usual way, unobservable individual-level heterogeneity is

¹By the term "integrability" we will understand that demand is generated from maximising a utility function subject to a linear budget constraint.

assumed to be captured in the additive, individual-specific error term, which is not taken into account when the integrability properties of the demand system are analysed. In recent years, however, it has become widely acknowledged that we must take unobservable heterogeneity into account when conducting empirical consumer demand analysis. In an early paper, Brown and Walker (1989) consider the additive demand model and show that in order for integrability to hold at both the individual and the average level, the individual error terms must be functionally dependent on prices and/or incomes. In other words: If we want integrability to hold both at the individual level and at the average level, we can not employ the usual modelling of demand systems with additive independent error terms. In contrast to the usual way of modelling demand systems, the random utility approach takes unobservable heterogeneity specifically into account by letting a random component either enter the individual utility maximization problem directly, or letting it enter the demand system. The idea is that each consumer has his or her own value of the random component (coefficient) and hence the distribution of the random coefficients represents the distribution of preference heterogeneity. Each consumer knows his or her own value of the random coefficient, but this value is unknown to the researcher. In a recent paper, Lewbel (2001) adopts the random utility approach to the usual additive model and derives conditions under which the conditional mean demand is integrable, assuming that individual demands are integrable. The conditions imply the result of Brown and Walker (1989). These integrability conditions are conditions on a matrix whose elements are covariances across individuals (households) between individual income responses and individual demands and hence they have testable implications for panel data: The conditions can be tested as properties of an estimate of the matrix of covariances. Because each element of the matrix is a covariance across individuals (households) between individual income responses and individual demands, this matrix can be estimated from panel data by a simple two step procedure. In a first step, the time series variation of the data can be used to estimate the individual income responses and the individual demands for all individuals. In a second step, the cross sectional variation in the data can be used to estimate the elements of the matrix of covariances as the sample covariances across households of the estimated individual effects.

In this paper, we exploit the unique long time series dimension of a Spanish panel data set on households expenditures to empirically test the integrability conditions developed in Lewbel (2001), following the two step procedure outlined above. The data is the Encuesta Permanente de Consumo (the ECP), which is a 6 year long data set with quarterly information on household expenditures, on prices and on demographics, collected by the Spanish National Bureau of Statistics in the period 1978-83 ². The ECP is to the best of our knowledge the longest real panel on households consumption covering a wide range of commodity groups we have available. The exceptional long time-series dimension allows us to estimate the individual income responses and individual demands that are needed for each household in order to construct an estimate of the matrix of covariances. It also allows us to be far more flexible in our specification of preference heterogeneity than other studies estimating demand systems with unobservable heterogeneity.

We take a semiparametric approach and model the demand of each household by an Almost Ideal

²The data was kindly provided by Lola Collado. Many thanks to her for answering numerous queries about the data. Also thanks to José M. Labeaga for help with this data.

Demand system. The Almost Ideal Demand system has linear Engel curves in log total expenditure. We introduce preference heterogeneity by allowing all the intercept and slope parameters of the Engel curves to be household-specific, but we impose no parametric restrictions on the distribution of preference parameters. Identification of such a flexible heterogeneity scheme is possible only because we have this long panel data. The idea is to view the time series dimension for each household as a repetition of the same thing. This is justified by the theory underlying demand systems: Demand systems are models of consumer behavior that describe how consumers allocate total expenditures to consumption goods *within* the period, given that consumers have already allocated a given amount of total expenditure to each period by solving an intra-temporal optimisation problem³. In other words, demand systems are static models and thus we have no dynamics in our model. We test for integrability and find that we can not reject integrability. Furthermore, this non-rejection is fairly significant with p -values as high as fifty percent. This finding implies that a set of completely heterogeneous Almost Ideal Demand systems generate an integrable conditional mean demand. The value of having integrable demands is that it facilitates welfare analysis. Suppose for example that the government increases the tax on Alcohol & Tobacco; from the demand system estimates alone we can assess quantitatively how demands will change following the tax increase, but in order to assess how the welfare of the consumers change, we need a utility framework. Given that integrability is not rejected, we know that the conditional mean demand is generated by utility maximisation of *some* utility function. But as mentioned above, we do not know what this utility function looks like (and it is beyond the scope of this paper to find out). We therefore calculate the welfare effects at the individual level for *each* household, utilising that each household is maximising an Almost Ideal utility function. We consider a tax increase that leads to the price of Food, Alcohol & Tobacco to increase by 20% and estimate the compensating variations for each household, using *that* household's particular utility function, conditional on income. We then compare the means of these conditional distributions of compensating variations to the compensating variation obtained from an AID model with no unobservable heterogeneity. This is interesting because the model with heterogeneity is the model that would usually be estimated from cross sectional data. We find that that compensating variations from the homogeneous model slightly undercompensates consumers as compared to the mean of welfare effects estimated from the heterogeneous model, but they are not very different. However, the substantial amount of heterogeneity in the welfare effects from the heterogeneous model highlights the importance of taking unobservable heterogeneity into account. Carrying out a similar comparison of demand system estimates of price and income elasticities estimated from the two models show that the homogeneous model severely overestimates income elasticities for luxury goods. This could be because the homogeneous model wrongly contributes all of the variation in budget shares to income, when in fact some of the variation is due to taste differences. This last result is in line with the findings of Christensen (2005).

Among our other findings are however some strong rejections of homogeneous consumer behavior: We strongly reject that different consumers have identical income effects. That is, we find strong evidence of preference heterogeneity in marginal propensities to spend. We also strongly reject that the intercepts are identical across consumers. Both these rejections of homogeneous consumer behavior are in stark

³Demand systems are thus the second stage of a two stage budgeting process, see Deaton and Muellbauer (1980).

contrast to what is usually assumed in demand system estimation: Usually, it is assumed that these behavioral coefficients are identical across consumers - at most it is assumed that income responses can vary with observable demographics (the classic example is to allow income responses to be different according to number of children in the household). But our finding suggests that this is not enough to explain the variation in budget shares. Our findings thus also adds to the growing body of research which shows the importance of taking unobservable heterogeneity into account. This is also evident when comparing our estimates of income elasticities for a model estimated the usual way (i.e. allowing no unobservable heterogeneity) to the estimates we obtain from the conditional mean demand generated from our heterogeneous demand system: The income elasticities for luxury goods are in magnitude far smaller in the heterogeneous model than in the model with no heterogeneity. This could be because some of the variation in the budget shares of these goods in the homogeneous model was wrongly contributed to income, when in fact it was a taste difference, thus leading to the income effects being overestimated in the homogeneous model.

Few other papers have looked at integrability of demand systems when accounting for unobservable heterogeneity. Hoderlein (2004) derives nonparametric tests of negativity and symmetry in a random utility setting which is more general than that of Lewbel (2001). Amongst other features, Hoderlein (2004) considers a model that is more general than the additive model. The framework of Hoderlein (2004) nests the framework of Lewbel (2001) and thus also nests our model. But none of these papers contain any empirical applications. Brown and Matzkin (1998) construct a random utility model by letting a random component enter the direct utility function and then derive the demand equations from the utility maximisation problem. Their paper contains no empirical application. Beckert (2005) estimates the demand for internet services, allowing for preference heterogeneity in a Cobb-Douglas utility framework. This model automatically generates a conditional mean demand which is integrable because individual preferences are Cobb-Douglas preferences. Calvet and Comon (2003) estimate an Almost Ideal Demand system with unobserved heterogeneity, but since they have cross sectional data, they are forced to being very restrictive in their heterogeneity specification; e.g. they can only identify a linear scheme with one heterogeneity parameter per individual and one per good, whereas panel data allows at least one heterogeneity parameter per individual per good. In our model, we furthermore allow heterogeneous price and income responses.

The integrability conditions in Lewbel (2001) bear a strong resemblance to the conditions for integrability of the unconditional average demand in Muellbauer (1975) and Mas-Colell (1985). The matrix of covariances that appear in Lewbel (2001) is roughly speaking a conditional version of the matrix of covariances that appear in Mas-Colell's work on aggregation. It is somewhat surprising that conditioning on observables like income does not provide more structure.

It is worth pointing out that while this paper deals with the question of obtaining integrability at the average level, while assuming integrability at the individual level prevails. A different type of question one could ask is "Can we obtain integrability - or just the Weak Axiom of Revealed Preference - at the average level without assuming integrability at the individual level, but instead by assumptions on the *distributions* of individual behavior?". This is the question asked in a strand of the theoretical demand aggregation literature by amongst others Hildenbrand (1994) and Grandmont (1992). The results of

Hildenbrand and Grandmont have an important feature in common with the conditions in Brown and Walker (1989) and Lewbel (2001). Namely that what is needed is what Hildenbrand and Grandmont denote *behavioral heterogeneity* in preferences, which are distributional assumptions on consumers' *behavior*, i.e. distributional assumptions on how consumers respond to changes in for example income. In other words, they also require that unobservable heterogeneity is functionally dependent on prices and/or incomes.

The rest of this paper is organised as follows. In Section 2 we present the integrability conditions and discuss what they imply for the ways in which we can introduce preference heterogeneity into demand systems in a way that is consistent with consumer theory. In Section 3 we formulate our theoretical model and provide a theoretical discussion of the integrability conditions in the context of our model. Section 4 presents the econometric model to be estimated and Section 5 presents the data. Section 6 describes the estimation method and Section 7 contains the results. In Section 8 we compare estimated demand system elasticities and welfare effects from our heterogeneous model to a model with no heterogeneity. Section 9 concludes.

2 The Integrability Conditions in the Presence of Unobservable Heterogeneity

In this section we present and discuss the theoretical conditions for integrability of a demand system with unobservable heterogeneity developed in Lewbel (2001).

2.1 Demand Systems and Integrability: The Usual Way

In order to put the integrability conditions for a demand system with unobservable heterogeneity into the right perspective, we first consider the usual way of modelling demands in empirical demand analysis. Usually, we specify an additive model: The demands are modelled as the sum of a function of observables (like prices and income) and an error term. The function of observables is common for all households, whereas the error term is household-specific and does not depend on observables. The error term then capture, among other things, unobservable heterogeneity and is typically assumed to have conditional mean zero. Let N denote the number of goods, let w_h denote the vector of budget shares for household h , let $p = (p_1, \dots, p_N)'$ denote the vector of prices for the N goods, let $\ln p$ denote the vector of log prices, let x_h denote total expenditure for household h , let z_h denote a K -dimensional vector of observable characteristics of household h (e.g. demographic characteristics) and let ε_h denote the error term specific to household h . Then the usual additive model can be written

$$w_h = G(\ln p, \ln x_h, z_h) + \varepsilon_h,$$

where

$$E[\varepsilon | \ln p, \ln x, z] = 0.$$

The function $G(\cdot)$ is then estimated from data. There is a vast literature on the estimation of demand systems, with both parametric and nonparametric specifications of $G(\cdot)$. One of the most well-known examples of a parametric form of $G(\cdot)$ is the Almost Ideal Demand system (hereafter denoted the AID system), introduced by Deaton and Muellbauer (1980b) and since used in numerous applications. The data used for demand system estimation is nowadays typically cross sectional household expenditure survey data, but until a couple of decades ago, household-level data were not that common and demand systems would be estimated from aggregate level data⁴. Regardless of whether data is aggregate level data or individual-level cross sectional data, and regardless of whether $G(\cdot)$ is specified parametrically or nonparametrically, the additive structure of the individual budget share function together with the zero mean condition implicitly defines $G(\cdot)$ as the *conditional mean* budget share function. This means that what is in fact being estimated is always an *average* (namely the conditional mean) budget share function.

Now, in order to be able to use the estimated price - and income elasticities from this demand system for welfare analysis, the estimated demand must be consistent with consumer theory, i.e. the conditional mean demand must be integrable. In order to perform welfare analysis at the individual level, we will also need that the behavior of each individual is consistent with consumer theory, i.e. that the individual demands $G + \varepsilon$ are integrable. Traditionally, empirical demand analysis conducts integrability analysis on G without taking the unobservable heterogeneity ε into account⁵. However, if one wants to explicitly interpret unobservable heterogeneity as containing preference heterogeneity, it seems natural to also require that individual demands are integrable. The first question that comes to mind is then whether it is possible that individual demands, $G + \varepsilon$, as well as the estimated demand, namely the conditional mean demand G , are integrable in the usual additive model? The answer to this question is no (Brown and Walker (1989)): If G as well as $G + \varepsilon$ are integrable, ε must be functionally dependent on prices and/or incomes (i.e. the additive error terms must be heteroskedastic)⁶. As a consequence of this result the literature on demand systems therefore turned to formulating demand systems that allow for unobservable heterogeneity to be functionally dependent on prices and incomes. The natural way to do this seems to be to adopt the random utility hypothesis as an approach for randomisation (e.g. Brown and Matzkin (1998), Beckert (2002)).

2.2 Demand Systems and Integrability in the Presence of Unobservable Heterogeneity

Lewbel (2001) adopts the random utility approach to the additive demand system model allowing for general heteroskedasticity of the error term and derive conditions under which the conditional mean demand is integrable, given that individual demands are integrable. These conditions have testable

⁴The AID system was, when it was first introduced in Deaton and Muellbauer (1980(b)), estimated on aggregate level data.

⁵Quoting Brown and Walker (1989): As Barten (1977) remarks, "disturbances are usually tacked on to demand equations as a kind of afterthought".

⁶With the exception of homothetic preferences. If $w_{ih} = \alpha_{ih} + \varepsilon_{ih}$, where $E[\varepsilon_i] = \alpha_{ih} - \bar{\alpha}_i = 0$, where $\bar{\alpha}_i$ is the mean of the α_{ih} 's, then both the individual and the conditional mean demand is integrable (because they are both Cobb-Douglas). However, there seems to be an overall consensus in the literature on demand systems that homothetic preferences are too restrictive to realistically describe consumer behavior.

implications for panel data which is what we utilize in this paper. In order to formulate Lewbel's conditions, consider a sample of H households, $h = 1, \dots, H$. Throughout, we assume independence across households. Let N , w , $\ln p$, $\ln x$ and z be as before. Let η denote an L -dimensional vector of unobservable characteristics with $L \geq N$ ⁷. Let g denote the individual budget share function of household h , and let $F(\eta \mid \ln x, \ln p, z)$ denote the conditional distribution of the unobserved characteristics in the population, conditional on observable characteristics. We can then write individual budget shares as

$$\begin{aligned} w &= g(\ln p, \ln x, z, \eta) \\ &= G(\ln p, \ln x, z) + \nu(\ln p, \ln x, z, \eta), \end{aligned} \tag{1}$$

where G is defined as the conditional mean:

$$\begin{aligned} G(\ln p, \ln x, z) &= E[w \mid \ln p, \ln x, z] \\ &= \int g(\ln p, \ln x, z, \eta) dF(\eta \mid \ln p, \ln x, z). \end{aligned} \tag{2}$$

The definition of G then implies that

$$E[\nu \mid \ln p, \ln x, z] = 0.$$

Notice that this formulation of g in itself imposes no restrictions on individual budget shares: We can choose g to be any budget share function, calculate G from (2) and then construct $\nu(\cdot)$ as the residual $g - G$. Obviously, this formulation nests the usual model.

Before turning to the integrability conditions, let us comment on how unobservable heterogeneity enters in this framework as compared to how it enters in the usual model. In the usual model, preference heterogeneity is implicitly assumed to be captured in the ε 's. Since ε is functionally independent of observables, price - and income effects are restricted to have the same functional form for all households: When differentiating the budget share function with respect to prices or income, there is no contribution from ε . This means that preference heterogeneity can only enter as level effects in the usual model. In other words, the usual model does not allow for unobservable heterogeneity in the marginal effects. The formulation in (1) has the household-specific error term as a function of both unobservables *and* observables (as was shown is necessary for integrability by Brown and Walker (1989)). This means that preference heterogeneity enters not just as level effects as in the usual model, but also as slope effects. For example, two households with identical income levels and identical observable characteristics can have different responses to a change in income. In other words, the result of Brown and Walker (1989) means that in order to ensure integrability both at the individual level of the conditional mean in models that allow for preference heterogeneity, it is necessary that preference heterogeneity enters not just as

⁷In order to ensure that the model produces a non-degenerate distribution of budget shares, it is necessary that there are at least as many unobservables per individual as there are goods (Beckert (2006)).

level effects (i.e. that some households persistently have a high budget share for some good and others a low budget share independently of prices and income levels), but also in the marginal effects (i.e. that different households respond differently to changes in prices or in their incomes, all other things being equal).

The error term being functionally independent on prices and/or incomes, however, only provides necessary conditions for integrability of the conditional mean demand. Sufficient conditions are provided in Lewbel (2001). Define the N by N matrices s and \tilde{s} by

$$s(\ln p, \ln x, z) = \frac{\partial g(\ln p, \ln x, z)}{\partial(\ln p)'} + \frac{\partial g(\ln p, \ln x, z)}{\partial(\ln x)} g(\ln p, \ln x, z)'$$

$$\tilde{s}(\ln p, \ln x, z) = s(\ln p, \ln x, z) + g(\ln p, \ln x, z)g(\ln p, \ln x, z)' - \text{diag}g(\ln p, \ln x, z).$$

These are the budget share analogs to the Slutsky matrix. From classical demand theory we know that the four conditions that ensure integrability of a continuously differentiable budget share function are: Adding up (budget shares add up to one), homogeneity (the budget share function is homogeneous of degree zero in prices and income), symmetry (that s is symmetric⁸) and negativity (that \tilde{s} is negative semidefinite). The corresponding budget share analogs to the Slutsky matrix for the conditional mean demand G are defined similarly

$$S(\ln p, \ln x, z) = \frac{\partial G(\ln p, \ln x, z)}{\partial(\ln p)'} + \frac{\partial G(\ln p, \ln x, z)}{\partial(\ln x)} G(\ln p, \ln x, z)'$$

$$\tilde{S}(\ln p, \ln x, z) = S(\ln p, \ln x, z) + G(\ln p, \ln x, z)G(\ln p, \ln x, z)' - \text{diag}G(\ln p, \ln x, z).$$

We assume that individual demands are integrable. In addition, the following independence assumption is invoked:

$$F_\eta \equiv F(\eta | \ln x, \ln p, z) = F(\eta | z), \quad (3)$$

which, roughly speaking, states that preferences are stochastically independent of prices and income. Under this independence assumption, \tilde{S} can be written

$$\tilde{S} = E[\tilde{s}] - M - \text{Var}[g], \quad (4)$$

where

$$M = \{M_{ij}\}_{i,j=1,\dots,N} = \left\{ \text{Cov} \left[\frac{\partial g_i}{\partial(\ln x)}, g_j \right] \right\}_{i,j=1,\dots,N}$$

and where $\text{Var}[g]$ is the variance-covariance matrix of g . For G to be integrable, G must satisfy adding up, homogeneity, symmetry and negativity. Adding up follows directly because adding up is satisfied at the individual level, and homogeneity follows from homogeneity at the individual level in conjunction with the independence assumption. Since a variance-covariance matrix is always symmetric and positive

⁸Note that \tilde{s} is symmetric if and only if s is symmetric.

definite, the negative of the variance matrix $-\text{Var}[g]$ is also symmetric and negative definite. $E[\tilde{s}]$ is symmetric and negative semidefinite because \tilde{s} is symmetric and negative semidefinite, which follows from integrability at the individual level. Therefore \tilde{S} is symmetric if and only if M is symmetric, and \tilde{S} is negative semidefinite if M is positive semidefinite. Note that the integrability conditions are sufficient conditions: Symmetry of M is necessary and sufficient, whereas the positive semidefiniteness of M is only sufficient.

2.2.1 The Matrix of Covariances

M is a matrix of covariances (not a variance-covariance matrix!). It expresses, roughly speaking, how income effects (i.e. marginal propensities to consume) vary with budget shares in response to changes in unobservables. Mas-Colell, Whinston and Green (1995) interpret the positive semidefiniteness of M as consumers with higher than average consumption of one commodity also tend to spend a higher than average fraction of their last unit of income on that commodity. The expression for the Slutsky matrix of the conditional mean demand, \tilde{S} , in (4) displays clearly how M comes about: The Slutsky matrix for the conditional mean demand is not equal to the conditional average of the Slutsky matrices for the individual demands. There is "something" left over, and this "something" is precisely the matrix of covariances M . Integrability at the individual level ensures that $E[\tilde{s}] - \text{Var}[g]$ is well-behaved (negative semidefinite), but imposes no structure whatsoever on M . This is very similar to the results found in the studies of demand aggregation. Here the averaging of demands is unconditional and thus aggregation happens over consumers with different incomes. But the same matrix of covariances (only now the covariances are unconditional) occurs in the special case where the income distribution is fixed (see the aggregation chapter in Mas-Colell, Whinston and Green (1995)) and also in the case of more general income distributions (see Mas-Colell (1985)). As Mas-Colell, Whinston and Green (1995) remark "the source of the aggregation problem rests squarely with the wealth effects on the consumption side". The result of Lewbel (2001) shows us that conditioning on income does not aid in getting rid of this problem caused by the income effects.

2.2.2 The Distribution of Unobservable Heterogeneity

There are two obvious cases in which the integrability conditions are met. One is the case where M is the zero matrix. This case is for example implied by Gorman aggregation: If all consumers have identical income effects for each commodity, then each element of M is the covariance between a constant and a random variable, which is always zero. Another case is where M is so close to being the zero matrix that the negative semidefiniteness of $E[\tilde{s}] - \text{Var}[g]$ is large enough to make \tilde{S} negative semidefinite. This case can be interpreted as budget shares and income effects varying very little with unobservables, i.e. that there is only little dispersion in preferences across consumers.

The integrability result presupposes that the distribution of unobservables conditional on observable demographics are independent of prices and incomes ((3)). This assumption has recently been examined empirically: Calvet and Comon (200?), Labeaga and Puig (2003), Browning and Collado (2006) and Christensen (200?) all contain empirical evidence suggesting that this assumption may not hold in the

data for all commodities. Calvet and Comon (2003) use the FES. The FES is a cross sectional data set and hence the authors are forced to rely on a restrictive identification scheme (they assume that preference heterogeneity can be described by one random parameter per good and one random parameter per individual) which may account for their strong finding; they find that the majority of the observed variation in budget shares is due to preference heterogeneity and hence that income effects can only explain very little of the observed differences in budget shares. Labeaga and Puig (2003), Browning and Collado (2006) and Christensen (2005) all use panel data (the same panel data set, namely the ECPF) and hence they can allow for more flexible specifications of unobservable heterogeneity than Calvet and Comon. Differing models and tests are employed in the three papers, but the findings are in concordance with each other, namely that for some, but not all, goods, there is evidence of correlated heterogeneity. When employing a similar test for the data set used in this paper we find no or very little evidence against the independence assumption.

3 How to Formulate Preference Heterogeneity and the Research Question

To estimate the matrix of covariances M , we must first specify the individual budget share function g and the distribution of preference heterogeneity F . g can be specified either parametrically or non-parametrically. We choose a parametric specification, because this allows us to model the unobservable household-specific characteristics η by the unknown parameters of the individual budget share functions, which facilitates identification. More precisely, we model the preference heterogeneity by taking g to be a parametric demand system and allowing all the intercept and slope parameters to be different across households. This amounts to specifying g as a variable-coefficient model and accordingly view the coefficients as random variables with a conditional distribution, conditional on observables⁹. This conditional distribution of the coefficients is then the distribution of preference heterogeneity. For example, if we took g to have the most simple Working-Leser form, that is, $g(\ln x) = \alpha + \beta \ln x$, we would view α and β as random variables with a conditional distribution across households such that each household h had its own intercept α_h and its own slope β_h . The values α_h and β_h would be known to household h , but unknown to the econometrician. In estimation, we can choose at one extreme to specify the distribution of coefficients completely nonparametrically, placing no restrictions on its form. In this case, we would estimate the distribution of coefficients as the empirical distribution of the estimated realizations of the random variable underlying the distribution of coefficients. Or, in other words, we would estimate each coefficient for each household. This approach is very general. The cost of this generality is that it involves estimating a large number of parameters; in the example above with g having the Working-Leser form, we would have to estimate one α and one β for each household. At the other extreme, we can choose to model the distribution of coefficients completely parametrically, for example by a normal distribution; in the example with g having the Working-Leser form, this would involve estimating only five parameters: The mean and variance of α , the mean and variance of β and the covariance of α with β . In between the

⁹We will refer to the parameters both as "coefficients" and as "parameters".

two extremes lies a whole range of (semi-parametric) possibilities, like for example assuming a mixture distribution for the coefficients a la Heckman and Singer (1984).

We choose the nonparametric extreme and thus place no restrictions on the distribution of coefficients. The reason for choosing the fully nonparametric approach to the modelling of the distribution of preference heterogeneity is that any restriction on this distribution would be completely ad hoc; this type of model has never been estimated on demand panel data for a complete set of goods before, and so there are no suggestions in the literature about what is reasonable to assume about the distribution of preference heterogeneity. Moreover, as we will show later in Section 3.4, distributional assumptions can actually imply that the matrix of covariances is symmetric and positive semidefinite. In other words, by imposing distributional assumptions, we risk imposing integrability of the conditional mean demand by assumption. Obviously, this would be highly undesirable. As mentioned earlier, the drawback of the fully nonparametric approach is that it involves a large number of parameters to be estimated, and therefore we will expect less precise estimates from this approach than from a parametric or a semiparametric approach.

Obviously, this flexible specification of preference heterogeneity is only feasible because our panel data set has large T . Our approach is similar to the idea underlying the mean-group estimator in Pesaran and Smith (1995): Like them, we also estimate a set of individual-specific parameters for each household (each group), but where Pesaran and Smith (1995) are interested in the average regression coefficient, we are interested in a different function of the estimated coefficients, namely in the matrix of covariances of the income effects with the budget shares, M .

In Section 3.1 we specify g and F , in Section 3.2 we calculate the conditional mean demand for our model, and in Section 3.3 we calculate the object of interest for the integrability test, the matrix of covariances M , for our model. Then we are finally able to state the research question in precise terms. In Section 3.4 we go beyond the model and give some examples of model specifications that in themselves lead to M being symmetric and positive semidefinite. Note that all that follows depend on the choice of g and F , since the matrix of covariances M is specific to these choices.

3.1 Individual Demands: An Almost Ideal Demand System with Household-Specific Parameters and a Nonparametric Distribution of Parameters

We choose g to be an AID system and introduce randomness in demands by making the intercept and slope parameters household-specific. Then the budget share equations for the N goods for household h are given by

$$w_{ih} = \alpha_{ih} + \beta_{ih} [\ln x - \ln P_h(p)] + \sum_j \gamma_{ij} \ln p_j, \quad i = 1, \dots, N, \quad (5)$$

where $P_h(p)$ is the price index given by¹⁰

$$\ln P_h(p) = \sum_k \alpha_{kh} \ln p_k + 0.5 \sum_k \sum_l \gamma_{kl} \ln p_k \ln p_l.$$

¹⁰Since the intercept parameter, usually denoted α_0 , in the price index is not identified we omit it without any loss of generality.

Since we assume that the demand of each individual is generated from utility maximisation, we have the usual restrictions on the parameters of an AID system, but now for α and β , they must hold for each h ¹¹:

$$\sum_{i=1}^n \alpha_{ih} = 1, \quad \sum_{i=1}^n \beta_{ih} = 0, \quad \sum_{i=1}^n \gamma_{ij} = 0 \text{ for all } j \quad (6)$$

$$\sum_{j=1}^n \gamma_{ij} = 0 \text{ for all } i \quad (7)$$

$$\gamma_{ij} = \gamma_{ji} \text{ for all } i, j. \quad (8)$$

Let θ denote the vector of random parameters in the demand system, i.e. $\theta = (\alpha, \beta)'$. The conditional distribution of preference heterogeneity is the conditional distribution of the unknown parameters θ , conditional on log total expenditure, log prices and demographics. Denote it by $F_\theta \equiv F(\alpha, \beta \mid \ln x, \ln p, z)$. The independence assumption (3) translates thus into that θ is conditionally independent of $\ln x$ and of $\ln p$ for all i, j , conditional on demographics z .

3.2 The Conditional Mean Demand

The conditional mean demand for commodity i is calculated from (2) as

$$\int g_i(\ln x, \ln p; \theta) dF = G_i(\ln x, \ln p).$$

The conditional mean budget share function G_i for commodity i is thus given as the conditional mean of the individual budget shares (5) with respect to the joint distribution of the α and β , conditional on $\ln x$, $\ln p$ and z . Let μ_{α_i} and μ_{β_i} denote the conditional means in the marginal distributions of α_i and β_i , $i = 1, \dots, N$, and write the parametric price index as $\ln P(p; \alpha)$ to remind ourselves that it depends α which is a vector of random parameter then using the independence assumption, we get from (5) that

$$\begin{aligned} G_i(\ln x, \ln p) &= \mu_{\alpha_i} + \mu_{\beta_i} \ln x - \int \beta_i \ln P(p; \alpha) dF + \sum_j \gamma_{ij} \ln p_j \quad (9) \\ &= \mu_{\alpha_i} + \mu_{\beta_i} \ln x - \sum_{k=1}^N \left(\int \beta_i \alpha_k dF \right) \ln p_k \\ &\quad - \mu_{\beta_i} 0.5 \sum_k \sum_l \gamma_{kl} \ln p_k \ln p_l + \sum_j \gamma_{ij} \ln p_j, \quad i = 1, \dots, N. \end{aligned}$$

Note that G in itself is not an AID system because of the terms $\int (\beta_i \alpha_k) dF$ stemming from the parametric price index. If it was not for this term, G would be an AID system with parameters given by the conditional mean of the corresponding parameters of the individual AID systems. But because β is not necessarily independent of α , the mean of the product of β_i with α_k is not necessarily equal to the product of the mean of β_i and the mean of α_k , and so G is not an AID system. From this observation it is clear that this must be a general point: When introducing unobservable heterogeneity in preferences by letting the parameters of a parametric demand system vary across individuals, the conditional mean demand will not have the generic form of the individual demands if the individual demands are nonlinear

¹¹See Deaton and Muellbauer (1980a) section 3.4 or Christensen (2005) Chapter 4 Appendix A.

in the parameters (i.e. are nonlinear in the unobservable preference heterogeneity). From this statement it is now clear that a sufficient condition for the conditional mean demand to be an AID system is that β_i is conditionally independent of α_k , $k = 1, \dots, N$, for all $i = 1, \dots, N$:

Example 1 Independent income effects

Assume that the income effect for commodity i , β_i , is independent of all the intercept parameters in the system for every commodity $i = 1, \dots, N$. Then the conditional mean demand G is itself an AID system:

$$G_i(\ln x, \ln p) = \mu_{\alpha_i} + \mu_{\beta_i} \ln x - \int \beta_i \ln P(p; \alpha) dF + \sum_j \gamma_{ij} \ln p_j,$$

and since β_i is independent of α ,

$$\begin{aligned} G_i(\ln x, \ln p) &= \mu_{\alpha_i} + \mu_{\beta_i} \ln x - \left(\int \beta_i dF \right) \sum_{k=1}^N \left(\int \alpha_k dF \right) \ln p_k \\ &\quad - \mu_{\beta_i} 0.5 \sum_k \sum_l \gamma_{kl} \ln p_k \ln p_l + \sum_j \gamma_{ij} \ln p_j \\ &= \mu_{\alpha_i} + \mu_{\beta_i} \left(\ln x - \left[\sum_{k=1}^N \mu_{\alpha_k} \ln p_k - 0.5 \sum_k \sum_l \gamma_{kl} \ln p_k \ln p_l \right] \right) \\ &\quad + \sum_j \gamma_{ij} \ln p_j, \quad i = 1, \dots, N. \end{aligned}$$

Thus, G is an AID system with parameters given as the conditional means of the corresponding individual parameters.

3.3 The Matrix of Covariances and the Research Question

The (i, j) 'th entry in the matrix of covariances, M , is the covariance between the partial derivative of the budget share function for commodity i with respect to $\ln x$ and the budget share for commodity j across households. Since in our model the partial derivative of the budget share function for household h for commodity i with respect to log total expenditure is β_{ih} , $h = 1, \dots, H$, the (i, j) 'th entry of M for our model is given by, where we use the independence assumption of parameters being conditionally independent of prices and total expenditures

$$\begin{aligned} M_{ij} &= \text{Cov} [\beta_i, g_j(\ln p, \ln x, z; \theta) | z] \\ &= \text{Cov} [\beta_i, \alpha_j | z] + \text{Cov} [\beta_i, \beta_j | z] \ln x - \text{Cov} [\beta_i, \sum_k \beta_j \alpha_k \ln p_k | z], \end{aligned} \tag{10}$$

where we have used that the covariance of a random variable with a constant is zero.

Now we are finally able to state the research question in precise terms. We have formulated a demand system with unobservable heterogeneity which is nested within the framework of Lewbel (2001). The integrability conditions in Lewbel (2001) then yield that the conditional mean demand for this demand system is integrable if the matrix of covariances M is symmetric and positive semidefinite. Our model is an AID system in which each household has its own intercept parameter and its own slope parameter.

The research question is thus whether a set of AID systems with all intercept and slope parameters being household-specific generates an integrable conditional mean demand. Taking a closer look at M , we see that without further assumptions, M is not necessarily symmetric. A sufficient condition for symmetry of M is that each of the three terms in (10) is symmetric. But only the second term, $\text{Cov}[\beta_i, \beta_j | z]$, is symmetric without further assumptions. Thus, G is not necessarily integrable.

3.4 Alternative assumptions on the distribution of preference heterogeneity

As mentioned earlier, making assumptions on the distribution of preference heterogeneity can have the unfortunate consequence that it makes M symmetric. In this section we give two examples of such assumptions:

Example 2 *Identical income effect coefficients in the AID system*

Let the budget share equations for household h be given by

$$g_i(\ln p, \ln x; \theta) = \alpha_{ih} + \beta_i \ln x - \beta_i \ln P_h(p) + \sum_j \gamma_{ij} \ln p_j, \quad i = 1, \dots, N,$$

for each $h = 1, \dots, H$. Then the partial derivative of the budget share function for commodity i with respect to log total expenditure is β_i for all households. Since M_{ij} is the covariance of the income effect for commodity i with the budget share of commodity j across households, and since the income effect is the same for all households, M_{ij} is the covariance between a constant and a random variable, so M_{ij} is zero for all i, j . This trivially implies that M is both symmetric and positive semidefinite.

This example shows that identical income effects for all households for all commodities in the AID system implies that the matrix of covariances is symmetric and positive semidefinite, i.e. identical income effects imply that G is integrable. This resembles the case of Gorman aggregation, where identical income effects (parallel individual Engel curves) implies that average demand (even unconditionally) is integrable.

Example 3 *Income effects independent of the other parameters of the demand system*

Suppose that β_i is conditionally independent of α_j , $j = 1, \dots, N$, conditional on $\ln x$, $\ln p$ and z for all i , and consider again M_{ij} . This independence assumption, together with all parameters being independent of $\ln x$ and $\ln p$, implies that the first term in M_{ij} is zero. i.e. $\text{Cov}[\beta_i, \alpha_j | z] = 0$. Furthermore, the third

term becomes

$$\begin{aligned}
Cov[\beta_i, \beta_j \ln P(p; \alpha) | z] &= \int \beta_i \beta_j \ln P(p; \alpha) dF - \left(\int \beta_i dF \right) \left(\int \beta_j \ln P(p; \alpha) \right) \\
&= \left(\int \beta_i \beta_j dF \right) \left(\int \ln P(p; \alpha) dF \right) \\
&\quad - \left(\int \beta_i dF \right) \left(\int \beta_j dF \right) \left(\int \ln P(p; \alpha) \right) \\
&= \left(\int \beta_i \beta_j dF - \left(\int \beta_i dF \right) \left(\int \beta_j dF \right) \right) \left(\int \ln P(p; \alpha) \right) \\
&= Cov[\beta_i, \beta_j | z] \left(\int \ln P(p; \alpha) \right),
\end{aligned}$$

i.e. M_{ij} reduces to

$$M_{ij} = Cov[\beta_i, \beta_j | z] \left(\ln x - \left(\int \ln P(p; \alpha) \right) \right),$$

i.e. M is symmetric. Furthermore, M is the product of the conditional variance-covariance matrix of $\beta = (\beta_1, \beta_2, \beta_3)$ conditional on z and the number $\ln x - \left(\int \ln P(p; \alpha) \right)$. The conditional variance-covariance matrix is positive definite which implies that M is positive semidefinite if $\ln x - \int \ln P(p; \alpha)$ is greater than or equal to zero. $P(p; \alpha)$ is a price index and thus it lies between 0 and 1, hence $\ln P(p; \alpha)$ is less than or equal to zero, hence the mean of $\ln P(p; \alpha)$ is also less than or equal to zero, which implies that $\ln x - \int \ln P(p; \alpha)$ is greater than zero for all values of total expenditure greater than 1. Since total expenditure is always much larger than 1, M is also positive semidefinite.

Note that these examples are specific to this particular model where g is an AID system. If choosing a different parametric form of g , one would have to re-examine which additional assumptions on the distribution of preference heterogeneity have unfortunate consequences for that choice of g ¹². We chose the AID system as our basic functional form, because the AID system is one of the most used parametric demand systems in the literature and because the income effects in that model are simply the β 's which simplifies estimations and interpretations.

The second example shows that if we were to choose a distributional form of the preference parameters, we need to allow for a general covariance-structure in that distribution (i.e. that we need to allow at least for the income effects to correlate with other parameters). Finding such a distribution would not be a big problem. However, a more complicated issue, which is independent of the choice of g , is to find a distribution from which the coefficients could be randomly drawn and yet always yield integrability at the individual level. It is not at all clear which distribution would ensure this. All these examples highlight the motivation for the chosen modelling strategy of not imposing any structure on the distribution of preference heterogeneity instead of imposing a random utility model across households.

¹²If for example choosing g is the QUAID system, identical income effects do not imply symmetry of M , see Lewbel(2001).

4 The Econometric Model

We model individual demand as an AID system with intercept and slope parameters being household-specific. The budget share equations for a given household h are thus given by

$$w_{iht} = \alpha_{ih} + \beta_{ih} [\ln x_{ht} - \ln P_{ht}(p)] + \sum_{j=1}^N \gamma_{ij} \ln p_{jt} + \sum_{k=1}^K \delta_{ik} z_{kht} + \varepsilon_{iht}, \quad (11)$$

$i = 1, \dots, N$, $t = 1, \dots, T$, with the price index given by

$$\ln P_{ht} = \sum_{k=1}^N \alpha_{hk} \ln p_{kt} + 0.5 \sum_{j=1}^N \sum_{k=1}^N \gamma_{kj} \ln p_{kt} \ln p_{jt},$$

for $t = 1, \dots, T$, and where z_{ht} is a K -dimensional vector of demographics for household h at time t and ε_{iht} is the idiosyncratic error term for household h , commodity i at time t . Note that α_i and β_i vary across households for all commodities, whereas the price coefficients γ_{ij} and the demographic coefficients δ_{ik} are restricted to being identical across households. We impose integrability at the individual level by imposing adding up, homogeneity and symmetry at the individual parameters. Adding up is satisfied by leaving out one commodity. We leave out the N 'th commodity and thus end up with a system of $N - 1$ equations. Homogeneity is imposed by using relative prices, relative to the left out commodity, i.e. we use the relative log prices $\ln \tilde{p}_{jt} = \ln p_{jt} - \ln p_{Nt}$, $j = 1, \dots, N - 1$. Symmetry is imposed by imposing $\gamma_{ij} = \gamma_{ji}$ for all $i, j = 1, \dots, N - 1$ in estimation. The error term structure of the demand system is given by $E[\varepsilon_{iht}] = 0$ for all i, t , $E[\varepsilon_{iht}\varepsilon_{jhs}] = \sigma_{ij}^2$ for all i, j , $t = s$ and $E[\varepsilon_{iht}\varepsilon_{jhs}] = 0$ for all i, j , $t \neq s$.

As a benchmark case we will be using an AID system with no heterogeneity, i.e.

$$w_{iht} = \alpha_i + \beta_i [\ln x_{ht} - \ln P_t(p)] + \sum_{j=1}^N \gamma_{ij} \ln p_{jt} + \sum_{k=1}^K \delta_{ik} z_{kht} + \varepsilon_{iht}, \quad (12)$$

$i = 1, \dots, N$, $t = 1, \dots, T$, with the price index given by

$$\ln P_{ht} = \sum_{k=1}^N \alpha_k \ln p_{kt} + 0.5 \sum_{j=1}^N \sum_{k=1}^N \gamma_{kj} \ln p_{kt} \ln p_{jt},$$

where integrability is now imposed on the homogenous parameters. When we compare the two models (11) and (12), we will refer to (11) as the heterogeneous model and to (12) as the homogeneous, or pooled, model.

5 Data

The data we use is a unique Spanish panel data sets on household expenditures, the Spanish Permanent Survey of Consumption (Encuesta Permanente de Consumo, hereafter the ECP), collected by the Spanish National Bureau of Statistics (Instituto Nacional de Estadística). It is a real panel of sizeable longitudinal length: the ECP covers the years 1978-83 with quarterly information on all recorded variables with households staying in the survey between 6 and 24 quarters. All households in the data set are headed by a married couple and may contain children or other adults cohabiting in the household. The data set

contains information on consumption expenditures for a wide range of commodity groups, price indices for these commodities as well as a variety of demographic variables such as labour market status, occupation, education level of the husband, the ages of the different household members and housing tenure. The version of the ECP we have available consists of 1641 households with more than 70 percent staying in the sample for at least 9 consecutive quarters. The total number of observations is 21.668. For complete lists of the variables recorded in the ECP, see Table A.1 and Table A.2 in Appendix A.

This Spanish household expenditure data set is exceptional in that it is a real panel with detailed information on a full range of commodities and prices, as well as on demographics. To the best of my knowledge, this is the longest real panels on household consumption expenditures covering a wide range of commodity groups, we have available. For comparison, the British Family Expenditure Survey (the FES) is a cross sectional data set and so does not provide the same possibility of taking account of unobservable heterogeneity. The American Consumer Expenditure Survey (the CEX) contains a rotating panel with information on several commodity groups, but it only has information on 4 quarters per household. Moreover, the consumption information in the panel part of the CEX consists of recall data only, and recent research shows that the best method for collecting accurate consumption expenditure data is a combination of diary and recall information, see Battistin (2003). In contrast, the Spanish data consists of a combination of diary and recall information with a grouping of commodities into which are recorded as diary information and which are recorded as recall information very close to the one recommended in Battistin (2003). The Panel Study of Income Dynamics (the PSID) is also a real panel and is of considerable length, but it only contains consumption expenditure information on food (food in and food out) and is therefore not well suited for analyzing demand choices. The British Household Panel Survey (the BHPS) is another real panel, but also this panel has insufficient information on consumption expenditures: The BHPS only records expenditures on household appliances and electronics and has one food question which asks how much the household approximately spends on food and groceries each week¹³. The European Community Household Panel Survey (the ECHP) has no records on consumption expenditures¹⁴. Finally, there exists a panel survey for Japan (The Japanese Panel Survey on Consumers, the JPSC) which contains consumption information. However, the JPSC asks all households about expenses only for the calendar month of September¹⁵, which makes it impossible to control for seasonal variation in demands.

In this paper, we use a subsample of the ECP. We select a sample consisting of the households that are observed for all 24 quarters in order to have as many observations per household as possible. In this sample, some households either experience changes to the husband's labor market status in form of unemployment spells or retirement, or the husband is retired in all 24 quarters¹⁶. This could be questionable because separability between the consumption of goods and labor supply is an underlying assumption of demand theory, and there seems to be some empirical evidence against it (e.g. Browning

¹³Browning, Crossley and Weber (2003) p. F563.

¹⁴The ECHP asks the household questions like "can you afford" various expenses, see Browning, Crossley and Weber (2003) p. F563, but this type of information is not useful for estimating demand systems.

¹⁵Browning, Crossley and Weber (2003) p. F563.

¹⁶The data contains no information about the wife's labor market status, but given the sample period - 1978-83 - it is likely that most of the wives are housewives. See also Christensen (2005).

and Meghir (1992)). But firstly, we need as large a sample with as much variation in total expenditures as we can get¹⁷. Secondly, the rejections of separability are much stronger for women than for men, and we only consider the labor supply of men. Thirdly, as can be seen from the table below, almost 70 percent of the husbands are employed throughout the survey period, and 22 percent are retired throughout or retires during the survey period. And since Christensen (2005) finds no rejections of separability between consumption and retirement in another representative Spanish panel data set (the ECPF) covering the years 1985-97, we are left with that only 10 percent of the sample experiences changes to their labor supply and hence 10 percent for which separability could be an issue. In the light of this, the sample selection does not seem too troublesome.

Labor market status	Number of households	Percent
Employed for all 24 quarters	168	67.50
Retired for all 24 quarters	31	12.42
Retires during the 24 quarters	25	10.04
Experiences spells of unemployment	25	10.04
Total	249	100.0

To get an idea of whether there could be selection problems stemming from selecting only the households that stay in the survey for the whole of the survey period, we compare the selected sample to the full data set. Summary statistics are presented in Table A.3 - A.5 in Appendix A. As can be seen, the characteristics of the households in our selected sample do not differ in any remarkable way from the full data set: The husbands' education levels are slightly lower, there is a slightly higher percentage of homeowners and the households spend on average less on every commodity group than in the full data set. None of the figures raise any concerns of sample selection problems, selected on observables.

To keep the integrability test as simple as possible, we construct four composite commodity groups from the 12 of the 14 commodities, leaving out Rent and Durables. The definition of the four commodity groups, together with their mean and standard deviations, are displayed in the table below.

Commodity group	Definition	Mean	Standard deviation
Food, alcohol & tobacco	Food at home + Foodout + Alcohol & tobacco	.60	.16
Clothing	Clothing	.10	.07
Utilities	Energy + Services + Medication + Transportation + Other	.24	.13
Leisure activities	Leisure + Holidays	.06	.07

The commodity group Food, Alcohol and Tobacco constitutes a larger percentage of the total budget

¹⁷Christensen (2007) carries out the integrability test on the sample in which all husbands are employed in a permanent job throughout the survey period. This sample, however, turns out to be too homogeneous for the test to have any power.

than is usually the case for food. It is abnormally large because the budget share for Food at home is abnormally large; an average budget share for Food of around 30% is what is usually observed in other expenditure surveys. The reason for this is that Spain went through a recession during the years where the ECP was collected¹⁸.

We construct the log prices for each of the composite commodity groups as the weighted average of the log prices for the goods in that particular commodity group, with the weights being the average budget share for the goods. When estimating the demand system, the left out good will be Utilities, and therefore the relative prices are the prices of Food, Alcohol&Tobacco, Clothing and Leisure&Holidays relative to the price of Utilities. The variation over the sample period in relative prices is shown in Figure 1 in Appendix A. As the graph shows there *is* independent variation between the three relative prices which means there is hope of some, or all, of the price coefficients being well identified.

We will estimate the matrix of covariances by a two-step procedure where we in the first step estimate the parameters of the demand system from the within variation in the data, and in the second step estimate the matrix of covariances from the between variation in the data. To get an idea of the amount of within and between variation in the data, we estimate a measure for each of the two sources of variations. A linear random effects estimation shows the fraction of variance in the error term which is due to the individual-specific part, which is a measure of the within-variation in the data. As can be seen from the table below, this fraction is between 25 and 50 percent of the total variation in the data, which seems high:

Commodity	Food, Alcohol and Tobacco	Clothing	Leisure and Holidays	Utilities
Within variation	.4331	.2597	.2869	.4423

By regressing each of the budget shares on the set of household dummies, we can use the R^2 s from these regressions as measures for how much of the variation in budget shares is explained by the between variation:

Commodity	Food, Alcohol and Tobacco	Clothing	Leisure and Holidays	Utilities
Between variation	.5389	.2526	.3324	.5104

Again, this fraction is between 25 and 50 percent of the total variation in budget shares, and we thus conclude that there seems to be sufficient variation in the data in both the within and the between dimension.

6 Estimation Procedure

The estimation of M is carried out in two steps. Recall that entry (i, j) in M is the covariance of the budget share for commodity i with the income effect for commodity j . In the first step, we estimate the demand system and get out the income effects and the predicted budget shares for each of the commodities. In the second step we calculate the covariance of the budget share for commodity i with the income effect for commodity j as the sample covariance across households of the estimated budget

¹⁸M. D. Collado (1998).

share for commodity i with the estimated income effect for commodity j .

6.1 Step 1: Estimation of the Demand System

We estimate the heterogeneous AID system in (11) on the sample of 249 households, allowing the intercept (the constant term) and the slope parameter (the income coefficient) to vary across households for each good. The left out good is Utilities, so the three composite goods are "Food, Alcohol & Tobacco", "Clothing" and "Leisure & Holidays". We control for seasonality with quarterly dummies, for the number of members in the household with the household size and for the husband's labor market status with a dummy for whether the husband is employed. Restricting the price coefficients, the coefficients on the quarterly dummies, the coefficients on household size and the coefficients on the employment dummy to being identical across households for each good introduces restrictions on the budget share equations across household for each good. Imposing symmetry at the individual level further introduces cross equation restrictions on the budget share equations across goods. Hence we need to estimate the AID system as one large system for all three goods and for all households simultaneously. We stack the budget shares for the three goods in one long vector on the left hand side and stack intercepts, log total expenditures, prices, quarterly dummies, household size and employment dummies according to the restrictions of identical coefficients and symmetry as discussed before, on the right hand side. Since we for each of the three goods have 249 households, each observed for 24 quarters, the system has $3 \cdot 249 \cdot 24 = 17\,928$ equations. Since intercepts and slopes vary across households, there are $3 \cdot 249 + 3 \cdot 249 = 1494$ household specific parameters to be estimated. Because of symmetry there are six price parameters to be estimated. And since there are three quarterly dummies and one household size variable and one employment dummy for each good, there are fifteen demographic parameters to be estimated. Thus, we have a total of 1515 parameters to be estimated from 17 928 observations.

The AID system is nonlinear in parameters, because the price index contains parameters. If there were no parameters to be estimated in the price index, the system would be linear in parameters. The usual way of estimating the AID is an iterative procedure which exploits this: In an initial step, the parametric price index is replaced by the Stone price index and the resulting linear model is estimated by least squares. The parameter estimates of this initial step are then used to calculate the parametric price index, which is then inserted in place of the Stone price index. A new set of parameters is estimated and the parametric price index is recalculated. This iterative procedure is continued until the parameter estimates converge (converge in the sense that they do not change from iteration to iteration). We estimate the variance-covariance matrix of the parameters using the least squares residuals from the final iteration.

As a benchmark case, we also estimate the corresponding homogeneous model (12) on the same sample, i.e. we estimate an AID system with all parameters being identical across households. This system has 27 parameters to be estimated from the same 17 928 observations. Estimation of the parameters is carried out by the iterative procedure described above, and standard errors are calculated using the least squares residuals from the final iteration, as above.

6.2 Step 2: Estimation of the matrix of covariances

The figures needed for estimating the matrix of covariances are the estimated income effects and the predicted budget shares for each of the three commodities for each of the 249 households. The estimated income effects are simply the estimates of the coefficients on log total expenditure, $\hat{\beta}_{ih}$, $i = 1, 2, 3$, $h = 1, \dots, 249$. For the predicted budget share for good j , household h , we use the the average over time of the predicted budget shares for good j for household h , i.e. $\hat{w}_{jh} = T^{-1} \sum_{t=1}^T \hat{w}_{iht}$, $i = 1, 2, 3$, $h = 1, \dots, 249$. The covariances of income effects with budget shares can now be calculated as the sample covariances across households of the estimated income effects with the estimated budget shares. Entry (i, j) of M is thus calculated as

$$\hat{M}_{ij} = H^{-1} \sum_{h=1}^H \hat{\beta}_{ih} \hat{w}_{jh} - \left(H^{-1} \sum_{h=1}^H \hat{\beta}_{ih} \right) \left(H^{-1} \sum_{h=1}^H \hat{w}_{jh} \right) \quad i, j = 1, 2, 3.$$

We bootstrap the standard errors of \hat{M} . We chose the bootstrap over the delta-method, since the latter often gives (more) biased results. We do a non-parametric bootstrap and sample from the data. We sample in clusters, that is, we sample households and for each sampled household, we sample all three budget shares. By sampling households we maintain the (true) covariance between time periods within a household. By sampling all three budget shares we ensure that the adding up restriction on budget shares holds in each bootstrap sample. We draw $B = 10.000$ bootstrap samples, estimate the household specific AID system and M for each bootstrap sample, and then calculate the variance-covariance matrix of the vectorized matrix of \hat{M} as the sample covariance of the bootstrap samples. To take an example, then the $(2, 3)$ 'th element in the estimate of the variance-covariance matrix of \hat{M} is calculated as

$$\text{Cov}(\hat{M}_{21}, \hat{M}_{31}) = (B - 1)^{-1} \sum_{b=1}^B \left(\hat{M}_{21}^b \hat{M}_{31}^b \right) - \left((B - 1)^{-1} \sum_{b=1}^B \hat{M}_{21}^b \right) \left((B - 1)^{-1} \sum_{b=1}^B \hat{M}_{31}^b \right).$$

7 Empirical Results

In this section we present the integrability test. We also present and discuss the household-specific demand system estimates resulting from the first stage estimation. Since the integrability test depends on these household-specific estimates - and since these are interesting in their own right - we start with them.

7.1 Results from the Estimation of the Heterogeneous Model

The output from the first stage estimation of the household specific AID system is 249 sets of intercept and slope estimates, one set for each household, together with six price coefficient estimates and fifteen demographic effects that are common across households. Firstly, we compare our estimation results with the estimates from the benchmark case, i.e. with the estimates from the homogenous AID system, to get an idea of whether our individual-specific estimates are reasonable. To this end, note that the demand for good i in the homogeneous AID system evaluated at unit prices (i.e. in the base year) is given by $\tilde{w}_i = \tilde{\alpha}_i + \tilde{\beta}_i \ln x$, $i = 1, 2, 3$, where $\tilde{\alpha}_i$ and $\tilde{\beta}_i$ are the homogenous intercept and slope parameters,

and that the conditional mean demand for good i in the household-specific AID system, evaluated at unit prices, is given by $G_i(\ln x) = \mu_{\alpha_i} + \mu_{\beta_i} \ln x$, where μ_{α_i} and μ_{β_i} are the means of the corresponding household-specific parameters α_{ih} and β_{ih} , $h = 1, \dots, H$. This means that the Engel curve for good i in the base year resulting from the estimation of the homogeneous AID system should be similar to the "Engel curve" for good i in the base year with intercept and slope coefficients given by the mean of the estimated household-specific intercepts and slopes. The estimates of the intercepts and slopes for the homogeneous model as well as the means of the household-specific estimates of intercepts and slopes and their standard errors are as follows (where (*) means the parameter is significant):

Coefficient	Homogeneous AID	Mean of household-specific AID
Intercept coefficient for good 1 (α_1)	1.9263 (.002936)*	1.2777 (0.038336)*
Intercept coefficient for good 2 (α_2)	.0856 (.002936)*	-.0565 (0.033828)
Intercept coefficient for good 3 (α_3)	-.4100 (.002938)*	-.3512 (0.037611)*
Slope coefficient for good 1 (β_1)	-.01104 (.002574)*	-.0515 (.003253)*
Slope coefficient for good 2 (β_2)	-.00015 (.002574)	.0132 (.002888)*
Slope coefficient for good 3 (β_3)	.0391 (.002575)*	.0314 (.003199)*

The estimates are similar for goods 1 and 3; for good 3 the Engel curves even coincide completely. For good 2 the point estimates suggest a difference, but as can be seen from the standard errors, then the slope coefficient in the homogeneous model is as the only parameter very imprecisely estimated and could thus just as well be positive. The remaining parameter estimates are directly comparable since they do not vary across households in the heterogeneous model. Below are the estimated price coefficients together with their standard errors (where (*) means the parameter is significant)¹⁹:

Coefficient	Homogeneous AID	Mean of household-specific AID
Price coefficient for good 1, price of good 1 (γ_{11})	-.1864 (.06649)*	-.0161 (.051443)
Price coefficient for good 1, price of good 2 (γ_{12})	.0107 (.02220)	-.0030 (.01724)
Price coefficient for good 1, price of good 3 (γ_{13})	-.0051 (.06805)	-.0633 (.05213)
Price coefficient for good 2, price of good 2 (γ_{22})	-.0300 (.01414)*	-.0394 (.01177)*
Price coefficient for good 2, price of good 3 (γ_{23})	.0147 (.04129)	-.0394 (.03150)
Price coefficient for good 3, price of good 3 (γ_{33})	.0481 (.15762)	.1209 (.1205)

As can be seen from the table, only two of the six price coefficients are significant in the homogeneous model and only one is significant in the heterogeneous model. This reflects a common finding, namely that it is often difficult to get significant price effects in demand system estimation.

Histograms of the parameter estimates are depicted in Appendix B. They show that all marginal distributions are unimodal, and that some are skewed. The latter justifies our choice of not modelling the distribution of preference heterogeneity as a normal distribution since that would have imposed symmetry of the marginal distributions. That the distributions are unimodal can be interpreted as there are no

¹⁹The estimates of the coefficients on the demographics are in Table ? in Appendix B

distinct "types" in the data; if for example the distribution of β_1 had had a bimodal distribution with most of its support in two points, this could be interpreted as if most of the households had either one β_1 or the other, i.e. that there are two "types" in the population. This situation would lend itself to a Heckman-Singer type approach (Heckman and Singer (1984)) in which the distribution of unobservables is modelled as a finite number of "types".

From the histograms it appears that there are quite large differences in the parameter estimates across households, suggesting a considerable amount of unobservable heterogeneity in the data. We test whether this is indeed the case by testing for identical income effects across households and identical intercepts across households, respectively, by an asymptotic F -type-test. Both test statistics are asymptotically chi-square with $3 \cdot (249 - 1) = 744$ degrees of freedom. For the case of the intercepts, we get a test statistic of $F = 2319$ and a p -value of $4.45 \cdot 10^{-161}$, so we strongly reject intercepts being identical across households. Similarly for the income effects, we get a test statistic of $F = 2378$ and a p -value of $7.59 \cdot 10^{-170}$, so we also strongly reject income effects being identical across households. The latter result means that if the matrix of covariances M turns out to be symmetric, it is *not* because the income effects are identical across households (which would imply M symmetric, see *Example 1*).

As a visual illustration of the heterogeneity in the data, we have picked a few households with the distribution of total expenditures roughly covering the support of total expenditures and depicted their Engel curves in Appendix B. Some households having upwards sloping Engel curves while other have downwards sloping Engel curves for the same good, in stark contrast to the homogeneous model which restricts the individual Engel curves for each good to exhibit identical slopes. This illustrates the substantial amount of heterogeneity already established by the tests above. Some Engel curves are steep, suggesting that the household is very sensitive to changes in income when it comes to consumption of that good, whereas others would hardly change their consumption of that good at all if their income changed. Is there a systematic relationship between having a great liking for a good and how sensitive to income changes one's demand for that good is? This question is answered by considering the correlation between intercepts and slopes. The table below shows the correlation between the estimates of α and β . The interesting part is the correlation between α_i 's and β_i 's for the same good²⁰: If a household has a high intercept for a certain good, this means that the household will always use a high proportion of their budget on that good, regardless of what their budget is. The correlation table shows that if this is the case, then the household has a low income effect for that good, since α_i and β_i are negatively correlated for all goods. This is what we would have expected: That a high, fixed taste for a good goes hand in hand with not wanting to change one's demand for that good in response to changes in one's income.

²⁰Obviously, the correlations between α 's and β 's with themselves are all negative, since the α 's sum to 1 and β 's sum to 0 by construction of the model.

	α_1	α_2	α_3	β_1	β_2	β_3
Intercept coefficient for good 1 (α_1)	1					
Intercept coefficient for good 2 (α_2)	-.38	1				
Intercept coefficient for good 3 (α_3)	-.31	-.04	1			
Slope coefficient for good 1 (β_1)	-.99	.38	.31	1		
Slope coefficient for good 2 (β_2)	.38	-.99	.04	-.38	1	
Slope coefficient for good 3 (β_3)	.32	.03	-.99	-.32	-.03	1

7.2 The Integrability Test

The estimate of the matrix of covariances M is²¹:

$$\hat{M} = \begin{bmatrix} 2.550 & .405 & -1.00 \\ (.459) & (.242) & (.209) \\ -.265 & .205 & -.059 \\ (.390) & (.228) & (.107) \\ -.537 & -.090 & .649 \\ (.450) & (.100) & (.196) \end{bmatrix}$$

As can be seen immediately, several of the elements of \hat{M} are not significantly different from zero. Below we report the χ^2 -statistics of the nine individual tests of

$$H_0 : M_{ij} = 0, \quad i, j = 1, \dots, 3,$$

which are each asymptotically chi square distributed with one degree of freedom. As can be seen, the three elements M_{11} , M_{13} and M_{33} are strongly significantly different from zero:

Element in \hat{M}	\hat{M}_{11}	\hat{M}_{21}	\hat{M}_{31}	\hat{M}_{12}	\hat{M}_{22}	\hat{M}_{32}	\hat{M}_{13}	\hat{M}_{23}	\hat{M}_{33}
χ^2 -statistic	30.94	.028	.23.07	.46	.807	.306	1.42	.792	11.00
p -value	0	.867	0	.498	.369	.580	.233	.373	.001

Hence, M is not the zero matrix and is thus not trivially symmetric. For completeness, we report a joint test of whether M is the zero matrix, since a joint test may have more power than the individual t type tests. The null hypothesis for this test can be written

$$H_0 : RM = 0,$$

where $R = I_9$ and M denotes the vectorized matrix of M . The Wald test statistic for H_0 is then given by $W_0 = [R\hat{M}]' [RV(\hat{M})R']^{-1} [R\hat{M}]$, which is asymptotically chi square with 9 degrees of freedom. We get $\hat{W}_0 = 104$, which results in a p -value of $2 \cdot 10^{-18}$, i.e. we strongly reject that M is the zero matrix. We therefore test whether M is symmetric by testing whether the off-diagonal elements are different from

²¹Row 1 is Food, Alcohol & Tobacco; row 2 is Clothing and row 3 is Leisure & Holidays. Bootstrapped standard errors are in parentheses. All elements are multiplied by 1000.

each other. The null hypothesis is:

$$H_0 : M_{ij} - M_{ji} = 0, \quad i \neq j,$$

and these test statistics must again be evaluated in the chi square distribution with one degree of freedom:

Element	$M_{12} - M_{21}$	$M_{13} - M_{31}$	$M_{23} - M_{32}$
χ^2 -statistic	.34	1.10	2.53
p -value	.559	29.4	11.2

As can be seen, we do not reject symmetry. Furthermore, the non-rejection is quite strong, with p -values between 11 and 56 percent. This means that this non-rejection occurs not because the estimate of M is too imprecise to reject that M is a matrix of zeroes, but because it is a real non-rejection. We therefore conclude that the conditional mean demand generated from a population of heterogeneous AID systems is in fact integrable and hence is consistent with utility maximisation. This means that there is some underlying utility function which generates the conditional mean demand. This in turn means that performing welfare analysis based on the estimated conditional mean is justified. However, it is not clear what this utility function looks like. From (9) in Section 3 we know what the conditional budget share function looks like and we know that it is not an AID system, hence the utility function underlying the conditional mean demand estimated from a population of heterogeneous AID systems is not the Almost Ideal utility. We therefore conclude that we can not calculate welfare effects directly based on the conditional mean demand function. What we can do, however, is to calculate welfare effects for each consumer, since each consumer has an Almost Ideal indirect utility function, and then consider the resulting distribution of welfare effects, conditional on income. This is done in the context of an example of a tax increase in the next section.

8 Comparing the Heterogeneous and Homogeneous Models

In this section we compare demand system estimates of price and income elasticities and welfare effects estimated from the heterogeneous model to those estimated from an AID model with no unobservable heterogeneity. This is interesting because the model with heterogeneity is the model that would usually be estimated from cross sectional data.

The value of having integrable demands is that it facilitates welfare analysis. Suppose for example that the government increases the tax on Alcohol & Tobacco; from the demand system estimates alone we can assess quantitatively how demands will change following the tax increase, but in order to assess how the welfare of the consumers change, we need a utility framework. Given that integrability is not rejected, we know that the conditional mean demand is generated by utility maximisation of *some* utility function. But as mentioned above, we do not know what this utility function looks like (and it is beyond the scope of this paper to find out). We therefore calculate the welfare effects at the individual level for *each* household, utilising that each household is maximising an Almost Ideal utility function. We consider a tax increase

that leads to the price of Food, Alcohol & Tobacco to increase by 20% and estimate the compensating variations for each household, using *that* household's particular utility function, conditional on income. We consider five percentiles in the income distribution, and hence the estimation in the heterogeneous model results in five conditional distributions each of 249 compensating variations. The homogeneous model assumes that all households have the same preferences, and hence this model produces only one compensating variation for each income percentile. We compare this one compensating variation to the mean of the 249 household-specific compensating variations for each income percentile. But first we compare estimates of price and income elasticities from the homogeneous and the heterogeneous model.

8.1 Positive Analysis: Price and Income Elasticities

Let e_i denote the income elasticity for good i , let e_{ij} denote the uncompensated price elasticity of good i with respect to the price of good j and let e_{ij}^c denote the compensated price elasticity. The price and income elasticities resulting from the homogeneous model are given by

$$\begin{aligned} e_i &= \frac{\tilde{\beta}_i}{w_i} + 1, \quad i = 1, \dots, N \\ e_{ii} &= \frac{1}{w_i} \left(\tilde{\gamma}_{ii} - \tilde{\beta}_i \tilde{\alpha}_i - \tilde{\beta}_i \sum_{k=1}^N \tilde{\gamma}_{ik} \ln p_k \right) - 1, \quad i = 1, \dots, N \\ e_{ij} &= \frac{1}{w_i} \left(\tilde{\gamma}_{ij} - \tilde{\beta}_i \tilde{\alpha}_j - \tilde{\beta}_i \sum_{k=1}^N \tilde{\gamma}_{jk} \ln p_k \right), \quad i, j = 1, \dots, N, \quad i \neq j \\ e_{ij}^c &= e_i w_j + e_{ij} \quad i = 1, \dots, N, \end{aligned}$$

where $\tilde{\alpha}$, $\tilde{\beta}$ and $\tilde{\gamma}$ are the parameters in the homogeneous model and w_i are the budget shares in the homogeneous model. We estimate the elasticities from the homogeneous model by inserting the parameter estimates from the homogeneous model of $\tilde{\alpha}$, $\tilde{\beta}$ and $\tilde{\gamma}$, as well as the predicted budget shares. We estimate an elasticity in each data point (i.e. for each household at each time period) and construct the final elasticity as a weighted average across households and time, the weights being the household's expenditure on the good in question relative to the total expenditure on that good.

Similarly, we derive the model price and income elasticities of the conditional mean demand resulting from the heterogeneous model. Let μ_{β_i} denote the mean in the marginal distribution of the slope coefficient for good i , let $E(\cdot)$ denote the mean with respect to the joint distribution of parameters and let G_i denote the predicted conditional mean budget share for good i , then

$$\begin{aligned} e_i &= \frac{\mu_{\beta_i}}{G_i} + 1, \quad i = 1, \dots, N \\ e_{ii} &= \frac{1}{G_i} \left(\gamma_{ii} - E(\beta_i \alpha_i) - \mu_{\beta_i} \sum_{k=1}^N \gamma_{ik} \ln p_k \right) - 1, \quad i = 1, \dots, N \\ e_{ij} &= \frac{1}{G_i} \left(\gamma_{ij} - E(\beta_i \alpha_j) - \mu_{\beta_i} \sum_{k=1}^N \gamma_{jk} \ln p_k \right), \quad i, j = 1, \dots, N, \quad i \neq j \\ e_{ij}^c &= e_i G_j + e_{ij} \quad i = 1, \dots, N. \end{aligned}$$

We estimate the elasticities from the heterogeneous model by inserting the sample mean of the esti-

mates of β_i , $H^{-1} \sum_{h=1}^H \hat{\beta}_{ih}$, for μ_{β_i} and the sample mean of the mean of the product of β_i with α_j , $H^{-1} \sum_{h=1}^H (\hat{\beta}_{ih} \hat{\alpha}_{jh})$, as well as the predicted budget shares from the conditional mean demand given in (9). As above, we estimate an elasticity for each household at each time period and construct the final elasticity as the weighted average as described above. The results are depicted in the tables below (bootstrapped standard errors in parantheses):

	<i>Income Elasticities</i>			
	Food, Alcohol&Tobacco	Clothing	Leisure&Holidays	Utilities
Homogeneous model	.8005 (.0735)	.9987 (.0970)	1.7858 (.1650)	1.2599 (.1158)
Heterogeneous model	.9073 (.1130)	1.1286 (.1085)	1.4169 (.1147)	1.0268 (.1205)

The income elasticities for the luxury goods Leisure&Holidays and Utilities are in magnitude far smaller in the heterogeneous model than in the model with no heterogeneity. This suggests that some of the variation in the budget shares of these goods in the homogeneous model was wrongly contributed to income, when in fact it was an income-independent taste difference, thus leading the income effects to being overestimated in the homogeneous model.

	<i>Compensated price elasticities: Homogeneous model</i>			
	Food, Alcohol&Tobacco	Clothing	Leisure&Holidays	Utilities
Food, Alcohol&Tobacco	-.5021 (.0509)	.1261 (.0432)	-.0436 (.0203)	.4399 (.0690)
Clothing	.6612 (.0202)	-1.1656 (.1105)	.1990 (.0124)	.3255 (.0243)
Leisure&Holidays	-.8130 (.2266)	.4323 (.0081)	.4066 (.0881)	-.2444 (.1818)
Utilities	.8676 (.0726)	.1308 (.0343)	-.0192 (.0632)	-.9488 (.1547)
	<i>Compensated price elasticities: Heterogeneous model</i>			
	Food, Alcohol&Tobacco	Clothing	Leisure&Holidays	Utilities
Food, Alcohol&Tobacco	-.1536 (.0411)	.0340 (.0817)	-.1121 (.0619)	.2408 (.0410)
Clothing	.1716 (.1290)	-.8945 (.4010)	-.2451 (.1100)	.9740 (.7501)
Leisure&Holidays	-1.051 (.1182)	-.4556 (.0862)	1.4560 (.3740)	-.2461 (.6702)
Utilities	.5213 (.0113)	.3968 (.0625)	-.0411 (.1294)	-.8663 (.2002)

	<i>Uncompensated price elasticities: Homogeneous model</i>			
	Food, Alcohol&Tobacco	Clothing	Leisure&Holidays	Utilities
Food, Alcohol&Tobacco	-.9508 (.0950)	.0382 (.0040)	-.0958 (.0171)	.2120 (.0358)
Clothing	.1015 (.0107)	-1.2754 (.1263)	.1339 (.0026)	.0413 (.0160)
Leisure&Holidays	-1.8140 (.2501)	.2360 (.0006)	.2902 (.1074)	-.7527 (.1915)
Utilities	.1615 (.0465)	-.0076 (.0042)	-.1012 (.0365)	-1.3074 (.2033)
	<i>Uncompensated price elasticities: Heterogeneous model</i>			
Food, Alcohol&Tobacco	-.6606 (.1067)	-.0602 (.0180)	-.1936 (.0038)	.0059 (.0221)
Clothing	-.4592 (.1935)	-1.0117 (.4854)	-.3464 (.1736)	.6818 (.6980)
Leisure&Holidays	-1.8434 (.1858)	.6028 (.1597)	1.3323 (.4391)	-.6131 (.6139)
Utilities	-.0528 (.0414)	.2901 (.0953)	-.1333 (.1583)	-1.1323 (.1731)

All own price elasticities, bot uncompensated and compensated, are negative with the exception of the own price elasticity for good 3. The reason e_{33} becomes positive is that the point estimate of γ_{33} is positive in both models, leading e_{33} to being positive as well. In addition to this, the predicted budget shares for good 3 are very small in both models (values ranging from very close to zero to .09, with a mean of .05) implying that the point estimate for e_{33} becomes large in magnitude. Looking closer at the parameter estimates of γ_{33} in both models, however, we see that the point estimates are highly insignificant in both models. Moreover, then the standard errors in both models on the estimate of γ_{33} are much larger than the standard errors on any of the other price coefficient estimates, suggesting that γ_{33} is insignificant also relative to all the other price coefficients. If we replace the estimate of γ_{33} in the heterogeneous with a negative number in the confidence interval of γ_{33} , which is [-0.1201,0.3619], then e_{33} is estimated to be negative. For example, if we use -.10 instead of the point estimate of γ_{33} , we get $e_{33} = -1.9$ (the compensated becomes -1.77), and if we insert -.001 we get $e_{33} = -.59$ (the compensated becomes -.49). Hence we conclude, that the positive point estimate of the own price elasticity for good is brought about by a highly insignificant positive point estimate of the price effect for good 3 and is thus insignificant in itself.

8.2 Welfare Effects: An Illustrative Example

Given that integrability of the conditional mean is not rejected we conclude the paper with an illustrative example that compares the estimated welfare effects from the model that allows no heterogeneity to the estimated welfare effects from the heterogeneous model following a tax increase which results in an increase of the price of Food, Alcohol & Tobacco of 20%. We calculate the compensating variations resulting from the heterogeneous model and from the benchmark case, homogeneous model, conditional on income percentiles. The homogeneous model provides one CV for each income percentile, since all consumers are assumed to have the same preferences. For the heterogeneous model, we calculate the compensating variation for *each* household at each percentile. The comparison thus consists of comparing

the conditional distribution of 249 compensating variations from the heterogeneous model with one single compensating variation provided by the homogeneous model for that income percentile.

The compensating variation (CV) is defined as the amount by which the consumer would have to be compensated after a price change to be as well off as before. Let p^0 denote the price before the price change and let p^1 denote the price after the price change. Let x^0 denote the (nominal) total expenditure level of the consumer and let ψ denote the indirect utility function, then the CV is given as

$$CV = \exp \{b(p^1) \ln \psi(x^0, p^0) + \ln a(p^1)\} - x^0,$$

where

$$\begin{aligned} \ln \psi(x, p) &= (\ln x - \ln a(p))/b(p) \\ \ln a(p) &= \sum_k \alpha_k \ln p_k + 0.5 \sum_k \sum_l \gamma_{kl} \ln p_k \ln p_l \\ b(p) &= \beta_0 \prod p_k^{\beta_k}. \end{aligned}$$

We consider the 10th, 25th, 50th, 75th and 90th percentiles of the total expenditure distribution. The results are below, all absolute numbers are in €. The first table provides the CV's resulting from the homogeneous model for each percentile together with the means of the individual specific CV's for each percentile. The difference between the homogeneous CV and the mean of the household-specific CV's is then taken relative to the total average budget in that particular percentile:

Percentile	CV in homogeneous model (€)	Means of CVs in heterogeneous model (€)	Difference between models (€)	Average total expenditure at percentile (€)	Difference relative to total (%)
10 th	44.2	46.4	2.2	339.4	.65 %
25 th	62.5	68.6	6.1	518.1	1.18 %
50 th	86.5	99.4	12.9	774.9	1.66 %
75 th	117.0	141.0	24.0	1132.1	2.12 %
90 th	152.5	192.8	40.3	1590.4	2.53 %

This table shows the compensation the homogeneous model prescribes is always lower than the mean of the heterogeneous CV's. The differences (undercompensations) between using the homogeneous model and the means from the heterogeneous model range from .65 to 2.53 percent of the total budget across the income distribution. This indicates that using the homogeneous model is not far off when interest lies in compensating the consumer with the average CV. If, however, focus is not on compensating the consumer with the average CV, the homogeneous model does not provide an answer. The conditional distributions of CV's, conditional on the different income percentiles considered, are depicted in the table below:

Income distribution	CV distribution						
	5 th	10 th	25 th	50 th	75 th	90 th	95 th
10 th	33	36	41	46	53	57	61
25 th	48	51	60	69	77	87	93
50 th	61	70	85	99	114	128	153
75 th	73	94	116	140	165	190	206
90 th	90	117	160	192	229	270	362

Take for example the consumer with median income who has a CV in the 95th percentile of the conditional distribution of the CV's. This consumer should be compensated with 153€ according to the heterogeneous model, as compared to the 86.5€ prescribed by the homogeneous model. This consumer would thus be undercompensated in the homogeneous model by 8.6 percent of his/her total budget, which is a quite a large undercompensation. On the other hand, a consumer with a CV in the 5th percentile of the CV distribution is overcompensated in the homogeneous model by 6.9 percent. In summary, it seems that consumers lose out on compensations ranging between .65 to 2.53 percent of their total budgets if policy recommendations were done following the homogeneous model, whereas the CV's calculated from the heterogeneous model show that some consumers would gain (and the government spending more money than needed) and others would lose if policy recommendations were done equally for all consumers using the calculations resulting from the homogeneous model. This example highlights the need for taking unobservable heterogeneity into account in empirical demand analysis.

9 Conclusions

In this paper we have carried out a panel data test of integrability of the conditional mean demand accounting for unobservable heterogeneity in preferences. A uniquely long panel data set on household expenditures allowed us to model individual demands as a set of heterogeneous Almost Ideal Demand systems by allowing all the intercept and slope parameters of the demand system to be household-specific. We do not reject integrability of the conditional mean demand: The tests for symmetry comes out with p -values between 11 and 56 percent in favour of symmetry, and furthermore the rejection of the matrix of covariances being the zero matrix showed that we have enough precision in our estimates to obtain a real non-rejection. Hence we conclude that a set of heterogeneous Almost Ideal Demand systems generates a conditional mean demand that is integrable. Given that integrability was not rejected, we compared the demand system estimates obtained from a model with no heterogeneity to our heterogeneous demand system estimates. This is interesting because the model with no heterogeneity is the model that has usually been estimated in the empirical demand literature. Our findings are that the homogeneous model severely overestimates income elasticities for luxury goods. This could be because the homogeneous model wrongly contributes all of the variation in budget shares to income, when in fact some of the variation is due to taste differences. An illustrative example of welfare effects following a tax increase show that compensating variations from the homogeneous model slightly undercompensates consumers as compared to the mean of welfare effects estimated from the heterogeneous model, but they are not very different. However, the substantial amount of heterogeneity in the welfare effects from the heterogeneous model

highlights the importance of taking unobservable heterogeneity into account.

This work could be extended by examining how sensitive the integrability test result is to the specification of individual demands. For example, individual demands could be modelled as an QUAID system, which would allow for curvature in individual Engel curves. Or one could attempt to estimate individual demands non-parametrically (for example by applying the non-parametric estimator for linear random coefficient models recently proposed in Hoderlein, Klemelä and Mammen (2007)); however, the 24 observations we have available per individual, which is a lot when it comes to parametric estimation, may prove not to be enough in practice to estimate individual demand non-parametrically.

References

- [1] Banks, J., R. Blundell and A. Lewbel (1997): Quadratic Engel Curves and Consumer Demand, *The Review of Economics and Statistics*, Nov, 4, 527-539.
- [2] Beckert, W. (2005): "Estimation of Heterogeneous Preferences, with an Application to Demand for Internet Services", *The Review of Economics and Statistics*, 87(3), 495-502.
- [3] Blundell, R., Chen, X. and Kristensen, D. (2007): "Nonparametric IV estimation of shape-invariant Engel curves", *Econometrica* forthcoming, CeMMAP Working Paper WP15/03.
- [4] Blundell, R., Duncan, A. and Pendakur, K. (1998): "Semiparametric Estimation and Consumer Demand", *Journal of Applied Econometrics*, Vol. 13, No. 5, 435-462.
- [5] Brown, B., and M. Walker (1989): The Random Utility Hypothesis and Inference in Demand Systems, *Econometrica*, Vol. 57, No. 4, 815-829.
- [6] Browning, M., Crossley, T. and Weber, G. (2003): "Asking Consumption Questions in General Purpose Surveys", *The Economic Journal*, 113, Nov, F540-F567.
- [7] Browning, M., and C. Meghir (1991): The Effects of Male and Female Labor Supply on Commodity Demands, *Econometrica*, Vol. 59, No. 4, 925-951.
- [8] Calvet, L. and Comon, E. (2003): "Behavioral Heterogeneity and the Income Effects", *The Review of Economics and Statistics*, Vol. 85(3), 653-669.
- [9] Christensen, M. (2005): Essays in Empirical Demand Analysis: Evidence from Panel Data, PhD Thesis, University of Copenhagen, Red Series No. 111, ISBN 87-91342-28-7.
- [10] Collado, M. D., (1998): Separability and Aggregate Shocks in the Life-Cycle Model of Consumption: Evidence from Spain, *Oxford Bulletin of Economics and Statistics*, Vol. 60, 227-247.
- [11] Deaton, A., and J. Muellbauer (1980a): Economics and Consumer Behavior, Cambridge University Press.
- [12] Deaton, A., and J. Muellbauer (1980b): "An Almost Ideal Demand System", *American Economic Review*, Vol. 70, No.3, 312-326.
- [13] Grandmont, J.-M., (1992): "Transformations of the Commodity Space, Behavioral Heterogeneity, and the Aggregation Problem", *Journal of Economic Theory*, Vol. 57, No. 1, 1-35.
- [14] Hardle, W. and Jerison, M. (1991): "Cross Section Engel Curves over Time", *Recherches Economiques de Louvain*, 57, 391-431.
- [15] Heckman J., and Singer, B. (1984): "A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data", *Econometrica*, Vol. 52, No. 2. pp. 271-320.

- [16] Hildenbrand, W., (1994): "Market Demand: Theory and Empirical Evidence", Princeton University Press.
- [17] Hoderlein, S., (2004): "Nonparametric Demand Systems, Instrumental Variables and a Heterogeneous Population", *Econometrica*, forthcoming.
- [18] Hoderlein, S., Klemelä, J. and Mammen, E. (2007): Reconsidering the Random Coefficient Model, mimeo.
- [19] Horowitz, J., (2000): The Bootstrap, Handbook of Econometrics, Vol. 5, J.J. Heckman and E.E. Leamer, eds., Elsevier Science B.V., 2001, Ch. 52, pp. 3159-3228.
- [20] Labeaga, J. M., and Puig (2003): "Demand system estimation with individual heterogeneity: an analysis using panel data on households", mimeo.
- [21] Lewbel, A. (1991): "The Rank of Demand Systems: Theory and Nonparametric Estimation", *Econometrica*, 59, 711-730.
- [22] Lewbel, A., (1999): "Consumer Demand Systems and Household Equivalence Scales", Handbook of Applied Econometrics, Blackwell Handbooks in Economics, H. Pesaran and M. Wickens (eds).
- [23] Lewbel A. (2001): Demand Systems with and without Errors, *American Economic Review*, Vol. 91 no. 3, 611-618.
- [24] Mas-Colell, A., (1985): "The Theory of General Economic Equilibrium: A Differentiable Approach", Econometric Society Monographs, Cambridge University Press.
- [25] Mas-Colell, A., M. D. Whinston and J. R. Green (1985): "Microeconomic Theory", Oxford University Press.
- [26] Muellbauer, J., (1975): Aggregation, Income Distribution and Consumer Demand, *Review of Economic Studies*, Vol. 62, 525-543.
- [27] Pesaran, H., and R. Smith (1995): "Estimating long-run relationships from dynamic heterogeneous panels", *Journal of Econometrics*, Vol. 68, 79-113.

A Data

Table A.1. - *Consumption Expenditures in the ECP*

Commodity group	Description
Food at home	Food at home
Alcohol & tobacco	Alcoholic drinks and tobacco
Clothing	Clothing
Rent	Renting (including the money paid to the owner and the water, but not electricity, heating etc.)
Energy	Electricity, heating and petrol
Services	Furniture and appliances repairing, products for cleaning, money paid to people for cleaning the house and other household services
Medication	Medical expenses
Transportation	Car repairing, public transportation and communications (phone, mail, etc)
Leisure	Books, cinemas and other entertainments
Education	Education
Foodout	Restaurants and cafeterias
Holiday	Holidays
Other	Hairdressing, non-durables for personal care (soap, cosmetics etc), pocket money given to children, other services
Durables	Durables (cars, furniture, tv, etc)

Table A.2. - *Demographic Variables in the ECP*

Variable name	Definition
hempl: Husband's employment status	1: employed 2: unemployed 3: military service having worked before 4: retired 5: living out of property rents 6: student (none) 7: housewife (none) 8: others
hgact: Husband's type of employment	0: missing 1: entrepreneurs or self-employed with employees 2: entrepreneurs or self-employed without employees 3: wage earners with a permanent job 4: wage earners with a temporary job 5: working in family business without salary 6: other
heduc: Husband's education	0: illiterate 1: primary school 2: secondary school, first level 3: secondary school, second level 4: secondary school, second level, professional studies 5: university degree (3 years) 6: university degree (5 years) and PhD's 7: less than 5 years of school
tenure: Housing tenure	1: renters 2: home owners 3: free accomodation 4: not documented, presumably missing
tenure2: Housing tenure of second house	1: renters 2: home owners 3: free accomodation 4: not documented, presumable missing 6: does not have a second house

Table A.3 - *Summary statistics in full data set and in sample*

Husband's employment status:	Full data set	Sample
Employed	77,97	79,47
Retired	18,66	17,00
Unemployed or out of labor force	3,37	3,53
Total	100	100
Husband's occupational status		
Wage earners with permanent job	63,56	61,67
Wage earners with temporary job	7,82	9,76
Self-employed	28,61	28,02
Nonpaid work	0,01	0,55
Total	100	100
Husband's education level:		
Illiterate	3,24	4,94
Less than 5 years of school	22,47	23,47
Primary school	60,05	62,00
Secondary school	10,45	6,93
University degree	3,79	2,66
Total	100	100
Housing tenure of main house:		
Renters	19,3	16,77
Home owners	76,18	78,97
Other (free accomodation or missing)	3,89	4,26
Total	100	100
Housing tenure of second house:		
Does not have a second house	90,18	86,93
Own second house	8,57	11,35
Rent or free accomodation or missing	1,25	1,72
Total	100	100

Table A.4 - *Sample means of expenditures and household composition in full data set and in sample* ²²

	Mean, full data set	Std.Dev, full data set	Mean, sample	Std.Dev, sample
Food at home	100.434	56.319	74.369	43.357
Alcohol & tobacco	8.640	9.713	6.585	7.403
Clothing	21.336	22.790	14.351	14.927
Rent	13.298	49.246	9.811	40.032
Energy	17.331	19.622	12.587	14.326
Services	4.711	7.447	3.041	4.125
Medication	4.552	10.666	3.132	7.465
Transportation	9.198	15.989	5.748	9.570
Leisure	7.635	12.820	5.696	10.505
Education	6.431	14.459	4.048	9.439
Foodout	2.202	13.591	1.388	9.860
Holidays	6.398	29.807	4.548	17.427
Other	17.702	29.807	12.767	20.829
Durables	32.006	91.541	22.006	68.002
Total Expenditure	182.365	134.273	180.076	132.589
Husband's age	51,24	11,89	51,78	11,15
Wife's age	48,61	11,95	49,15	11,26
Number of children	1,27	1,36	1,32	1,39
Number of adults	2,55	0,80	2,72	0,90
Total household size	3,97	1,54	4,04	1,60

²²These are pooled means and standard deviations, based on all respectively 21668 and 5976 observations.




Figure 1: The variation in relative prices during the sample period.

B Empirical Results

Table B.1 - Parameter estimates from a homogeneous AID and means of household-specific estimates

	N4H249 Homogeneous	N4H29 Means
δ_1^{dq1}	-.0137 (.00357)	-.0072 (.002745)
δ_1^{dq2}	-.0172 (.00354)	-.0158 (.00273)
δ_1^{dq3}	.0062 (.00363)	.0043 (.00279)
δ_2^{dq1}	-.0052 (.00354)	-.0006 (.00272)
δ_2^{dq2}	-.0140 (.00352)	-.0137 (.00270)
δ_2^{dq3}	.0046 (.00355)	.0033 (.00272)
δ_3^{dq1}	.0016 (.00356)	.0091 (.00276)
δ_3^{dq2}	-.0032 (.00351)	-.0316 (.00272)
δ_3^{dq3}	.0100 (.00357)	.0075 (.00277)
δ_1^{fam}	.0077 (.0008857)	-.0150 (.00298)
δ_2^{fam}	.002955 (.0008857)	-.00007 (.00298)
δ_3^{fam}	-.0017 (.0008858)	.0054 (.00298)
δ_1^{emp}	-.0572 (.0033339)	-.0180 (.00544)
δ_2^{emp}	.0113 (.0033339)	.0027 (.00544)
δ_3^{emp}	-.0011 (.0033334)	.0062 (.00544)




Figure 2: Intercepts: Food, Alcohol & Tobacco




Figure 3: Intercepts: Clothing.




Figure 4: Intercepts: Leisure & Holidays




Figure 5: Slope coefficients: Food, Alcohol & Tobacco




Figure 6: Slope coefficients: Clothing.




Figure 7: Slope coefficients: Leisure & Holidays




Figure 8: Individual Engel curves for Food, Alcohol & Tobacco.




Figure 9: Individual Engel curves for Clothing.




Figure 10: Individual Engel curves for Leisure & Holidays.