HEDG Health, Econometrics and Data Group

THE UNIVERSITY of York

HEDG Working Paper 09/18

Quantile Regression Analysis of the Rational Addiction Model: Making unobservable heterogeneity observable

Audrey Laporte

Alfia Karimova

Brian Ferguson

July 2009

ISSN 1751-1976

york.ac.uk/res/herc/hedgwp

Quantile Regression Analysis of the Rational Addiction Model: Making unobservable heterogeneity observable

Audrey Laporte, PhD*
Department of Health Policy, Management and Evaluation
University of Toronto
Toronto, Ontario, Canada
e-mail: audrey.laporte@utoronto.ca

Alfia Karimova, MA
Department of Economics
University of Guelph
Guelph, Ontario, Canada
e-mail: akarimov@uoguelph.ca

Brian Ferguson, PhD
Department of Economics
University of Guelph
Guelph, Ontario, Canada
e-mail: brianfer@uoguelpoh.ca

***corresponding author**

I. Introduction

The Rational Addiction (RA) model assumes that individual decisions about the consumption of harmful and addictive commodities are made on a rational basis (Becker and Murphy, 1988). In this context, rational means forward looking, i.e. a tendency to take account of future consequences of current consumption decisions. Different individuals may well attach different weights to the present relative to the future. The degree to which an individual is forward looking in her consumption decisions is revealed not by her current consumption level but rather by the time path of her consumption of an addictive commodity. Hence, the need to estimate a forward looking second order difference equation (SODE) as part of the process of testing the RA model. Most studies using micro level data estimate a single SODE for the whole sample. This involves estimating an average propensity to be forward looking for the entire sample, even when it is believed that different fully rational individuals in the same sample may have different propensities to be forward looking. Forward looking behaviour is an aspect of treating the consumption of an addictive commodity as part of an inter-temporal optimization problem. Inter-temporal optimization is characterized by what are known as saddle point dynamics and the information about an individual's propensity to be forward looking is contained in what are known as the characteristic roots of the equation (Ferguson, 2003).

In a sample of heterogeneous individuals we expect propensity to be forward looking to differ across individuals and the best way to identify these differences is by looking at the dynamic behaviour of the individual consumption paths. Estimating a common SODE for everyone hides this key difference.

In this paper, we make the argument that the best place to look for differences in individual propensities to be forward looking is in dynamic behaviour considered at different points in the distribution of the consumption of an addictive commodity. To do this we adopt techniques of Quantile Regression, (QR) estimating RA type difference equations in consumption across quantiles of cigarette consumption. We use panel data to ensure that we are examining the behaviour of individuals across time. Our hypothesis is that we will find differences in the degree of forward looking behaviour characterizing the time paths of consumption across quantiles in the micro-level data.

II. Theoretical Context

The Becker-Murphy model of rational addiction is an individual-level model of intertemporal optimization, where the commodity being consumed happens to be addictive to some degree or another (Becker and Murphy, 1988). As with virtually all such models, its solution involves saddlepoint dynamics - i.e. the consumption trajectory for the addictive commodity is driven by two characteristic roots, one stable and one unstable. Empirically, the RA model is implemented as some variant on the basic form

$$(1) \qquad C_t = \alpha_0 + \alpha_1 C_{t-1} + \alpha_2 C_{t+1} + \alpha_3 P_t + \varepsilon_t$$

While the theoretical RA model is a model of individual, not market, behaviour, it is probably still safe to say that the majority of the empirical literature has involved estimating RA models on aggregate level data of some sort - national level, state level or provincial level, because to estimate it on individual level data requires that one has panel data, and until recently in most countries, suitable panel data sets were scarce. Increasingly, however, official health and expenditure surveys are taking longitudinal form, so that data are becoming available at a level of observation which corresponds to the level of observation of the theory. Recent papers which have used individual level panel data include Baltagi and Geishecker (2006) and Jones and Labeaga (2003).

Estimated on longitudinal individual level data, RA econometrics clearly falls under the heading of dynamic panel data (DPD) econometrics (for a general overview, see Arellano 2003). DPD issues are typically discussed in the context of a backward looking first order difference equation such as:

$$(2) \qquad C_t = \alpha_0 + \alpha_1 C_{t-1} + \alpha_3 P_t + \varepsilon_t$$

It is typically assumed that there are differences between individuals' tastes (i.e. differences between individuals' utility functions) which mean that, when faced with exactly the same set of prices, different individuals will make different decisions about the quantity of any given commodity which they want to consume. That is the reason that OLS regression, when run on individual level data (without a lagged dependent variable), can yield very low R^2 values and very high t-statistics. In such a case, all individuals might be strongly responsive to changes in price, say, yielding high t-statistics, but differences in preferences mean that individual consumption levels are scattered widely around the mean, and since the OLS line fits the conditional mean, the R^2 will tend to be low.

In fitting dynamic relations on individual data an additional consideration arises. Expositions of DPD typically assume that individual heterogeneity in tastes can be represented by a fixed term γ_i, where i indexes the individual. Then (2) becomes

$$(3) \qquad C_{it} = \alpha_0 + \alpha_1 C_{it-1} + \alpha_3 P_{it} + \gamma_i + \varepsilon_t$$

where we would leave the i subscript off the P term if we are dealing with the case where all individuals face the same price in period t. It is typically assumed that the taste heterogeneity term is not independently observable, especially in the case where the number of individuals in the panel, N, is very large relative to the number of observations on each individual, T, so that it is not feasible to include individual intercept terms. Neglecting it, however, creates an omitted variable bias problem for dynamic modeling. If we assume that γ is positively related to consumption of C, so that high values of γ indicate a preference for C and low values a dislike for C, and we assume that γ is indeed constant over time - people's fundamental preferences do not change during the period spanned by the longitudinal data set, then a high γ today means that γ was also high yesterday, and will have caused both C_{it} and C_{it-1} to be high (i.e. individual i will have tended to consume large quantities of C) while someone with a low γ will have consumed low quantities of C both yesterday and today. Thus even if $\alpha_1 = 0$, so that there is no habit formation - an increase in C_{it-1} does not automatically lead to an increase in C_{it} in the data set, because of the effect of γ_i, there will tend to be a positive association between C_{it} and C_{it-1} across individuals, so that the estimate of the coefficient α_1 will tend to be upward biased.

Clearly, in the RA model with unobservable heterogeneity in tastes,

$$(4) \qquad C_{it} = \alpha_0 + \alpha_1 C_{it-1} + \alpha_2 C_{it+1} + \alpha_3 P_t + \gamma_i + \varepsilon_{it}$$

the same issue arises, but this time a high value of γ which does not change over time for the individual will tend to push all three of C_{it}, C_{it-1} and C_{it+1} up while a low value will tend to push all three down, with the result that we would expect the omitted variable bias effect to bias the estimates of both α_1 and α_2 upward. In that case, unobservable heterogeneity would tend to make a variable whose current consumption level was completely independent of past and future consumption appear to be fitting the forward looking, rational addiction pattern. Presumably the same effect would tend to bias upward the coefficients on lead and lag consumption for a commodity which was rationally addictive.

In the DPD literature on first order difference equations, it is typically assumed that individual heterogeneity, at least with regards to taste for commodity C, can adequately be represented by an individual-specific constant γ_i term. Given that assumption, the most common DPD approach is probably to take the equation explaining consumption of C in period t-1:

$$(5) \qquad C_{it-1} = \alpha_0 + \alpha_1 C_{it-2} + \alpha_3 P_{it-1} + \gamma_i + \varepsilon_{t-1}$$

and subtract it from equation (3) above, yielding

$$(6) \qquad \Delta C_{it} = \alpha_1 \Delta C_{it-1} + \alpha_3 \Delta P_{it} + \Delta\varepsilon_t$$

where the overall intercept and the unobserved heterogeneity term are removed by the differencing. This approach raises problems of its own, however, since $\Delta C_{it-1} = C_{it-1} - C_{it-2}$ and $\Delta\varepsilon_t = \varepsilon_{it} - \varepsilon_{t-1}$. Since, from equation (5), C_{it-1} depends on ε_{t-1}, the moving average error term in (6) is correlated with one of the RHS explanatory variables, and we have an endogeneity problem. DPD analysis then proceeds to instrument ΔC_{it-1}, usually with lagged values of ΔC_{it} and of C_{it} which do not overlap the MA error term. In the RA case, differencing would give

$$(7) \qquad \Delta C_{it} = \alpha_1 \Delta C_{it-1} + \alpha_2 \Delta C_{it+1} + \alpha_3 \Delta P_{it} + \Delta\varepsilon_t$$

where, since $\Delta C_{it+1} = C_{it+1} - C_{it}$, and C_{it} depends on ε_{it} and therefore on $\Delta\varepsilon_t$ we have two RHS variables which are correlated with the MA error term, although in different ways -

ΔC_{it+1} is positively correlated with the ε_{it} part of $\Delta\varepsilon_t$ while ΔC_{it-1} is negatively correlated with the ε_{t-1} part. Since the endogenous RHS variables in this transformed equation are simply differently-timed versions of each other, we would presumably use the same variables to instrument both. Jones and Labeaga (2003) apply this approach to a Spanish data set.

As stated earlier, it is common, when running OLS equations on large micro data sets, to obtain high t-statistics and very low R^2 values. Again, this is because an OLS equation fits the conditional mean of the data, but individual observations may be scattered widely about the conditional mean. Thus in Figure 1 below we have a set of individuals all facing the same prices and income, and hence the same budget line representing the choice between two commodities, C and X, but choosing very different optimal consumption points along that budget line. An OLS equation representing the demand for one of the two commodities on this diagram, estimated on panel data from this set of individuals, would not necessarily do a particularly good job of predicting the behaviour of any single individual in response, say, to changes in the price they all face, but could still do a good job representing the behaviour of the mean of the distribution of their consumption choices.

Each individual in Figure 1 will be consuming at a point of tangency between the common budget line and her own, individual, indifference curve between the two commodities. While this means that, at each individual's consumption point, her indifference curve must have a slope equal to the ratio of the prices of the two commodities, the fact that each has chosen a different (C, X) combination along that budget line means that there are significant differences in their marginal rates of substitution between the two commodities, measured at a common (C, X) point. This, of course, means that there will be differences between the marginal utilities (MU) they derive from consumption of the two goods.

In the case of most commodities, differences in individual's MUs simply reflect different degrees of satisfaction derived from consumption of a certain number of each of the two goods.

Figure 1. Individual Preferences

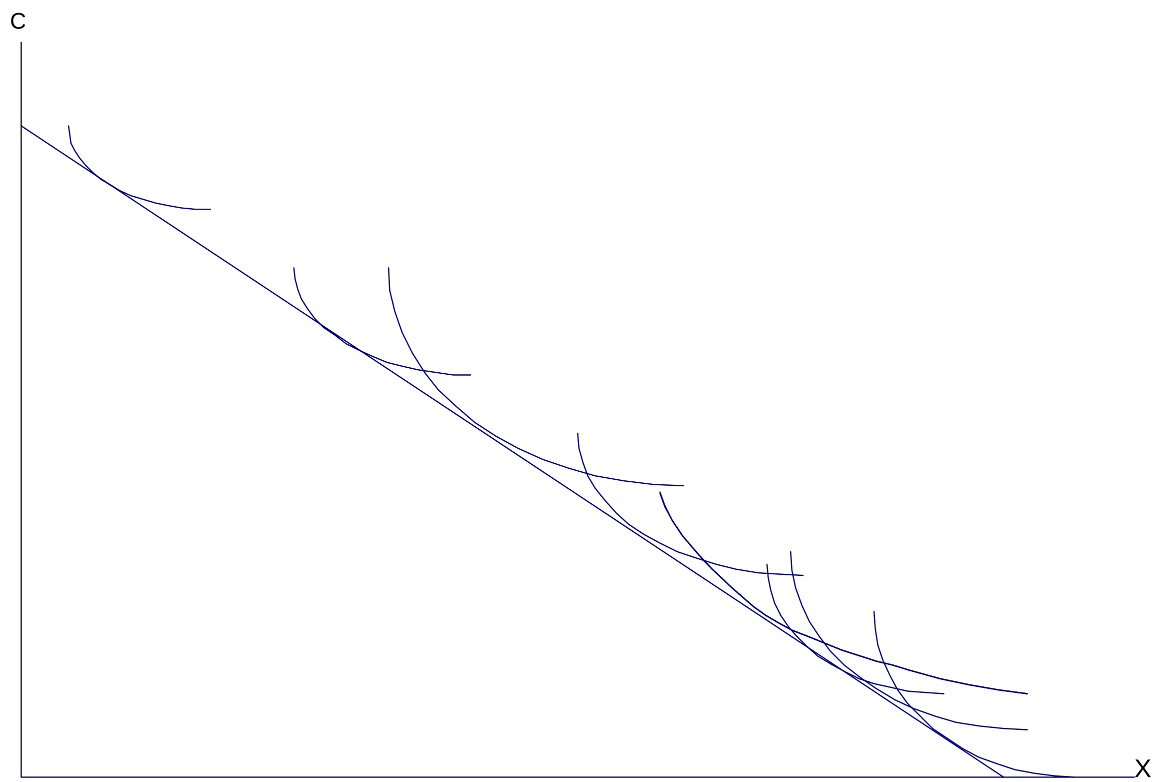

The matter is slightly different, however, in the case of a commodity like cigarettes which has the property that current consumption results in future damage to the individual's health, so that the choice of current consumption involves a tradeoff between the utility derived from current consumption and the disutility derived from future consequences.

When the commodity is, like cigarettes, addictive, there is the additional consideration that increased consumption today will tend to lead to further increased consumption in the future, with consequences for further future damage to health. When using the single period indifference curve diagram, then, we can think of the marginal utility of the consumption of cigarettes today as containing two components - the satisfaction derived from the current consumption adjusted by a measure of the current utility value of the future disutility which today's consumption will lead to. If we consider two individuals, each of whom derives the same utility from units of commodity X and each of whom derives the same immediate satisfaction from the act of consuming C, they will still choose different optimal points along the budget line if they have different assessments of the disutility they will derive from future health damage, of the probability of suffering health damage (since, for example, by no means do all smokers develop lung cancer), of their individual susceptibility to addiction, or if they simply have different rates of time preference. And so, even if they agree on all of the other aspects of the future consequences of their current choices, they may put different weights on future events. All of these factors enter into the degree to which their consumption of cigarettes will tend to be forward looking, and we can generally say that more forward looking individuals will place greater weight on any or all of these factors and will therefore have a lower current net MU from smoking than would a less forward looking individual who derives exactly the same satisfaction from the current act of consuming cigarettes. It is worth emphasizing that differences in the degree to which an individual is forward looking, or more precisely differences in the degree to which they net future disutility out of current utility from smoking, depends on more than simply differences in individual discount rates - differences between individual assessments of the riskiness of smoking, broadly defined, also enter the calculation. In any event, a more forward looking individual will have a lower MU from smoking today than will a less forward looking individual, will thus have a larger MRS: MU(X)/MU(C) compared to a less forward looking individual at identical values of C and X and will therefore choose a consumption pair involving more X and less C than will a less forward looking individual. When we are considering individual decisions about smoking, then, individual heterogeneity includes individual differences in the degree of forward looking behaviour.

III. Methodological Approach

The standard explanation of DPD analysis works best when we are concerned with the possibility that individual heterogeneity might make a commodity which involves no habit formation at all look as if current consumption does in fact depend on past consumption. If the true value of the autoregressive coefficient is zero, it makes sense to try and sweep individual heterogeneity out of the equation, when estimating the effect of, say, prices. It is also a reasonable approach when we believe that there is habit formation and the autoregressive term is roughly the same for everyone, so that the mean value of the AR term adequately represents all of the individuals in the sample. It should be remembered, though, that sweeping out individual differences is not an end in itself, rather it is an approach to compensating for the omitted variable bias which results from the fact that the individual heterogeneity is unobservable. An alternative approach would be to try to make the unobservable heterogeneity observable.

If we had a larger T relative to N than the typical panel data set has, we would often tackle this simply by including individual intercepts. Since this is not possible in our case (as it is not in most panel data studies) we propose an alternative approach. We take as our starting point the proposition that individual heterogeneity is the primary reason for the scatter of tangency points along the budget line in Figure 1, and that individual heterogeneity represents fundamental differences in individuals' preference structures, so the scatter will be long lasting. Since we cannot introduce individual fixed effects given our data set, we instead analyze the distribution of cigarette consumption using QR (see Koenker 2005**)** rather than OLS-type regression. QR has been used before in the addictions literature, by Manning, Blumberg and Moulton (1995) to investigate the question of whether the price elasticity of demand for alcohol varies with the quantity consumed. While OLS regression yields an estimate of the equation explaining the conditional mean of the distribution of consumption, QR yields estimates of equations characterizing the conditional quantiles of the distribution of consumption as functions of the explanatory variables. This allows us to investigate whether the response to changes in explanatory variables at a high quantile of the distribution of consumption differs significantly from behaviour associated with a low quantile, as seems quite plausible. It also lets us test whether the distribution of consumption changes shape, or simply shifts, as the values of the explanatory variables change[1].

It is sometimes suggested that the results of QR should be characterized as explaining the shape of the distribution and not as characterizing the behaviour of individuals at various points on the distribution. It is not clear, however, why this should

[1]Much of the literature dealing with QR in a panel data context focuses on methods of estimating fixed effects(Koenker, 2004**)**. Our quantile equation, being derived from the RA model, is a forward looking second order difference equation. The theoretical literature on QR of autoregressive processes (see, for example, Koenker and Xiao, 2006) deals with first order difference equations.

be any more of a problem for QR than for OLS, which we tend to think of as characterizing the behaviour of the average individual, or as characterizing average behaviour when it actually characterizes the mean of the distribution of individuals. If there are systematic differences between the equations associated with different quantiles, we would infer that there are significant differences between the behaviour and therefore the tastes of the individuals who are located at those quantiles. It is certainly true that if the mix of individuals in the sample changes significantly the interpretation of the estimated conditional quantile equations may have to be changed. Suppose, for example, that the size of the data set was increased by the entry of a large number of people whose daily cigarette consumption was extremely low. Then the location of the quantiles of the distribution would change - most of the quantiles would now be located at lower levels of cigarette consumption, and a smoker who had been at the median of the distribution, and whose behaviour had been adequately characterized by the conditional median equation, would now be at one of the higher quantiles, not because her behaviour had changed but because the location of the median of the distribution had slid out from underneath her. In this case, though, the location of the conditional mean would also be expected to have changed. This is not a concern with our data set, however, since we use only observations on individuals who were present in all of the data periods. In the case of the demand for cigarettes, changing attitudes towards smoking would also change the shape of the distribution of consumption over time and thus cause the meaning of the quantiles - their location in terms of absolute number of cigarettes smoked - to change. This would be a concern if we had, for example, data running back into the 1960s and 70s, but our time span (2000/01-2006/07) is too short for that kind of mass change in tastes to be a significant consideration. The assumption, then, is that individuals' fundamental preferences remain constant through our data period.

The advantage of QR, as we noted above, is that it gives us information about the impact of changes in explanatory variables on the consumption behaviour of people who have chosen to place themselves at various points in the distribution of cigarette consumption. It is not uncommon in the literature for micro-econometric studies of cigarette demand to separate the data into subsets consisting of heavy and light smokers and estimate demand functions on those subsets separately, and we noted above the Manning et al. study which considered price elasticities of demand across quantiles of alcohol consumption. We are more interested, however, in a different factor - whether the degree of forward looking behaviour changes over the quantiles of consumption.

Standard DPD analysis aims to sweep out the unobservable individual heterogeneity which might be biasing the AR coefficient in a first order difference equation and yield an unbiased estimate of that coefficient. This presumes that people at different points along the distribution of consumption have intertemporal consumption trajectories which obey the same difference equations. While this might be an acceptable assumption in the case of a first order difference equation, in the forward looking RA structure as we have already noted, the point at which an individual's indifference curve is tangent to her budget line in any period will depend in part on the degree to which she is forward looking. A more forward looking individual will tend to be lower down, along the C axis, and further out, along the X axis, than will a less forward looking individual.

Consider the RA second order difference equation $C_{it} = \alpha_0 + \alpha_1 C_{it-1} + \alpha_2 C_{it+1}$. We write the characteristic equation for this SODE as $-\alpha_2 \lambda^2 + \lambda - \alpha_1 = 0$, which we rearrange to read $\lambda^2 - 1/\alpha_2 \lambda + \alpha_1 / \alpha_2 = 0$.[2] Since we are dealing with an optimization problem, the SODE will display saddlepoint behaviour, with one unstable (i.e. larger than one) characteristic root and one stable (i.e. less than one) root (both roots will be real and positive). One approach to solving a saddlepoint difference equation involves applying what is known as the backward solution approach to the stable root and the forward solution approach to the unstable root[3]. In the backward solution approach, the stable root is applied to past shocks which affect the current value of the dependent variable while in the forward solution process, the inverse of the unstable root is applied to future values of the explanatory variables. Thus it is said that the inverse of the unstable root reflects the effect of future shocks on current consumption while the stable root reflects the effect of past shocks[4]. Jones and Labeaga (2003) refer to the backward looking effect as the strength of the addiction effect and the forward looking term as strength of the forward looking element in consumption. The advantage of QR is that it allows for an investigation of whether the strength of the addictive effect, and of the forward looking effect, vary across quantiles of consumption.

In addition to using QR methodology to allow for individual heterogeneity, we also make use of the fact that decisions about consumption of individual commodities are not made in isolation from decisions about consumption of other commodities. The individual's preference structure will be reflected in their consumption decisions about all of the commodities in the consumption basket, not just consumption of the commodity

[2] The roots calculated from this form of characteristic equation are the inverse of those often reported in the literature. Essentially the difference is between the way econometricians calculate roots and the way economic theorists calculate them. This is why the econometrics literature often refers to stability requiring both roots be outside the unit circle while the theoretical literature requires both to be inside the unit circle. It is not a fundamental issue, but can cause confusion.

[3] See, for example, Gandolfo(1997)

[4] See Frank Chaloupka (1991) bearing in mind that his calculated roots are the inverses of ours, so that what he refers to as the smaller of the two roots is the inverse of the larger of our roots.

whose demand function is being estimated. In standard textbook discussions of consumption theory, this is reflected in the fact that the demand function for a single commodity is written as a function of the prices of the entire set of consumption commodities. In empirical applications, we tend to simplify by replacing the vector of prices of other commodities by a single price index, with the CPI or some other share-weighted price index. There is, however a body of empirical work (Pollack and Wales, 1992) which uses information on the entire vector of an individual's purchases (or at least a significant subset, invoking multi-stage budgeting arguments) to determine the form of her consumption preferences. In the expenditure share literature, Seemingly Unrelated Regression techniques are used to estimate an interrelated system of equations where the share of the consumer's expenditure which is allocated to each good is estimated as a function of the prices all of the commodities in the budget set. System estimation methods have been applied to addictive commodities such as alcohol and tobacco, and Pollack and Wales discuss variants of the demand system approach which allow for habit formation, although not for the forward looking element of the RA model.

One drawback to using the demand system approach in the case of longitudinal micro data is that it is quite possible that all of the individuals living in a single geographic area - a Canadian province, for example - may face the same price for some commodities or the data may be limited to a single, province wide price index for a category of commodities so that while we can see how the price of a commodity changes over time within a province, we cannot observe regional differences in price within a province. Even with panel data if T is small relative to N and if the number of geographic areas is small - again in the Canadian case, the subjects will be living in one of only ten provinces - we may have very little variation in price. Moreover, it is not price or relative prices which give us information about individual preferences, it is the set of consumption choices made by individuals who all face the same price. In other words, individual heterogeneity is represented not by the slope of the budget line but by the point on identical budget lines at which different people choose to consume. It seems reasonable, therefore, to hypothesize that by including quantities of other commodities consumed as explanatory variables in a dynamic model to be estimated using individual panel data we may be able to make the heterogeneity observable[5]. Clearly this raises its own endogeneity issues, since the quantities of the other consumption variables will be chosen simultaneously with, and subject to the same budget constraint as, the dependent variable, but at the very least it raises the possibility that we will have a larger set of instrumenting variables available. In the estimated equations, then, we include values of other elements of the individual's consumption bundle as explanatory variables (i.e. fruits and vegetables). Since we will be estimating the demand for cigarettes, we also include variables representing relevant aspects of the individual's state of health, since these can be seen as factors affecting her individual preferences for healthy and unhealthy commodities.

The basic structure, then, is a Becker-Murphy RA model augmented with variables characterizing the individuals' health status, their educational attainment, their family circumstances, their age, whether they face workplace restrictions on smoking, whether anyone in their household smokes in the house, their income, their education level and their choices with regard to purchases of fruits and vegetables.

---

[5] Jones and Labeaga (2003) also make use of individual spending on other commodities – in their case they use expenditure data to personalize share-weighted consumer price indices for deflation purposes.

IV. Data

The data are drawn from several cycles of the Canadian National Population Health Survey (NPHS). The NPHS is taken every two years, and has been longitudinal since the 1994/95 cycle. We initially sub-set the data to ensure that we only had people who remained in the panel for all of the cycles, then reduced it further to remove absolute non-smokers, on the assumption that their individual preferences were fundamentally different from those of people who had at one time been smokers. While the smoking equation was asked through all of the cycles of the survey, we were forced to reduce the time span because other questions which we wanted to use to derive explanatory variables from - questions about consumption of certain other commodities, for example - were only asked in a few of the cycles. Ultimately we were able to use only the 2002/3 and 2004/5 cycles as the core of our analysis (i.e. as sources for C_t data), with C_{t-1} data also being drawn from the 2000/01 cycle and C_{t+1} data from the 2006/7 cycle.

Within this reduced sample, we excluded anyone who reported zero consumption in all of the four cycles from 2000/1 to 2006/7, using only individuals who reported positive consumption of cigarettes in at least one cycle. This left 4148 observations, two values of C_t on each individual, of which a significant number were still zero. As a consequence, we include dummy variables for whether the individual has attempted to quit within the last 6 months and whether they classify themselves as former daily (frmrdaily) or former occasional (frmrocc) smokers. We include these last two to control for the effect of people who have chosen to go cold turkey, rather than to follow a smooth intertemporal consumption path. We also include a dummy variable for people who identify themselves as nonsmokers (nonsmoker) but who report positive consumption in any of the data periods.

The 25th percentile of the unweighted distribution of cigarette consumption in the data set was one cigarette. Figure 2 below shows the number of daily cigarettes associated with the various quantiles of the distribution of consumption.

Figure 2

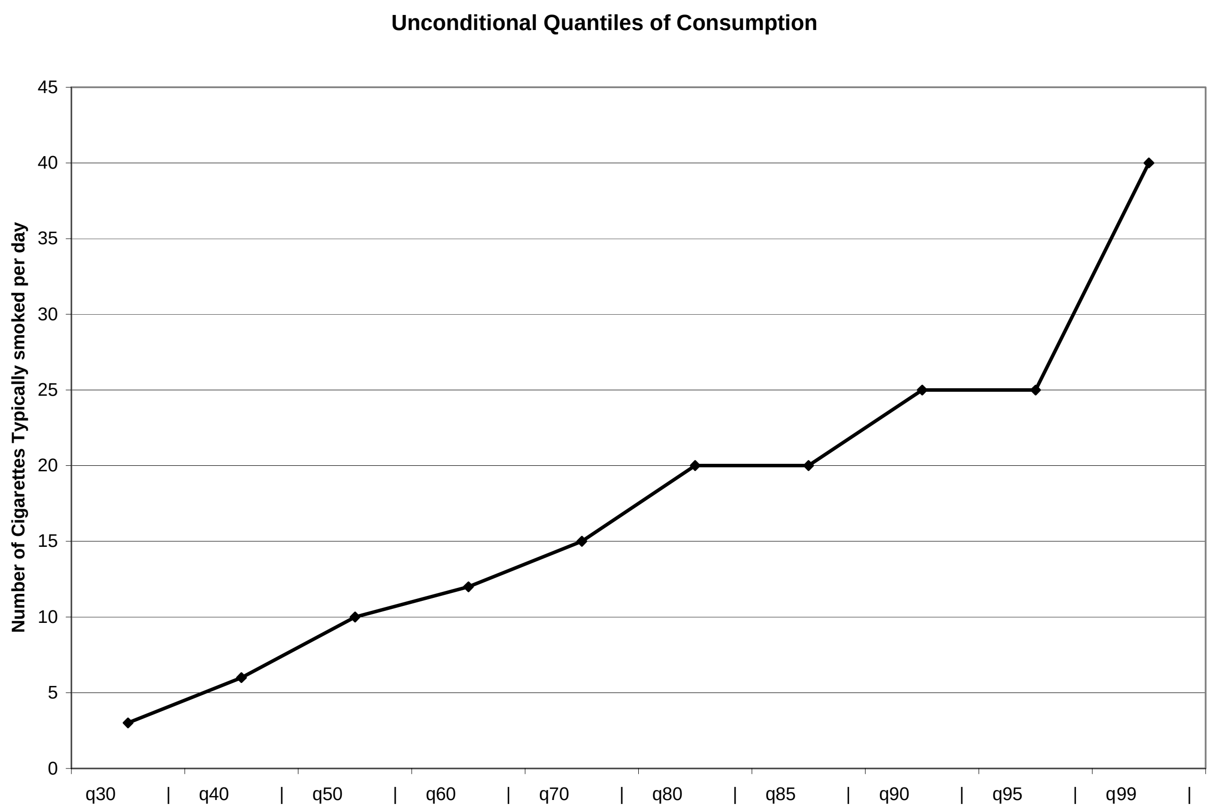

The NPHS is not a diary survey, so individuals were not asked to report actual consumption of cigarettes on any particular day. Rather they were asked how many cigarettes they would smoke on a typical day. In essence, their answer is their expected daily cigarette consumption. Because of the nature of the cigarette question and the fact that the observations on an individual's typical daily consumption are taken two years apart, it seemed unlikely that there would be a statistical link between the error term in the C_t equation and the value of C_{t+1} for any individual, so we did not instrument future consumption.

The price data were based on the nominal price of a carton of 200 cigarettes, drawn from the Non-Smokers Rights Association of Canada[6], combined with provincial level data on the Consumer Price Index for cigarettes for the individual provinces, since nominal price data were not available for all of the periods in the sample. The dynamic structure of the model is represented by the presence of lead and lag cigarette consumption. We also include, in most of the equations, lead (leadpt) and lag (lagpt) cigarette prices along with the current price (pt)[7].

Among other explanatory variables we included dummy variables for level of education attained (highsc, postsec, univcoll), dummies for age group (age 1519, age2039, age4059 (base), age6079, age80pl), and dummies for ranges of income (inclt20, inc2039, inc4079(base), inc80pl). We also included data on family circumstances - married and the presence of children under 12 years in the household (kids) - and a dummy for the sex of the respondent (male). We included dummies for a number of indicators of health status - whether the respondent uses asthma medication (astmed), whether they are on blood pressure medication (bpmeds), whether they have had a heart attack (attack), whether they have diabetes (diab) and whether they find life stressful (stress), and a variable indicating whether, overall, their self-assessed health status is poor (healthpoor). We include a dummy variable for whether they have attempted to quit smoking within the last 6 months(quit), another for whether they face restrictions on

[6] http://www.nsra-adnf.ca/cms/file/pdf/cigarette_prices_Canada_17_April_2009.pdf

[7] This is consistent with the earliest version of the RA model, although it is more common in the literature for researchers to include only the current price. See Becker, Grossman and Murphy (1990).

smoking in the workplace (workres) and one for whether any member of their household smokes in the house (smhouse).

Among other consumption variables we include variables for weekly consumption of fruits, juice, salad, potatoes, carrots and vegetables generally, and a variable indicating whether the respondent assesses his dietary habits to be poor (eatpoor)[8]. We also include a variable for the quantity of alcohol which the individual consumes (drinks). In the NPHS, this variable was measured in a manner which came close to being a diary method, as respondents were asked how much alcohol they had consumed on each of the previous seven days. We aggregated this variable into a weekly alcohol consumption variable. Again, had the NPHS been a diary survey we would have faced issues of endogeneity between expenditure (and hence quantity) of alcohol and the various food items. Because the cigarette consumption question asked about typical, rather than precise daily, consumption, however, it seemed unlikely that there would be a statistical endogeneity problem with regard to the food and drink questions, so we did not instrument those variables.

V. Results

Table 1 below reports the results of OLS estimation of the full RA model. In this case, the lead and lag consumption coefficients satisfy the sum and product conditions for real roots and saddlepoint behaviour, and the roots of the difference equation are 2.67 and .371. In solving a saddlepoint equation, we apply what is known as the forward solution process to the unstable root and the backward solution process to the stable root. In the

[8]In some cycles the NPHS asks about consumption of milk and fish and of soft drinks. Unfortunately those questions were not asked in the cycles we were using.

forward solution process, the inverse of the unstable root is applied to future values of the explanatory variables. Thus it is said that the inverse of the unstable root reflects the effect of future shocks on current consumption while the stable root reflects the effect of past shocks. The inverse of the unstable root from Table 1 is 0.375 and the stable root is 0.371, so by Chaloupka's measure, the addiction and forward looking effects are of similar magnitude. The hypothesis that the coefficient on lead consumption will be less than that on lagged consumption is rejected, but since this is generally the case in the RA literature, our results are at least not outliers. It is also the case, as we noted above, that a range of different elements of the individuals preferences and views about the probability of different outcomes, and not just pure time preference, determine the degree to which the individual is forward looking, so it is not clear to what degree this should be taken as a test of the argument that cigarette consumption decisions are forward looking.

Looking at the other variables in Table 1, we see that the coefficient on the current price is negative, that on lag price is positive and that on lead price is negative. The sum of the three is positive. Among the other explanatory variables, the fact that someone in the family (and this may refer to the respondent) smokes in the house has a significant positive effect on current cigarette consumption, having attempted to quit in the past is associated with a lower cigarette consumption, facing restrictions on smoking at work reduces consumption and being male significantly increases consumption. The coefficients on the education variables are consistent with expectations, with higher education being associated with smoking less, but only the "university and college" dummy (univcoll) attains significance. Interestingly, both younger and older people smoke less than do the reference group, the 40-59 age group; this may reflect a combination of a generational effect in the younger age groups (the middle age group would have entered the smoking ages when smoking was still socially acceptable in Canada) and health effects in the older groups (with possible reverse causality of non-smokers live longer than smokers).

Interestingly enough, the health variables are generally non-significant, although they do tend to have negative coefficients. The only one which comes close to significant is being diabetic, and that has a positive coefficient. The income variables are non-significant, a result which carries over into the QR estimation, and, oddly, the "kids" variable has a positive coefficient, although it falls short of significance.

Among the dietary choice variables, alcohol consumption tends to be positively associated with the quantity smoked while increased consumption of fruit juice and carrots are negatively associated with average quantity smoked. Oddly enough, increased consumption of potatoes is associated with increase consumption of cigarettes[9]. The variable indicating that the respondent said that they had a poor diet is positively and significantly associated with the quantity of cigarettes smoked, possibly indicating a degree of self-awareness which would not be inconsistent with consumption choices being deliberately made.

Next we proceed to QR analysis of the RA model. Because QR yields an estimating equation, and standard errors, for each quantile investigated, the results of QR are usually presented in graphical form. For the most part we will follow that convention. However, Table 2 below reports the QR coefficients, bootstrap standard errors and t-statistics for lead and lag consumption and current, lead and lag prices.

---

[9]The NPHS survey instructs respondents not to count French fries or potato chips as servings of potatoes, but there is no indication as to how closely that instruction is followed.

Because the bottom twenty percent of the distribution of consumption is zeroes (the twenty-fifth percentile is one cigarette) we report results from the 25^{th} percentile up.

We note that the price effects in the tables above do not seem to be indicating downward sloping demand curves, and when we test for significance of the long run price effects (not reported here) we generally cannot reject the null that they are zero. We will return to this point below.

Figure 3 shows the estimated intercepts for the quantile equations. In this figure, as in the other coefficient graphs, the solid centre line shows the coefficient estimates while the dashed lines on either side of the centre line show the 95% confidence interval for the estimates. Here we see the estimated intercepts rising as we go from the lowest to the highest quantiles, as we would expect. The horizontal scale of the diagram is slightly, though not seriously, misleading since we generally work in increments of ten percentiles, but after the 80^{th} percentile report the 85^{th}, 90^{th}, 95^{th} and 99^{th} percentiles. In addition, Figure 3 includes a horizontal dotted line, which shows the single intercept estimated in the OLS equations reported above. As we would expect, the OLS intercept is above those for the lower quantiles and below those for the upper quantiles.

Figures 4 and 5 below show the coefficients on lag and lead consumption. We note that the coefficient on lagged consumption tends to increase over the quantiles while that on lead consumption tends to decrease as we move to higher quantiles. Figure 6 below shows the calculated characteristic roots from the OLS and QR versions of the equations, while Figure 7 shows the smaller root and the inverse of the larger root from the QR equations. Following Chaloupka (1991) and Jones and Labeaga (2003) in interpreting these values as the strengths of the addiction effect and of the forward looking effect respectively, we see that the forward looking effect is strongest in the lower quantiles of cigarette consumption and decreases as we move to higher quantiles while the addiction effect is weaker in the lower quantiles, becoming stronger at higher quantiles of cigarette consumption. This would mean that the dynamics described by the SODE derived from the RA model are consistent with intuition, suggesting that heavier smokers are less forward looking than are lighter smokers.

Figure 8 shows the coefficients for the alcohol consumption variable “drinks”. Interestingly enough, given that cigarettes and alcohol are often taken to be complementary commodities, the QR coefficients alternate between being positive and negative and are never statistically significantly different from zero, while the OLS coefficient is positive and, as seen in the table above, would be significant at the 5% level in a 1 tail test, and is significant at the 10% level in a two-tailed test. Figure 9 depicts the coefficient on juice consumption, showing that at all quantiles an increase in regular consumption of juice is associated with reduced cigarette consumption, and that the effect is statistically significant from the 60th percentile up. Figure 10 shows the coefficient on consumption of potatoes to be positive and significant at several quantiles (and nearly so at the rest - significant were we looking at one-tailed tests instead of two-tailed tests). Figure 11 depicts the effect of a self-assessed poor diet to be positive and generally significant or nearly so, virtually constant across quantiles and hence little different from the corresponding OLS coefficient. Figure 12 shows that the effect of being male is non-significant at the bottom quantiles, but that it increases as we move to the higher quantiles of cigarette consumption, becoming significant at all but the top quantile.

The health-state variables had less impact on consumption than we had expected:

Figure 13 demonstrates that the effect of being in self-assessed poor health is generally negative but never significant, while Figure 14 shows odd, but generally non-significant effects across quantiles from having had a heart attack. Figure 15 portrays the coefficient on a variable indicating that life is stressful: while it rises over the quantiles it is generally non-significant.

Figure 16 shows that reporting having attempted to quit is associated with a reduction in the number of cigarettes smoked, although the effect is significant only in the lower quantiles, while Figure 17 shows that facing workplace restrictions on smoking has a negative effect, but that it is significant only at the upper quantiles, quite possibly because smokers in the lower quantiles were not smoking at work in any event.

Figure 18 below shows the effect of having a university or college education on cigarette consumption across quantiles. Consistent with the OLS results, this coefficient was negative and significant at the median and below.

VI. Discussion

The QR results show a mixed set of effects of explanatory variables, in a few cases quite different from the effects shown by OLS regression. The workplace restrictions variable, for example, which has a negative and significant coefficient in the OLS regression, shows a negative effect across all but the lowest quantile, statistically significant only from the 70th percentile up, but the QR results show that workplace restrictions have a much larger effect on the upper quantiles of consumption than the OLS coefficient would suggest. The health variables proved unexpectedly weak in the regressions, and income had a non-significant effect throughout. While the individual fruit and vegetable variables did not do much in any of the equations, the self-assessed poor diet variable was associated with higher cigarette consumption, with an effect that was roughly constant across quantiles. It is possible that this variable is more informative about people's general dietary habits than are variables that relate to a few specific items (e.g. carrots, salad).

As we noted above, we included in our equations dummy variables for individuals who made discrete jumps from being daily or occasional smokers to zero consumption – cold turkey dummies. The coefficients on these variables were, not surprisingly, consistently negative and statistically significant. We did this on the assumption that an individual who went cold turkey was different from one who followed a smooth trajectory, even a trajectory which was heading towards zero consumption. Such an individual might, for example, have discovered that they had unexpected health problems and had gone cold turkey on their doctor's advice. We dummied these individuals out because we wanted the coefficients on lead and lag consumption, and therefore the roots of the SODE characterizing their behaviour, to reflect the behaviour of individuals who were following a trajectory which represented the solution to their forward looking optimal control problem – the type of trajectory which the cold turkey individuals might have followed had they not, for example, had a health shock. Since going cold turkey is likely to be associated with having a health shock, we investigated the relation between these dummies and our health status variables, and found that there was indeed a correlation, and that when we experimented with dropping these dummies from the equations, several of the health variables became significant. Our preferred specification, however, includes the cold turkey dummies, so the health variables lose significance in it.

We noted above that the price variable had a persistent tendency to have the wrong sign and that the long run price effect was generally non-significant. This is, of course, inconsistent with market level studies, which generally find price to have a strong negative effect on cigarette consumption. There are a number of possible explanations for this - one is the lack of variability in the price series. We were forced to assume that everyone living in any one of Canada's ten provinces in a given year faced the same price which, especially if true, means that there is not much variation in the price series within the data set. This is a common problem in Canadian micro-econometric cigarette studies[10]. Another problem may be associated with the self-assessed nature of the consumption data. People show a tendency to respond in "round" numbers: numbers ending in 0 or 5. Combining that with the fact that cigarettes can only be bought in packs, not individually, at least not on the legal market for which we have price data, this would tend to reduce the measured effect of price on consumption. We also note that, while we excluded from our data set "never smokers" and people who quit smoking before the four NPHS cycles from which we drew our data, there are still a significant number of zeros in our data set, associated with people who quit smoking during the span of years covered by the cycles which we are using. This may also help explain the non-significance of the price variable.

We also found less effect from the education variables than we might have anticipated. Having a university or college (univcoll) education has a negative but weak effect across the quantiles, despite having been significant in the OLS equation. One possible explanation is that education has its effect on the decision whether or not to be a

---

[10] See, for example, Beatty (2008)

non-smoker, which we have not modeled here. Another is that the effect of having a college education requires a sample of 4000 to detect, but if that is the case then it must be a very weak effect.

An alternative possibility, which might be worth exploring, is that, as some of the smoking literature has suggested, education has only a weak effect on smoking behaviour, and that the observed negative association between smoking and education is a result of the smoking decision and the education decision both being functions of the degree to which an individual is forward looking. A more forward looking individual would tend to smoke less (either be a non-smoker or be in the bottom quantiles of smoking) and also to invest more in higher education. If the direct effect of education is weak, then most of the effect of education in individual-level smoking equations may capture the effect of differences in the degree to which individuals are forward looking – i.e. proxying that variable. In an OLS regression, which estimates only one set of roots for the entire sample, variations in education might proxy for differences in individual forward looking propensities from the average. In the QR equations, we allow the propensity to be forward looking to vary across equations, which may leave the education variables with only their direct effect to reflect, which would tend to weaken their effect.

The results on the strength of the addiction effect and of the propensity to be forward looking seem reasonable, with the lower quantiles showing a tendency to put more weight on the future than do the upper quantiles. If this result holds up in further research, it may help to explain some of the odder results found in the empirical RA literature[11].

[11] We refer here to the well known tendency of RA models to yield implausible values for the discount factor.

The failure of some of the explanatory variables to have a significant effect raises the question of whether QR alone might be sufficient to control for individual heterogeneity, without the need to add other variables. To investigate this, we ran a bare bones RA model, regressing current quantity of cigarettes only on lag and lead quantity and on current price. We calculated the characteristic roots for both the OLS and QR cases, and report the addiction and forward looking effects in Table 3 below:

Table 3

Quantile	OLS stable root	OLS unstable root	QR stable root	QR unstable root
Q25	0.472	1.55	0.206	1.44
Q30	0.472	1.55	0.225	1.304
Q40	0.472	1.55	0.365	1.157
Q50	0.472	1.55	0.6	1
Q60	0.472	1.55	complex roots	complex roots
Q70	0.472	1.55	0.634	1.456
Q80	0.472	1.55	0.46	1.88
Q85	0.472	1.55	0.756	1.57
Q90	0.472	1.55	0.734	1.71
Q95	0.472	1.55	0.846	1.45
Q99	0.472	1.55	1	1.25

with the corresponding addiction and forward looking effects in Table 4:

Table 4

Quantile	OLS addiction effect	OLS forward effect	QR addiction effect	QR forward effect
Q25	0.472	0.645	0.224	0.692
Q30	0.472	0.645	0.225	0.767
Q40	0.472	0.645	0.365	0.864
Q50	0.472	0.645	0.6	1
Q60	0.472	0.645	complex roots	complex roots
Q70	0.472	0.645	0.634	0.687
Q80	0.472	0.645	0.46	0.532
Q85	0.472	0.645	0.756	0.635
Q90	0.472	0.645	0.734	0.583
Q95	0.472	0.645	0.846	0.688
Q99	0.472	0.645	1	0.8

The estimated roots and effects show patterns similar to those from the full models, but are less well behaved. In the case of the 60th quantile, the roots are a complex conjugate pair, at the fiftieth quantile we have a unit root as we do at the 99th percentile, but in the latter case it is the smaller root which is the unit root. While not a formal test, these results seem to suggest that omitting the other explanatory variables may well bias the estimates of the roots.

Our results, and those of Baltagi and Geishecker (2006) and Jones and Labeaga (2003) suggest that RA models estimated on micro panel data yield quite different results from models estimated on market or aggregate level time series data. In particular, the micro level studies seem to show much larger differences between the stable and unstable roots, and much more definitive evidence of saddlepoint behaviour. Aggregate level studies are much more likely to manifest unit root behaviour (Laporte, 2006) and seem to suffer from all of the problems which are associated with unit root variables in the macrodynamic literature. It has been customary in the literature to test for RA simply by looking at whether the coefficients on lead and lag consumption are positive. While this is a necessary condition, it is not sufficient - in the case in Table 3 above where the roots are complex, for example, the coefficients on lead and lag consumption are both positive. Optimization at the individual level implies saddlepoint behaviour, and testing for saddlepoint behaviour requires at the very least testing that the sum of the coefficients on lead and lag consumption is significantly less than one and for preference, calculating the roots. Those aggregate level studies which find one unit root and one stable root probably should not, despite having positive coefficients on lead and lag consumption, be taken as evidence of RA behaviour.

If the difference between the results of micro and aggregate level panel studies holds up under further investigation, it should be interpreted as meaning that a representative agent approach is not informative when it comes to modeling market-level cigarette data. While an understanding of the intertemporal behaviour of the agents who make up the market is necessary for an understanding of market behaviour - it gives us an idea as to what variables should be included in a market-level cointegrating equation, for example - we cannot assume that the dynamic behaviour of the market reflects the dynamic behaviour of each individual in the market. The difference between individual and market level dynamic results suggests that aggregation conditions and problems must be taken seriously if we want to understand how individual level behaviour drives market level behaviour.

Overall, the results are consistent with the RA model, in that we find evidence of saddlepoint dynamics, and that the forward looking effect is weaker among the higher quantiles of cigarette consumption. Clearly some of the results call for further investigation, most notably our failure to find a significant negative price effect. Nevertheless, we believe the results do suggest that QR is a useful tool for understanding

RA dynamics at the level of the individual.

VII. Conclusion

In this paper we attempt to illustrate the application of QR to the RA model using micro-level panel data. We show that QR can be used to deal with unobserved heterogeneity. In particular, differences in the degree to which individual consumers of cigarettes are forward looking. QR methods offer a fruitful alternative to DPD methods when it comes to investigating RA models because of the information they provide about differences in the degree of forward looking behaviour and forward looking behaviour is the key to rationality in the consumption of addictive commodities.

References

Arellano, Manuel (2003): Panel data econometrics: Oxford University Press, Oxford and New York

Baltagi, Badi H; Geishecker, Ingo (2006): “Rational Alcohol Addiction: Evidence from the Russian Longitudinal Monitoring Survey” *Health Economics*, 15 (9) September, pp. 893-914

Beatty, Timothy K. M (2008): “Expenditure Dispersion and Dietary Quality: Evidence From Canada” *Health Economics* 17: 1001-1014

Becker, Gary S; Murphy, Kevin M (1988) A Theory of Rational Addiction, *Journal of Political Economy*, vol. 96, no. 4, pp. 675-700.

Becker, Gary S., Grossman, Michael and Murphy, Kevin M. (1990): "An empirical analysis of Cigarette Addiction", *National Bureau of Economic Research Working Paper* No. 3322, April

Chaloupka Frank (1991) “Rational Addictive Behavior and Cigarette Smoking” *Journal of Political Economy* 99(4), 722-742.

Ferguson, Brian S. and G.C. Lim (2003): Dynamic Economic Models in Discrete Time: Theory and Empirical Applications, Routledge, London

Gandolfo, Giancarlo (1997): Economic Dynamics Springer-Verlag Berlin-Heidelberg

Jones, Andrew M; Labeaga, Jose M (2003): “Individual Heterogeneity and Censoring in Panel Data Estimates of Tobacco Expenditure” *Journal of Applied Econometrics*, 18 (2) March-April pp. 157-77

Koenker, Roger (2004): “Quantile regression for longitudinal data”, *Journal of Multivariate Analysis,* 91 (1).

Koenker, Roger (2005): Quantile Regression, Econometric Society Monographs, no. 38. Cambridge University Press, Cambridge and New York:

Koenker, Roger; Xiao, Zhijie (2006): “Quantile Autoregression” *Journal of the American Statistical Association*, 101 (475) September pp. 980-90

Laporte, Audrey (2006) “Price responsiveness of demand for cigarettes: Does rationality matter?”, *Substance Use and Misuse*, Vol. 41(4):511-531.

Manning, Willard G; Blumberg, Linda; Moulton, Lawrence H (1995): “The Demand for Alcohol: The Differential Response to Price” *Journal of Health Economics*, 14 (2), June

Pollak, Robert A. and Terence J. Wales: Demand system specification and estimation
Oxford University Press, Oxford, 1992

Table 1: OLS RA Model

Variable	Coef.	Std. Err.	t	P>|t|	[95% Conf.	Interval]
lagct	0.326298	0.010384	31.42	0	0.305939	0.346656
leadct	0.3291	0.010779	30.53	0	0.307968	0.350232
pt	-0.03669	0.0230887	-1.59	0.112	-0.08195	0.008576
lagpt	0.061229	0.014929	4.10	0	0.031961	0.090497
leadpt	-0.01335	0.014676	-0.91	0.363	-0.04213	0.01542
smhouse	2.436083	0.192794	12.64	0	2.058102	2.814063
quit	-0.9622	0.188635	-5.1	0	-1.33203	-0.59238
workres	-0.41437	0.19568	-2.12	0.034	-0.79801	-0.03074
drinks	0.026913	0.015726	1.71	0.087	-0.00392	0.057744
juice	-0.31082	0.100769	-3.08	0.002	-0.50838	-0.11326
fruit	-0.07094	0.099403	-0.71	0.475	-0.26583	0.123939
salad	-0.24641	0.233399	-1.06	0.291	-0.704	0.211177
potato	0.518319	0.251221	2.06	0.039	0.02579	1.010848
carrot	-0.44698	0.265884	-1.68	0.093	-0.968254	0.0743
veg	-0.01751	0.107816	-0.16	0.871	-0.22888	0.193872
eatpoor	0.622377	0.20454	3.04	0.002	0.221368	1.023385
kids	0.164988	0.119839	1.38	0.169	-0.06996	0.399937
married	-0.09615	0.184112	-0.52	0.602	-0.45711	0.264808
male	0.3645	0.186155	1.96	0.05	-0.00046	0.729465
highsc	0.246902	0.201567	1.22	0.221	-0.14828	0.642083
postsec	-0.3027	0.215381	-1.41	0.16	-0.72496	0.119568
univcoll	-0.50356	0.220498	-2.28	0.022	-0.93585	-0.07126
age1519	-0.97932	0.449629	-2.18	0.029	-1.86084	-0.0978
age2039	-0.69821	0.202534	-3.45	0.001	-1.09528	-0.30113
age6079	-0.70479	0.286331	-2.46	0.014	-1.26616	-0.14343
age80pl	-2.07138	1.041958	-1.99	0.047	-4.11419	-0.02858
inclt20	-0.30504	0.254746	-1.2	0.231	-0.80448	0.194405
inc2039	-0.23718	0.223657	-1.06	0.289	-0.67567	0.201305
inc80pl	0.138723	0.430305	0.32	0.747	-0.70491	0.982355
healthpoor	-0.35753	0.27054	-1.32	0.186	-0.88793	0.172876
astmed	-0.30942	0.335345	-0.92	0.356	-0.96688	0.348037
bpmeds	-0.24916	0.290444	-0.86	0.391	-0.81859	0.320268
diab	0.695119	0.418905	1.66	0.097	--0.12616	1.516399
attack	-0.55662	0.625854	-0.89	0.374	-1.78363	0.670395
stress	0.212177	0.193574	1.1	0.273	-0.16733	0.591687
frmrdaily	-8.42488	0.246723	-34.15	0	-8.90859	-7.94117
frmrocc	-5.54235	0.631549	-8.78	0	-6.78053	-4.30417
nonsmoker	-5.49602	1.193733	-4.6	0	-7.83638	-3.15566
_cons	5.485306	0.83267	6.59	0	3.852821	7.117791
R^2 =0.7116	F(38, 4109) =266.80					

Table 2 Price and Consumption results from complete Quantile RA model

Quantile	Variable	Coefficient	Bootstrap Std. Error	t-statistic
q25	lagct	0.252629	0.013747	18.38
	leadct	0.350578	0.021512	16.30
	pt	-0.04119	0.023007	-1.79
	lagpt	0.049704	0.014027	3.54
	leadpt	0.01702	0.013969	1.22
q30	lagct	0.262825	0.017598	14.93
	leadct	0.376771	0.018381	20.50
	pt	-0.05144	0.02297	-2.24
	lagpt	0.053002	0.013481	3.93
	leadpt	0.019736	0.075916	1.18
q40	lagct	0.290092	0.015636	18.55
	leadct	0.397003	0.018386	21.59
	pt	-0.03115	0.020835	-1.5
	lagpt	0.043503	0.013052	3.33
	leadpt	0.003302	0.013807	0.24
q50	lagct	0.313533	0.016393	19.13
	leadct	0.417866	0.018195	22.97
	pt	-0.01698	0.019532	-0.87
	lagpt	0.048056	0.013465	3.57
	leadpt	-0.0142	0.013409	-1.06
q60	lagct	0.322044	0.018708	17.21
	leadct	0.415264	0.02022	20.54
	pt	-0.02258	0.020085	-1.12
	lagpt	0.050002	0.013996	3.57
	leadpt	-0.02429	0.013421	-1.81
q70	lagct	0.338853	0.022175	15.28
	leadct	0.377617	0.023202	16.27
	pt	-0.02442	0.021481	-1.14
	lagpt	0.052283	0.015902	3.29
	leadpt	-0.02444	0.01498	-1.63

Quantile	Variable	Coefficient	Bootstrap Std. Error	t-statistic
q80	lagct	0.364286	0.023441	15.54
	leadct	0.321665	0.024054	13.37
	pt	-0.02731	0.026018	-1.05
	lagpt	0.055858	0.016782	3.33
	leadpt	-0.03122	0.016733	-1.87
q85	lagct	0.369309	0.025081	14.72
	leadct	0.315925	0.023126	13.66
	pt	-0.02203	0.031816	-0.69
	lagpt	0.045414	0.021023	2.16
	leadpt	-0.02892	0.019318	-1.5
q90	lagct	0.3885092	0.0308919	12.58
	leadct	0.312717	0.024299	12.87
	pt	-0.00357	0.039162	-0.09
	lagpt	0.042795	0.024161	1.77
	leadpt	-0.03697	0.025475	-1.45
q95	lagct	0.401627	0.031354	12.81
	leadct	0.310551	0.027381	11.34
	pt	-0.00784	0.053104	-0.15
	lagpt	0.041798	0.037343	1.12
	leadpt	-0.04011	0.034792	-1.15
q99	lagct	0.4962556	0.061814	8.03
	leadct	0.315605	0.077296	4.08
	pt	-0.0005	0.119415	-0.5
	lagpt	0.013005	0.071058	0.18
	leadpt	0.02484	0.077502	0.32

Figure 3: Intercepts

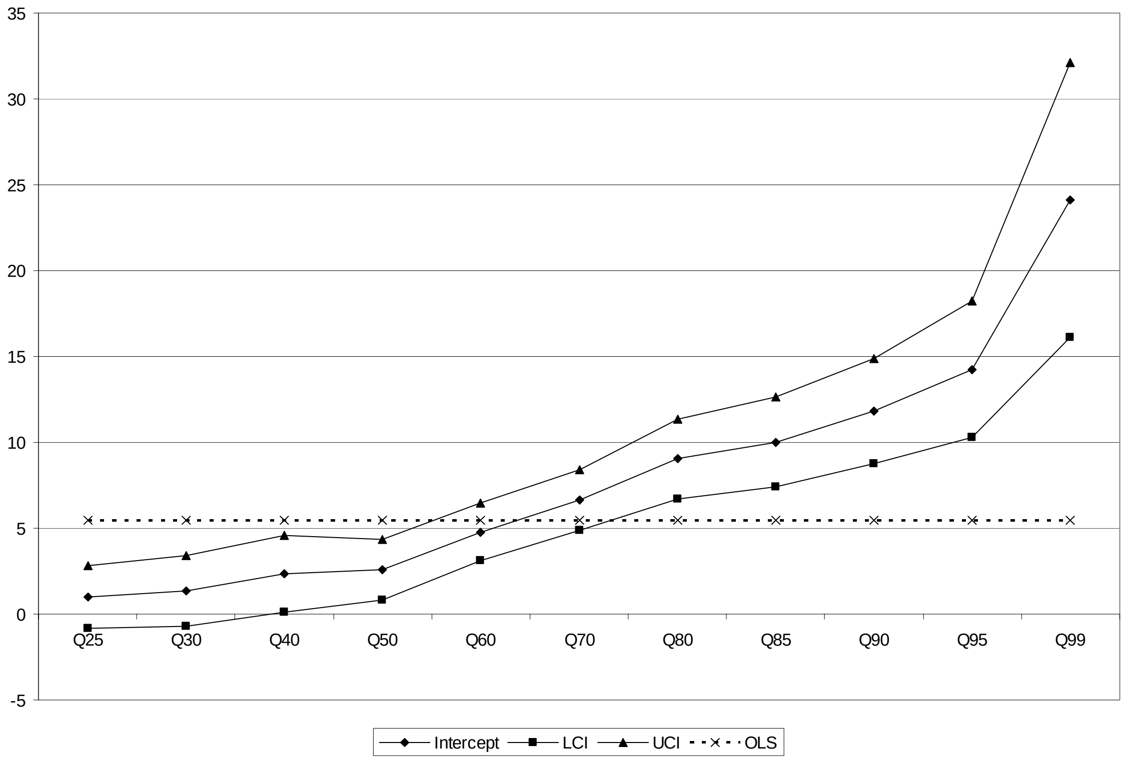

Figure 4: Full RA Lag C Coefficient

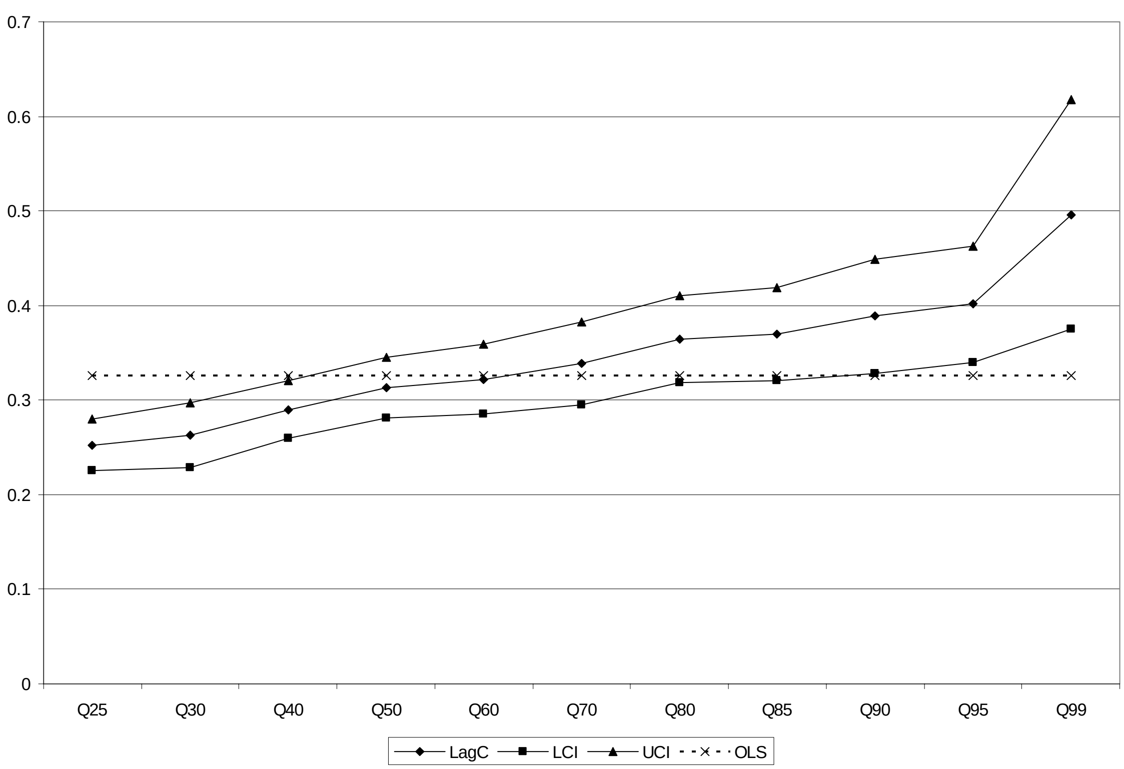

Figure 5: Full RA Lead C Coefficient

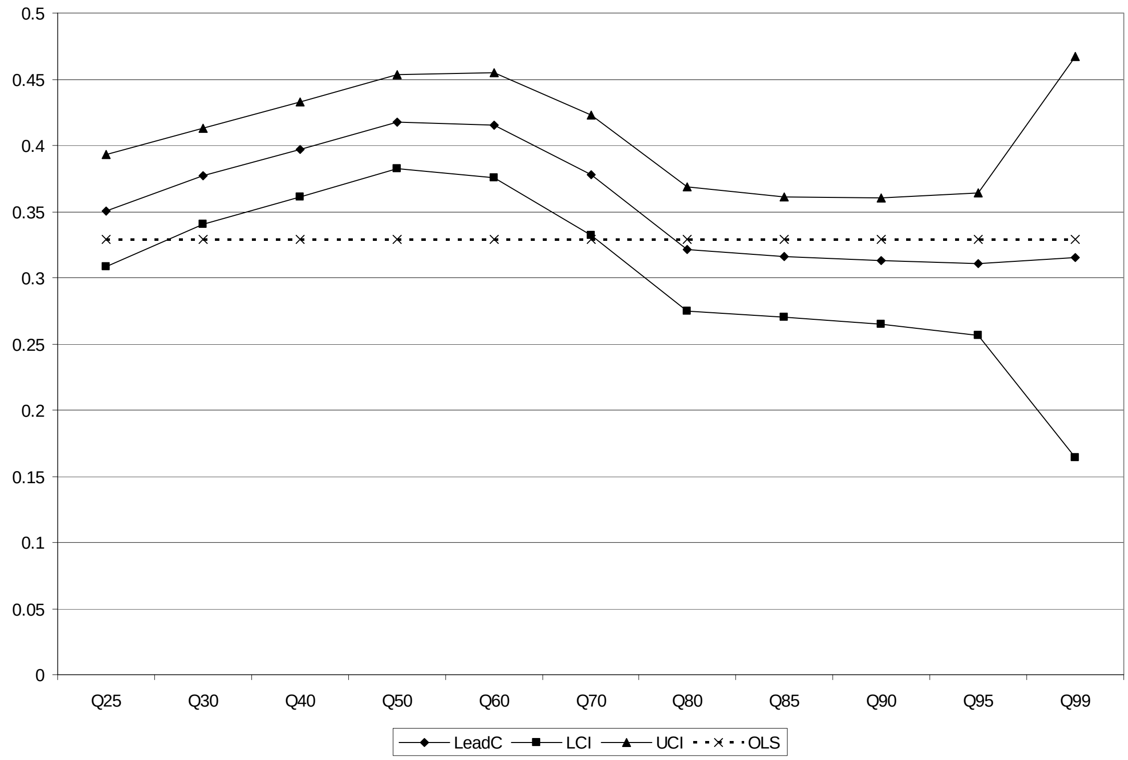

Figure 6: Full RA characteristic roots

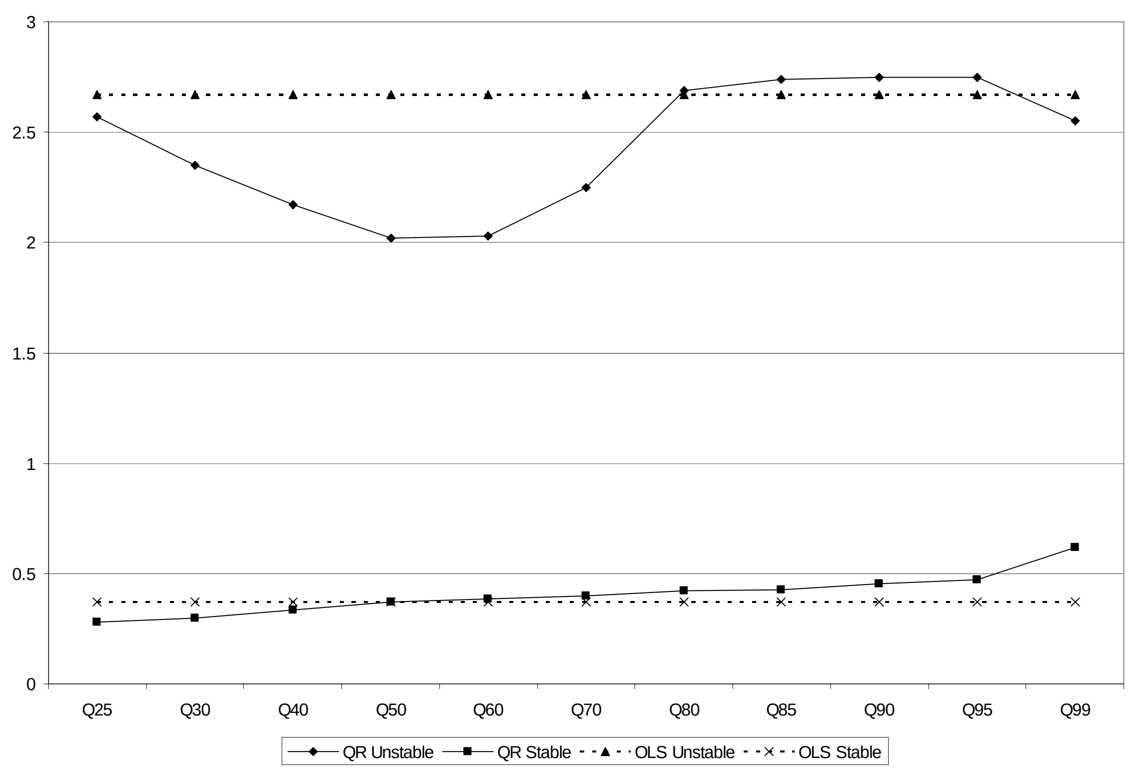

Figure 7: Full RA addictive and forward looking effects

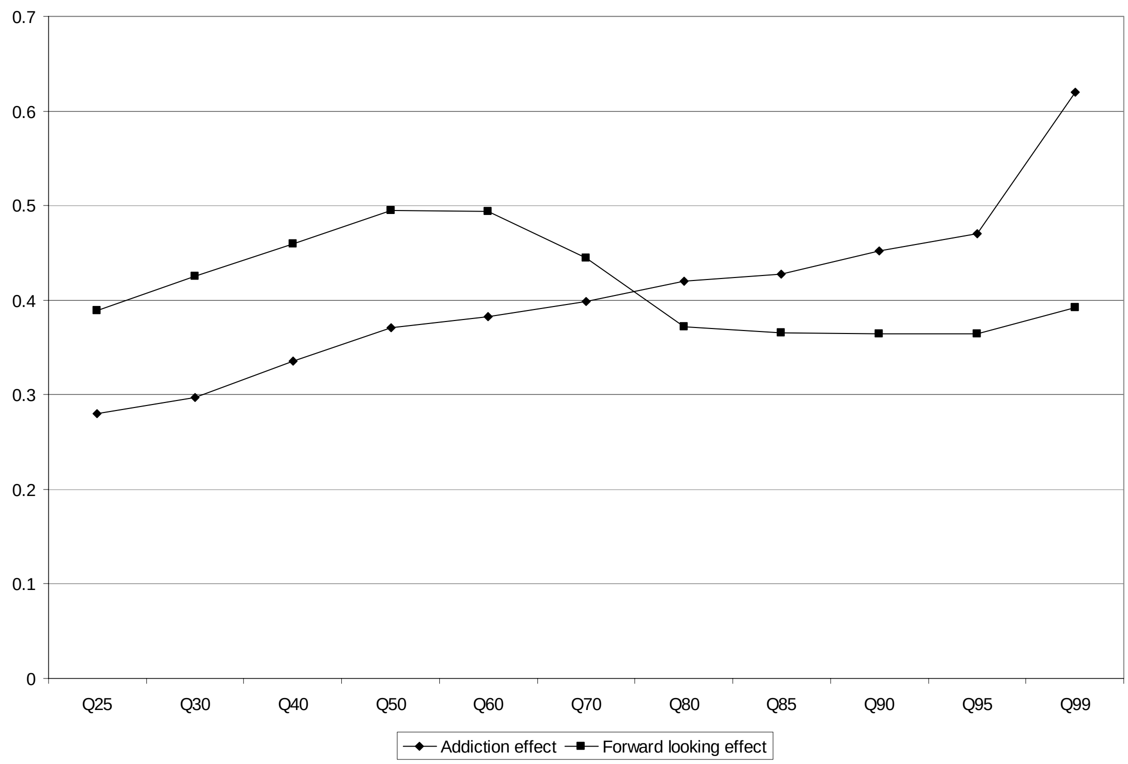

Figure 8: Drinks

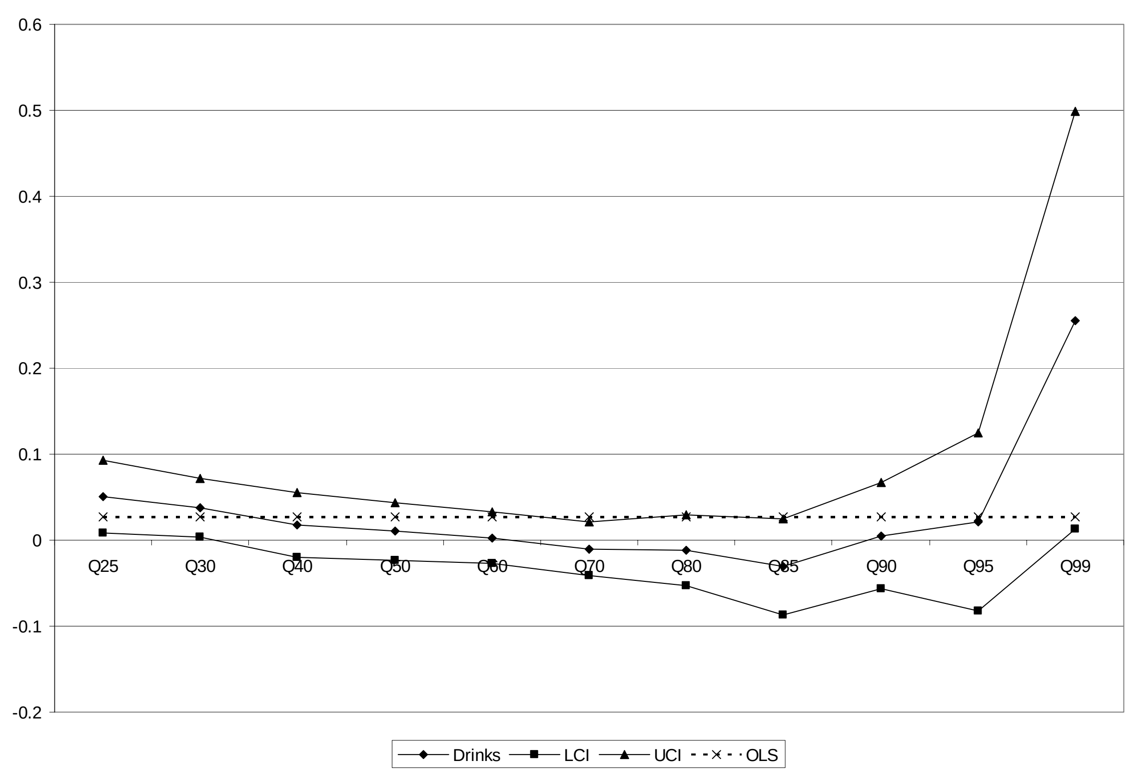

Figure 9: Juice

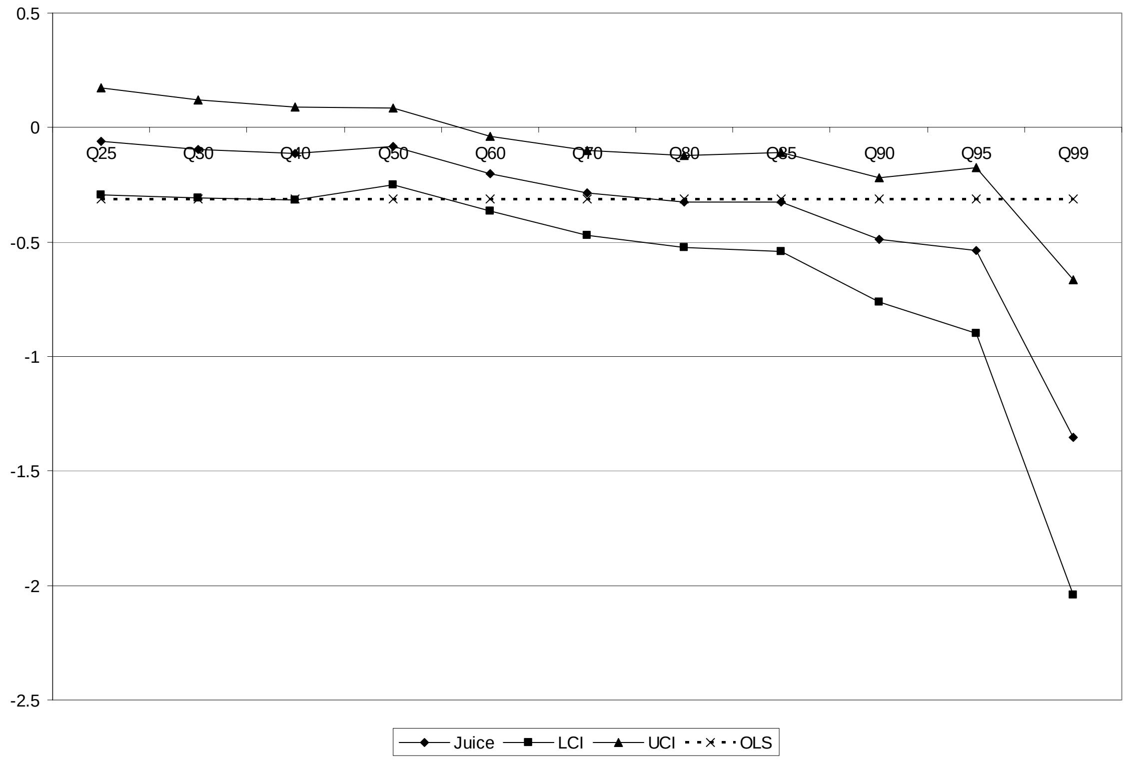

Figure 10: Potatoes

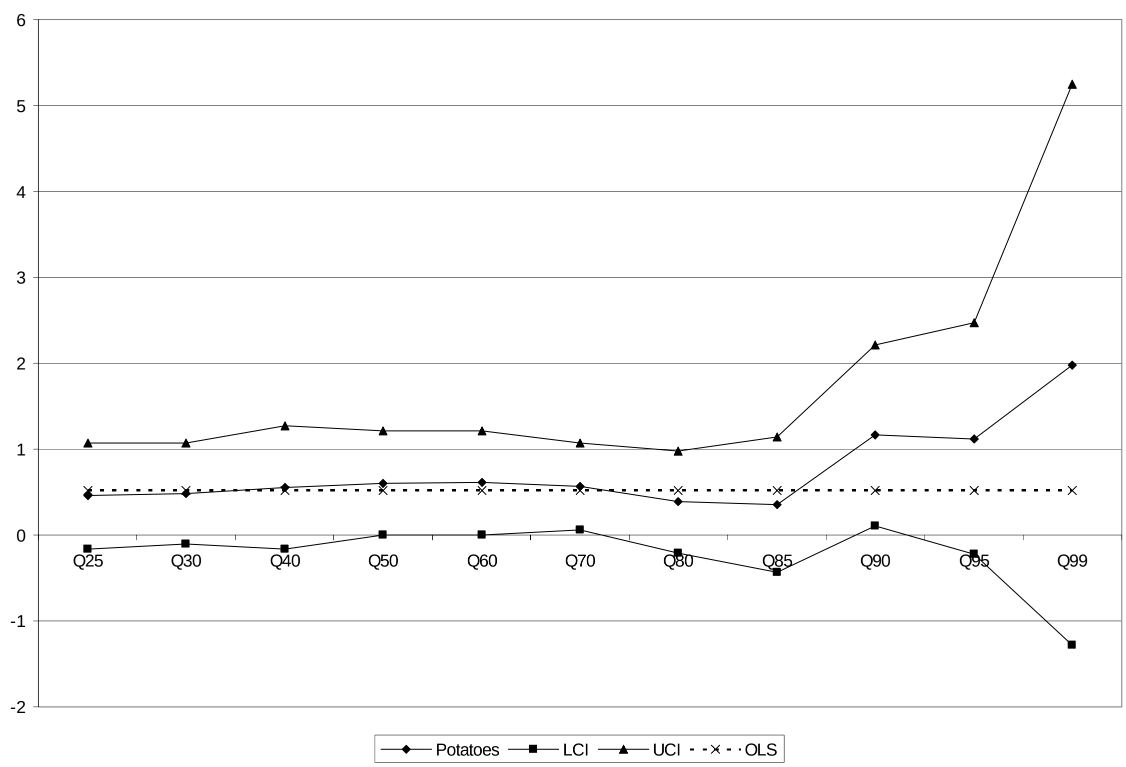

Figure 11: Self-assessed poor diet

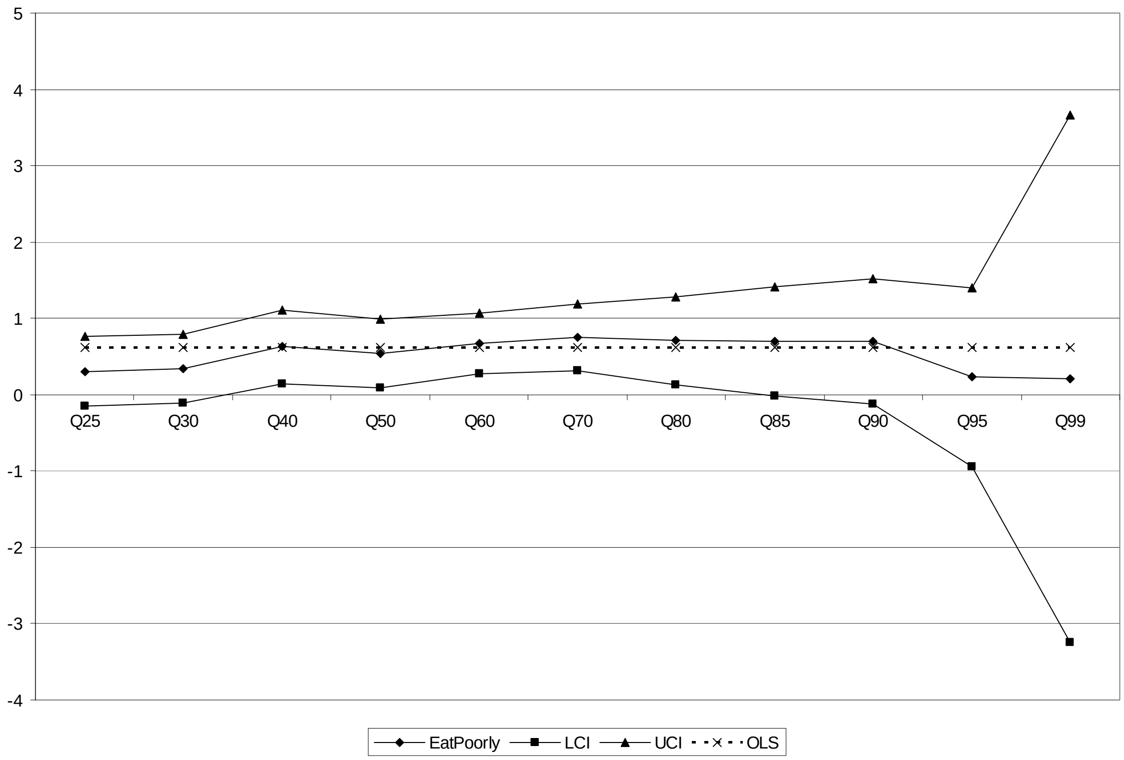

Figure 12: Male

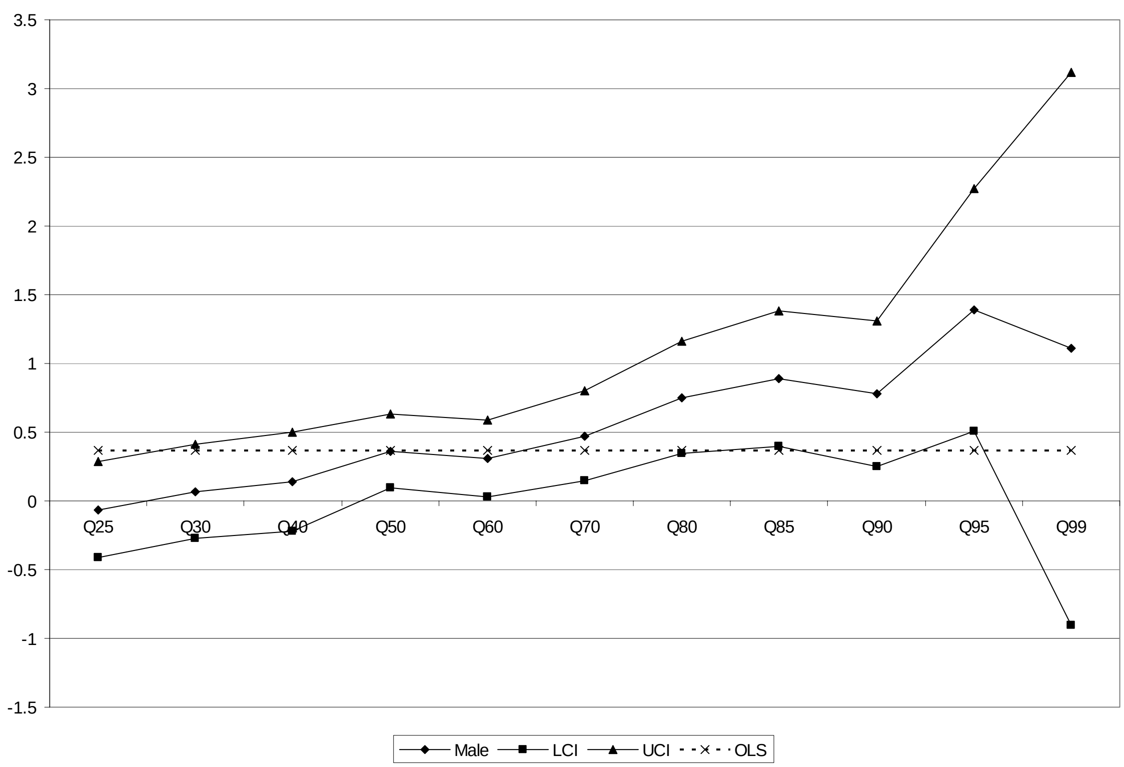

Figure 13: In poor health

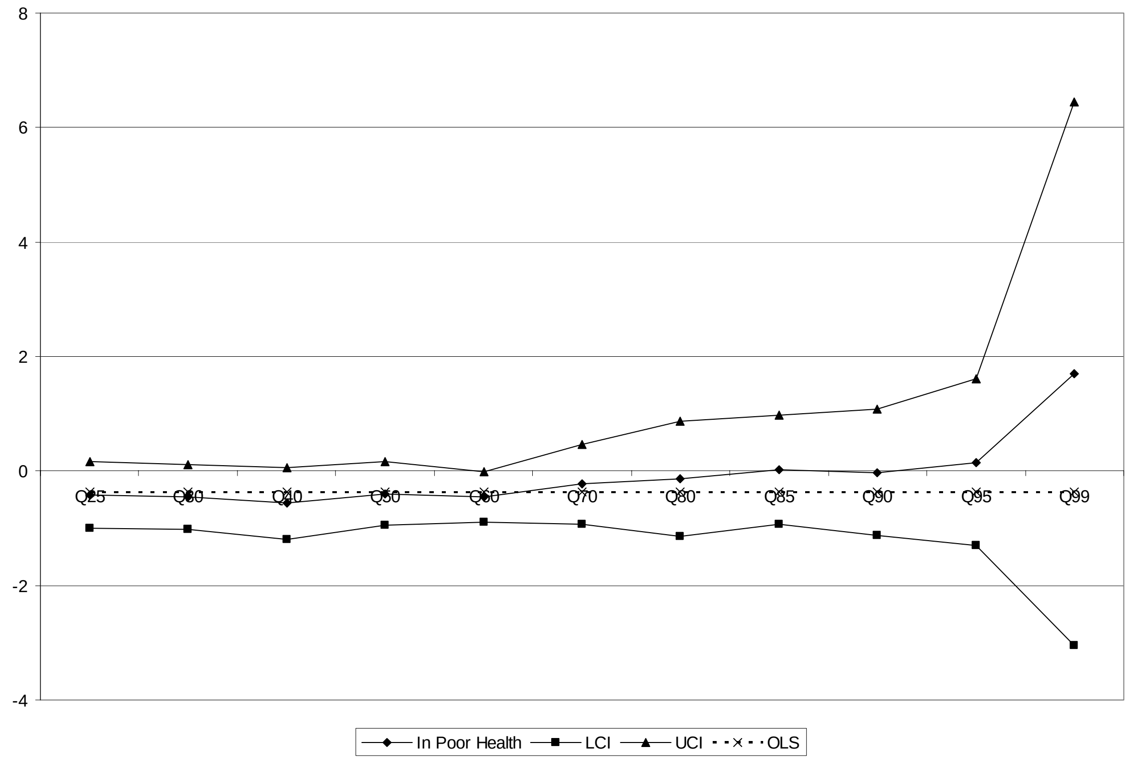

Figure 14: Had a heart attack

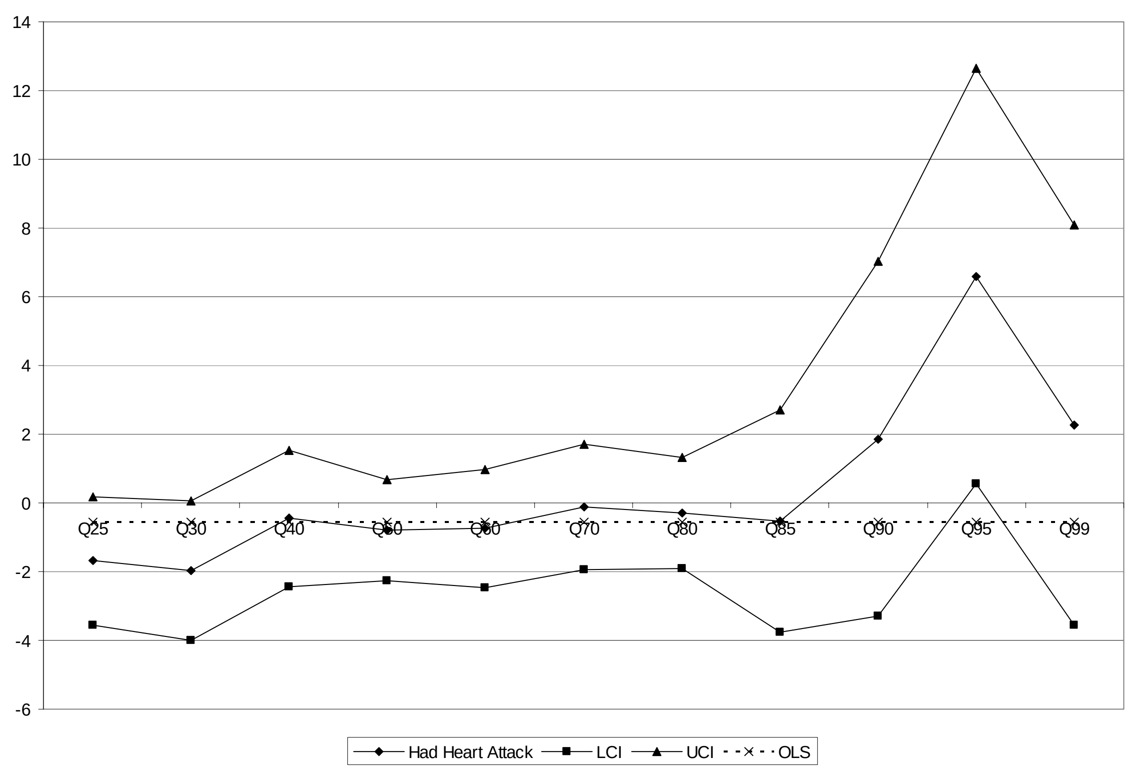

Figure 15: Life stressful

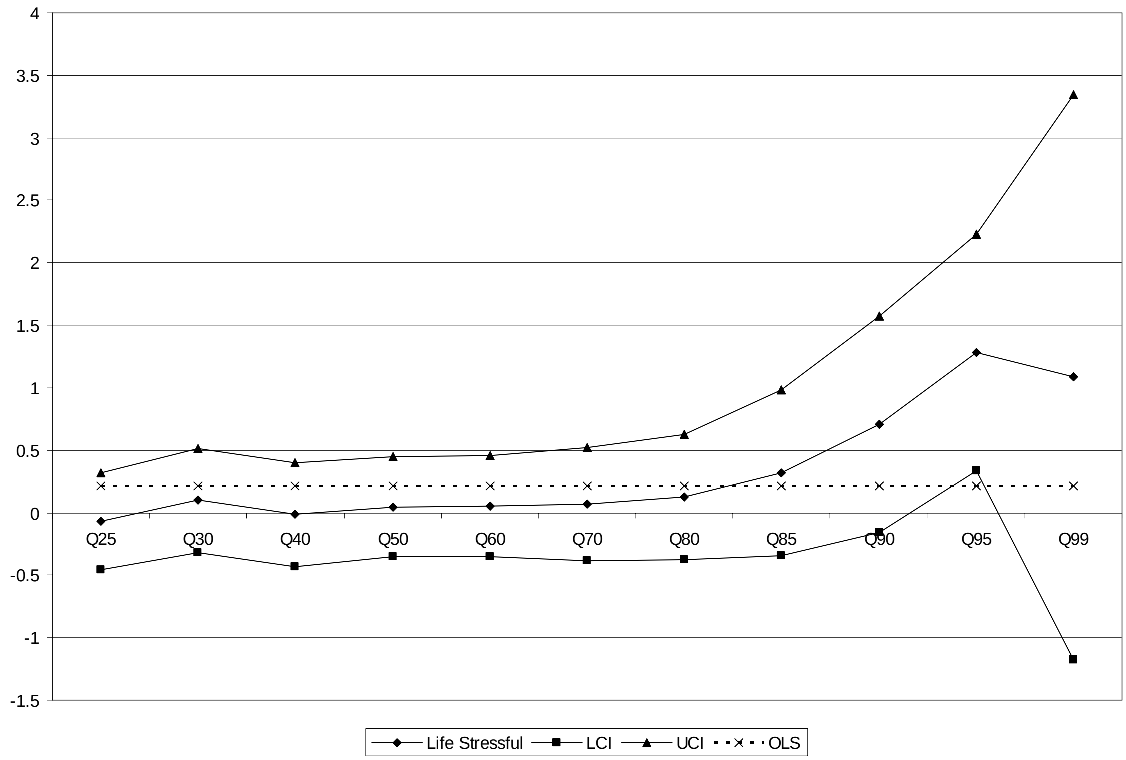

Figure 16: Tried to quit

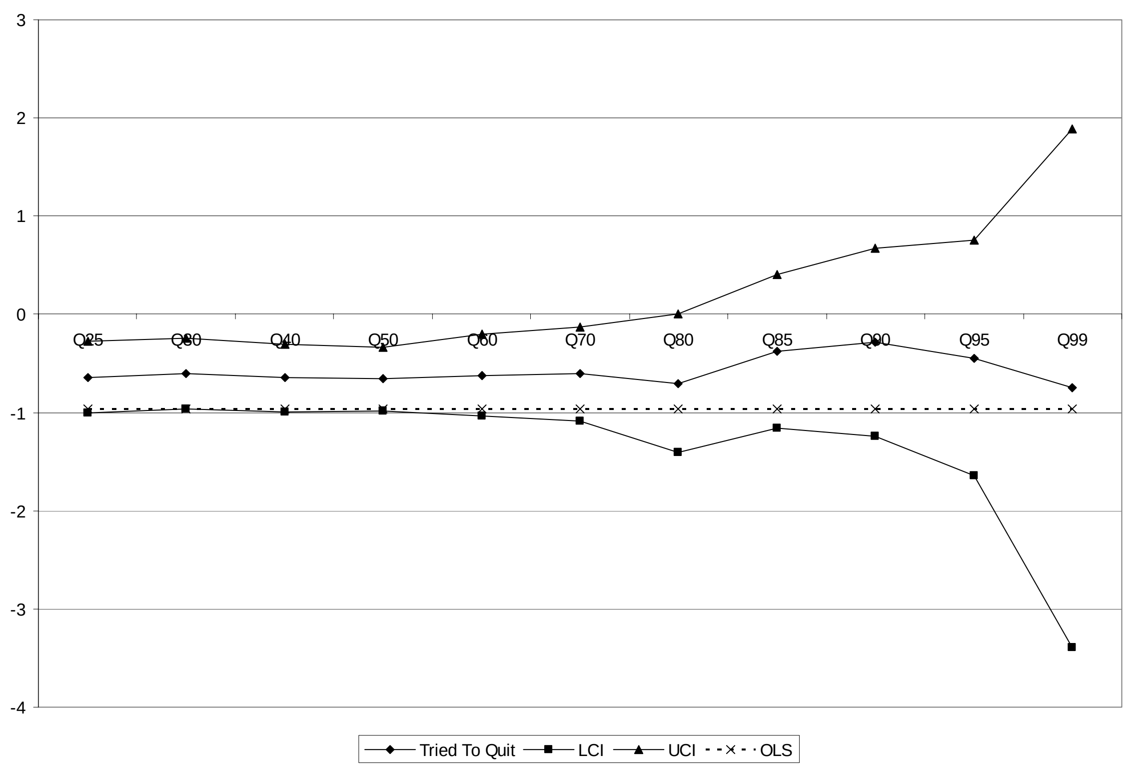

Figure 17: Faces workplace smoking restrictions

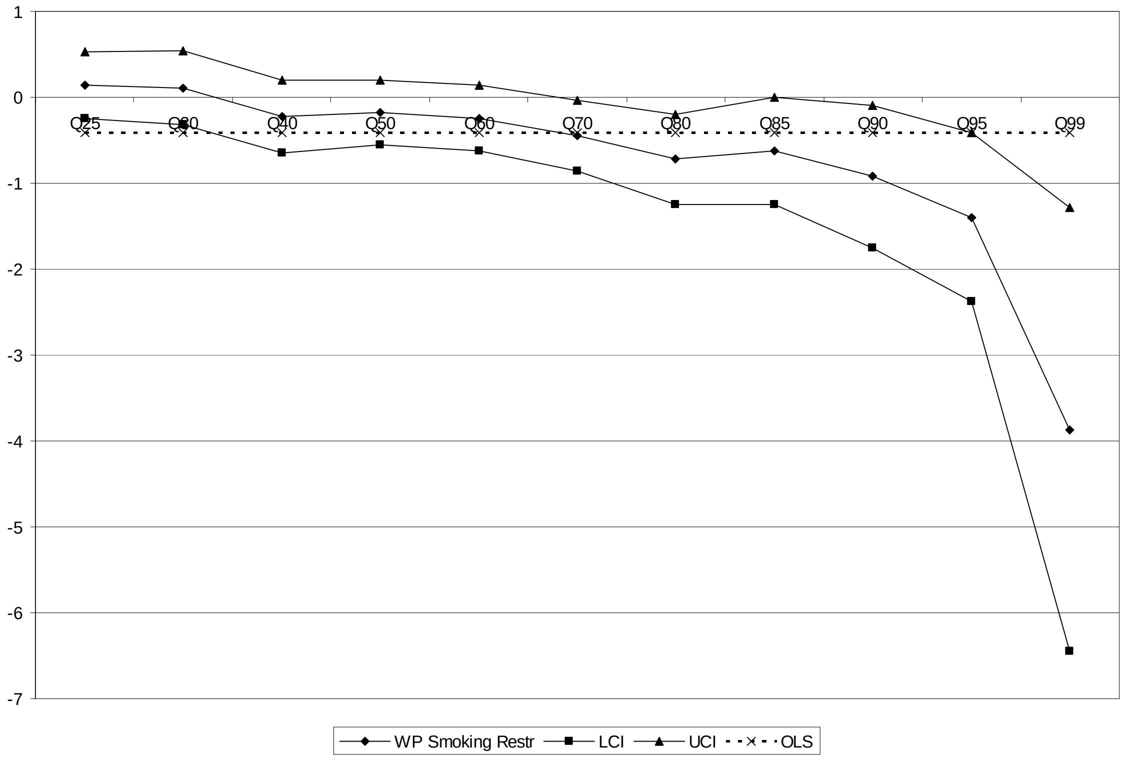

Figure 18: Has a University or College Education