

Revealed Preferences in a Heterogeneous Population

Stefan Hoderlein* Jörg Stoye†

Brown University NYU

August 13, 2009

Abstract

This paper explores the empirical content of the weak axiom of revealed preference (WARP) for repeated cross-sectional data or for panel data where individuals experience preference shocks. Specifically, in a heterogeneous population, think of the fraction of consumers violating WARP as the parameter of interest. This parameter depends on the joint distribution of choices over different budget sets. Repeated cross-sections do not reveal this distribution but only its marginals. Thus, the parameter is not point identified but can be bounded.

We frame this as a copula problem and use copula techniques to analyze it. The bounds, as well as some nonparametric refinements of them, correspond to intuitive behavioral assumptions in the two goods case. With three or more goods, these intuitions break down, and plausible assumptions can have counterintuitive implications. Inference on the bounds is an application of partial identification through moment inequalities. We implement our analysis with the British Family Expenditure Survey (FES) data. Upper bounds are frequently positive but lower bounds not significantly so, hence FES data are consistent with WARP in a heterogeneous population.

Keywords: Revealed Preference, Weak Axiom, Heterogeneity, Partial Identification, Moment Inequalities.

*Stefan Hoderlein, Brown University, Department of Economics, Robinson Hall #302C, Providence, RI 02912, USA, email: stefan_hoderlein@yahoo.com.

†Jörg Stoye, Department of Economics, New York University, 19 W 4th St, New York, NY 10012, USA, email: j.stoye@nyu.edu. Excellent research assistance by Sonya Mihaleva is gratefully acknowledged. Parts of the paper were written while Stoye was visiting Mannheim respectively the Cowles Foundation, whose hospitality is gratefully acknowledged. We also thank Don Andrews, Dan McFadden, and a conference audience at Brown University for comments. Any and all errors are ours.

1 Introduction

Motivation. The weak axiom of revealed preference (WARP) is among the core elements of the theory of rational consumer behavior. In a heterogeneous population, agents differ in their specific demand behavior, yet economic theory predicts that they individually obey the weak axiom. This paper explores the empirical content of this prediction for repeated cross-sectional data, i.e. the ability of such data to reject, be consistent with, or even imply (up to sampling uncertainty) the weak axiom. We approach this question as an exercise in bounds or *partial identification*. In particular, we are interested in the fraction of consumers violating WARP. This parameter should be zero according to economic theory. Its empirical value depends on the joint distribution of choices over different budget sets. Repeated cross-sections do not reveal this joint distribution but do reveal the marginal distribution of demand on every single budget set. Bounds on the fraction of consumers who violate WARP are implied; indeed, they are closely related to the classic Fréchet-Hoeffding bounds. We develop these bounds, refine them using nonparametric assumptions, and apply them to the U.K. Family Expenditure Survey. One motivation for this exercise is to provide a complement to the nonparametric estimation of “revealed preference” bounds on behavior derived from similar data sets (e.g., Blundell, Browning, and Crawford (2003, 2008)). We provide some insight as to how much mileage can be gained from strict revealed preference assumptions alone, without additional aggregation assumptions and only invoking weak assumptions on the dependence structures.

To see the gist of the identification problem, suppose one is interested in the joint distribution of demand on just two different (intersecting) budget lines, but one only knows the marginal distribution of demand on each of these budget lines. Our motivating application is that one faces repeated cross-sections of the population of interest, a setting that corresponds to many practical applications and data sets. The scenario also applies, however, if one sees many choices from the two budgets by one and the same agent, yet this agent follows a random utility model, i.e., there is a random component in the individual’s utility function. Finally, random utility and cross-sectional variation could be compounded: Data could be panel data from a heterogeneous population but one might also be worried about randomly varying preferences of the individuals and therefore use a random utility model anyway. Whatever the interpretation, the problem is that one knows the marginal but not the joint distribution of choices; in other words, the aspect of the relevant distribution that is not identified is precisely the copula.

- figure 1 about here -.

The problem is illustrated in figure 1, which displays two intersecting budget lines and (as

shaded areas) the marginal distributions of consumers on those lines. WARP is violated by those consumers whose choices lie in the emphasized segment of each budget line. The proportion of these consumers in the population depends on the displayed marginal distributions but also on the copula linking them. This copula has an intuitive interpretation: Consumers can be thought of as being ordered with respect to their revealed preference for good 2 in any given period. The copula describes how much realignment with respect to this ordering occurs as consumers get relocated from one budget to the other one. Two copulas stand out as extreme: The “best behaved” population might be one where this ordering is precisely maintained; the “worst behaved” (or at least most heterogeneous) one might correspond to its complete reversal. For this simple example and if distributions are continuous, these two dependence structures do indeed generate upper and lower worst-case bounds on $\mathbb{P}(\text{WARP violated})$. Furthermore, the two copulas just described correspond to the Fréchet-Hoeffding upper and lower limit copulas. In the continuous case, the problem thus becomes an application of a classic finding.

Our contribution is to observe this connection to the copula literature and to exploit it in numerous ways. First, we develop the result for mixed discrete-continuous distributions, with the above case as corollary, and also show how the resulting bounds can (under assumptions) be integrated over budget sets to bound $\mathbb{P}(\text{WARP violated})$ for populations that face heterogeneous budgets. Second, we use the existing literature on copulas, but also some novel ideas, to refine bounds from above and below. In particular, we impose some nonparametric dependence structure between demand in different budgets, i.e. we nonparametrically constrain unobserved heterogeneity, leading to tighter bounds. In the two-good case, it turns out that some such assumptions are both intuitively meaningful (and perhaps reasonable) and qualitatively affect bounds in the way that one might have expected. Third, we generalize the analysis to three and more goods. This generalization has some unpleasant features: While Fréchet-Hoeffding bounds still apply, the according worst-case copulas do not correspond to plausible, or at least easily comprehensible, restrictions on heterogeneity in the population. What is more, mathematically natural generalizations of the aforementioned, partially identifying assumptions fail to have clear intuitions any more and may have unexpected effects on the bounds. As one particular example, many assumptions which seemingly force the population to be well-behaved can actually induce spurious violations of WARP, that is, they can refine the lower bound of $\mathbb{P}(\text{WARP violated})$ away from zero for data that were generated by a rational population. We illustrate this by considering the triangular structure often advocated in nonseparable systems of equations.

We finally bring the analysis to a practical application, estimating the bounds on data from the British Household Expenditure Survey. Inference on the bounds is an application of moment inequalities, a recently burgeoning literature in econometric theory that we apply and adapt. The

empirical result is that point estimators of bounds indicate occasional violations of WARP but these are far from statistically significant. The data are consistent with WARP, either because consumers exhibit minimal rationality or because WARP is just too weak (or, of course, both).

Related Literature. This paper touches on a number of distinct issues, including the integrability of stochastic demand functions, the theory of copulas, and inference on parameters that are partially identified by moment inequalities. Consequently, there are points of contact to numerous literatures, some of which might be called classic and some of which are currently developing.

This is primarily a paper about revealed preferences. The revealed preference approach to consumer demand was introduced and popularized by Afriat (1978), and Varian (1982, 1983). This work lays the economic foundations for our approach. However, all empirical applications at the time considered one (usually representative) consumer, and the very notion of unobserved heterogeneity does not arise. The closest predecessor to our identification analysis is the literature on integrability of stochastic demand; see, in particular, the classic analysis in McFadden and Richter (1990) and the synthesis and extension in McFadden (2005). Special cases of some of our results could alternatively (if clumsily) be derived from there. One main difference is that we explicitly attack the problem as one of partial identification and consider upper and lower bounds on the fraction of rational consumers rather than just asking whether the lower bound is zero.¹ Also, we do not restrict attention to continuous demand. Perhaps more importantly, by considering WARP rather than SARP we test a substantially more fundamental requirement of rationality and turn the problem into one that is precisely suited to the tools developed in the literature on copulas, as well as in the literature on moment inequalities. In particular, we provide an economic interpretation involving the behavior of individuals in a heterogeneous population, show how this relates to refinements proposed in the copula literature, and establish the sense in which these intuitions break down in the high dimensional case. As already mentioned, our work is also related to applications of revealed preference to consumer demand, in particular by Blundell, Browning, and Crawford (2003, 2008); see Cherchye, Crawford, de Rock, and Vermeulen (2009) for an overview. This literature tests revealed preference theory, and uses it to derive bounds on demand regression. However, their stochastic models of unobserved heterogeneity are limited. For instance, Blundell, Browning, and Crawford (2003) focus on revealed preference analysis using the mean regression, which comes close to imposing a representative agent assumption because additive deviations from the conditional mean cannot in general be generated by a structural

¹McFadden (2005) could be extended in this direction, however. As the formal identification problem would also occur in a random utility setting, our work is furthermore related to Marschak (1960), Manski (2007), and to Bandhopadhyay et al. (2002, 2004), more on which later.

model (due to well known aggregation problems of WARP; see, e.g., Mas Colell, Whinston, and Green (1995), p.110). Our contribution complements this line of research by adding unobserved heterogeneity. Since we are employing a nonseparable model in a consumer demand setup, our contribution is also nested within the wider econometric literature on nonparametric identification of economic hypotheses using nonseparable models; see Matzkin (2006) for a lucid overview.

Formally, the present identification problem is also related to one that received attention in the treatment effects literature. Consider learning about the distribution of a treatment effect, $\Delta \equiv Y_1 - Y_0$ (or variations thereof), when a randomized experiment identifies the marginal distributions of potential outcomes Y_0 and Y_1 . Clearly, a similar partial identification problem to ours emerges, namely, marginals are perfectly but copulas are not at all identified. The issue is commonly avoided by focusing attention on the expected value of the treatment, which does not depend on the copula. Researchers genuinely interested in the distribution of the treatment effect have, however, brought Fréchet-Hoeffding's as well as related bounds to this problem (Heckman, Smith, and Clemens (1997), Manski (1997), Fan and Park (2009)). While motivated by a very different question, this literature has some formal similarities to what we are doing. The technical difference is that we are interested in features of the joint distribution, notably $\mathbb{P}(\text{WARP violated})$, that do not correspond to interesting aspects of the distribution of Δ , thus the detail of our identification analysis is quite different. Also, both Heckman, Smith, and Clemens (1997) and Manski (1997) recognize that inference on the resulting bounds is nonstandard but do not focus on it; Fan and Park's (2009) results on inference do not apply here.

Finally, inference on our bounds is an application of moment inequalities, a currently very active literature. While we do not provide a conceptual innovation to this field, it is interesting to note that mechanical application of existing approaches, in particular of Andrews and Soares (2007), can be improved upon by exploiting the specific structure of our bounds. We expect that the same will hold true for many other applications of moment inequalities, and that this paper might accordingly be of interest as case study of such an application.

Structure of the Paper. The remainder of this paper is structured as follows. Section 2 is devoted to identification analysis: We describe and solve the identification problem, that is, we find bounds on the fraction of consumers that violate WARP under the assumption that all observable features of population distributions are known. We provide worst-case bounds as well as bounds that use partially identifying assumptions and conduct this analysis in two as well as more dimensions, with the latter analysis having a qualitatively different message. Section 3 develops the necessary tools for inference. Section 4 contains our empirical application. Section 5 concludes, and an appendix collects all proofs.

2 Identification Analysis

This section analyzes identification of the fraction of a population who violate WARP. Thus, we discover what we could learn about this fraction if the population distribution of observables were known. Estimation and inference will be considered later.

Consider, therefore, a population of agents who face an income process $(Y_t)_{t=1,\dots,T}$ and a consumption set \mathbb{R}_+^m , where $m \geq 2$ denotes the number of distinct goods. Individual demand is given by the time invariant function $Q(y_t, \mathbf{p}_t, \mathbf{z}, \mathbf{a}) : \mathbb{R} \times \mathbb{R}_+^m \times \mathcal{Z} \times \mathcal{A} \rightarrow \mathbb{R}_+^m$, $t = 1, \dots, T$, where $Q(y_t, \mathbf{p}_t, \mathbf{z}, \mathbf{a}) \in \{\mathbf{x} \in \mathbb{R}_+^m : \mathbf{p}_t \mathbf{x} \leq y_t\}$, $\mathbf{z} \in \mathcal{Z} \subseteq \mathbb{R}^l$ denotes observable covariates (assumed time invariant for simplicity) and $\mathbf{a} \in \mathcal{A}$ denotes time invariant unobservable covariates, including what one might want to think of as preferences. Note that Q is nonstochastic and constant across consumers; without loss of generality, heterogeneity is absorbed by \mathbf{A} . In the repeated cross section scenario that constitutes our leading application, one would think of $\mathbf{A} = \mathbf{a}_l$ as a consumer with preference ordering \mathbf{a}_l . However, l could also denote preferences in situation l , and as already mentioned, the analysis could equally be applied to a random utility model. In any case, the distribution of $(Y_t, Q(Y_t, \mathbf{p}_t, \mathbf{Z}, \mathbf{A}), \mathbf{Z})$ is identified for every t in the sample.² The sequence $(\mathbf{p}_t)_{t=1,\dots,T}$ is considered nonstochastic.

To obtain testable implications, we have to assume some structure across time periods.

Assumption 1. *For any time periods s and t , $\Delta Y_{st} \equiv Y_s - Y_t$ is independent of \mathbf{A} conditional on $(\mathbf{Z}, \mathbf{Y}_t)$.*

This type of assumption is standard in the related literature on nonseparable models (Matzkin (2006)). It states that preferences for the goods in question and income changes are independent conditional on current income and household characteristics. For an intuition, suppose there are two types of income shocks, say positive and negative, where the size depends on covariates $(\mathbf{Z}, \mathbf{Y}_t)$ (think of this conditioning as allocating an individual to a cell defined by values $z_0, y_{t,0}$). Suppose further that for a good k , there are two types of individuals, \mathbf{a}_k and $\mathbf{a}_{k'}$ say, where type \mathbf{a}_k idiosyncratically likes good k and type $\mathbf{a}_{k'}$ does not. Then assumption 1 states that, conditional on covariates having a certain value and for both positive and negative income shocks, there must be equal proportions of \mathbf{a}_k and $\mathbf{a}_{k'}$ in the population. This implies identification of the distribution of $(Y_s, Y_t, Q(Y_s, \mathbf{p}_s, \mathbf{Z}, \mathbf{A}), \mathbf{Z})$ as well as $(Y_s, Y_t, \mathbf{Q}(Y_t, \mathbf{p}_t, \mathbf{Z}, \mathbf{A}), \mathbf{Z})$; conditional on covariates, we can identify the joint distribution of consumption in the respective period, and income across periods, from the respective marginals. The assumption is *not* sufficient to identify the joint distribution of consumption across periods, i.e. of $(Y_s, Y_t, \mathbf{Q}(Y_s, \mathbf{p}_s, \mathbf{Z}, \mathbf{A}), \mathbf{Q}(Y_t, \mathbf{p}_t, \mathbf{Z}, \mathbf{A}), \mathbf{Z})$, from its marginals.

²We use the following conventions: Large letters denote random variables and small letters denote realizations (as well as nonrandom variables). Vectors are identified by bold typeface.

Fix any two time periods s and t and initially condition on a realization of (Y_s, Y_t, \mathbf{Z}) ; integration of the resulting bounds will be considered in a second step. Recall that the accordingly conditioned distributions of $\mathbf{Q}_t \equiv \mathbf{Q}(Y_t, \mathbf{p}_t, \mathbf{Z}, \mathbf{A})$ and $\mathbf{Q}_s \equiv \mathbf{Q}(Y_s, \mathbf{p}_s, \mathbf{Z}, \mathbf{A})$ are identified. A given consumer's choices violate WARP if one choice would have been strictly affordable given the other budget, that is, if $\mathbf{p}'_s \mathbf{Q}_t \leq y_s$ and $\mathbf{p}'_t \mathbf{Q}_s \leq y_t$, with at least one inequality being strict. The fraction of consumers who violate WARP (or, equivalently from an identification point of view, the population probability of violating it) is

$$\mathbb{P}_{yz}(\text{WARP violated}) = \mathbb{P}_{yz}((\mathbf{p}'_s \mathbf{Q}_t \leq y_s, \mathbf{p}'_t \mathbf{Q}_s < y_t) \vee (\mathbf{p}'_s \mathbf{Q}_t < y_s, \mathbf{p}'_t \mathbf{Q}_s \leq y_t)), \quad (2.1)$$

for all $s, t \in \{1, \dots, T\}$, where $\mathbb{P}_{yz}(\cdot) \equiv \mathbb{P}(\cdot | Y_s = y_s, Y_t = y_t, \mathbf{Z} = \mathbf{z})$. This probability is a feature of the joint distribution of $(\mathbf{Q}_t, \mathbf{Q}_s | Y_t, Y_s, \mathbf{Z})$ and hence, is not identified under our assumptions. We will initially develop bounds on it for the two-good case. This case turns out to be characterized by a tight relation between bounds and meaningful (if not necessarily reasonable) assumptions about the evolution of demand in the population. We then generalize the analysis to three and more goods, illustrating all concepts with an example in the three goods case, where a graphical intuition is still available. The multiple goods case qualitatively differs from the two goods one. Bounds are easily derived by generalizing previous concepts, but plausible conditions on individual behavior are harder to find. Conversely, natural generalizations of the behavioral interpretation of the two dimensional case will fail to provide reasonable bounds. In particular, we show that depending on the distribution of individuals, the frequently advocated triangular structure of nonseparable models may achieve upper as well as lower worst-case bounds and may spuriously indicate violations of WARP. The reason for this may be found in the difficulty of finding a natural ordering of goods.

2.1 The Two Good Case

2.1.1 The General Result

Fix again any two time periods s and t . The following identification result applies no matter whether $(\mathbf{Q}_s, \mathbf{Q}_t | Y_s, Y_t, \mathbf{Z})$ is distributed continuously, discretely, or as a mixture of the two. We begin by setting $m = 2$, thus the consumers' problem is characterized by two time periods and two budget lines \mathcal{B}_s and \mathcal{B}_t in \mathbb{R}_+^2 . Normalizing $\mathbf{p}_t = (1, p_t)$, and using the budget equation to write $\mathbf{Q}_t = (y_t - p_t Q_t, Q_t)$, these budget lines can be expressed as $\mathcal{B}_s \equiv \{q \geq 0 : p'_t q = y_t\}$ and $\mathcal{B}_t \equiv \{q \geq 0 : p'_s q = y_s\}$, and their intersection occurs at $Q_s = Q_t = (y_t - y_s)/(p_t - p_s)$. Assume w.l.o.g. that $y_s/p_s > y_t/p_t$, i.e. \mathcal{B}_s has the larger vertical intercept, then WARP is violated iff $Q_t \geq (y_t - y_s)/(p_t - p_s) > Q_s$ or $Q_t > (y_t - y_s)/(p_t - p_s) \geq Q_s$. (See again figure 1, where the

distribution of Q_s corresponds to the lighter shaded probability.) The probability of this event is constrained by the marginal distributions of Q_t and Q_s but also depends on how consumers are re-ordered along the budget lines between periods s and t . It can be bounded as follows.

Proposition 1. *Suppose the model of individual demand as outlined above holds. Let assumption A1 be true and assume that the conditional distributions of \mathbf{Q}_s and \mathbf{Q}_t are known. Finally, assume that all probabilities are well defined. Then,*

$$\begin{aligned} \max \left\{ \begin{array}{l} \mathbb{P}_{yz}(Q_s \leq (y_t - y_s)/(p_t - p_s)) - \mathbb{P}_{yz}(Q_t < (y_t - y_s)/(p_t - p_s)) \\ - \min \{ \mathbb{P}_{yz}(Q_s = (y_t - y_s)/(p_t - p_s)), \mathbb{P}_{yz}(Q_t = (y_t - y_s)/(p_t - p_s)) \}, 0 \end{array} \right\} \\ \leq \mathbb{P}_{yz}(\text{WARP violated}) \leq \\ \min \left\{ \begin{array}{l} \max \{ \mathbb{P}_{yz}(Q_s < (y_t - y_s)/(p_t - p_s)), \mathbb{P}_{yz}(Q_t > (y_t - y_s)/(p_t - p_s)) \}, \\ \mathbb{P}_{yz}(Q_s \leq (y_t - y_s)/(p_t - p_s)), \mathbb{P}_{yz}(Q_t \geq (y_t - y_s)/(p_t - p_s)) \end{array} \right\}. \end{aligned}$$

Moreover, these bounds are tight, i.e. best possible given the available information.

This result provides tight bounds for the parameter of interest that can be determined from the marginals, using only observable, and hence estimable quantities. It holds for any value of y_t, y_s, p_t, p_s and z , and it specifically provides bounds for the subpopulation defined by y, z . Other than regularity conditions needed to ensure that all probabilities exist and are well defined, it only requires assumption 1, which restricts the income process. Note that this assumption could be weakened provided one has access to instruments, which define controls such that the conditional independence holds. Hence we consider this a rather weak assumption, and thus proposition 1 a rather general result. Its proof can be found in the appendix.

2.1.2 Specialization to Continuous Demand

We now specialize proposition 1 to the case where $(\mathbf{Q}_s, \mathbf{Q}_t | Y_s, Y_t, \mathbf{Z})$ is distributed continuously. We will also work with this case, which leads to a rather simple and intuitive result, later on. Whether the continuity assumption is realistic depends on one's perspective. If one thinks of all British consumers as the population of interest, then the true population distribution is of course discrete, albeit so finely grained that the simplification gained from assuming continuity may be worth the price. If one thinks even of all consumers as a very large sample from the data generating process of true interest, then continuity is a substantively appropriate assumption. We also duly note that the difference between our above result and this subsection's corollary is driven by the necessity to take care of mass points that sit at very specific locations; in that sense, it can be thought of as accommodating nongeneric complications. An important exception to these

considerations occurs if the demand function is expected to have genuine mass points, for example at corner solutions.

In the continuous case, we can simplify (2.1) as follows:

$$\mathbb{P}_{yz}(\text{WARP violated}) = \mathbb{P}_{yz}(p'_s Q_t < y_s, p'_t Q_s < y_t). \quad (2.2)$$

This event has a simple geometric interpretation. Denote the boundary of the time t budget constraint by

$$\mathcal{B}_t \equiv \{q \geq 0 : p'_t q = y_t\},$$

then the set of time t choices that are affordable given the time s budget is

$$\underline{\mathcal{B}}_{t,s} \equiv \{q \geq 0 : p'_t q = y_t, p'_s q \leq y_s\},$$

the intersection of \mathcal{B}_t with the half-space below (the hyperplane containing) \mathcal{B}_s . Hence,

$$\mathbb{P}_{yz}(\text{WARP violated}) = \mathbb{P}_{yz}((Q_t, Q_s) \in \underline{\mathcal{B}}_{s,t} \times \underline{\mathcal{B}}_{t,s}),$$

the probability that both Q_t and Q_s are contained in the respective half-space; consider again figure 1. Given that the marginal probabilities ($\mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{s,t}), \mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{t,s})$) are known, the problem of bounding $\mathbb{P}_{yz}((Q_t, Q_s) \in \underline{\mathcal{B}}_{s,t} \times \underline{\mathcal{B}}_{t,s})$ is easily recognized as the original Fréchet-Hoeffding bounding problem, and it can indeed be verified that our result simplifies to this classic result.³

Corollary 2. (Fréchet-Hoeffding Bounds) *Let the conditions of Proposition 1 be satisfied and assume in addition that the conditional distributions of Q_s and Q_t be continuous. Then*

$$\begin{aligned} \max\{\mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}) + \mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s}) - 1, 0\} &\leq \mathbb{P}_{yz}(\text{WARP violated}) \\ &\leq \min\{\mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}), \mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s})\}. \end{aligned}$$

These bounds are again tight.

In the two-dimensional case, the crucial probabilities can be written in terms of $(Q_s, Q_t, p_s, p_t, y_s, y_t)$. Recalling that s is normalized to correspond to the budget line with higher intercept, we have

$$\begin{aligned} Q_s \in \underline{\mathcal{B}}_{s,t} &\iff Q_s < (y_t - y_s)/(p_t - p_s) \\ Q_t \in \underline{\mathcal{B}}_{t,s} &\iff Q_t > (y_t - y_s)/(p_t - p_s). \end{aligned}$$

³These bounds can furthermore be derived as solutions to linear programming problems along the lines of McFadden (2005, theorem 3.1). We thank Dan McFadden for pointing this out.

Notice in particular that the lower half-space $\underline{\mathcal{B}}_{t,s}$ is an *upper* contour set of Q_t . The Fréchet-Hoeffding Bounds then become

$$\begin{aligned} & \max \{ \mathbb{P}_{yz}(Q_s < (y_t - y_s)/(p_t - p_s)) - \mathbb{P}_{yz}(Q_t < (y_t - y_s)/(p_t - p_s)), 0 \} \\ & \leq \mathbb{P}_{yz}(\text{WARP violated}) \leq \\ & \min \{ \mathbb{P}_{yz}(Q_s < (y_t - y_s)/(p_t - p_s)), 1 - \mathbb{P}_{yz}(Q_t < (y_t - y_s)/(p_t - p_s)) \}, \end{aligned}$$

which is the expression we will use for estimation, inference and in the application.

We will restrict attention to continuous distributions henceforth. Before turning to extensions of our result, we elaborate a further on an informal corollary as well as a generalization.

(i) Empirical Content of WARP A corollary of our result is to identify the empirical content of WARP. Specifically, choice probabilities are consistent with WARP iff the lower bound on $\mathbb{P}_{yz}(\text{WARP violated})$ equals zero. This is the case if

$$\begin{aligned} & \mathbb{P}_{yz}(Q_s \leq (y_t - y_s)/(p_t - p_s)) - \mathbb{P}_{yz}(Q_t < (y_t - y_s)/(p_t - p_s)) \\ & - \min \{ \mathbb{P}_{yz}(Q_s = (y_t - y_s)/(p_t - p_s)), \mathbb{P}_{yz}(Q_t = (y_t - y_s)/(p_t - p_s)) \} \leq 0, \end{aligned}$$

which in the continuous case simplifies to

$$\begin{aligned} & \mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}) + \mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s}) \leq 1 \\ & \iff \mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}) \leq \mathbb{P}_{yz}(Q_t \notin \underline{\mathcal{B}}_{t,s}) \\ & \iff \mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s}) \leq \mathbb{P}_{yz}(Q_s \notin \underline{\mathcal{B}}_{s,t}). \end{aligned}$$

This corollary fully applies to higher dimensions (with notational adaptations indicated later). Intuitively, it means that if a budget plane is rotated, the probability mass on the part of the plane that is rotated toward the origin must weakly shrink, whereas the mass on the part that is rotated away from the origin must weakly expand. There is no empirical content if the “before” and “after” budgets fail to intersect.⁴

⁴Apart from an alternative proof via McFadden (2005), this finding for the continuous case could be deduced by combining previous results of Bandyopadhyay et al. (2002, 2004). The inequalities correspond to what these authors call “stochastic substitutability,” and in Bandyopadhyay et al. (2004), they show that under our auxiliary assumptions, stochastic substitutability is equivalent to their “Weak Axiom of Stochastic Revealed Preference” (WASRP). Bandyopadhyay et al. (2002) establish that WARP for individual demand implies WASRP for aggregate demand. These results jointly imply the inequalities.

Of course, our results are significantly stronger than the implication from WARP toward stochastic substitutability. While Bandyopadhyay et al. (2002) do not use the auxiliary assumptions, one might wonder whether the present analysis can be adapted to analogously strengthen their original result. We show in ongoing work that this is indeed so.

These restrictions certainly do not appear very strong. Furthermore, they are implied not only by WARP but also by a host of other restrictions on individual behavior. For example, they are easily derived from the assumption that consumers choose independently (across consumers and time periods) from uniform distributions over budget planes.⁵ Nonparametric, assumption-free tests of WARP from cross-sectional data will have accordingly limited power, but this is simply due to the limited empirical content that WARP has on its own.

(ii) Unconditional Bounds We conclude this section by providing expressions for unconditional bounds. Suppose one also knows (or can estimate) the joint distribution of $(Y_1, \dots, Y_T, \mathbf{Z})$; this is a realistic assumption as panel income data sets exist for many countries. Under assumption 1, one can then generate unconditional bounds on $\mathbb{P}(\text{WARP is violated})$ by integrating the preceding bounds over (Y_s, Y_t, \mathbf{Z}) . For the worst-case bounds this means the following.⁶

Lemma 2.1. *Let the conditions of Proposition 1 hold and suppose that the distribution of (Y_s, Y_t, \mathbf{Z}) is known. Then*

$$\begin{aligned} \int \max\{\mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}) + \mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s}) - 1, 0\} F(d(y_s, y_t, \mathbf{z})) &\leq \mathbb{P}(\text{WARP is violated}) \\ &\leq \int \min\{\mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}), \mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s})\} F(d(y_s, y_t, \mathbf{z})). \end{aligned}$$

These bounds are tight.

2.1.3 Nonparametric Refinements

One upshot of the preceding section is that the identification problem is really about copulas, more specifically, about the copula connecting Q_s and Q_t . Recognizing this allows one to refine the above bounds by importing results about copulas. We now present some such assumptions and their exact implications.

The lower and upper bounds on $\mathbb{P}_{yz}(\text{WARP violated})$ correspond to measures of dependence between Q_s and Q_t that are extremal in an obvious sense; namely, they impose perfectly positively respectively negatively dependence. Many nonparametric measures of dependence interpolate between these extremes. Restrictions on any of them may induce narrower bounds.

One nonparametric dependence concept that has gained popularity in the copula literature is quadrant dependence:

⁵This observation resembles a classic discussion by Becker (1962).

⁶To keep expressions simple, we here abuse notation: $\underline{\mathcal{B}}_{s,t}$ and $\underline{\mathcal{B}}_{t,s}$ depend on (Y_s, Y_t) and therefore vary over the integrals.

Definition 3. *The copula linking Q_s and Q_t exhibits positive [negative] quadrant dependence if $\mathbb{P}_{yz}((Q_s, Q_t) \geq (q_s, q_t)) \geq [\leq] \mathbb{P}_{yz}(Q_s \geq q_s) \mathbb{P}_{yz}(Q_t \geq q_t)$ for all scalars q_s, q_t .*

Positive quadrant dependence means that large (and small) values of the individual variables coincide more often than would be expected under independence. In our application, this is to say that consumers who reveal strong taste for good 2 at time s tend to do the same at time t . Negative quadrant dependence is the intuitive opposite.

Imposing quadrant dependence leads to the following refinement.

Lemma 2.2. *(i) Let the conditions of Proposition 1 hold and assume positive quadrant dependence, then*

$$\max\{\mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}) + \mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s}) - 1, 0\} \leq \mathbb{P}_{yz}(\text{WARP is violated}) \leq \mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}) \mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s}).$$

(ii) Let the conditions of Proposition 1 hold and assume negative quadrant dependence, then

$$\mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}) \mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s}) \leq \mathbb{P}_{yz}(\text{WARP is violated}) \leq \min\{\mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}), \mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s})\}.$$

These bounds are tight.

Positive respectively negative quadrant dependence therefore neatly separate the worst-case bounds into two regions, one that is associated with positive and one that is associated with negative dependence. The boundary between the regions corresponds to independence. Substantively, it is certainly positive rather than negative quadrant dependence that we mean to suggest as interesting restriction on behavior across choice situations.

Numerous nonparametric measures of dependence can be used to strengthen positive quadrant dependence. In particular, one could impose that the copula exhibit tail monotonicity, stochastic monotonicity, corner set monotonicity, or likelihood ratio dependence (known as affiliation in the auctions literature). See, for example, Nelsen (2006) for definitions of all of these, which are listed roughly in order of increasing stringency. Imposing any of them would lead to the same bounds identified above: All of them imply quadrant dependence, so the bounds cannot be wider; but all of them also allow for independence as boundary case as well as for the relevant one of the original worst-case bounds, so the bounds do not become tighter. Within this family, quadrant dependence therefore stands out as the weakest restriction that generates the above refinement.⁷

⁷Comparisons of nonparametric concepts of positive dependence in de Castro (2007) and Yanamigoto (1972) support the same conclusion, i.e. quadrant dependence is weakest among large classes of such concepts, none of which would lead to tighter bounds. Two concepts that are insufficient to generate the above bounds are positive correlation and a positive value of Kendall's τ .

We therefore now proceed to a different class of restrictions. Positive dependence can also be modelled by presuming that consumers do not “move too much” in terms of quantiles. Letting F_t denote the c.d.f. of Q_t , this idea could be formalized as follows:

Definition 4. *The copula linking Q_s and Q_t exhibits κ -limited quantile changes if $|F_s(Q_s) - F_t(Q_t)| \leq \kappa$ a.s.*

The assumption has not, to our knowledge, been previously entertained for copulas but is close in spirit to “sensitivity assumptions” that are occasionally used in econometrics and statistics; compare, for example, the parameter λ in Horowitz and Manski (1995) or Γ in Rosenbaum (2002). In particular, $\kappa \in [0, 1]$ is a user-defined parameter that determines how much deviation from perfectly positive dependence is allowed. By changing it, one can scale the assumption from implying perfect dependence ($\kappa = 0$) to being completely vacuous ($\kappa = 1$). Whether bounds are affected depends on how tightly κ is chosen. The precise result is as follows.

Lemma 2.3. *Let the conditions of Proposition 1 hold and let the copula linking the conditional distributions of Q_s and Q_t exhibit κ -limited quantile changes, then*

$$\begin{aligned} & \max\{\mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}) + \mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s}) - 1, 0\} \leq \mathbb{P}_{yz}(\text{WARP is violated}) \\ & \leq \max\{\min\{\mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}), \mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s}), \kappa + \mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}) + \mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s}) - 1\}, 0\}. \end{aligned}$$

These bounds are again tight. In particular, the assumption leads to a refinement of bounds iff $\kappa < 1 - \max\{\mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}), \mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s})\}$, and it implies that WARP is not violated at all iff $\kappa \leq 1 - \mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}) - \mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s})$.

An observation of special interest is that the upper bound on $\mathbb{P}_{yz}(\text{WARP is violated})$ is zero for a sufficiently restrictive choice of κ , namely when $\kappa \leq 1 - \mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}) - \mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s})$. This choice is feasible whenever this inequality’s r.h.s. is positive, i.e. iff the lower bound is zero. By the same token, whenever the lower bound is zero, the threshold value $\kappa^* \equiv 1 - \mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}) - \mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s})$ identifies just how much quantile movement is needed to generate a violation of WARP. The higher this number, the more contrived are the quantile movements that are needed to contradict WARP. One could, therefore, think of κ^* as an indicator of just how suggestive the data are in favor of WARP.

The approach can be extended by importing other suggestions from the literature on copulas. For example, one could bound dependence between choice on \mathcal{B}_s and choice on \mathcal{B}_t in terms of the medial correlation coefficient (Blomquist’s β), the rank correlation coefficient (Spearman’s ρ), or Kendall’s τ . The resulting bounds on joint c.d.f.’s, and hence on $\mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}, Q_t \in \underline{\mathcal{B}}_{t,s})$,

then follow from known results (Nelsen et al. (2001), Nelsen and Úbeda-Flores (2004)). We omit elaborations because these bounds involve much more algebra than the preceding ones, but notice that numerical evaluation would in all cases be easy, and also that such restrictions can be used to tighten bounds from below as well as above if one is willing to make according assumptions.

2.2 The Multiple Goods Case

We now analyze the multiple goods case, emphasizing differences to the two goods one. Integration of bounds, for one thing, is completely unchanged, and we therefore condition on (y_s, y_t, \mathbf{z}) throughout this section. We also continue to restrict attention to continuous distributions. As a result, the Fréchet-Hoeffding bounds from corollary 2 still apply. Adapting the notation to higher dimensions, we have

$$\begin{aligned} \max\{\mathbb{P}_{yz}(\mathbf{Q}_s \in \underline{\mathcal{B}}_{s,t}) + \mathbb{P}_{yz}(\mathbf{Q}_t \in \underline{\mathcal{B}}_{t,s}) - 1, 0\} &\leq \mathbb{P}_{yz}(\text{WARP violated}) \\ &\leq \min\{\mathbb{P}_{yz}(\mathbf{Q}_s \in \underline{\mathcal{B}}_{s,t}), \mathbb{P}_{yz}(\mathbf{Q}_t \in \underline{\mathcal{B}}_{t,s})\}, \end{aligned}$$

where

$$\mathcal{B}_t \equiv \{\mathbf{q} : \mathbf{p}'_t \mathbf{q} = y_t, \mathbf{q} \geq \mathbf{0}\}$$

and

$$\underline{\mathcal{B}}_{t,s} \equiv \{\mathbf{q} : \mathbf{p}'_t \mathbf{q} = y_t, \mathbf{p}'_s \mathbf{q} \leq y_s, \mathbf{q} \geq \mathbf{0}\}.$$

The analysis could be generalized to mixed continuous-discrete distributions, but the necessary bookkeeping regarding point masses becomes very tedious. The worst-case bounds are illustrated in the following example, which we will revisit later.

Example 1. *Let there be three goods, let $\mathbf{p}_1 = (10, 6, 5)$, $\mathbf{p}_2 = (5, 10, 6)$, and $y_1 = y_2 = 30$. The according budget triangles are most easily described by their corners: \mathcal{B}_1 is spanned by $((3, 0, 0), (0, 5, 0), (0, 0, 6))$, \mathcal{B}_2 is spanned by $((6, 0, 0), (0, 3, 0), (0, 0, 5))$. Assume that \mathbf{Q}_1 is supported on $(\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3) = ((5/7, 0, 32/7), (1, 0, 28/7), (9/7, 20/7, 0))$, that \mathbf{Q}_2 is supported on $(\mathbf{b}_1, \mathbf{b}_2) = ((0, 0, 5), (6, 0, 0))$, and that the joint distribution of $(\mathbf{Q}_1, \mathbf{Q}_2)$ is characterized by the following population-level contingency table, where the bold row and column indicate marginal distributions.*

	\mathbf{b}_1	\mathbf{b}_2	
\mathbf{a}_1	1/4	0	1/4
\mathbf{a}_2	0	1/4	1/4
\mathbf{a}_3	1/4	1/4	1/2
	1/2	1/2	

It is easily calculated that of the support of \mathbf{Q}_1 , only \mathbf{a}_2 lies on $\underline{\mathcal{B}}_{1,2}$. On the support of \mathbf{Q}_2 , it is \mathbf{b}_1 that lies on $\underline{\mathcal{B}}_{2,1}$. Thus, there is no violation of WARP. Using only the marginals, one would find $\mathbb{P}_{yz}(\mathbf{Q}_1 \in \underline{\mathcal{B}}_{1,2}) = 0.25$, $\mathbb{P}_{yz}(\mathbf{Q}_2 \in \underline{\mathcal{B}}_{2,1}) = 0.5$, and

$$0 \leq \mathbb{P}_{yz}(\text{WARP violated}) \leq 0.25.$$

Tightness of bounds follows from known results, but it is instructive to inspect the joint distributions that achieve the bounds. To do so, order elements of \mathcal{B}_s increasingly according to $\mathbf{p}'_t \mathbf{q}$, i.e. according to how expensive they would be given time t prices. Similarly, order elements of \mathcal{B}_t according to $\mathbf{p}'_s \mathbf{q}$ but in *decreasing* order. These orderings have in common that their level sets are parallel to $\mathcal{B}_s \cap \mathcal{B}_t$; furthermore, they identify $\underline{\mathcal{B}}_{s,t}$ as a lower and $\underline{\mathcal{B}}_{t,s}$ as an upper contour set.⁸ Then the lower Fréchet-Hoeffding bound is achieved by assuming that all consumers maintain their quantile position with respect to these orderings, and the upper bounds are achieved by inversion of the orderings.

In two dimensions, these orderings are the ones we identified before, and they have clear intuitions. In higher dimensions, the case is more complicated. Quite literally, the consumers are ordered according to how much one would have to pay them (or tax them) so that they could just afford their previous consumption bundle. This ordering does not have much economic significance.⁹ In particular, it is hard to see how an ordering of subjects on \mathcal{B}_s that depends on time t prices (and vice versa) would arise from natural restrictions on behavior. Therefore, while the mechanism generating the bounds is instructive, we do not claim that it corresponds to compelling behavioral assumptions. This raises the question whether interesting results can be generated either without ordering \mathcal{B}_s and \mathcal{B}_t beforehand or by ordering them in a more natural way.

We now explore the idea of weakly maintaining, with varying degrees of stringency, the ordering of consumers with respect to the orthants, i.e. with respect to “taste for good k .” Mathematically, these orderings may appear arbitrary because they privilege the standard basis over other possible bases of Euclidean space. Substantively, however, there may be reason to do just that: Co-

⁸Tie-breaking rules of the orderings matter only if distributions have mass points, which were assumed away.

⁹To see a vestige of economic meaning, assume for the moment that preferences exist and consider the problem of bounding the compensating variation for a consumer who moves from \mathcal{B}_s to \mathcal{B}_t . With the few assumptions imposed here, a lower bound on this CV will always be zero because the consumer could be indifferent across all bundles. The upper bound is given by $(p'_t q_s - y_t)$, the payment needed so that the consumer can afford her previous bundle. The ordering over \mathcal{B}_s therefore accords with the upper bound on compensating variation as consumers move from \mathcal{B}_s to \mathcal{B}_t , whereas a similar reasoning reveals that the ordering on \mathcal{B}_t accords with an upper bound on equivalent variation. Perfectly positive dependence then means that both orderings coincide, perfectly negative dependence maximizes their disagreement.

ordinatewise orderings map onto an interesting dimension of heterogeneity if goods are substitutes and complements only to a limited degree, so that a high demand for good k indeed indicates taste for it. Whether this assumption is defensible depends on how the orthants are defined. The favorable case would be that they correspond to well separated groups of nondurables. We also note that none of the below assumptions induce bounds that are easily expressed in closed form; in particular, although they are adaptations of quantile constancy and its weakenings, they do lead not to analogs of the preceding section's lemmas.

The strongest possible generalization of quantile constancy would be to maintain that all consumers keep their quantile position with respect to every good:

Definition 5. *Let $(\bar{\mathbf{Q}}_s, \bar{\mathbf{Q}}_t)$ and $(\tilde{\mathbf{Q}}_s, \tilde{\mathbf{Q}}_t)$ be distributed independently according to the distribution of $(\mathbf{Q}_s, \mathbf{Q}_t)$. The copula joining \mathbf{Q}_s and \mathbf{Q}_t exhibits co-ordinatewise quantile constancy if*

$$\left(\bar{\mathbf{Q}}_s - \tilde{\mathbf{Q}}_s\right) \odot \left(\bar{\mathbf{Q}}_t - \tilde{\mathbf{Q}}_t\right) \geq \mathbf{0}$$

a.s., where \odot denotes the componentwise product. In words, any matching components of $(\bar{\mathbf{Q}}_s - \tilde{\mathbf{Q}}_s)$ and $(\bar{\mathbf{Q}}_t - \tilde{\mathbf{Q}}_t)$ are concordant (their product is non-negative) a.s.

Co-ordinatewise quantile constancy is perhaps the most natural generalization of quantile constancy. It is obviously quite strong; indeed, it is not only point identifying but testable in the sense of generating cross-marginal restrictions that might be violated in the data.¹⁰ A less obvious observation is that co-ordinatewise quantile constancy generalizes the intuition behind quantile constancy but has a quite different impact. In particular, it need not lead to the lower worst-case bounds, as is illustrated in our example.

Example 2. (continued) *The fictitious population distributions are compatible with co-ordinatewise quantile constancy. They then imply that $\mathbb{P}(\text{WARP violated}) = 0.25$. If the location of \mathbf{b}_2 were changed to $(0, 3, 0)$, co-ordinatewise quantile constancy would have to be violated.*

Co-ordinatewise quantile constancy turns out to generate the *upper* worst-case bounds (admittedly due to careful rigging of the example). Intuitively, what happens here is that we maintain an ordering of consumers on budget planes that might be substantively plausible but that disagrees with the ordering underlying the lower worst-case bounds.

¹⁰Hence, despite our wording, it is not strictly an assumption about copulas only and accordingly not known in the copulas literature. Heterogeneous Cobb-Douglas preferences constitute an important special case where the assumption is fulfilled.

One might complain that co-ordinatewise quantile constancy is too strong an assumption, especially when it is incompatible with the data but even otherwise. We therefore turn to weakenings of it. All of these are assumptions about copulas proper and, therefore, not testable from observation of marginals.

To begin, one could impose quantile constancy not along all orthants simultaneously, but first along one, then (conditionally) another one and so forth.

Definition 6. *The copula joining \mathbf{Q}_s and \mathbf{Q}_t exhibits triangular quantile constancy with respect to an ordering (i_1, \dots, i_m) of orthants if for independent random variables $(\bar{\mathbf{Q}}_s, \bar{\mathbf{Q}}_t)$ and $(\tilde{\mathbf{Q}}_s, \tilde{\mathbf{Q}}_t)$ distributed as $(\mathbf{Q}_s, \mathbf{Q}_t)$,*

$$\begin{aligned} \bar{Q}_s^{i_1} \geq \tilde{Q}_s^{i_1} &\iff \bar{Q}_t^{i_1} \geq \tilde{Q}_t^{i_1} \\ \left[\bar{Q}_s^{i_1} = \tilde{Q}_s^{i_1} \right] &\implies \left[\bar{Q}_s^{i_2} \geq \tilde{Q}_s^{i_2} \iff \bar{Q}_t^{i_2} \geq \tilde{Q}_t^{i_2} \right] \\ &\text{etc.,} \end{aligned}$$

where $\bar{Q}_s^{i_1}$ is the i_1 -component of $\bar{\mathbf{Q}}_s$ etc.

Triangular quantile constancy weakens co-ordinatewise quantile constancy by specifying a tie-breaking rule if maintenance of quantile positions along different orthants would encounter a contradiction. Note that the assumption's impact generally depends on how one orders the orthants; it is invariant to interchange of orthants, and at the same time coincides with co-ordinatewise quantile constancy, if (and only if) the latter is compatible with the data. While not previously known in the copulas literature, triangular quantile constancy resembles assumptions made in the literature on nonparametric demand estimation. An obvious problem is that the ordering of orthants will typically be arbitrary, thus point identification is arguably achieved at a high cost in terms of credibility.

We therefore also propose some assumptions that are symmetric with respect to permutations of the orthants, hence they do not depend on how the orthants are ordered. The assumptions weaken co-ordinatewise quantile constancy in a different way, namely by merely requiring that quantile positions are maintained in some stochastic sense.

Definition 7. *The copula linking \mathbf{Q}_s and \mathbf{Q}_t exhibits co-ordinatewise positive quadrant dependence if $\mathbb{P}_{yz}(Q_s^i \geq a, Q_t^i \geq b) \geq \mathbb{P}_{yz}(Q_s^i \geq a) \mathbb{P}_{yz}(Q_t^i \geq b)$ for all scalars a, b and orthants $i \leq m$.*

Definition 8. *The copula linking \mathbf{Q}_s and \mathbf{Q}_t exhibits association if $\mathbb{P}_{yz}(\mathbf{Q}_s \in A, \mathbf{Q}_t \in B) \geq \mathbb{P}_{yz}(\mathbf{Q}_s \in A) \mathbb{P}_{yz}(\mathbf{Q}_t \in B)$ for all upper contour sets A and B .*

Both definitions generalize the idea behind quadrant dependence, the former one by applying it separately along each orthant, the latter one by allowing for general upper contour sets. Association is strictly stronger; it is also the assumption that previously appeared in the literature (Esary et al. (1967)).

On the face of it, the assumptions enforce some consistency of tastes across budget problems and hence, appear optimistic in the sense of limiting the probability of violating WARP. This impression is not quite accurate for two reasons. The less important one is that both assumptions are fulfilled by the independent copula as a boundary case, hence neither of them will exclude the possibility that $\mathbb{P}_{yz}(\mathbf{Q}_s \in \underline{\mathcal{B}}_{s,t}, \mathbf{Q}_t \in \underline{\mathcal{B}}_{t,s}) = \mathbb{P}_{yz}(\mathbf{Q}_s \in \underline{\mathcal{B}}_{s,t})\mathbb{P}_{yz}(\mathbf{Q}_t \in \underline{\mathcal{B}}_{t,s})$. In particular, neither will reduce the upper bound on $\mathbb{P}_{yz}(\text{WARP violated})$ to zero unless it was zero to begin with. More importantly, just as with co-ordinatewise quantile constancy, the assumptions can actually be pessimistic in the sense of refining bounds from below but not from above.

Example 3. (continued) *Imposing triangular quantile constancy, with the first orthant having highest priority, implies $\mathbb{P}_{yz}(\text{WARP violated}) = 0.25$. Imposing association refines the bounds to $0.125 \leq \mathbb{P}_{yz}(\text{WARP violated}) \leq 0.25$.¹¹ Merely imposing co-ordinatewise positive quadrant dependence does not affect the bounds.*

Our example indicates that the potential tension between apparently reasonable assumptions and WARP remains; thus, it contains a piece of cautionary advice for users who wish to use specifications of this type in nonparametric analysis of demand. Triangular quantile constancy point identifies $\mathbb{P}_{yz}(\text{WARP violated})$ at its highest possible level, and association refines away the *lower* half of the identified set. In particular, both assumptions imply spurious violations of WARP – the lower bound is strictly positive even though the data were generated by rational consumers. Intuitively plausible restriction on copulas in higher dimensional commodity spaces may in general (although certainly not for every particular data set) be logically inconsistent with presuming consumer rationality. The intuitive reason why this can happen is again the same: The ordering of consumers that is maintained – at least in some stochastic sense – is simply not relevant for the Fréchet-Hoeffding problem.

3 Hypothesis Tests and Confidence Intervals

Estimation of bounds presents a relatively routine nonparametric estimation problem, and we will give details on our procedure in the empirical section. Inference on partially identified parameters, however, raises a number of conceptual and technical issues that are the subject of a currently

¹¹This claim is established in the appendix.

active literature, notably (for our purposes) Andrews and Soares (2007), Imbens and Manski (2004), and Stoye (2009). To tackle these, we continue to assume continuity of relevant population distributions, and we also focus on worst-case Fréchet bounds, that is, we use (2.1.2) to construct confidence regions for parameters of interest

$$\theta_{yz} \in \Theta_{yz} = [\max\{\pi_{yz} - \psi_{yz}, 0\}, \min\{\pi_{yz}, 1 - \psi_{yz}\}],$$

where

$$\begin{aligned}\pi_{yz} &= \mathbb{P}_{yz}(Q_s < (y_t - y_s)/(p_t - p_s)) \\ \psi_{yz} &= \mathbb{P}_{yz}(Q_t < (y_t - y_s)/(p_t - p_s)).\end{aligned}$$

We will assume the following:

Assumption 2. *There exist estimators*

$$\begin{aligned}\hat{\pi}_{yz} &= \hat{\mathbb{P}}_{yz}(Q_s < (y_t - y_s)/(p_t - p_s)) \\ \hat{\psi}_{yz} &= \hat{\mathbb{P}}_{yz}(Q_t < (y_t - y_s)/(p_t - p_s))\end{aligned}$$

such that

$$\sqrt{\phi_n} \begin{bmatrix} \hat{\pi}_{yz} - \pi_{yz} \\ \hat{\psi}_{yz} - \psi_{yz} \end{bmatrix} \xrightarrow{d} N \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_{\pi_{yz}}^2 & \rho_{yz} \sigma_{\pi_{yz}} \sigma_{\psi_{yz}} \\ \rho_{yz} \sigma_{\pi_{yz}} \sigma_{\psi_{yz}} & \sigma_{\psi_{yz}}^2 \end{bmatrix} \right),$$

where $\phi_n = nh^d$, h being the kernel bandwidth and d the dimensionality of (Y, Z) , and where estimators $(\hat{\sigma}_{\pi_{yz}}, \hat{\sigma}_{\psi_{yz}}, \hat{\rho}_{yz})$ are available.

In our application, we estimate $\hat{\sigma}_{\pi_{yz}}$ and $\hat{\sigma}_{\psi_{yz}}$ and set $\rho = 0$ since samples are drawn from different time periods. The latter leads to algebraic simplification in several of the results reported below; we exploit these simplifications in our implementation but omit them here. Also, all results apply equally if the distribution of $\sqrt{\phi_n} (\hat{\pi}_{yz} - \pi_{yz}, \hat{\psi}_{yz} - \psi_{yz})$ is approximated by a bootstrap distribution $F_{\pi_{yz}^* - \hat{\pi}_{yz}, \psi_{yz}^* - \hat{\psi}_{yz}}$, provided of course that this approximation can be justified. It is the bootstrap version that we implement in the next section.

Two major issues regarding inference in this setting are the following. First, should a confidence region cover the identified set, i.e. the relevant coverage probability is $\mathbb{P}(\Theta_{yz} \subseteq CI_{1-\alpha}(\Theta))$, or should it cover the partially identified parameter, i.e. the relevant coverage probability is $\mathbb{P}(\theta_{yz} \in CI_{1-\alpha}(\theta))$? The answer plainly depends on what is conceived as the quantity of interest. We will leave the choice to the reader and develop both types of inference.

Second, some simple confidence regions will display undesirable behavior because they fail to account for non-pivotality of underlying test statistics. More specifically, observe that identification through Fréchet bounds can be seen as identification through moment inequalities:

$$\max\{\pi_{yz} - \psi_{yz}, 0\} \leq \theta_{yz} \leq \min\{\pi_{yz}, 1 - \psi_{yz}\}$$

is equivalent to the conjunction of

$$\pi_{yz} - \psi_{yz} - \theta_{yz} \leq 0 \tag{3.1a}$$

$$-\theta_{yz} \leq 0 \tag{3.1b}$$

$$\theta_{yz} - \pi_{yz} \leq 0 \tag{3.1c}$$

$$\theta_{yz} - 1 + \psi_{yz} \leq 0. \tag{3.1d}$$

Confidence intervals will be lower contour sets of a test statistic that aggregates violations of sample versions of these inequalities. The distribution of this statistic depends on which inequalities bind, and it is not possible to pre-test for the binding inequalities with sufficient precision in boundary cases where one or more inequalities are close to binding. In line with the recent literature on moment inequalities, we resolve this by using conservative pre-tests. Conceptually, this method is by now well understood, including results for very general settings. By exploiting the specific structure of the present problem, we can however improve on mechanical implementation of available methods. In particular, we know that (3.1a) and (3.1b) as well as (3.1c) and (3.1d) potentially compete, and we will accordingly pre-test for the comparison of those pairs of inequalities. What is more, (3.1b) is nonstochastic, simplifying some developments. In a final step, one should generally also pre-test for the possibility that the upper and lower bound simultaneously bind at the true parameter value, that is, there is point identification. Simple algebra shows that in the present case, this obtains if $\min\{\pi_{yz}, 1 - \pi_{yz}, \psi_{yz}, 1 - \psi_{yz}\} = 0$. These cases are frequently simple to pre-test because they involve degenerate probabilities; furthermore, all data points that we consider below are far away from them. We will therefore ignore this aspect by assuming the following:

Assumption 3. $\varepsilon \leq \pi_{yz}, \psi_{yz} \leq 1 - \varepsilon$ for some $\varepsilon > 0$.

Assumption 3 greatly simplifies matters because it implies that Θ_{yz} is asymptotically large relative to standard errors.

Many of the ideas presented in this section can be extended to inference on Fréchet-Hoeffding bounded probabilities more generally and without assumption 3; we do this in ongoing research. We also reiterate that the ideas can be equivalently used if the above, normal approximation is replaced with a bootstrap approximation. It is important to understand, however, that use of the bootstrap does *not* automatically resolve the pre-testing problems analyzed here. That is, valid inference requires that one bootstrap the above sampling distribution but then use conservative pre-tests as detailed below. Also, we keep the treatment simple by formulating simple pre-tests in which certain nulls are either rejected or not. Smooth pre-tests (corresponding to smooth

shrinkage, rather than thresholding, for certain test statistics) might seem conceptually appealing, but experience with inference on moment inequalities suggests that gains from such refinements are at best moderate (Andrews and Jia (2008)). In our application, empirical results are clear enough that we desist from complicating the analysis by such refinements.

Unlike papers cited above, we will not establish uniform (in parameter space) validity of our confidence regions. However, this is only because we would then require uniform validity of asymptotic joint normality respectively the bootstrap, and we do not wish to get into the required regularity conditions. The pre-tests inherent in our constructions do avoid the aforementioned, non-pivotality driven uniformity failures.

3.1 One-Sided Confidence Regions

We first construct one-sided confidence regions for the upper and lower bound on θ_{yz} . Let $\widehat{\sigma}_{sum}^2 = \widehat{\sigma}_{\pi_{yz}}^2 + \widehat{\sigma}_{\psi_{yz}}^2 - 2\widehat{\rho}_{yz}\widehat{\sigma}_{\pi_{yz}}\widehat{\sigma}_{\psi_{yz}}$ and

$$c_{1-\alpha}^l = \Phi^{-1}(1 - \alpha)\widehat{\sigma}_{sum},$$

where Φ is the standard normal c.d.f. Then we have:

Lemma 3.1. *Let assumption 2 hold. Then*

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(\max\{\pi_{yz} - \psi_{yz}, 0\} \geq \max \left\{ \widehat{\pi}_{yz} - \widehat{\psi}_{yz} - \phi_n^{-1/2} c_{1-\alpha}^l, 0 \right\} \right) = 1 - \alpha.$$

This one-sided confidence interval will be conservative if $\pi_{yz} - \psi_{yz} < 0$, reflecting that the underlying testing problem is then degenerate.

One-sided testing for the upper bound presents additional problems because $\phi_n^{1/2} \left(\min \left\{ \widehat{\pi}_{yz}, 1 - \widehat{\psi}_{yz} \right\} - \min \left\{ \pi_{yz}, 1 - \psi_{yz} \right\} \right)$ is not an asymptotic pivot. We resolve this issue by implicitly shrinking $\phi_n^{1/2} (\pi_{yz} + \psi_{yz} - 1)$ toward zero, making for a test that is pointwise exact and uniformly valid, but conservative along certain local parameter sequences. Thus, let b_n be a user-defined sequence s.t. $b_n \rightarrow 0$ but $\phi_n^{1/2} b_n \rightarrow \infty$. If $\left| \widehat{\pi}_{yz} + \widehat{\psi}_{yz} - 1 \right| \leq b_n$, let $c_{1-\alpha}^u$ be s.t.

$$\Phi_{BVN} \left(c_{1-\alpha}^u, c_{1-\alpha}^u; 0, 0, \widehat{\sigma}_{\pi_{yz}}, \widehat{\sigma}_{\psi_{yz}}, -\widehat{\rho}_{yz} \right) = 1 - \alpha,$$

where $\Phi_{BVN}(z_1, z_2; \mu_1, \mu_2, \sigma_1, \sigma_2, \rho)$ denotes the bivariate normal c.d.f. with the indicated parameters, evaluated at (z_1, z_2) . If $\widehat{\pi}_{yz} + \widehat{\psi}_{yz} - 1 < b_n$, let $c_{1-\alpha}^u = \widehat{\sigma}_{\pi_{yz}} \Phi^{-1}(1 - \alpha)$; if $\widehat{\pi}_{yz} + \widehat{\psi}_{yz} - 1 > b_n$, let $c_{1-\alpha}^u = \widehat{\sigma}_{\psi_{yz}} \Phi^{-1}(1 - \alpha)$. Then we have:

Lemma 3.2. *Let assumption 2 hold. Then*

$$\begin{aligned} & \lim_{n \rightarrow \infty} \mathbb{P} \left(\min \{ \pi_{yz}, 1 - \psi_{yz} \} \leq \min \left\{ \widehat{\pi}_{yz}, 1 - \widehat{\psi}_{yz} \right\} + \phi_n^{-1/2} c_{1-\alpha}^u \right) \\ &= \lim_{n \rightarrow \infty} \mathbb{P} \left(\min \{ \pi_{yz}, 1 - \psi_{yz} \} \geq \min \left\{ \widehat{\pi}_{yz}, 1 - \widehat{\psi}_{yz} \right\} - \phi_n^{-1/2} c_{1-\alpha}^u \right) \\ &= 1 - \alpha. \end{aligned}$$

Here, the first claim will be used in the next subsection, and the second one is relevant to test statistical significance of strictly positive upper bounds, as we will do later. Note that if $\widehat{\pi}_{yz}$ and $1 - \widehat{\psi}_{yz}$ are similar with similar estimated standard errors and we furthermore have $\rho_{yz} = 0$, then calibration of $c_{1-\alpha}^u$ from the bivariate standard normal distribution is essentially equivalent to setting $c_{1-\alpha}^u = \widehat{\sigma}_{\psi_{yz}} \Phi^{-1}((1 - \alpha)^{1/2}) \approx \widehat{\sigma}_{\psi_{yz}} \Phi^{-1}(1 - \alpha/2)$, where the last step presumes a reasonably small α . This simplification will hold in our application.

3.2 Inference on the True Parameter

We now construct confidence regions for the partially identified parameter θ_{yz} . An according confidence region $CI(\theta)$ must fulfil

$$\lim_{n \rightarrow \infty} \inf_{\theta_{yz} \in \Theta_{yz}} \mathbb{P}(\theta_{yz} \in CI_{1-\alpha}(\theta)) = 1 - \alpha.$$

Here, the inf-operator reflects the fact that coverage probabilities cannot in general be approached uniformly over Θ_{yz} . For example, if Θ_{yz} has an interior, then it is typically the case that any reasonable estimator of it is a 100%-confidence region for any θ_{yz} in this interior. The coverage probability of $1 - \alpha$ therefore pertains to a least favorable choice of θ_{yz} within Θ_{yz} .

A valid construction is here straightforwardly defined as

$$CI_{1-\alpha}(\theta) = \left[\max \left\{ \widehat{\pi}_{yz} - \widehat{\psi}_{yz} - \phi_n^{-1/2} c_{1-\alpha}^l, 0 \right\}, \min \left\{ \widehat{\pi}_{yz}, 1 - \widehat{\psi}_{yz} \right\} + \phi_n^{-1/2} c_{1-\alpha}^u \right],$$

i.e. by intersecting the one-sided $(1 - \alpha)$ -confidence intervals for the upper and lower bound. Intuitively, this works because under our assumptions, the length of the identified set is asymptotically large relative to standard errors, thus the corresponding hypothesis testing problem is either one-sided (if θ_{yz} is on the boundary of Θ_{yz}) or degenerate (if it is interior). Formally, we observe the following.

Lemma 3.3. *Let assumptions 2 and 3 hold. Then*

$$\lim_{n \rightarrow \infty} \inf_{\theta_{yz} \in \Theta_{yz}} \mathbb{P}(\theta_{yz} \in CI_{1-\alpha}(\theta)) = 1 - \alpha.$$

3.3 Inference on the Identified Set

A confidence region for Θ_{yz} is constructed by projecting a simultaneous confidence set for the upper and lower bound. Thus, let

$$(z_1, z_2) \sim N \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \hat{\sigma}_{\pi yz}^2 & \hat{\rho}_{yz} \hat{\sigma}_{\pi yz} \hat{\sigma}_{\psi yz} \\ \hat{\rho}_{yz} \hat{\sigma}_{\pi yz} \hat{\sigma}_{\psi yz} & \hat{\sigma}_{\psi yz}^2 \end{bmatrix} \right).$$

Let the sequence b_n be as before and let (c^l, c^u) minimize $(c^l + c^u)$ (and thus the length of the resultant confidence region) subject to the constraint that

$$\mathbb{P}(z_1 \geq -c^u, z_2 \leq c^u, z_1 - z_2 \leq c^l) \geq 1 - \alpha$$

$$\text{if } \left| \hat{\pi}_{yz} + \hat{\psi}_{yz} - 1 \right| \leq b_n,$$

$$\mathbb{P}(z_1 \geq -c^u, z_1 - z_2 \leq c^l) \geq 1 - \alpha$$

$$\text{if } \hat{\pi}_{yz} < 1 - \hat{\psi}_{yz}, \text{ and}$$

$$\mathbb{P}(z_2 \leq c^u, z_1 - z_2 \leq c^l) \geq 1 - \alpha$$

otherwise. Let

$$CI_{1-\alpha}(\Theta) = \left[\max \left\{ \hat{\pi}_{yz} - \hat{\psi}_{yz} - \phi_n^{-1/2} c^l, 0 \right\}, \min \left\{ \hat{\pi}_{yz}, 1 - \hat{\psi}_{yz} \right\} + \phi_n^{-1/2} c^u \right]$$

$$\text{if } \hat{\pi}_{yz} - \hat{\psi}_{yz} > -b_n \text{ and}$$

$$CI_{1-\alpha}(\Theta) = \left[0, \min \left\{ \hat{\pi}_{yz}, 1 - \hat{\psi}_{yz} \right\} + \phi_n^{-1/2} c_{1-\alpha}^u \right]$$

otherwise, where $c_{1-\alpha}^u$ is as in section 3.1. Intuitively, the final case distinction exploits the fact that if $\pi_{yz} - \psi_{yz}$ is known to be negative, then noncoverage risk is exclusively incurred at the interval's upper end. Then:

Lemma 3.4. *Let assumption 2 hold. Then*

$$\lim_{n \rightarrow \infty} \mathbb{P}(\Theta_{yz} \subseteq CI_{1-\alpha}(\Theta)) = 1 - \alpha.$$

3.4 Testing for Violations of WARP

Recall that a corollary of bounds on $\mathbb{P}_{yz}(\text{WARP violated})$ is a delineation of the empirical content of WARP in repeated cross-sectional data: these are compatible with WARP as long as the lower bound is zero. It is, therefore, of special interest to test whether the lower bound is nonzero for at least some (y, z) . To do so, it is not appropriate to just collect values of (y, z) for which the preceding confidence region fails to include 0 – simultaneously applied to many such values, this

approach would tend to generate false positives. One must therefore control for familywise error rates. Given the large number of values (y, z) , a Bonferroni correction would, on the other hand, be hopelessly conservative. We rather propose to use current developments in the econometrics of partial identification. Specifically, the null hypothesis that the lower bound on \mathbb{P}_{yz} (WARP violated) equals zero for all (y, z) can be written as

$$H_0 : (\pi_{yz} - \psi_{yz} \leq 0), \forall (y, z) \iff \max_{y,z} \{\pi_{yz} - \psi_{yz}\} \leq 0.$$

The second way of writing the hypothesis reveals that testing it would be an application of current work by Chernozhukov, Lee, and Rosen (2009). Specifically, they propose to test this hypothesis by the following algorithm:

- Pre-select a conservative (that is, guaranteed to be too large in the limit) estimator \widehat{V} of $V = \arg \max_{yz} \{\pi_{yz} - \psi_{yz}\}$.
- Letting $\theta(yz)$ be the population objective function, $\widehat{\theta}(yz)$ an estimator, and $\widehat{\sigma}(yz)$ an estimator of the pointwise standard error of $\widehat{\theta}$, simulate the distribution of the studentized error process $\tau(yz) = (\widehat{\theta}(yz) - \theta(yz)) / \widehat{\sigma}(yz)$. Let τ^* be the $(1 - \alpha)$ -quantile of $\sup_{yz} \tau(yz)$.
- Reject the null if $\sup_{yz} \{\widehat{\theta}(yz) - \tau^* \widehat{\sigma}(yz)\} > 0$.

Results in Chernozhukov, Lee, and Rosen (2009) imply that the method could be implemented here, provided we impose enough assumptions to ensure uniform consistency of estimators $(\widehat{\pi}_{yz}, \widehat{\psi}_{yz})$. We desist from such an implementation because our empirical results imply that the result will be acceptance of H_0 . The method might be important, though, in other data sets that may generate (potentially spurious) indications of violations of WARP. We also note that essentially the same remarks apply to testing whether the upper bound is ever significantly positive, which is equivalent to testing for rejection of

$$H'_0 : (\min\{\pi_{yz}, 1 - \psi_{yz}\} \leq 0), \forall (y, z) \iff \max_{y,z} \{\min\{\pi_{yz}, 1 - \psi_{yz}\}\} \leq 0.$$

4 Empirical Application: WARP and the British Household Expenditure Survey

In this section, we analyze a real world data set with the framework we propose in this paper. The data come from the Family Expenditure Survey (FES) and were the basis for successful, recent applications of revealed preference approaches (Blundell, Browning, Crawford (2003, 2008)). This section is structured as follows: We first provide a description of the data we use. Then we present some econometric details. Finally, we display the empirical results.

4.1 Description of the Data

The FES reports a yearly cross section of labor income, expenditures, demographic composition and other characteristics of about 7,000 households. We use the years 1974-1993, but exclude the respective Christmas periods as they contain too much irregular behavior. As is standard in the demand system literature, we focus on the subpopulation of two person households where both are adults, at least one is working, and the head of household is a white collar worker. This is to reduce the impact of measurement error; see Lewbel (1999) for a discussion. We provide a summary statistic of our data in table 1 in the appendix.

We form several expenditure categories. The first category is related to food consumption and consists of the subcategories food bought, food out (catering) and tobacco. The second category contains housing expenditures, namely rent or mortgage payments and household goods and services, excluding furniture. The last group consists of motoring and fuel expenditures. For brevity, we call these categories food, housing and energy. These broader categories are formed since more detailed accounts suffer from infrequent purchases (recall that the recording period is 14 days) and are thus often underreported. These three categories account for 20-30% of total expenditure on average, leaving a fourth residual category. Results actually displayed were generated by considering consumption of food versus nonfood items, but similar analyses were performed for all of the goods, and with similar results. We removed outliers by excluding the upper and lower 2.5% of the population in the three groups.

For the pairwise comparisons, we normalize prices by dividing all variables by the general price index excluding the good into consideration (in particular, for food we consider the price of food vs. the price of all nondurable goods except food). This removes both general inflation and transforms all prices to be relative to the price index. Quantities are defined by dividing the normalized expenditures by the respective normalized price, e.g. food by the food price index. We also divide total expenditure by the price index.

To account for possible endogeneities, i.e. violations of assumption 1, we use labor income as an instrument. It is constructed as in the household below average income study (HBAI), that is, it is roughly defined as labor income after taxes and transfers. We include the remaining household covariates as regressors. Specifically, we use principal components to reduce the vector of remaining household characteristics to a few orthogonal, approximately continuous components, mainly because we require continuous covariates for nonparametric estimation. Since we already condition on a lot of household information by using the specific subgroup, we only use the first principal component. While this is arguably ad hoc, we perform some robustness checks like alternating the component or adding several others, and results do not change appreciably.

4.2 Econometric Specification and Empirical Results

We estimate conditional probabilities (π_{yz}, ψ_{yz}) via a locally linear estimator with a standard Epanechnikov kernel. The bandwidth is selected by cross validation. We checked the sensitivity of our results by varying the bandwidth; there was no material effect on results. Sampling distributions of $\phi_n^{1/2}(\hat{\psi}_{yz} - \psi_{yz})$ and $\phi_n^{1/2}(\hat{\pi}_{yz} - \pi_{yz})$ were simulated by a hybrid bootstrap (Shao and Tu (1995)). Specifically, we use the inverted c.d.f. of $F_{\hat{\psi}_{yz}^* - \hat{\psi}_{yz}}$, where $\hat{\psi}_{yz}^*$ denotes the bootstrap estimator and $\hat{\psi}_{yz}$ the original estimator, to derive a consistent estimator for $F_{\hat{\psi}_{yz} - \psi_{yz}}$, and similarly for $\hat{\pi}_{yz}$. As elaborated before, we also assume that $\hat{\psi}_{yz}$ and $\hat{\pi}_{yz}$ are generated from independent samples.

When applying our local polynomial estimators to the choice data, we proceed as follows: We first group the population into “bands” of three years, e.g., we collect all people surveyed in the years 1974-1976 into one group. This is done to increase the number of observations. As a consequence, our cross sections actually comprise 3 years, and we assume the individuals to face the mean price in this period.¹² We then make pairwise comparisons for all pairs of groups that are not too distant in time. This is done because for groups that are many years apart, apparent violations of WARP could plausibly be driven by changing preferences. Specifically, we compare every cross section with the two adjacent ones only.

Our first important finding is that for most such comparisons, the income change swamps the price effect, leading to upper bounds of zero. This is easy to explain: The order of magnitude of the relative price change is -0.05 , while most quantities are around 10. Thus, the overall effect of a price change on quantities is $\Xi_{ts} = (p_t - p_s) Q_s \approx 0.5$. Figure 2 provides a graphical representation of the density of this effect in the second year (so the variable is $\Xi_{22} = (p_3 - p_2) Q_2$, corresponding to the years of 1978 and 1981).

Fig. 2 approx. here

The probability mass is highly concentrated between -1 and 0 . In contrast, mean real income increased in the same two periods from 50.4 to 54.3, and median income increased less dramatically from 43.5 to 45.5. Still, the typical case is that $\pi_{yz} = 1 - \psi_{yz} = 0$, so that $\mathbb{P}_{yz}(\text{WARP violated})$ is point identified at 0.

Of course, this observation might not attest to the population’s rationality but rather to the income and price processes being uninformative. Hence, we focus on regions in the data that are at least potentially informative, as operationalized by a nonzero upper bound on $\mathbb{P}_{yz}(\text{WARP}$

¹²This induces a measurement error. Compared to the already incurred measurement error, and in light of the fact that all that matters is the change in prices, we feel that this is a minor issue.

violated). For instance, given the distribution of Ξ_{22} , we focus on income changes that are between 0 and -1 , i.e. individuals who become marginally poorer. The following density shows the distribution of the upper bound on a 10×10 grid that combines each of the (9, 18, ..., 90)-quantiles of the income distribution in the period (1977-1982) with all of the income changes in $(-1, -0.9, \dots, -1)$. Hence, we focus on consumers who experienced small income losses but look at varying positions of the income distribution; see figure 2. This choice of subsample is certainly ad hoc; we leave to future research a more systematic treatment of the choice of region for which the data is informative. Figures 3 and 4 display our results. We start with the distribution of the point estimate for the upper bound in our subsample.

Fig. 3 approx. here

The point estimate of the upper bound exceeds 1% in 93 out of the 100 points of support of the regressors; it exceeds 5% (the threshold indicated by the vertical line) in 80 cases. Thus, the data appear potentially informative about WARP. To check whether the positive values are statistically significant, we use lemma 3.2 to compute a lower confidence bound for the upper bound. That is, we pre-test for equality of π_{yz} and $(1 - \psi_{yz})$ at significance level $\tau_n = 0.01$ (of course, the idea is that $\tau_n \rightarrow 0$ as $n \rightarrow \infty$). If equality is rejected, we form a standard 95% confidence band for whichever of π_{yz} and $(1 - \psi_{yz})$ appears smaller; if it is not rejected, we form joint confidence sets, which are effectively 97.5% for either parameter because $\pi_{yz} \approx (1 - \psi_{yz})$ in these cases. (See the elaborations after lemma 3.2 for a more detailed explanation.) We reiterate that this inference is locally conservative if $\pi_{yz} \approx (1 - \psi_{yz})$, but that this feature is unavoidable unless one is willing to sacrifice validity of confidence regions for these local cases.

Many of the positive upper bounds are statistically significant. In 84 out of 100 cases, the lower confidence interval is above 0.01. The mean and median of the upper bound are 0.177 respectively 0.146. More than a third of cases are above 0.25, while the 95% bootstrap CI usually has length between 0.10 and 0.15. Figure 4 shows the distribution of the upper bound and that of the lower end of the 95% bootstrap CI around the upper bound.

Fig. 4 approx. here

While these numbers should be interpreted with care due to our failure to control familywise error rates, they do suggest that the data are reasonably informative – in the sense of allowing for potential violation of WARP – in this selected subpopulation. Still, figure 5 illustrates that we find hardly any evidence against WARP.

Fig. 5 approx. here

The lower bound is typically close to zero even in this informative subpopulation; it exceeds 0.05 in only 1 out of 100 instances. What is more, one-sided 95% bootstrap confidence intervals for $\pi_{yz} - \psi_{yz}$, constructed in accordance with lemma 3.1, include zero at 97 of 100 positions on our grid. Hence, we cannot statistically distinguish the positive lower bounds from zero even with confidence regions that fail to control familywise error rates.

We now showcase the effect of one of the refinements discussed in section 2. Specifically, we impose positive quadrant dependence (PQD; see section 2.1.3). The refined bounds are illustrated in figure 6, and figure 7 shows through a comparison the effect of the introduction of PQD on the distribution of the upper bound.

Figs. 6,7 approx. here

As explained in section 2, PQD will not induce refined upper bounds of zero unless worst-case upper bounds were zero, however the upper bounds are much reduced, typically by 50 – 75%. The highest possible proportion of “violators” is substantially reduced and exceeds 20% only at very select data points. Having said that, most of the positive bounds are statistically significantly so, with many p-values being small enough that controlling for familywise error rates would not overturn this conclusion. Hence, the data are still consistent with some violations of WARP. Lower bounds on $\mathbb{P}_{yz}(\text{WARP violated})$ remain zero because in the two-dimensional case, PQD refines worst-case bounds from above but not below; recall that this is not true for natural generalizations of it in the higher dimensional case. We interpret this finding as being supportive of rationality: While WARP in isolation may be a rather weak restriction, combined with the plausible PQD property imposed on the population, the restriction becomes more stringent. And while the data are still potentially informative about this restriction (as indicated by the significantly positive upper bound), we still do not detect violations.

None of these results change appreciably if we include a measure of household characteristics and/or correct for endogeneity using a control function approach. Moreover, they are stable across the large groups of goods we consider, for pairwise comparisons (e.g., energy vs. non-energy). In summary, we tend to think that at least as a reasonable approximation to behavior, WARP is more corroborated than questioned by these data, but we would like to emphasize the need for further research with other data.

5 Conclusion

This paper investigated exactly what power revealed preference assumptions have under realistic data constraints. The leading question was to bound the fraction of a population that violates

WARP given repeated cross-section data. Side results were to elucidate the exact empirical content of WARP and to carry out an inference exercise that applies recent insights about inference under partial identification. The empirical result with respect to the U.K. Family Expenditure Survey is that even for those observations where budget planes meaningfully overlap, as reflected by large upper bounds on the probability of violating WARP, lower bounds are not significantly positive, i.e. WARP cannot be rejected. Furthermore, imposing a very weak, nonparametric limitation on heterogeneity (namely, positive quadrant dependence with respect to budget shares spent on different goods) leads to uniformly rather small, though not uniformly zero, upper bounds.

The core difference between this paper and existing work that estimates demand for applied purposes is that we consider the revealed preference paradigm on individual level in isolation, being careful to impose no or very weak homogeneity assumptions. This, of course, leads to less conclusive results. While the data may be interpreted to be mildly supportive of WARP, this could certainly be due not to the population being substantively rational, but to the weak axiom being, well, weak. To be sure, we do not mean to implicitly criticize other papers, but rather to augment them by showing how much mileage can be gained from revealed preference assumptions proper. Thus, our motivation is somewhat similar to early papers on partial identification of treatment effects, which frequently stress the conceptual value of understanding just how much one could learn from the data without identifying assumptions. Insofar as the result is somewhat negative, the substantive upshot might well be a corroboration of approaches that use stronger assumptions. We hope, however, to illuminate just how much sharper conclusions than ours will depend on using sharper assumptions, whether or not these assumptions are formally semi- or nonparametric.

Appendix I - Summary Statistics of Data: Household Characteristics, Income and Normalized Expenditures

Variable	Minimum	1st Quartile	Median	Mean	3rd Quartile	maximum
number of female	0	1	1	1.073	1	2
number of retired	0	0	0	0.051	0	1
number of earners	0	1	2	1.692	2	2
Age of HHhead	19	31	49	46	58	90
Fridge	0	1	1	0.987	1	1
Washing Machine	0	1	1	0.882	1	1
Centr. Heating	0	1	1	0.804	1	1
TV	0	1	1	0.874	1	1
Video	0	0	0	0.407	1	1
PC	0	0	0	0.792	0	1
number of cars	0	1	1	1.351	2	10
number of rooms	1	4	5	5.455	6	26
HHincome	6.653	37.550	52.210	61.820	73.920	3981.000
Food	0	5.565	7.346	7.867	9.602	52.519
Housing	0	4.052	7.859	9.715	12.910	375.486
Energy	0	1.271	1.812	2.121	2.509	34.103

Appendix II - Proofs

Proposition 1 Recall the Fréchet-Hoeffding bounds: For any two random variables X_1 and X_2 and events (in the relevant algebras) A_1 and A_2 , one has the tight bounds

$$\max\{\mathbb{P}(X_1 \in A_1) + \mathbb{P}(X_2 \in A_2) - 1, 0\} \leq \mathbb{P}(X_1 \in A_1, X_2 \in A_2) \leq \min\{\mathbb{P}(X_1 \in A_1), \mathbb{P}(X_2 \in A_2)\}.$$

To see validity of the lower bound, note that

$$\begin{aligned} & \mathbb{P}_{yz}(\text{WARP violated}) \\ &= \mathbb{P}_{yz}(Q_s \leq (y_t - y_s)/(p_t - p_s), Q_t \geq (y_t - y_s)/(p_t - p_s), Q_s \neq Q_t) \\ &= \mathbb{P}_{yz}(Q_s \leq (y_t - y_s)/(p_t - p_s), Q_t \geq (y_t - y_s)/(p_t - p_s)) - \\ & \quad \mathbb{P}_{yz}(Q_s \leq (y_t - y_s)/(p_t - p_s), Q_t \geq (y_t - y_s)/(p_t - p_s), Q_s = Q_t) \\ &\geq \mathbb{P}_{yz}(Q_s \leq (y_t - y_s)/(p_t - p_s), Q_t \geq (y_t - y_s)/(p_t - p_s)) - \mathbb{P}_{yz}(Q_s = Q_t = (y_t - y_s)/(p_t - p_s)) \\ &\geq \max\{\mathbb{P}_{yz}(Q_s \leq (y_t - y_s)/(p_t - p_s)) - \mathbb{P}_{yz}(Q_t < (y_t - y_s)/(p_t - p_s)), 0\} - \\ & \quad \min\{\mathbb{P}_{yz}(Q_s = (y_t - y_s)/(p_t - p_s)), \mathbb{P}_{yz}(Q_t = (y_t - y_s)/(p_t - p_s))\}, \end{aligned}$$

where the first equality spells out the event that WARP is violated, the next two steps use basic probability calculus, and the last step uses the *lower* Fréchet-Hoeffding bound on $\mathbb{P}_{yz}(Q_s \leq (y_t - y_s)/(p_t - p_s), Q_t \geq (y_t - y_s)/(p_t - p_s))$ as well as the *upper* Fréchet-Hoeffding bound on $\mathbb{P}_{yz}(Q_s = Q_t = (y_t - y_s)/(p_t - p_s))$. The expression in the lemma is generated by taking the maximum between the last expression and zero, observing that this renders redundant the max-operator in the preceding display.

To see that the bound is tight, consider the joint distribution of (Q_s, Q_t) characterized as follows: (i) a probability of $\min\{\mathbb{P}_{yz}(Q_s = (y_t - y_s)/(p_t - p_s)), \mathbb{P}_{yz}(Q_t = (y_t - y_s)/(p_t - p_s))\}$ is assigned to the event $(Q_s = Q_t = (y_t - y_s)/(p_t - p_s))$, (ii) the remaining probability mass of Q_s and Q_t is linked by the Fréchet-Hoeffding lower bound (perfectly positive dependence) copula.

To see validity of the upper bound, note that

$$\mathbb{P}_{yz}(\text{WARP violated}) \leq \min\{\mathbb{P}_{yz}(Q_s \leq (y_t - y_s)/(p_t - p_s)), \mathbb{P}_{yz}(Q_t \geq (y_t - y_s)/(p_t - p_s))\}$$

by the upper Fréchet-Hoeffding bounds and furthermore that

$$\begin{aligned}
& \mathbb{P}_{yz}(\text{WARP violated}) \\
&= \mathbb{P}_{yz}((Q_s \leq (y_t - y_s)/(p_t - p_s), Q_t > (y_t - y_s)/(p_t - p_s)) \\
&\quad \text{or } (Q_s < (y_t - y_s)/(p_t - p_s), Q_t \geq (y_t - y_s)/(p_t - p_s))) \\
&= 1 - \mathbb{P}_{yz}(Q_s \geq (y_t - y_s)/(p_t - p_s), Q_t \leq (y_t - y_s)/(p_t - p_s)) \\
&\leq 1 - \min \{ \mathbb{P}_{yz}(Q_s \geq (y_t - y_s)/(p_t - p_s)), \mathbb{P}_{yz}(Q_t \leq (y_t - y_s)/(p_t - p_s)) \} \\
&= \max \{ \mathbb{P}_{yz}(Q_s < (y_t - y_s)/(p_t - p_s)), \mathbb{P}_{yz}(Q_t > (y_t - y_s)/(p_t - p_s)) \},
\end{aligned}$$

where all equalities use basic probability calculus and the inequality utilizes a lower Fréchet-Hoeffding bound on $\mathbb{P}_{yz}(Q_s \geq (y_t - y_s)/(p_t - p_s), Q_t \leq (y_t - y_s)/(p_t - p_s))$.

To see that the bound is tight, note that it is achieved by the Fréchet-Hoeffding upper bound (perfectly negative dependence) copula.

Lemma 2.1 Validity of the bounds is immediate. To see that they are tight, note that it is possible for lower [upper] bounds to be achieved simultaneously for every realization of (Y_s, Y_t, Z) .

Lemma 2.2 Assume positive quadrant dependence. Recall that $\underline{\mathcal{B}}_{s,t}$ is a lower contour set and $\underline{\mathcal{B}}_{t,s}$ an upper one, hence

$$\mathbb{P}_{yz}(Q_s \notin \underline{\mathcal{B}}_{s,t}, Q_t \in \underline{\mathcal{B}}_{t,s}) \geq \mathbb{P}_{yz}(Q_s \notin \underline{\mathcal{B}}_{s,t})\mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s}),$$

hence

$$\begin{aligned}
\mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}, Q_t \in \underline{\mathcal{B}}_{t,s}) &= \mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s}) - \mathbb{P}_{yz}(Q_s \notin \underline{\mathcal{B}}_{s,t}, Q_t \in \underline{\mathcal{B}}_{t,s}) \\
&\leq \mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s}) - \mathbb{P}_{yz}(Q_s \notin \underline{\mathcal{B}}_{s,t})\mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s}) \\
&= (1 - \mathbb{P}_{yz}(Q_s \notin \underline{\mathcal{B}}_{s,t}))\mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s}) = \mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t})\mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s}).
\end{aligned}$$

The refined lower bound for (ii) is established similarly. The old lower and upper bounds are tight because the distributions that generate them are consistent with positive respectively negative quadrant dependence. The bounds at $\mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t})\mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s})$ are tight because independence of Q_s and Q_t cannot be excluded.

Lemma 2.3 The main claim – shown in the next paragraph – is that $\mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}, Q_t \in \underline{\mathcal{B}}_{t,s}) \leq \kappa + \mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}) + \mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s}) - 1$, which in conjunction with previous results implies that $\mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}, Q_t \in \underline{\mathcal{B}}_{t,s})$ is bounded above by the r.h.s. minimum. The additional

max-operator is needed because $\mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}, Q_t \in \underline{\mathcal{B}}_{t,s}) \geq 0$, yet the minimum will be negative if $\kappa \leq 1 - \mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}) - \mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s})$. This observation also establishes the very last claim.

To see the main claim, observe that $\underline{\mathcal{B}}_{t,s}$ corresponds to quantile positions $[1 - \mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s}), 1]$ on \mathcal{B}_t . It follows that under κ -limited quantile movements, consumers located on $\underline{\mathcal{B}}_{t,s}$ can occupy quantile positions on \mathcal{B}_s not lower than $[1 - \mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s}) - \kappa, 1 - \kappa]$. On the other hand, $\underline{\mathcal{B}}_{s,t}$ corresponds to quantile positions $[0, \mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t})]$ on \mathcal{B}_s . This implies the claim. A refinement is obtained iff $\kappa + \mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}) + \mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s}) - 1 < \min\{\mathbb{P}_{yz}(Q_s \in \underline{\mathcal{B}}_{s,t}), \mathbb{P}_{yz}(Q_t \in \underline{\mathcal{B}}_{t,s})\}$, which is equivalent to the inequality given in the lemma. Tightness obtains because whenever the new upper bound binds, the quantile movements constructed to prove it are feasible. Finally, as suggested by the bounds' symmetry in the arguments, the same bound is obtained by projecting the set of consumers who violate WARP onto \mathcal{B}_t .

Example 3 The first and third claim are easy to see, we will establish the one regarding association. To see the lower bound, let $U \equiv \{\mathbf{q} : (1, 2, 1) \cdot \mathbf{q} \geq 5.1\}$, then U contains \mathbf{a}_1 , \mathbf{a}_2 , and \mathbf{b}_2 but not \mathbf{a}_2 or \mathbf{b}_1 . Now write

$$\begin{aligned} \mathbb{P}_{yz}(\mathbf{Q}_1 \in U, \mathbf{Q}_2 \in U) &\geq \mathbb{P}_{yz}(\mathbf{Q}_1 \in U)\mathbb{P}_{yz}(\mathbf{Q}_2 \in U) \\ \iff \mathbb{P}_{yz}(\mathbf{Q}_2 \in U | \mathbf{Q}_1 \in U) &\geq \mathbb{P}_{yz}(\mathbf{Q}_2 \in U) \\ \iff \mathbb{P}_{yz}(\mathbf{Q}_2 \notin U | \mathbf{Q}_1 \in U) &\leq \mathbb{P}_{yz}(\mathbf{Q}_2 \notin U) \\ \iff \mathbb{P}_{yz}(\mathbf{Q}_2 \notin U | \mathbf{Q}_1 \notin U) &\geq \mathbb{P}_{yz}(\mathbf{Q}_2 \notin U) \\ \iff \mathbb{P}_{yz}(\mathbf{Q}_2 = \mathbf{b}_1 | \mathbf{Q}_1 = \mathbf{a}_2) &\geq \mathbb{P}_{yz}(\mathbf{Q}_2 = \mathbf{b}_1) = 1/2, \end{aligned}$$

implying the claim. The bound is tight because association allows for independence.

Lemma 3.1 Our assumptions immediately imply that

$$\lim_{n \rightarrow \infty} \mathbb{P}\left(\pi_{yz} - \psi_{yz} \geq \widehat{\pi}_{yz} - \widehat{\psi}_{yz} - \Phi^{-1}(1 - \alpha)\widehat{\sigma}_{sum}\phi_n^{-1/2}\right) = 1 - \alpha.$$

Let $\pi_{yz} - \psi_{yz} \geq 0$, then

$$\begin{aligned} &\lim_{n \rightarrow \infty} \mathbb{P}\left(\max\{\pi_{yz} - \psi_{yz}, 0\} \geq \max\left\{\widehat{\pi}_{yz} - \widehat{\psi}_{yz} - \phi_n^{-1/2}c_{1-\alpha}^l, 0\right\}\right) \\ &= \lim_{n \rightarrow \infty} \mathbb{P}\left(\pi_{yz} - \psi_{yz} \geq \widehat{\pi}_{yz} - \widehat{\psi}_{yz} - \phi_n^{-1/2}c_{1-\alpha}^l\right) \\ &= 1 - \alpha. \end{aligned}$$

If $\pi_{yz} - \psi_{yz} < 0$, then

$$\begin{aligned}
& \lim_{n \rightarrow \infty} \mathbb{P} \left(\max\{\pi_{yz} - \psi_{yz}, 0\} \geq \max\left\{\widehat{\pi}_{yz} - \widehat{\psi}_{yz} - \phi_n^{-1/2} c_{1-\alpha}^l, 0\right\} \right) \\
&= \lim_{n \rightarrow \infty} \mathbb{P} \left(\widehat{\pi}_{yz} - \widehat{\psi}_{yz} - \phi_n^{-1/2} c_{1-\alpha}^l \leq 0 \right) \\
&> \lim_{n \rightarrow \infty} \mathbb{P} \left(\widehat{\pi}_{yz} - \widehat{\psi}_{yz} - \phi_n^{-1/2} c_{1-\alpha}^l \leq \pi_{yz} - \psi_{yz} \right) \\
&= 1 - \alpha.
\end{aligned}$$

Lemma 3.2 We only show the first claim, the proof for the second one is very similar. We distinguish between three cases.

Case 1: $\pi_{yz} = 1 - \psi_{yz}$.

We then have

$$\begin{aligned}
& \lim_{n \rightarrow \infty} \mathbb{P} \left(\min\{\pi_{yz}, 1 - \psi_{yz}\} \leq \min\left\{\widehat{\pi}_{yz}, 1 - \widehat{\psi}_{yz}\right\} + \phi_n^{-1/2} c_{1-\alpha}^u \right) \\
&= \lim_{n \rightarrow \infty} \mathbb{P} \left(\min\{\widehat{\pi}_{yz} - \pi_{yz}, 1 - \widehat{\psi}_{yz} - (1 - \psi_{yz})\} \geq -\phi_n^{-1/2} c_{1-\alpha}^u \right) \\
&= \lim_{n \rightarrow \infty} \mathbb{P} \left(\phi_n^{1/2} \min\{\widehat{\pi}_{yz} - \pi_{yz}, 1 - \widehat{\psi}_{yz} - (1 - \psi_{yz})\} \geq -c_{1-\alpha}^u \right) \\
&= \lim_{n \rightarrow \infty} \mathbb{P} \left(\phi_n^{1/2} \max\{\widehat{\pi}_{yz} - \pi_{yz}, \psi_{yz} - \widehat{\psi}_{yz}\} \leq c_{1-\alpha}^u \right) \\
&= 1 - \alpha,
\end{aligned}$$

where the first step uses that $\pi_{yz} = 1 - \psi_{yz}$, and the last steps use distributional assumptions as before.

Case 2: $\pi_{yz} < 1 - \psi_{yz}$.

In this case, we have $c_{1-\alpha}^u = \widehat{\sigma}_{\pi_{yz}} \Phi^{-1}(1 - \alpha)$ with probability approaching one, and

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(\min\{\pi_{yz}, 1 - \psi_{yz}\} \leq \min\left\{\widehat{\pi}_{yz}, 1 - \widehat{\psi}_{yz}\right\} + \phi_n^{-1/2} c_{1-\alpha}^u \right) = \lim_{n \rightarrow \infty} \mathbb{P} \left(\pi_{yz} \leq \widehat{\pi}_{yz} + \phi_n^{-1/2} c_{1-\alpha}^u \right) = 1 - \alpha.$$

Case 3: $\pi_{yz} > 1 - \psi_{yz}$.

This case is entirely similar to case 2.

Note finally that the pre-test will not be consistent along local alternatives where $\pi_{yz} + \psi_{yz} - 1 = O(\phi_n^{-1/2})$, but it is easy to see that the confidence regions will in fact be conservative in those cases.

Lemma 3.3 Coverage properties of $CI_{1-\alpha}(\theta)$ depend on where θ_{yz} lies in Θ_{yz} . Let $\theta_{yz} = \max\{\pi_{yz} - \psi_{yz}, 0\}$, thus θ_{yz} assumes its lower bound, then

$$\begin{aligned}
& \lim_{n \rightarrow \infty} \mathbb{P}(\theta_{yz} \in CI_{1-\alpha}(\theta)) \\
&= \lim_{n \rightarrow \infty} \mathbb{P}\left(\max\left\{\widehat{\pi}_{yz} - \widehat{\psi}_{yz} - \phi_n^{-1/2}c_{1-\alpha}^l, 0\right\} \leq \max\{\pi_{yz} - \psi_{yz}, 0\} \leq \min\left\{\widehat{\pi}_{yz}, 1 - \widehat{\psi}_{yz}\right\} + \phi_n^{-1/2}c_{1-\alpha}^u\right) \\
&\geq \lim_{n \rightarrow \infty} \mathbb{P}\left(\max\left\{\widehat{\pi}_{yz} - \widehat{\psi}_{yz} - \phi_n^{-1/2}c_{1-\alpha}^l, 0\right\} \leq \max\{\pi_{yz} - \psi_{yz}, 0\} \leq \min\left\{\widehat{\pi}_{yz}, 1 - \widehat{\psi}_{yz}\right\}\right) \\
&\geq \lim_{n \rightarrow \infty} \mathbb{P}\left(\begin{array}{l} \max\left\{\widehat{\pi}_{yz} - \widehat{\psi}_{yz} - \phi_n^{-1/2}c_{1-\alpha}^l, 0\right\} \leq \max\{\pi_{yz} - \psi_{yz}, 0\}, \\ \max\{\pi_{yz} - \psi_{yz}, 0\} \leq \max\{\pi_{yz} - \psi_{yz}, 0\} + \varepsilon + \min\left\{\widehat{\pi}_{yz}, 1 - \widehat{\psi}_{yz}\right\} - \min\{\pi_{yz}, 1 - \psi_{yz}\} \end{array}\right) \\
&= \lim_{n \rightarrow \infty} \mathbb{P}\left(\begin{array}{l} \max\left\{\widehat{\pi}_{yz} - \widehat{\psi}_{yz} - \phi_n^{-1/2}c_{1-\alpha}^l, 0\right\} \leq \max\{\pi_{yz} - \psi_{yz}, 0\}, \\ \min\{\pi_{yz}, 1 - \psi_{yz}\} - \min\left\{\widehat{\pi}_{yz}, 1 - \widehat{\psi}_{yz}\right\} \leq \varepsilon \end{array}\right) \\
&= \lim_{n \rightarrow \infty} \mathbb{P}\left(\max\left\{\widehat{\pi}_{yz} - \widehat{\psi}_{yz} - \phi_n^{-1/2}c_{1-\alpha}^l, 0\right\} \leq \max\{\pi_{yz} - \psi_{yz}, 0\}\right) \\
&\geq 1 - \alpha.
\end{aligned}$$

Here, the first inequality uses that $c_{1-\alpha}^u > 0$. To see the second inequality, write $\min\left\{\widehat{\pi}_{yz}, 1 - \widehat{\psi}_{yz}\right\} = \min\{\pi_{yz}, 1 - \psi_{yz}\} + \min\left\{\widehat{\pi}_{yz}, 1 - \widehat{\psi}_{yz}\right\} - \min\{\pi_{yz}, 1 - \psi_{yz}\}$ and observe that

$$\min\{\pi_{yz}, 1 - \psi_{yz}\} - \max\{\pi_{yz} - \psi_{yz}, 0\} = \min\{\pi_{yz}, \psi_{yz}, 1 - \pi_{yz}, 1 - \psi_{yz}\} \geq \varepsilon,$$

thus $\min\{\pi_{yz}, 1 - \psi_{yz}\} \geq \max\{\pi_{yz} - \psi_{yz}, 0\} + \varepsilon$ by assumption. The very last step was established before.

Lemma 3.4 Let $\pi_{yz} - \psi_{yz} \geq 0$, then $\widehat{\pi}_{yz} - \widehat{\psi}_{yz} > -b_n$ with probability approaching one, and

$$\begin{aligned}
& \lim_{n \rightarrow \infty} \mathbb{P}(\Theta_{yz} \subseteq CI_{1-\alpha}(\Theta)) \\
&= \lim_{n \rightarrow \infty} \mathbb{P}\left(\max\left\{\widehat{\pi}_{yz} - \widehat{\psi}_{yz} - \phi_n^{-1/2}c^l, 0\right\} \leq \pi_{yz} - \psi_{yz}, \min\{\pi_{yz}, 1 - \psi_{yz}\} \leq \min\left\{\widehat{\pi}_{yz}, 1 - \widehat{\psi}_{yz}\right\} + \phi_n^{-1/2}c^u\right) \\
&= \lim_{n \rightarrow \infty} \mathbb{P}\left(\widehat{\pi}_{yz} - \widehat{\psi}_{yz} - \phi_n^{-1/2}c^l \leq \pi_{yz} - \psi_{yz}, \min\{\pi_{yz}, 1 - \psi_{yz}\} \leq \min\left\{\widehat{\pi}_{yz}, 1 - \widehat{\psi}_{yz}\right\} + \phi_n^{-1/2}c^u\right) \\
&= \lim_{n \rightarrow \infty} \mathbb{P}\left(\phi_n^{1/2}\left(\widehat{\pi}_{yz} - \widehat{\psi}_{yz} - (\pi_{yz} - \psi_{yz})\right) \leq c^l, \phi_n^{1/2}\left(\min\{\widehat{\pi}_{yz}, 1 - \widehat{\psi}_{yz}\} - \min\{\pi_{yz}, 1 - \psi_{yz}\}\right) \geq -c^u\right) \\
&\geq \lim_{n \rightarrow \infty} \mathbb{P}\left(\phi_n^{1/2}\left(\widehat{\pi}_{yz} - \widehat{\psi}_{yz} - (\pi_{yz} - \psi_{yz})\right) \leq c^l, \phi_n^{1/2}\left(\min\{\widehat{\pi}_{yz} - \pi_{yz}, \psi_{yz} - \widehat{\psi}_{yz}\}\right) \geq -c^u\right) \\
&= 1 - \alpha.
\end{aligned}$$

where the steps are either algebraic or justified in earlier proofs.

Let $\pi_{yz} - \psi_{yz} < 0$, then $\hat{\pi}_{yz} - \hat{\psi}_{yz} \leq -b_n$ with probability approaching one, and

$$\begin{aligned} & \lim_{n \rightarrow \infty} \mathbb{P}(\Theta_{yz} \subseteq CI_{1-\alpha}(\Theta)) \\ &= \lim_{n \rightarrow \infty} \mathbb{P}\left(\min\{\pi_{yz}, 1 - \psi_{yz}\} \leq \min\{\hat{\pi}_{yz}, 1 - \hat{\psi}_{yz}\} + \phi_n^{-1/2} c_{1-\alpha}^u\right) \\ &\geq 1 - \alpha, \end{aligned}$$

where the last step was shown before.

References

- [1] Afriat, S.N., 1973. On a System of Inequalities in Demand Analysis: An Extension of the Classical Method, *International Economic Review*, 14, 460-472.
- [2] Andrews, D.W.K., and P. Jia, 2008. Inference for Parameters Defined by Moment Inequalities: A Recommended Moment Selection Procedure. Cowles Foundation Discussion Paper 1676.
- [3] Andrews, D.W.K., and G. Soares, 2007. Inference for Parameters Defined by Moment Inequalities Using Generalized Moment Selection. Cowles Foundation Discussion Paper 1631.
- [4] Bandyopadhyay, T., I. Dasgupta, and P.K. Pattanaik, 2002. Demand Aggregation and the Weak Axiom of Stochastic Revealed Preference. *Journal of Economic Theory*, 107, 483-489.
- [5] —, 2004. A General Revealed Preference Theorem for Stochastic Demand Behavior. *Economic Theory*, 23, 589-599.
- [6] Becker, G.S., 1962. Irrational Behavior and Economic Theory. *Journal of Political Economy*, 70, 1-13.
- [7] Blundell, R., M. Browning, and I. Crawford, 2003. Nonparametric Engel Curves and Revealed Preference. *Econometrica*, 71, 205-240.
- [8] Blundell, R., M. Browning, and I. Crawford, 2008. Best Nonparametric Bounds on Demand Responses. *Econometrica*, 76, 1227-1262.
- [9] Chernozhukov, V., S. Lee, and A. Rosen, 2009. Intersection Bounds: Estimation and Inference. CEMMAP Working Paper 19/09.
- [10] Cherchye, L., Crawford, I., De Rock, B., and F. Vermeulen, 2009. The Revealed Preference Approach to Demand. In D. Slottje (ed.), *Quantifying Consumer Preferences*. Emerald Books.

- [11] Deaton, A., and J. Muellbauer, 1980. An Almost Ideal Demand System. *American Economic Review*, 70, 312-26.
- [12] de Castro, L.I., 2007. Affiliation and Positive Dependence in Auctions. Discussion paper, University of Illinois, Urbana-Champaign.
- [13] Esary, J.D., F. Proschan, and D. W. Walkup, 1967. Association of Random Variables, with Applications. *Annals of Mathematical Statistics*, 38, 1466-1474.
- [14] Fan, Y. and S.S. Park, 2009. Sharp Bounds on the Distribution of Treatment Effects and their Statistical Inference, *Econometric Theory*, forthcoming.
- [15] Heckman, J.J., Smith, J.A., and N. Clements, 1997. Making the Most Out of Programme Evaluations and Social Experiments: Accounting for Heterogeneity in Programme Impacts. *Review of Economic Studies*, 64, 487–536
- [16] Horowitz, J.L., and C.F. Manski, 1995. Identification and Robustness with Contaminated and Corrupted Data. *Econometrica*, 63, 281-302.
- [17] Imbens, G., and C.F. Manski, 2004. Confidence Intervals for Partially Identified Parameters. *Econometrica*, 72, 1845-1857.
- [18] Lewbel, A., 1999. Consumer Demand Systems and Household Expenditure. In H. Pesaran and M. Wickens (eds.), *Handbook of Applied Econometrics*. Basil Blackwell.
- [19] Manski, C.F., 1997. The Mixing Problem in Programme Evaluation. *Review of Economic Studies*, 64, 537–553.
- [20] —, 2003. *Partial Identification of Probability Distributions*. Springer Verlag.
- [21] —, 2007. Partial Identification of Choice Probabilities. *International Economic Review*, 48, 1393-1410.
- [22] Matzkin, R., 2006. Heterogeneous Choice. In R. Blundell, W. Newey, and T. Persson (eds.) *Advances in Econometrics: Proceedings of the 9th World Congress*. Cambridge University Press.
- [23] Marschak, J., 1960. Binary Choice Constraints on Random Utility Indicators. In K. Arrow (ed.) *Stanford Symposium on Mathematical Methods in the Social Sciences*. Stanford University Press.

- [24] McFadden, D., 2005. Revealed Stochastic Preference: A Synthesis. *Economic Theory*, 26, 245-264.
- [25] McFadden, D., and M. Richter, 1990. Stochastic Rationality and Revealed Stochastic Preference. In J. Chipman, D. McFadden, and M. Richter (eds.) *Preferences, Uncertainty and Optimality: Essays in Honour of Leo Hurwicz*. Westview Press.
- [26] Nelsen, R.B., 2006. *An Introduction to Copulas (2nd Edition)*. Springer Verlag.
- [27] Nelsen, R.B., J.J. Quesada-Molina, J.A. Rodríguez-Lallena, and M. Úbeda-Flores, 2001. Bounds on Bivariate Distribution Functions with Given Margins and Measures of Association. *Communications in Statistics – Theory and Methods*, 30, 1155-1162.
- [28] Nelsen, R.B. and M. Úbeda-Flores, 2004. A Comparison of Bounds on Sets of Joint Distribution Functions Derived from Various Measures of Association. *Communications in Statistics – Theory and Methods*, 33, 2299-2305.
- [29] Rosenbaum, P.R., 2002. *Observational Studies (2nd Edition)*. Springer Verlag.
- [30] Shao, J., and D. Tu, 1995. *The Jackknife and Bootstrap*. Springer Verlag.
- [31] Stoye, J., 2009. More on Confidence Intervals for Partially Identified Parameters. *Econometrica*, 77, 1299-1315.
- [32] Varian, H., 1982. The Nonparametric Approach to Demand Analysis , *Econometrica*, 50, 945-974.
- [33] —, 1983. Nonparametric Tests of Consumer Behaviour, *Review of Economic Studies*, 50, 99-110.
- [34] Yanagimoto, T., 1972. Families of Positively Dependent Random Variables. *Annals of the Institute of Statistical Mathematics*, 24, 559-573.

Marginal Distributions of Preferences Across Budget Sets




Figure 1

Distribution of Quantity-Weighted Price Changes




Figure 2

Distribution of Upper Bound




Figure 3

Upper Bound and Associated Lower Conf. Int.




Figure 4

Distribution of Lower Bound




Figure 5

Distribution of Upper Bound with PQD

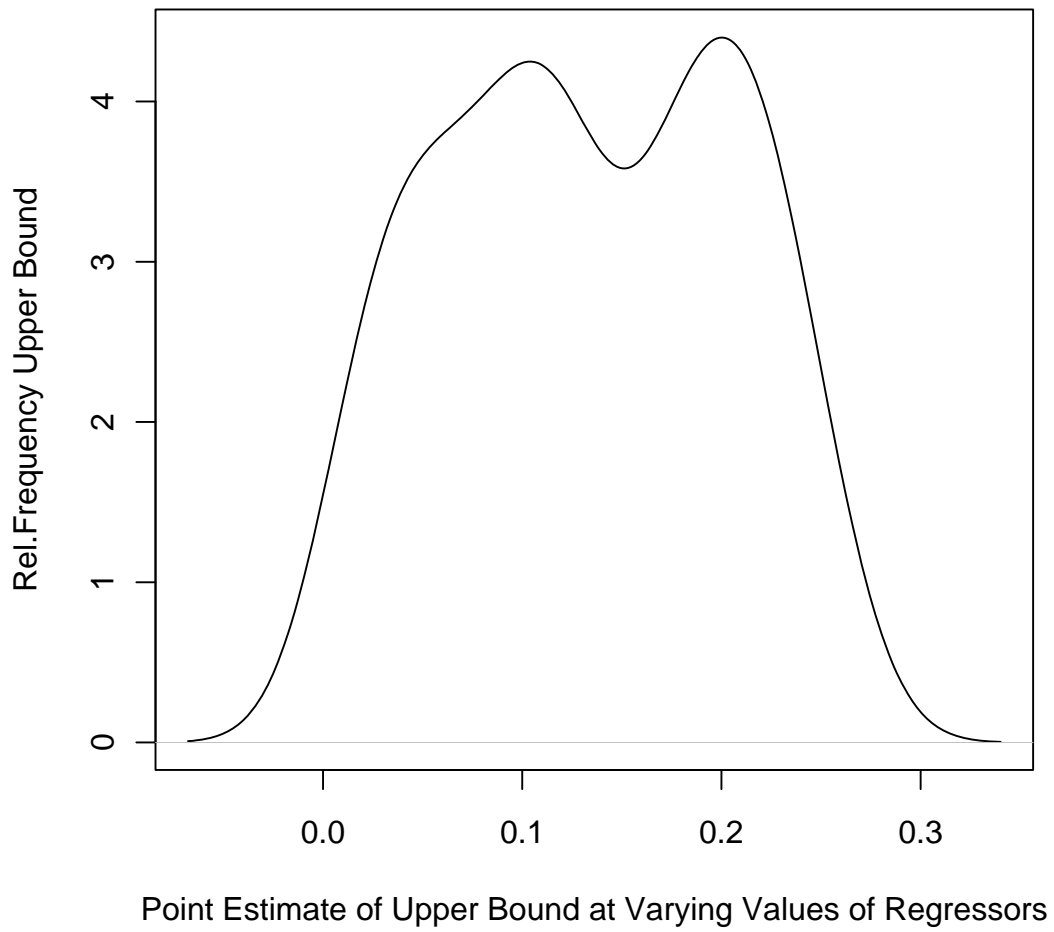


Figure 6

Distribution of Upper Bound - Comparison w/o PQD

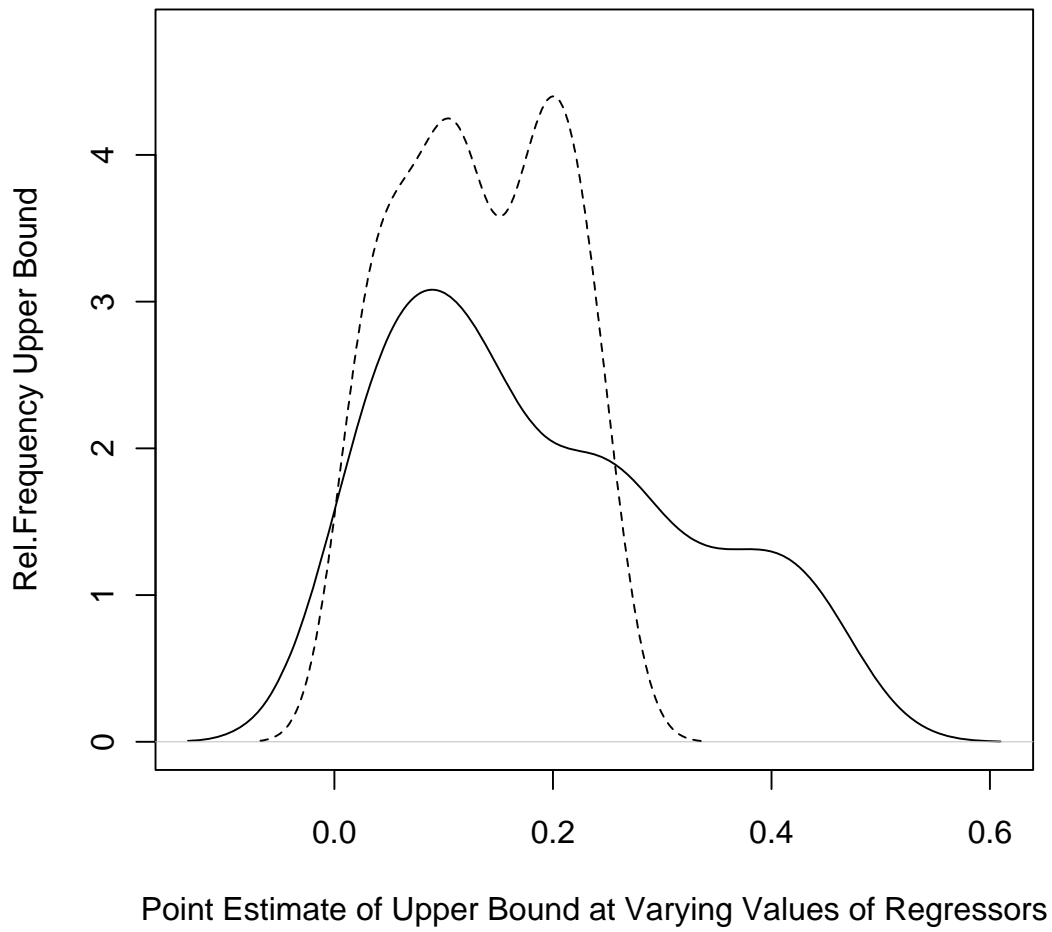


Figure 7