Journal of Monetary Economics 57 (2010) 988–999

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme

Measuring consumption smoothing in CEX data

Martin Gervais *, Paul Klein

University of Southampton, United Kingdom

ARTICLE INFO

Article history:
Received 9 January 2008
Received in revised form
22 August 2010
Accepted 1 September 2010
Available online 21 September 2010

ABSTRACT

A new method of measuring the degree of consumption smoothing is proposed and implemented using data from the Consumer Expenditure Survey. The structure of this Survey is such that estimators previously used in the literature are inconsistent, simply because income is measured annually and consumption is measured quarterly. An AR(1) structure is imposed on the income process to obtain a proxy for quarterly income through a projection on annual income. By construction, this proxy gives rise to a measurement error which is orthogonal to the proxy itself—as opposed to the unobserved regressor—leading to a consistent estimator. Our estimates are contrasted with the output of two estimators used in the literature. This comparison shows that while the first (OLS) estimator tends to overstate the degree of risk sharing, the second (IV) estimator grossly understates it.

1. Introduction

A new method of measuring the degree of consumption smoothing is proposed and implemented using data from the U.S. Consumer Expenditure Survey (CEX).

This project is motivated by a large and growing literature on models with heterogeneous agents. In order for this heterogeneity to matter, so that agents are not just scaled versions of one another, these models typically feature some kind of friction that prevents agents from perfectly sharing idiosyncratic risk. The question is what specification is the most realistic one. One useful selection criterion, as suggested in Krueger and Perri (2004), is the extent to which models replicate the degree of risk sharing measured in the data. Properly measuring the degree of risk sharing in the data is therefore important in evaluating models with various kinds of frictions.

The empirically most obvious aspect of risk sharing—indeed it is also the one that Krueger and Perri (2004) focus on—is the extent to which household consumption responds to idiosyncratic earnings shocks. As in Dynarski and Gruber (1997), this notion is operationalized to mean the regression coefficient of the quarterly idiosyncratic consumption change on the quarterly idiosyncratic change in earnings. Since this is a temporal concept, it may as well be called the degree of *consumption smoothing*.[1] For the purpose of interpretation, it is useful to keep in mind that autarky implies that this coefficient is equal to one. On the other hand, if insurance markets are perfect and consumption and leisure are separable, then this coefficient is zero. Intermediate values can then be interpreted as measuring the *degree* of risk sharing or consumption smoothing.

Mace (1991) was the first to use CEX data to estimate the degree to which households smooth consumption in the presence of variable earnings or employment status. Her OLS estimates lend support to the full consumption insurance or

* Corresponding author.
E-mail addresses: gervais@soton.ac.uk (M. Gervais), pklein2@gmail.com (P. Klein).

[1] Clearly the value of this regression coefficient may depend on the length of the period. Presumably it would be smaller at the monthly level and bigger at the annual level. Limitations on the data, however, prevent us from looking either at the monthly or annual level. Although data on monthly consumption is in principle available in the CEX, the survey is conducted in such a way that significant components of reported consumption are not allowed to change across months within a three-month period. Therefore, in this paper, the focus is entirely on the quarterly frequency.

0304-3932/$ - see front matter
doi:10.1016/j.jmoneco.2010.08.009

full risk sharing proposition. However, Nelson (1994) points out potential problems with Mace's methodology, which leads Nelson to reject complete risk sharing.[2] Similarly, Dynarski and Gruber (1997) claim that the OLS estimates used by both Mace (1991) and Nelson (1994) suffer from measurement error and propose an IV method to deal with this problem. Meanwhile, our work shows that the structure of the CEX gives rise to non-classical measurement error that renders this IV approach invalid in this context, that is, as a measure of consumption smoothing as defined above.

Unlike the papers cited above, our analysis deals with a major problem with the CEX, which is that consumption and income are not observed for coincident periods. Specifically, the structure of the data is as follows. A household reports information regarding seven quarters (21 months). Consumption expenditure data are available for the last four of these quarters, and income data are available for two 12-month periods, one covering the first 12 months and the other covering the last 12 months. The main contribution here is to develop and implement an estimation strategy that is appropriate given the structure of the CEX.[3]

It is worth stressing in this context that our approach is designed to address the problem created by what might be called the asynchronicity of consumption and income observations in the CEX. It is not designed to address the problem of inaccurate reporting by households, which is the usual measurement error problem. Without denying that misreporting of income and/or consumption may be an important problem, this paper focuses elsewhere. From now on, to be clear on this matter, when we specifically talk about the measurement error that arises from asynchronicity, we will not refer to it as measurement error but synchronization error.

The estimator developed here uses a proxy for the true regressor and is based on the following simple result. As pointed out in Wooldridge (2002), if the proxy for true income is such that the implied measurement error is orthogonal to the proxy itself, then the OLS estimator using the proxy as the regressor is consistent.[4] The challenge, of course, is to find a proxy with the required orthogonality property. This is achieved by replacing the unobserved regressor by a linear projection on observable variables. In order to compute the linear projection of income contemporaneous with our consumption measure, an AR(1) structure is imposed on monthly income whose parameters are estimated using the generalized method of moments (GMM).

The main result is that the null hypothesis of perfect consumption risk sharing can be rejected, but that the degree of risk sharing is quite high. In particular, it is much higher than it would appear if measured using the method of Dynarski and Gruber (1997). Having rejected the null, it would of course be very interesting to consider alternative hypotheses about how consumption is smoothed in the data. Several recent papers speak to this point, including Kaplan and Violante (2009) and Broer (2009); in fact, this literature is growing fast. No attempt is made here to add to this literature; instead attention is confined to establishing facts about consumption smoothing.

An additional result of potential practical importance is that it appears that using food as a measure of consumption—historically the only measure available in PSID data—is not a bad approximation. As one would expect, durable goods purchases are found to be more responsive to changes in income than purchases of non-durable goods. More surprisingly, it turns out that households with a lot of financial assets do not smooth consumption as much as households with less financial assets, consistent with recent findings by Guvenen (2007). Similarly, households with relatively high income smooth consumption to a lesser extent than households with relatively low income. Finally, married households consistently seem to smooth consumption to a lesser extent than their non-married counterparts.[5]

The paper is organized as follows. In Section 2 the CEX data and its structure are described, with details contained in the online supplement. In Section 3 the instrumental variables approach of Dynarski and Gruber (1997) is described and shown to be invalid in this context. A solution to the asynchronicity problem is presented in Section 4. Section 5 presents some simulation results, contrasting the properties of existing estimators with our own. Section 6 presents some estimation results from the CEX data. Section 7 concludes.

2. Description of the data

The problem with measuring the degree to which households smooth consumption in the presence of income variability has always been the scarcity of reliable consumption data. It is well known that both the CPS and the PSID contain high quality income data. However, the CPS provides virtually no consumption data and the PSID, at least historically, only has information on food and housing consumption, which is clearly not ideal. CEX is the only survey in the

[2] In particular, Nelson (1994) points out that changes in monthly expenditures, which is the measure of consumption used by Mace (1991), are likely to overstate changes in consumption as they may not reflect changes in service flows. Accordingly, Nelson (1994) uses changes in quarterly instead of monthly consumption.

[3] Instead of confronting this inherent problem with the structure of CEX data, Blundell et al. (2008) impute consumption to households in the PSID using information from the demand for food, which is available both in CEX and PSID data. See Blundell et al. (2005) for details on the imputation procedure. Note, however, that PSID data are also subject to temporal alignment issues, in particular because food consumption corresponds to usual weekly spending, while income is measured on an annual basis.

[4] This is in contrast to classical measurement error, where the measurement error is orthogonal to the true regressor.

[5] Since this result is perhaps counterintuitive at first sight, it deserves comment. Presumably a married person's consumption is less sensitive to shocks to his or her own individual income than a single person's consumption is. However, this is not what is measured here. Rather, it is the response of *household* consumption to shocks to *household* income. In this sense, there is no particular reason to believe that married couples are better at smoothing than single-person households.

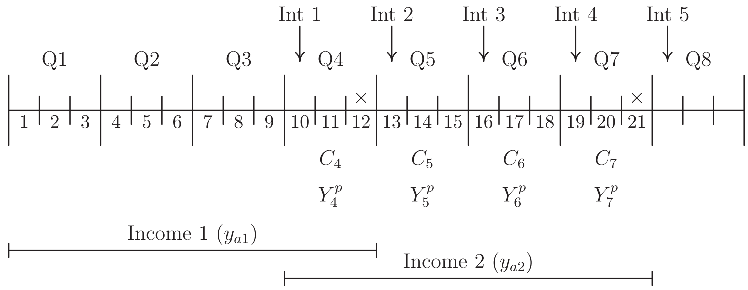

Fig. 1. Life of an interviewee in the CEX. Note: int 1–5 refer to interviews 1 though 5. Q1 to Q8 refer to quarters, but not necessarily calendar quarters. The numbers 1–21 refer to months, and represent the 21 months over which consumer units in the CEX report information. The $\times$'s for month 12 and 21 indicate that data on the last pay check is available for these two months. C_4 to C_7 represent four quarterly observations of consumption. Y_4^p to Y_7^p represent projected income for these quarters. y_{a1} and y_{a2} represent annual observations of income.

U.S. which collects detailed consumption data.[6] In addition, the CEX collects income data, which essentially makes the CEX the only dataset suitable for this purpose.[7] Another desirable feature of the CEX is that it provides a (short) panel for both consumption and income. However, for the purpose of estimating the response of consumption to income, the structure of the Survey is far from ideal.

Each household or consumer unit (CU) in the Survey is interviewed 5 times.[8] Interviews occur every three months, but the interview month varies across CUs: if the first interview occurs in January of a given year, the last interview takes place in January of the following year. The first interview is only used to collect data on various characteristics (such as race and education) of the CU and its members as well as information on some durable goods. Whereas consumption data pertaining to each of the last three months is collected at each subsequent interview, income data are only collected during the second and fifth interviews and pertains to the last 12 months.

Fig. 1 illustrates the life of a household in the Survey. In this figure quarters, labelled Q1 to Q8, do not necessarily refer to calendar quarters: they are calendar quarters only for households whose interview month corresponds to the first month of a calendar quarter. Notice that while the five interviews only span 12 months, the data collected span seven quarters. Since annual income data (y_{a1} and y_{a2}) are collected at the second and fifth interviews, they overlap for three months and thus only span 21 months. The $\times$'s in the figure (at months 12 and 21) indicate that information regarding the last paycheck received by members of the household is collected for these months, denoted by y_{12}^m and y_{21}^m. respectively.[9] Since consumption information is collected at each of the last four interviews, there are four consumption observations, labelled C_4 though C_7. Finally, the Y^p's refer to projected income, which will be explained in detail in Section 4.

2.1. Income data

Our definition of income consists of after-tax labor earnings plus government-mandated deductions and benefits.[10] More precisely, for every CU, yearly income is defined as wages and salary (including all compensations from the employer) plus a fraction (0.864) of self-employment (farm and non-farm) income. Also included are the following government transfers: social security benefits, unemployment compensation, public assistance and welfare payments, as well as other transfers. From that amount, total taxes paid (federal, state and local, including property taxes, all net of refunds) as well as social security and (government and railroad) retirement contributions are deducted. Total income is then deflated by the CPI for the relevant 12 months. See the online supplement for more details.

2.2. Consumption data

Quarterly consumption is defined as the expenditure on all items purchased by the household during the quarter (except vehicles and houses) plus imputed consumption flow values for owner-occupied housing and vehicles. Imputed

[6] The consumption data used in this paper are from the CEX Interview Survey. Battistin (2003) and Attanasio et al. (2007) document differences between the interview and the diary components of the CEX.

[7] Attanasio and Davis (1996) combine high quality consumption data from the CEX to high quality income data from the CPS. While this strategy is suitable to study risk sharing across groups of households, it masks the degree of risk sharing at the household level as idiosyncratic risk washes out in the aggregation procedure.

[8] The primary sampling unit in the CEX is called a consumer unit. A consumer unit consists of individuals who are either related or share their income to make joint expenditures. The CEX makes a subtle distinction between households and consumer units, but here the terms household and CU are used interchangeably.

[9] This last paycheck need not refer only to that month, nor does it need to cover the entire month. Fortunately, information on the period of time the pay check covers is also collected.

[10] The reason why government-mandated contributions and benefits are included in our concept of income is that we are ultimately interested in how well the private sector is able to share risk, as opposed to society as a whole.

Table 1
Benchmark sample selection.

Selection criterion	Observations deleted	Remaining observations
Original dataset		164,507
Incomplete income data	97,554	66,953
Inconsistent race	527	66,426
Inconsistent sex	1,351	65,075
Inconsistent age	2,106	62,969
Inconsistent education	1,078	61,891
Non-positive annual income(s)	1,432	60,459
Zero or missing food consumption	2,310	58,149
Income (y_{a2}) less than total food consumption	1,498	56,651
Negative medical care expenditures	1,537	
Our benchmark sample		**55,114**

Note: Original data from Consumer Expenditure Survey, 1980Q1–2002Q1.

housing services for homeowners is defined as the *rental equivalent*, i.e. the amount that the respondent expects the property to fetch in the rental market. For vehicles, the purchase price of vehicles is imputed following the procedure outlined in Cutler and Katz (1991) in order to impute the purchase price of vehicles owned by the household; this number is then divided by 32 to get the quarterly flow value. To this any expenditure on maintaining and repairing the vehicle(s) is added. For all other durables, no attempt is made to impute any service flows from the stocks: instead, expenditures during the quarter are simply added to consumption in that quarter. All observations are deflated by the relevant CPI. The availability of CPI categories thus dictates our choice of consumption categories.[11]

2.3. Idiosyncratic income and consumption

To extract the idiosyncratic component of log annual incomes, denoted by y^i_{a1} and y^i_{a2}, log incomes are regressed on a constant, a cubic in age, aggregate log GDP per head (not seasonally adjusted), the number of earners in the household, the number of individuals in the household, the number of members below the age of 18, the number of members above the age of 64, and dummies representing marital status, education and race; the residuals are then retained. Since it is consumption *changes* that are regressed on income *changes*, any individual fixed effects are automatically eliminated.

Idiosyncratic consumption is extracted in the same way as income, except that a seasonal dummy is introduced in addition to the other controls in order to take care of the possibility that consumption varies systematically with the time of year.

2.4. Sample selection

Our sample runs from the first quarter of 1980 to the first quarter of 2002. Table 1 summarizes the benchmark sample selection. Consumer units whose income in either year is considered incomplete are excluded.[12] Also excluded are households for whom the characteristics of the reference person are inconsistent over time, either because the reference person grows younger or ages by more than one year from one quarter to the next, gets less or more educated too fast, or undergoes a sex or race change. Because of a coding mistake in the CEX data, all households whose interviews span the years 1981 and 1982 are dropped because of changes in characteristics. Households with at least one non-positive yearly income observation and those with missing or zero food consumption are also excluded. Next, households whose income in year 2 is insufficient to cover their total food consumption for that year are excluded. This criterion is meant to eliminate households whose income is badly measured.[13] Finally, households whose consumption of medical care is negative are dropped from the sample. It should be noted that while these criteria are sufficient to guarantee good consumption data for interviews 2 and 5, they do not guarantee that consumption is available for interviews 3 and 4. There are, however, fewer than 50 households for which consumption in those interviews is missing.

With this sample as a benchmark, the sensitivity of the results to many other sample definitions is explored below. In particular, the results are shown to be robust to eliminating self-employed households, who may be in a position to

[11] Our categories are similar to those of Hobijn and Lagakos (2005). See the online supplement for details.

[12] CUs are considered *complete* income reporters if they list any major source of income, independent of other income responses—see Cutler and Katz (1991) for details.

[13] The income of the poor in CEX data is known to be poorly measured (see Meyer and Sullivan, 2003). This criterion is also used by Blundell et al. (2008) for their CEX sample.

manipulate their income in response to consumption expenditures.[14] The results are also robust to restricting the sample to the working-age population (between 21 and 64), as well as to the elimination of households who live in rural areas.

3. Synchronization error and the IV solution

The fundamental problem with estimating the degree to which households can smooth consumption in the presence of income risk from CEX data is the fact that consumption and income are not observed for the same periods of time. Having no observations of income corresponding to consumption, a natural though problematic strategy is to use the change in annual income as a proxy for the income change between quarters Q7 and Q4. Doing so, however, introduces a discrepancy between the true regressor and the proxy, which can be described as measurement error (here we call it synchronization error), though it has nothing to do with misreporting. Of course, the change in consumption from Q4 to Q7 is not a quarterly change, as the change from Q4 to Q5 is. However, for the purpose of making the argument of this section, let us assume that the objects of interest are changes from Q4 to Q7.

Now consider the problem of estimating β in the regression equation

$$C = \beta Y + \varepsilon, \tag{1}$$

where C measures the change in the idiosyncratic component of consumption ($C_7 - C_4$), Y measures the change in the idiosyncratic component of income ($Y_7 - Y_4$), and $\mathsf{E}[\varepsilon] = \mathsf{E}[Y\varepsilon] = 0$. Since Y is unobserved, it needs to be replaced by a proxy $\hat{Y}$. In the existing literature, the proxy is simply the difference between the two yearly observations of income ($y_{a2} - y_{a1}$). The difference between the true regressor and this proxy will be denoted by η, defined via $\hat{Y} = Y + \eta$. The OLS estimator in this case is not likely to be consistent. If one assumes that $\mathsf{E}[\eta\varepsilon] = 0$, then

$$\operatorname*{plim}_{N\to\infty} \beta_N^{\mathrm{OLS}} = \beta \frac{\mathsf{E}[Y\hat{Y}]}{\mathsf{E}[\hat{Y}^2]}.$$

If the measurement error were "classical," i.e. $\mathsf{E}[Y\eta] = 0$, one would have

$$\frac{\mathsf{E}[Y\hat{Y}]}{\mathsf{E}[\hat{Y}^2]} = \frac{\mathsf{E}[Y^2]}{\mathsf{E}[Y^2] + \mathsf{E}[\eta^2]} < 1,$$

so that the OLS estimator would be asymptotically biased towards zero, thereby overstating the degree of risk sharing.

In order to deal with this possible asymptotic bias, Dynarski and Gruber (1997) use an instrumental variable (IV) approach with a second measure of the income change as an instrument. As shown in Fig. 1, the CEX provides information about the amount of the CUs last paycheck, as well as its frequency, which we refer to as (log) monthly incomes y_{12}^m and y_{21}^m. This is then used to define the instrument via $Z = y_{21}^m - y_{12}^m$. This instrument is invalid because the measurement error arises from asynchronicity, as opposed to misreporting of income, rendering it non-classical. To see this, notice that

$$\operatorname*{plim}_{N\to\infty} \beta_N^{\mathrm{IV}} = \frac{\mathsf{E}[CZ]}{\mathsf{E}[\hat{Y}Z]}.$$

Assuming that $\mathsf{E}[\eta\varepsilon] = 0$, one obtains

$$\operatorname*{plim}_{N\to\infty} \beta_N^{\mathrm{IV}} = \beta \frac{\mathsf{E}[YZ]}{\mathsf{E}[YZ] + \mathsf{E}[Z\eta]}. \tag{2}$$

This instrumental variable strategy will thus be valid if and only if the synchronization error in the change in annual income is uncorrelated with the instrument, $\mathsf{E}[Z\eta] = 0$.

There are very strong reasons to believe that this condition is violated in this context, simply because of the structure of the CEX as illustrated in Fig. 1. For the sake of argument, suppose that income in the second year (y_{a2}) is high relative to what it is trying to measure, i.e. relative to income in the seventh quarter (Y_7). This results in a positive synchronization error η. Now if y_{a2} is high relative to Y_7, it is very likely that income during quarters Q4 to Q6 was above average. In particular, income in the third month of Q4, i.e. y_{12}^m, is also likely to be high. But if y_{12}^m is above average, it is also likely to be greater than y_{21}^m. This means that the instrument, $y_{21}^m - y_{12}^m$, is likely to be negative. For such reasons, one would expect the synchronization error to be negatively correlated with the instrument. Of course the converse assumption—that y_{a2} is low—would lead to the same conclusion. As Eq. (2) shows, the IV estimator in that case is biased upward (provided of course that the instrument is positively correlated with the true regressor), that is, this estimator tends to understate the degree of risk sharing.

[14] A household is defined here as self-employed if it receives at least 50% of its total income from self-employment in either of the two yearly observations of income. Around 6% of our sample satisfies this definition. Note, however, that only 62% of those who are self-employed according to the first observation of income remain self-employed in the second observation of income.

4. Projection-based estimation

It has already been established that when the regressor is measured with error, the OLS estimator is asymptotically biased by the factor

$$\frac{\mathsf{E}[Y\hat{Y}]}{\mathsf{E}[\hat{Y}^2]}.$$

Furthermore, if measurement error is classical, i.e. $\mathsf{E}[Y\eta]=0$, then this ratio is strictly less than one. However, if $\mathsf{E}[\hat{Y}\eta]=0$, then $\mathsf{E}[\hat{Y}^2]=\mathsf{E}[\hat{Y}Y]$ and hence

$$\frac{\mathsf{E}[Y\hat{Y}]}{\mathsf{E}[\hat{Y}^2]}=1.$$

Thus consistency is achieved if the measurement error is orthogonal not to the regressor but to its proxy.[15]

One proxy for the regressor that will certainly have the desired orthogonality property is the linear projection of the unobserved regressor on something that can be observed, such as the two observations of annual income. By definition of the projection, the projection error is orthogonal to the projection itself. The only remaining problem is to construct that projection. To do so, let W be defined as

$$W=\begin{bmatrix} y^i_{a1} \\ y^i_{a2} \end{bmatrix}.$$

Denote the actual (unobserved) idiosyncratic income change by Y and its linear projection on W by Y^p. Then

$$Y^p=\alpha W,$$

where

$$\alpha=\mathsf{E}[YW']\mathsf{E}[WW']^{-1}.$$

Thus in order to compute the projection one needs to estimate the covariance matrix of idiosyncratic annual incomes ($\mathsf{E}[WW']$) and the covariance between quarterly income and annual income ($\mathsf{E}[YW']$). For the latter it is necessary to impose some structure on the autocovariance function, as described in the next section.

4.1. Constructing the projection

Our strategy is as follows. First, a parameterized structure is imposed on the data generating process of income. Then, GMM is used to estimate the parameters of that structure. The estimated parameter values can then be used to compute the desired covariance matrix.

Let y^i_t denote (log) monthly income for household i. Assume that the stochastic process governing y^i_t is given by

$$y^i_t=\rho y^i_{t-1}+\varepsilon^i_t, \tag{3}$$

where ε^i_t is the idiosyncratic shock received by household i in period t and ρ measures the persistence of income.[16] Recall that income data consist of two annual observations of income, denoted y^i_{a1} and y^i_{a2}, which overlap for exactly three months. What is needed, however, is a measure of quarterly income, which will be constructed from estimates of monthly income y^i_t, t=1,…,21. First note that given y_1,

$$y^i_t=\rho^{t-1}y^i_1+\sum_{k=2}^{t}\rho^{t-k}\varepsilon^i_k.$$

One can then express the first annual income observation in terms of monthly income;

$$y^i_{a1}=\ln\left(\frac{1}{12}\sum_{t=1}^{12}\exp\left\{\rho^{t-1}y^i_1+\sum_{k=2}^{t}\rho^{t-k}\varepsilon^i_k\right\}\right).$$

Similarly, one can express the second annual income observation in terms of monthly income;

$$y^i_{a2}=\ln\left(\frac{1}{12}\sum_{t=10}^{21}\exp\left\{\rho^{t-1}y^i_1+\sum_{k=2}^{t}\rho^{t-k}\varepsilon^i_k\right\}\right).$$

[15] This result is related to the unobserved variable problem discussed in Zellner (1970), Goldberger (1972) and Pagan (1984).

[16] It should be emphasized that although this may not be the "true" income process, all that matters from the present point of view is that the implied structure for the autocovariance function is a good approximation of reality.

The following moments for GMM are used: $\mathsf{E}[y^i_{a1}\ y^i_{a1}]$; $\mathsf{E}[y^i_{a2}\ y^i_{a2}]$; and $\mathsf{E}[y^i_{a1}\ y^i_{a2}]$. Note that these moments are the elements of the covariance matrix of idiosyncratic annual incomes, $\mathsf{E}[WW']$. The parameters to be estimated are the variance of y^i_1, denoted by $\sigma^2_{y_1}$; the variance of ε^i_t, denoted by σ^2_ε; and the persistence parameter ρ. Using this notation, one obtains the following *approximate* results[17]:

$$\mathsf{E}[y^i_{a1}y^i_{a1}] \approx \left(\frac{1-\rho^{12}}{1-\rho}\right)^2 \sigma^2_{y_1} + \frac{1}{(1-\rho)^2}\left(11-2\rho\frac{1-\rho^{11}}{1-\rho}+\rho^2\frac{1-\rho^{22}}{1-\rho^2}\right)\sigma^2_\varepsilon,$$

$$\mathsf{E}[y^i_{a2}y^i_{a2}] \approx \rho^{18}\left(\frac{1-\rho^{12}}{1-\rho}\right)^2 \sigma^2_{y_1} + \left[\left(\frac{1-\rho^{18}}{1-\rho^2}\right)\left(\frac{1-\rho^{12}}{1-\rho}\right)^2 + \frac{1}{(1-\rho)^2}\left(11-2\rho\frac{1-\rho^{11}}{1-\rho}+\rho^2\frac{1-\rho^{22}}{1-\rho^2}\right)\right]\sigma^2_\varepsilon,$$

$$\mathsf{E}[y^i_{a1}y^i_{a2}] \approx \frac{\rho^9(1-\rho^{12})^2}{(1-\rho)^2}\sigma^2_{y_1} + \left[\frac{1-\rho^{12}}{(1-\rho)^2}\left(\frac{1-\rho^9}{1-\rho}+\frac{\rho^{21}-\rho^3}{1-\rho^2}\right)+\frac{(1+\rho)(1-\rho^{11})}{1-\rho}+\frac{(1-\rho^{10})}{1-\rho}\right]\sigma^2_\varepsilon.$$

The estimated process has a persistence parameter $\rho=0.87$, which is much lower than most estimates from the literature. For example, using PSID data, Storesletten et al. (2004a) find the persistent component of income to be close to a unit root process. The comparatively low estimate of ρ appears to emanate from the use of a richer set of covariates than is common in the literature. These covariates, which are highly predictable, account for a large fraction (49.5%) of the cross-sectional variation of income. On the other hand, following Deaton and Paxson (1994), Storesletten et al. (2004a) only use age and year of birth as controls, which potentially leaves their income process highly correlated with characteristics (such as marital status and education) which are predominantly deterministic.[18]

The variance of the idiosyncratic shock (σ^2_ε) is estimated to equal 0.18, while the variance of first month's income ($\sigma^2_{y_1}$), is estimated to be 0.86. Note that stationarity of the income process would imply a variance of first month income equal to 0.69. In a separate estimation, stationarity is rejected at any reasonable significance level. Given the values of σ^2_ε, σ^2_y and ρ, the covariance between quarterly and annual income can be computed, $\mathsf{E}[YW']$, and hence so can α and ultimately projected income Y^p_i, i=4,5,6,7 referred to in Fig. 1.

4.2. Estimation

Given the projections of quarterly income, there are four observations of both consumption and income for each household. Thus there are three sets of observations on log-differences of consumption and income from which to estimate the single regression coefficient β. The first set refers to the change in income and consumption from the first to the second quarter of the sample life of an interviewee, the second set to the change from the second to the third interview, etc.[19] The orthogonality between the disturbance term and the regressor in each of these equations gives rise to three moment conditions, which can be written as follows:

$$\mathsf{E}\left[\frac{1}{n_k}\sum_{i=1}^{n_k}\mathcal{X}_{k,i}(\mathcal{Y}_{k,i}-\beta\mathcal{X}_{k,i}-\gamma_k)\right]=0$$

for k=1,2,3, where $\mathcal{X}_{k,i}$ is the ith observation of the log difference of imputed income in the kth set of observations. Similarly, $\mathcal{Y}_{k,i}$ is the ith observation of the log difference of consumption in the kth set of observations. These moment conditions are then used as the basis for a GMM estimation, where the weighting matrix is chosen optimally. In practice, the need to allow for non-zero constants γ_k is eliminated by removing the sample-specific mean from each variable in each of the three samples. This procedure should remove any contamination of the results by any possible systematic tendency for reported consumption to rise or fall across interviews.

5. Simulations

To illustrate the properties of our estimator relative to the OLS and IV estimators, artificial monthly data from month 1 to month 21 are generated according to the process for income implied by the estimates from the previous section. The true value of β is assumed to be 0.1. The data generating process for consumption is thus given by $C^i_t=\beta Y^i_t+\xi^i_t$, where Y^i_t denotes income of household i in month t for t=1,2,…,21. It is also assumed that ξ^i_t is identically and independently distributed (over time and across households) normally with zero mean and variance 0.2.[20] One hundred samples of size 500,000 are simulated.

[17] The approximation errors that emanate from Jensen's inequality and are fairly small. The largest absolute deviation between the true moment and the approximated moment is about 0.02; the formula above understates the covariance between y^i_{a1} and y^i_{a2} by approximately that amount. A derivation of these approximate results is available upon request.

[18] In Storesletten et al. (2004b), where they use year dummies, a cubic in age and education as controls, they report a R^2 equal to 0.23.

[19] Recall that, for instance, the "first quarter" has nothing to do with the first quarter of a calendar year (January–March). Rather, it is the quarter that the household is asked about in its second interview. This could be any three-month period (e.g. May–July), and not necessarily a calendar quarter at all.

[20] The results are not sensitive to the variance of ξ^i_t.

Table 2
Simulation results.

	β^{OLS}	β^{IV}	β^{PRO}
Estimate	0.073	0.415	0.100

Note: These regression coefficients are computed on the basis of 21 simulated monthly observations of consumption and income.

The consumption observations are aggregated into seven three-month periods or quarters. The income observations are aggregated into two overlapping 12-month periods or years, the first covering months 1–12 and the second months 10–21. The OLS regression coefficient β^{OLS} is defined as the regression coefficient of the change in consumption from quarter 1 to quarter 4 on the income change from year 1 to year 2. Meanwhile, β^{IV} is defined as the IV regression coefficient of the quarterly change in consumption as defined above on the yearly change in income as defined above, using the change in income from month 12 to month 21 as an instrument. Finally, we compute β^{PRO} as the regression coefficient on quarterly consumption changes on quarterly income changes, where quarterly income is computed as the projection of quarterly income on yearly income using the technique developed above.

The results, shown in Table 2, indicate that the IV estimate is more than 5 times higher than the OLS estimate, as one would expect given the discussion in Section 3. While the estimate from our proposed projection method, labelled β^{PRO}, is equal to the true value of β, the IV estimate understates the degree of risk sharing and the OLS estimate overstates it.

5.1. Robustness

The structure that we impose on the autocovariance function of idiosyncratic income (AR(1)) may be important for the results. In order to check the robustness of our procedure, income data were generated according to a process that features a permanent, a persistent, and a transitory component, and projections were computed as if income had been generated by an AR(1) process. Specifically, following Storesletten et al. (2004a), assume that earnings can be represented as follows:

$$y_{i,t} = \zeta_i + z_{i,t} + x_{i,t},$$

where ζ_i is a permanent component (a fixed effect), $z_{i,t}$ is an AR(1) process, and $x_{i,t}$, the purely transitory component, is i.i.d. over time and households.

Having generated income according to this stochastic process, the exercise is repeated by first estimating the coefficient of an AR(1) process for income, and then following the exact same procedure as above to estimate the same three parameters. First, the simulation results suggest that the orthogonality property sufficient for our estimator to be consistent, discussed at the beginning of Section 4, holds to a high degree of precision. Second, the coefficient estimate is 0.101, obviously very close to that reported in the third column of Table 2, suggesting that our estimator is robust to different specifications of the earnings process.

6. Estimation results

Table 3 shows our estimate of the degree of risk sharing together with the OLS and IV estimates. The standard errors in that table were computed using a bootstrap strategy, with 1000 samples of the same number of observations as in the original sample (55,114).[21] As expected, the OLS estimate is low and the IV estimate is high relative to our proposed projection method estimate. Given the size of our sample, all these estimates are fairly precise and thus all statistically different from zero.

Table 4 shows estimates of the degree of consumption smoothing across broad categories of goods. The sample size for each of these categories (shown in parentheses in Table 4) varies as more households can have non-positive consumption of a particular category and still have positive total consumption. For example, around 32,000 households in the sample report positive expenditures on alcohol and tobacco. As one would expect, all our estimates lie in between the OLS and the IV estimates. It is interesting to note that the estimate for food is not substantially different from that for total consumption. This suggests that using food—historically the only component of consumption available in PSID data—as a proxy for consumption might not be such a bad idea when answering questions related to risk sharing.[22] This result, however, is somewhat sensitive to the sample. In a very homogeneous sample of non-self-employed, white, married, working-age households living in urban areas, the estimate for food consumption is 30% lower than that for total consumption, mainly because the estimate for food at home is lower than in the benchmark sample (see Table 5).

Table 4 also suggests that household durables are used as a buffer stock as suggested in Browning and Crossley (2009). Note also that housing consumption appears very smooth. This should not be surprising given how housing consumption is measured, but also because households with substantial changes in housing consumption are lost since the CEX does not

[21] Note that some observations still need to be discarded either because consumption in interviews 3 or 4 is non-positive, or, in the IV case, because the instrument (monthly income) is missing.

[22] Limited data on food expenditures are also available in the CPS Food Security Supplement, which was first administered in 1995.

Table 3
Estimates.

	β^{OLS}	β^{IV}	β^{PRO}
Estimate	0.0666	0.1909	0.1278
	(0.0032)	(0.0196)	(0.0145)
Sample size	55,114	34,379	55,083

Note: Authors' calculations based on CEX data from 1980Q1 to 2002Q1. Standard errors are in parentheses.

Table 4
Estimates across categories of goods.

	β^{OLS}	β^{IV}	β^{PRO}
Total consumption	0.067	0.191	0.128
	(55114)	(34379)	(55083)
Consumption less cars and housing	0.066	0.198	0.154
	(55114)	(34379)	(55114)
Food	0.054	0.188	0.125
	(55006)	(34346)	(55015)
Food at home	0.040	0.160	0.077
	(54736)	(34184)	(54761)
Food away from home	0.089	0.375	0.197
	(43043)	(29650)	(43306)
Alcohol and tobacco	0.048	0.136	0.106
	(31942)	(22549)	(32236)
Housing[a]	0.031	0.075	0.050
	(55054)	(34346)	(55080)
Household durables[b]	0.132	0.541	0.389
	(50081)	(32611)	(50065)
Transportation[c]	0.132	0.300	0.227
	(52001)	(33842)	(52105)
Education	0.061	0.161	0.122
	(38431)	(25601)	(38810)

Note: Authors' calculations based on CEX data from 1980Q1 to 2002Q1. Sample sizes are in parentheses.
[a] Housing includes rental and imputed rents.
[b] Household durables consist of furniture, household operations and apparel.
[c] Transportation includes all private and public expenditures on transportation, as well as the imputed service flow of privately owned vehicles.

Table 5
Projection estimates across categories of goods for different samples.

	Benchmark	Working age	Without rural	Without self-employed
Total consumption	0.128	0.124	0.117	0.114
Total less cars and housing	0.154	0.152	0.140	0.137
Food	0.125	0.106	0.083	0.077
Food at home	0.077	0.056	0.040	0.045
Food away from home	0.197	0.224	0.216	0.202
Alcohol and tobacco	0.106	0.095	0.096	0.115
Housing	0.050	0.045	0.044	0.043
Household durables	0.389	0.415	0.395	0.370
Transportation	0.227	0.206	0.213	0.211
Education	0.122	0.134	0.136	0.143

Note: Authors' calculations based on CEX data from 1980Q1 to 2002Q1. The second column reproduces results obtained under the benchmark sample defined in Table 1. In the third column, the sample is restricted to working age (21–64) reference persons, reducing the sample size to 41,604 observations. In the fourth column, households who live in rural areas are also dropped, further reducing the sample to 37,608 observations. In the final column, self-employed households have been dropped as well, leaving a sample of 35,017 observations.

follow households once they have moved. A related result is that consumption less cars and housing is less smooth than total consumption. Table 5 verifies that these results are robust to restricting the sample to working age (21–64) population, urban households, as well as to the non-self-employed.

Table 6
Conditional estimates.

	β^{OLS}	β^{IV}	β^{PRO}
Total sample	0.067	0.191	0.128
Results by financial asset holdings			
Fin assets < median	0.066	0.192	0.120
Fin assets > median	0.068	0.187	0.149
Results by income			
Income < median	0.053	0.159	0.122
Income > median	0.087	0.222	0.161
Results by type of employment			
Self-employed	0.053	0.158	0.159
Not self-employed	0.069	0.193	0.124
Results by race			
White	0.067	0.191	0.138
Black	0.057	0.211	0.087
Results by marital status			
Married	0.078	0.202	0.145
Not married	0.045	0.137	0.102
Results by education			
Less than high school	0.065	0.253	0.106
High school	0.063	0.187	0.129
Some college	0.070	0.179	0.130
College graduate	0.069	0.169	0.145
Results by sign of shocks			
Negative shocks	0.085	0.272	0.231
Positive shocks	0.030	0.269	0.051

Note: Authors' calculations based on CEX data from 1980Q1 to 2002Q1.

Table 7
Conditional projection estimates for different samples.

	Benchmark	Working age	Without rural	Without self-employed
Total sample	0.128	0.124	0.117	0.114
Results by financial asset holdings				
Fin assets < median	0.120	0.117	0.111	0.103
Fin assets > median	0.149	0.147	0.143	0.149
Results by income				
Income < median	0.122	0.118	0.111	0.102
Income > median	0.161	0.157	0.146	0.150
Results by type of employment				
Self-employed	0.159	0.152	0.142	
Not self-employed	0.124	0.120	0.112	
Results by race				
White	0.138	0.133	0.122	0.123
Black	0.087	0.100	0.111	0.106
Results by marital status				
Married	0.145	0.142	0.125	0.128
Not married	0.102	0.094	0.100	0.085
Results by education				
Less than high school	0.106	0.161	0.136	0.158
High school	0.129	0.107	0.110	0.091
Some college	0.130	0.124	0.113	0.112
College graduate	0.145	0.123	0.113	0.110

Note: Authors' calculations based on CEX data from 1980Q1 to 2002Q1. The second column reproduces results obtained under the benchmark sample defined in Table 1. In the third column, the sample is restricted to working age (21–64) reference persons, reducing the sample size to 41,604 observations. In the fourth column, households who live in rural areas are also dropped, further reducing the sample to 37,608 observations. In the final column, self-employed households have been dropped as well, leaving a sample of 35,017 observations.

Table 6 reports a set of estimates conditional on particular characteristics of the household. First, the CEX reports data on financial assets collected in the last interview.[23] Surprisingly, households with financial wealth below the median seem to smooth consumption more than those whose wealth is above the median. This result, which at first appears

[23] Financial assets consist of savings accounts, checking and brokerage accounts, savings bonds, as well as securities (stocks and bonds). It should be noted that ideally one would like to have access to financial assets at the time of the first interview, as the amount of financial assets at the end of the sample life can

Table 8
Conditional estimates for consumption less cars and housing.

	β^{OLS}	β^{IV}	β^{PRO}
Total sample	0.066	0.191	0.128
Results by financial asset holdings			
Fin assets < median	0.069	0.200	0.143
Fin assets > median	0.064	0.192	0.178
Results by income			
Income < median	0.057	0.160	0.144
Income > median	0.082	0.234	0.188
Results by type of employment			
Self-employed	0.049	0.165	0.189
Not self-employed	0.069	0.200	0.150
Results by race			
White	0.067	0.198	0.165
Black	0.056	0.247	0.088
Results by marital status			
Married	0.076	0.203	0.173
Not married	0.048	0.160	0.124
Results by education			
Less than high school	0.069	0.259	0.127
High school	0.065	0.210	0.156
Some college	0.070	0.203	0.160
College graduate	0.063	0.147	0.169

Note: Authors' calculations based on CEX data from 1980Q1 to 2002Q1.

counterintuitive, is consistent with recent findings in Guvenen (2007). It is worth emphasizing that these results were obtained through entirely different methods, and using different datasets—Guvenen (2007) uses the PSID and thus food consumption to estimate a structural model of stock market participation. He suggests that this result is due to the higher fraction of entrepreneurs among high wealth households. Table 6 shows that self-employed households indeed do not smooth consumption to the same extent as other households. However, this may be due to the fact that self-employed households can manipulate their income in response to expenditures. Furthermore, the "counterintuitive" result still holds in a sample that excludes self-employed households (see Table 7).

In line with the results for financial wealth, high income households smooth consumption to a lesser extent than low income households. Moreover, consumption responds to a much greater extent to a decrease in income than to an increase in income. Meanwhile, married households do not smooth consumption as much as singles. The results by race indicate that whites do not smooth consumption as much as blacks. However, the large difference between the OLS and IV estimates for blacks suggests that their income may be particularly badly measured.[24]

Table 7 shows that while the above-mentioned results are robust to alternative sample selections, such is not the case for education. *A priori*, Table 6 seems to suggest that more education is associated with less consumption smoothing. However, Table 7 shows the estimate for households whose reference person has less than a high school degree increases as the sample becomes more homogeneous. Furthermore, the difference between the estimates for high school and college changes with the sample, and effectively disappears once the sample is restricted to non-self-employed, white, married, working-age households living in urban areas (not shown). Finally, note that the patterns that emerge from Table 6 for total consumption also holds for consumption less cars and housing, as shown in Table 8.

7. Concluding remarks

The focus in this paper is on the structure of the Consumer Expenditure Survey and the problems it raises for obtaining consistent estimates of the extent to which households smooth consumption in the presence of variable earnings. Careful consideration of these problems leads one to reject both the OLS and IV approaches previously used in this literature: while the OLS estimator tends to overstate the degree of risk sharing, the IV estimator tends to understate it.

In response to this problem, an estimation strategy that is appropriate given the structure of the CEX is proposed. This approach uses the fact that an OLS estimator is consistent if the proxy variable used in the regression is orthogonal to its measurement error. A proxy variable with that property is constructed by projecting the unobserved regressor on observables.

(*footnote continued*)
be influenced by income shocks received in the previous periods. Unfortunately, the CEX only asks questions related to financial assets during the last interview.

[24] The sample size of blacks is also small (5699) relative to that of whites (47,490).

This estimation strategy is used to estimate the extent to which households smooth consumption across broad categories of goods and for different groups of the population. The main result is that while full risk sharing can formally be rejected, the degree of risk sharing is nevertheless quite high, lying in between the OLS and IV estimates. Another finding is that food consumption may be a reasonable proxy for total consumption. Moreover, the results are consistent with the notion that household durables are used as a buffer stock as suggested in Browning and Crossley (2009). Perhaps more surprisingly, households with a lot of financial assets are found to smooth consumption to a lesser extent than households with fewer financial assets; this is consistent with recent findings by Guvenen (2007).

Acknowledgments

We are grateful to the Editor (Robert King), the Associate Editor (Martin Eichenbaum), and especially an anonymous referee for their helpful comments and suggestions. We thank seminar participants at many conferences and institutions as well as Orazio Attanasio, Gadi Barlevy, Lance Lochner, Lars Osberg, Fabrizio Perri, Chris Robinson, David Strömberg, Todd Stinebrickner and Lars E. O. Svensson for useful comments. Michio Suzuki provided excellent research assistance and Laura Paszkiewicz from the Bureau of Labor Statistics was extremely helpful with CEX data. Both authors gratefully acknowledge financial support from the Social Sciences and Humanities Research Council of Canada.

Appendix A. Supplementary data

Supplementary data associated with this article can be found in the online version at doi:10.1016/j.jmoneco.2010.08.009.

References

Attanasio, O., Battistin, E., Ichimura, H., 2007. What really happened to consumption inequality in the US? In: Berndt, E., Hulten, C. (Eds.), Hard-to-Measure Goods and Services: Essays in Honor of Zvi Griliches. University of Chicago Press.

Attanasio, O., Davis, S.J., 1996. Relative wage movements and the distribution of consumption. Journal of Political Economy 104 (6), 1227–1262.

Battistin, E., 2003. Errors in survey reports of consumption expenditures. Institute for Fiscal Studies Working Paper W03/07.

Blundell, R., Pistaferri, L., Preston, I., 2005. Imputing consumption in the PSID using food demand estimates from the CEX. Institute for Fiscal Studies WP04/27.

Blundell, R., Pistaferri, L., Preston, I., 2008. Consumption inequality and partial insurance. American Economic Review 98 (5), 1887–1921.

Broer, T., 2009. Stationary equilibrium distributions in economies with limited commitment. Unpublished manuscript.

Browning, M., Crossley, T.F., 2009. Shocks, stocks and socks: smoothing consumption over a temporary income loss. Journal of the European Economic Association 7 (6), 1169–1192.

Cutler, D.M., Katz, L.F., 1991. Macroeconomic performance and the disadvantaged. Brookings Papers on Economic Activity 1991 (2), 1–74.

Deaton, A., Paxson, C., 1994. Intertemporal choice and inequality. Journal of Political Economy 102 (3), 437–467.

Dynarski, S., Gruber, J., 1997. Can families smooth variable earnings? Brookings Papers on Economic Activity 8 (1), 229–303

Goldberger, A., 1972. Maximum-likelihood estimation of regressions containing unobservable independent variables. International Economic Review 13 (1), 1–15.

Guvenen, F., 2007. Do stockholders share risk more effectively than non-stockholders? Review of Economics and Statistics 89 (2), 275–288

Hobijn, B., Lagakos, D., 2005. Inflation inequality in the united states. Review of Income and Wealth 51 (4), 581–606.

Kaplan, G., Violante, G.L., 2009. How much insurance in Bewley models? Unpublished manuscript.

Krueger, D., Perri, F., 2004. Understanding consumption smoothing: evidence from US consumer expenditure data. Journal of the European Economic Association 3 (2–3), 340–350.

Mace, B.J., 1991. Full insurance in the presence of uncertainty. Journal of Political Economy 99 (5), 928–956.

Meyer, B.D., Sullivan, J.X., 2003. Measuring the well-being of the poor using income and consumption. Journal of Human Resources 38, 1220–1880.

Nelson, J.A., 1994. On testing for full insurance using consumer expenditure survey data. Journal of Political Economy 102 (2), 384–394.

Pagan, A., 1984. Econometric issues in the analysis of regressions with generated regressors. International Economic Review 25 (1), 221–247.

Storesletten, K., Telmer, C.I., Yaron, A., 2004a. Consumption and risk sharing over the life cycle. Journal of Monetary Economics 51 (3), 609–633.

Storesletten, K., Telmer, C.I., Yaron, A., 2004b. Cyclical dynamics in idiosyncratic labor-market risk. Journal of Political Economy 112 (3), 695–717.

Wooldridge, J.M., 2002. Econometric Analysis of Cross Section and Panel Data. MIT Press, Cambridge, MA.

Zellner, A., 1970. Estimation of regression relationships containing unobservable independent variables. International Economic Review 11 (3), 441–454.