



Contents lists available at ScienceDirect

## European Economic Review

journal homepage: [www.elsevier.com/locate/eer](http://www.elsevier.com/locate/eer)Consumption in the shadow of unemployment<sup>☆</sup>Rodolfo G. Campos<sup>a,\*</sup>, Iliana Reggio<sup>b</sup><sup>a</sup> Department of Economics, IESE Business School, Camino del Cerro del Águila 3, 28023 Madrid, Spain<sup>b</sup> Department of Economics, Universidad Carlos III, Calle Madrid 126, 28903 Getafe, Madrid, Spain

## ARTICLE INFO

## Article history:

Received 14 September 2014

Accepted 25 April 2015

Available online 14 May 2015

## JEL classification:

E21

E24

D12

D84

## Keywords:

Consumption

Unemployment

Life-cycle models

Spain

Great Recession

## ABSTRACT

By how much do employed households reduce their consumption when the aggregate unemployment rate rises? In Spain during the Great Recession a one point increase in the unemployment rate was related to a strong reduction in household consumption of more than 0.7% per equivalent adult. This reduction is consistent with forward-looking agents responding to downward revisions of their expectations on future income growth rates: the shadow of unemployment. Using consumption panel data that include information on physical quantities we show that the drop in consumption expenditure was truly a reduction in quantities, and not a switch to cheaper alternatives.

© 2015 Elsevier B.V. All rights reserved.

## 1. Introduction

In a typical country, between half and two-thirds of GDP corresponds to household consumption. Understanding the process of how households adjust their consumption expenditure in the face of worsening labor market conditions is important for research on the dynamics of aggregate consumption, business cycles in general, and public policy. For example, the evolution of consumption and unemployment affects fiscal policy because it directly impacts government revenues and expenses through taxes and transfers. A fall in consumption in tandem with an increase in the unemployment rate can therefore severely strain the budget balance, especially during periods of depressed economic activity.

During recessions, some households experience unemployment directly but a relatively large fraction of households remains relatively unaffected. In most households, the person who contributes the largest share to a household's labor income—the primary earner—remains employed. Previous research has extensively focused on the consumption decisions of households that experience unemployment (e.g., Gruber, 1997; Browning and Crossley, 2001). However, and despite of its aggregate implications, the question of how the large fraction of households who are not directly hit by unemployment individually react to a rising unemployment rate is not definitely settled. Do these households reduce their consumption in response to a rise in the aggregate unemployment rate? By how much, and, if so, why?

<sup>☆</sup> Part of this research was carried out while Iliana Reggio was visiting Banco de España.

\* Corresponding author.

E-mail addresses: [rcampos@iese.edu](mailto:rcampos@iese.edu) (R.G. Campos), [ireggio@eco.uc3m.es](mailto:ireggio@eco.uc3m.es) (I. Reggio).

To answer these questions, we use a representative sample of Spanish households. Using a sub-sample of households whose primary earner was not hit by unemployment, we estimate whether a rising unemployment rate had a sizable effect on their consumption. We choose to focus on Spain in the Great Recession because, in addition to having high quality consumption data at the household level over the period of interest, the country experienced a rapid and unexpected rise in the unemployment rate. Over the period 2006–2011, the average Spanish unemployment rate increased by more than 13 percentage points, from less than 8.5% to 21.6%.<sup>1</sup>

We find that Spanish households responded strongly to the rising unemployment rate. Households in which the primary earner stayed employed reduced their consumption per equivalent adult by more than 0.7% per point increase in the unemployment rate. What makes these households respond so strongly to the aggregate unemployment rate? We find that the explanation does not lie in a drop in contemporaneous household income. Instead, we argue that aggregate unemployment casts a shadow on future income growth rates. A rise in the unemployment rate indicates lower future income growth. In turn, forward-looking behavior implies that households adjust consumption downwards.

A challenge in our estimation is how to distinguish the effect of the aggregate unemployment from other time-varying variables at the macro level. Our empirical strategy is to identify the effect of the unemployment rate on consumption from the variation of the unemployment rate calculated for groups of households of similar demographic characteristics. The intuitive idea behind it is that economic agents will predominantly respond to labor market conditions of population groups that are most similar to them. We form groups based on age and education attainment of the primary earner. Because education is largely predetermined for the age groups we consider, both categories can be assumed to be exogenous to the consumption decision.

Our empirical strategy requires the availability of an appropriate data set—one in which the effect the unemployment rate exhibits sufficient variation for successful identification, especially on the cross-sectional dimension. The Spanish experience during the Great Recession is ideally suited for this task. Not only did the average unemployment rate unexpectedly skyrocket over the period 2006–2011, this increase in the unemployment rate was not homogeneous across population groups. In particular, the rise in the unemployment rate was relatively larger for younger and less educated people.

Our focus on a sample of employed households has a number of advantages. First, using the restricted subset allows us to tie the unemployment rate to a particular time horizon over which income expectations are formed. We use the theoretical implications of a standard life cycle/permanent income hypothesis (LC/PIH) model to distinguish between the effect on consumption of changes in innovations to current and future income growth. Because income by employed households is not immediately affected by a rise in the unemployment rate (unlike what happens for households transitioning into unemployment), the consumption response of employed households to the unemployment rate measures adjustments to changes in expectations about future, not current, income growth.

The second advantage of focusing exclusively on employed workers is that it removes some of the caveats that arise when using a composite consumption good. Consumption responses of unemployed workers are likely to differ from employed workers along several dimensions—some of them unobserved. As argued by [Browning and Crossley \(2009\)](#), there may be changes in the composition of expenditures when a household becomes unemployed. Expenditures related to work, such transport and clothing, are likely to be differentially reduced by unemployed workers. Using a composite consumption measure would not be appropriate in a mixed sample that also includes unemployed households.

Third, the restriction to households who remain employed removes potential pitfalls in the estimation due to unobservable missing variables. For example, [Carroll et al. \(2003, p. 587\)](#) argue that unemployed and employed workers differ in their response to an increase in unemployment risk if they have accumulated savings for precautionary reasons. Upon experiencing unemployment a household with accumulated savings will run down these savings, despite a worsening labor environment, whereas an employed household will not. If the stock of savings is unobserved, as it usually is in consumption surveys, the response by unemployed workers will have an additional layer of unobserved heterogeneity relative to that of employed workers.

Research closely related to ours includes that of [Stephens \(2004\)](#) and [Benito \(2006\)](#), who studied the response of consumption to changes in subjective job-loss probabilities in the US and in the UK. Not many surveys contain information both on job-loss probabilities and consumption. In fact, [Stephens \(2004\)](#) and [Benito \(2006\)](#) are forced to restrict their analysis to food items. In comparison, the use of unemployment rates allows us to use more comprehensive definitions of consumption.<sup>2</sup> Because [Stephens \(2004\)](#) also uses a LC/PIH framework for interpretation, his findings are of particular interest because they are easily comparable to our results. In fact, our findings are complementary to his.

[Stephens \(2004\)](#) finds that changes in subjective job-loss probabilities do not have an effect on consumption of employed workers. Because he—like us—frames his results using a LC/PIH model, his results can be given a precise structural interpretation. The timing of the variables he uses implies that he tests whether changes in expectations about current income affect consumption; the answer is negative. In comparison, we test whether changes in expectations to future income affect consumption; this time the answer turns out to be positive. Taken together, these results imply that bad news

<sup>1</sup> In comparison, in the US the unemployment rate averaged 4.8% in 2007, and peaked at 9.6% in 2010.

<sup>2</sup> An additional problem with subjective measures uncovered by [Stephens \(2004\)](#) is that, although relevant to predict future job losses, these probabilities tend to be overly pessimistic. Also, they exhibit extreme bunching at focal probabilities such as 0%, 50%, and a 100%.

about income growth contained in labor market indicators, the shadow of unemployment, affect consumption primarily through expectations about periods that lie in the future.

More generally, our research question is related to an extensive economic literature that relates consumption behavior to unemployment and, more comprehensively, to income shocks.<sup>3</sup> In particular, our results are consistent with previous results on forward-looking behavior using data from the US. Stephens (2001) finds that US households reduce their food consumption in years prior to a job loss. Nalewaik (2006) shows that US households alter nondurable consumption in response to income changes as far as six years into the future. Whereas Stephens (2001) and Nalewaik (2006) do not propose a mechanism through which households anticipate an increased probability of a future job loss or income drop, our results suggest that the unemployment rate may be the carrier of such information.

In the last part of our paper we establish whether the drop in consumption expenditure reflects a change in actual consumption or prices paid by households. The implications for welfare of whether a rising unemployment rate is related to drops in quantities or prices are quite different. Recent research by Aguiar and Hurst (2005, 2007) has highlighted the potential for home production to explain drops in expenditure, particularly those that occur at retirement. For Spain, Luengo-Prado and Sevilla (2013) found that part of the drop of expenditure at retirement is explained by a switch to home production and a drop in prices, not by a reduction of actual consumption. We adapt the methodology of Aguiar and Hurst (2007) and use quantity data that is available in the Spanish consumption survey. Whereas previous findings had uncovered that households exiting the labor force pay lower prices, our results indicate that the increase in the unemployment rate reduces the consumption by employed households with no discernible effect on prices.

## 2. Data and empirical strategy

To estimate the influence of the aggregate unemployment rate on consumption by households who are individually unaffected by unemployment we first classify households into demographic groups according to the level of education and age of the primary earner in the household. Then, we assign to each household the unemployment rate of its specific demographic group and estimate whether its change, which we denote with  $\Delta U_t$ , has a discernible effect on a household's consumption growth rate, which we denote with  $\Delta c_t$ .

### 2.1. Data sources and sample selection

We use the Spanish labor force survey, *Encuesta de Población Activa* (EPA), to calculate unemployment rates corresponding to the education–age groups. The unemployment rate evolved differently across groups. This is best seen in Fig. 1, in which we plot the evolution of the unemployment rate for each of the age–education groups over the period 2006–2011. Whereas unemployment rates for all groups were contained in the range 2.9%–14.6% in 2006, they fanned out, increasingly so after 2007, to cover a range three times as large, 7.2%–43.6%, in 2011. The lowest unemployment rates correspond to the population with higher education aged 45–54 and 55–64. The highest unemployment rate, on the other side of the spectrum, corresponds to the group of those in the lowest education category and who are aged less than 30.


Over the period 2006–2011, there was a fair amount of variation in group unemployment rates. The cross-sectional and longitudinal variation were of roughly the same size; indeed, slightly larger in the cross-section. Across groups, and over the whole period, the unemployment rate was on average 14.1% with an overall standard deviation of 9.3%. The between standard deviation stood at 6.9% and the within standard deviation at 6.5%. According to our empirical strategy, the sizable amount of cross-sectional variation exhibited by the Spanish data is consequential for identifying the effect of the unemployment rate on consumption.

Our household consumption data is obtained from the Spanish consumer expenditure survey EPF (*Encuesta de Presupuestos Familiares. Base 2006*), a survey that provides detailed information on consumption, unemployment, and socioeconomic characteristics at the household level for the period 2006–2011. Whereas the Consumer Expenditure Survey (CEX), the consumption survey commonly used for the US, has been shown to have some shortcomings, Spanish consumption data has fared better in validation studies. It is well known that consumption measured in the CEX has important discrepancies with Personal Consumption Expenditure (PCE), the aggregate consumption series in the US (Slesnick, 1992; Garner et al., 2006; Heathcote et al., 2010). The two series exhibit a large gap. In contrast, consumption measured in the EPF accounts for 87% of Spanish aggregate consumption during the period 2006–2011. In this respect the EPF also compares favorably to prior consumption expenditure surveys in Spain. Because the Spanish statistics agency made several enhancements aimed at increasing coverage when it transitioned to the new EPF in 2006, the coverage ratio in the EPF is higher than in the preceding surveys, the ECPF-85 (which ran from 1985 to 1996) and the ECPF-97 (which ran from 1997 to 2005).<sup>4</sup>

Consumption measured in the CEX is also less pro-cyclical than aggregate consumption. Campos et al. (2013) show that consumption measured from the CEX underestimates the cyclical correlation of aggregate consumption (PCE) with GDP by

<sup>3</sup> Jappelli and Pistaferri (2010) provide a useful survey of this literature in which they distinguish between anticipated and unanticipated income shocks.

<sup>4</sup> Pou and Alegre (2002) calculated that consumption measured in the ECPF-85 over the whole period 1985–1996 accounted for 80% of aggregate consumption. Luengo-Prado and Sevilla (2013) report a coverage ratio in the ECPF-97 of 85%, on average, over the period they consider.



**Fig. 1.** The unemployment rate across education and age groups. The figure plots the evolution of the unemployment rate for each of the age-education groups. Group unemployment rates were calculated from the Spanish labor force survey (EPA). Data plotted in this figure are tabulated in Appendix A (Table 11).

40%. It is therefore likely that consumption measured from the CEX underestimates its co-movement with unemployment, which is a very cyclical variable. For our purposes, an additional shortcoming of the CEX is that consumption and income are observed asynchronously, which may lead to inference problems. For example, [Gervais and Klein \(2010\)](#) show that this causes the estimators of risk-sharing tests that rely on data from the CEX to be inconsistent. In contrast, in the EPF consumption and income belong to the same period.

The EPF samples households on a yearly basis and provides population weights that can be used to obtain results that are representative of the Spanish population. Consumption items are classified using the COICOP/HBS classification. Our measure of consumption is defined as expenditure on nondurable consumption goods and services. We obtain real household consumption expenditure by adjusting for inflation using the Spanish price index (*IPC Base 2006*). We calculate consumption by equivalent adult by adjusting household consumption with the OECD equivalence scale to take into account possible economies of scale in consumption. In the Appendix we show that using another equalization procedure or using per-capita variables does not make a difference.

The EPF consists of a panel that follows households for two periods. Because we first difference data we are effectively working with a cross-section. We focus on the working-age population and restrict the sample to those households in which the primary earner is aged 25–64. We start with a sample of 30,036 households. In our baseline specification we consider households in which the primary earner is employed in both waves (dropping 6727 households).<sup>5</sup> We drop households in which the primary earner is self-employed (dropping 4944 households). We keep only those households that report positive values on food consumption (dropping 70 households). Finally, we eliminate households in which the identity of the primary earner changes from one year to the other (we drop 1101 households). Missing values of some controls led to dropping 9 additional households. Our final sample consists of 17,182 households.

## 2.2. Empirical strategy

### 2.2.1. Consumption and unemployment

Household consumption  $C$  is the sum of expenditure on nondurables and services. To take into account possible economies of scale we express it as consumption per equivalent adult using the OECD scale. Households are interviewed twice, in two consecutive years. Our variable of interest is yearly household consumption growth  $\Delta C_t \equiv \log C_t - \log C_{t-1}$ .

We group households according to the education level and age of the primary earner (in the second interview) and calculate a group-specific unemployment rate as the ratio of the unemployed to the labor force in each one of these groups. The four education groups are: less than 1st cycle in high school, completed 1st cycle in high school, completed 2nd cycle in high school, higher education. Age groups are also four: less than 30, 30–44, 45–54, and 55–64.<sup>6</sup> The change from one year to the other of the group unemployment rate is  $\Delta U_t \equiv U_t - U_{t-1}$ .

<sup>5</sup> There is no way to be certain that an employed primary earner did not experience unemployment spells in between. However, the definition of primary earner applies to the individual with the highest earnings in the household, therefore reducing the likelihood of unobserved unemployment spells. Because we include income changes and a temporary contract dummy as additional controls, the potential selection problem is further ameliorated. We also estimated our baseline equation excluding households with temporary contracts and found similar results.

<sup>6</sup> This classification follows the standard classification used by INE when cross-tabulating data from its labor force survey according to education and age. In the Appendix we show that our results are robust to a finer classification of age, using 5-year age bins, and also to an alternative definition of groups according to whether the primary earner's job is skilled or unskilled.

### 2.2.2. Specification and estimation

Letting  $i$  index households,  $g$  groups, and  $t$  stand for time, we regress household consumption growth on group unemployment using the following specification:

$$\Delta c_{igt} = \gamma \Delta U_{gt} + \beta' X_{igt} + \alpha_t + \varepsilon_{igt}. \quad (1)$$

The covariates in  $X$  vary at the household level whereas  $\Delta U$  is measured only at the group level;  $\alpha$  is a time dummy, and  $\varepsilon$  represents the estimation error. The partial effect of the unemployment rate of the group on consumption growth is captured by the coefficient  $\gamma$ . An increase by one percentage point in the unemployment rate implies a change in  $\gamma$  log-percentage points in household consumption. We use the population weights available in the consumption survey to estimate Eq. (1) by weighted least squares.

The fact that  $\Delta U$  varies at the group level may lead to within-group correlation and pose a potential problem in obtaining standard errors for  $\hat{\gamma}$ , as illustrated by Moulton (1990). We therefore cluster standard errors by groups and report these adjusted standard errors in all tables. Because the number of groups in our data is relatively small, an issue discussed by Wooldridge (2003) and Cameron et al. (2008), we apply finite sample bias correction and compute the  $p$ -values for the null hypotheses that the coefficients on  $\Delta U$  are zero using a  $t$ -distribution, instead of a standard normal distribution. We performed several robustness checks on the estimation of these standard errors. We computed bootstrapped standard errors and found similar significance levels. We also ran estimations by group and obtained even stronger results.

### 2.2.3. Preference shifters and additional controls

Our specification includes time dummies to capture changes in the macroeconomic environment, as well as variables that are likely to affect the level of consumption and the consumption profile. Variables that affect the level of consumption enter the specification in first differences, and those that influence the consumption profile enter the specification in levels. For the level of consumption, our variables include household size and the number of kids below 16 or dependents below 25. Variables that affect the consumption profile consist of dummies for educational attainment, a dummy for work in a skilled occupation, gender, marital status, and age dummies for the primary earner to capture life-cycle effects. We include regional dummies to control for systematic differences in the evolution of consumption across regions. We define regions as the *Comunidades Autónomas* (CCAA), the first-level political and administrative division of Spain.<sup>7</sup>

We introduce a set of additional controls. A first concern might be that consumption changes may depend on wealth, or access to credit. We use home ownership as a proxy for wealth and include dummies for owning a primary home or a secondary home as additional controls. We also control for whether the household has a mortgage. Having a mortgage proxies for access to credit and, since households with a mortgage are a subset of those who own a home, it also allows for a more flexible relationship between home ownership and consumption.

A second concern is that changes in consumption might reflect differentials in the availability of income sources unrelated to the labor market. Because our sample includes urban and rural households, we include a dummy for rural households who may, in principle, derive a substantial part of their income from agricultural activities.

Our sample consists of primary earners who are continuously employed. However, the unemployment rate faced by a household's primary earner is potentially correlated with the labor market situation of other household members. For this reason we also include the change in the number of adults employed other than the primary earner and the change in the number of unemployed in the household as additional controls. An added benefit from including employment variables of household members is that they may be controlling for potential effects of leisure within the household on the marginal utility of consumption.<sup>8</sup>

To address the dual nature of the Spanish labor market we include a dummy for when the primary earner has a temporary labor contract. Previous research has argued that the dual structure of labor contracts is a key characteristic of the Spanish labor market and that it explains the evolution of unemployment (Costain et al., 2010; Bentolila et al., 2012).<sup>9</sup>

### 2.2.4. Income

Because increases in the unemployment rate may be correlated with drops in current household income across groups it is convenient to include measures of income growth as additional regressors. A first approach is to include income growth at the household level, defined as the growth rate of the logarithm of household income  $\Delta y_t \equiv \log Y_t - \log Y_{t-1}$ , as an additional regressor in  $X$  on the right hand side of (1).

<sup>7</sup> There are 18 regions: 15 on the Iberian peninsula, 2 corresponding to the Balearic and Canary Islands, and 1 region which consists of Ceuta and Melilla, two autonomous cities on the African continent which are bundled together in the data.

<sup>8</sup> In our data we cannot distinguish between full-time and part-time-work. However, in Spain part-time jobs are not very prevalent. According to data from EUROSTAT, over the period 2006–2011, the percentage of workers in part-time jobs in Spain was, on average, 12.6. In comparison, on average over the same period, this percentage was 18.7 for the European Union (EU27) and 19.9 for the Euro area (EA17).

<sup>9</sup> Costain et al. (2010) show that labor market duality, i.e., the coexistence of temporary contracts with low firing costs and permanent contracts with high firing costs, increases the response of the unemployment rate to macroeconomic shocks. Bentolila et al. (2012) use a calibrated search model and argue that Spain could have avoided about 45% of its unemployment surge during the Great Recession had it adopted the French employment protection legislation.

The inclusion of income growth on the right hand side of (1) can be contentious. It is well known that survey data on income is likely to contain considerable measurement error (Altonji and Siow, 1987), and may therefore potentially bias our results. We have countered this potential threat in a number of ways. As a first check, we have run all our regressions with and without income growth and found that the coefficient on  $\Delta U_t$  is not much affected. Second, we have tested what happens if household income growth is replaced by average income growth  $\overline{\Delta y}_t$ , taking the average over households that belong to the same education–age group. Because we are averaging, part of the measurement error is likely to be averaged out. Again, results are not much affected.

The group average of income growth is also of interest for reasons unrelated to measurement error. Because  $\Delta U_t$  is measured as a group average, constructing  $\overline{\Delta y}_t$  in the same way provides a natural counterpart to it. If  $\overline{\Delta y}_t$  is included as a right hand side variable, the change in the unemployment rate effectively measures the effect of the unemployment rate through channels other than income changes common to the group.

An additional approach is to separate household income growth into  $\overline{\Delta y}$ , the part explained by the demographic group, and the part left unexplained:  $\Delta y - \overline{\Delta y}$ . The interpretation of the remainder  $\Delta y - \overline{\Delta y}$  is that it is the part of a household's income growth that is idiosyncratic to the household, i.e., not explained by the group average. Because the group growth rate and the idiosyncratic growth rate sum, by definition, to household income growth, this can be interpreted as a more flexible version of the first approach which included  $\Delta y$ . As we discuss in Section 3.2, this separation also aids the interpretation in terms of the effect of unemployment on current and future income growth.

### 3. Results

In Spain during the Great Recession households that did not themselves experience unemployment significantly reduced their consumption in response to a rise in the unemployment rate. A one point increase in the average unemployment rate implied a drop of more than 0.7% in household consumption per equivalent adult. This result is robust to the inclusion of different income measures.

In the Appendix we include several robustness checks. We consider the use of alternative equivalence scales and alternative definitions of the group unemployment rate. To exclude the possibility that outliers are driving our results we

**Table 1**  
Consumption response to the unemployment rate and income measures.

Variables	(1) Without income	(2) Total income change	(3) Group average income change	(4) Average + idiosync. income change
$\Delta U$	-0.752*** (0.233)	-0.737*** (0.227)	-0.736*** (0.234)	-0.735*** (0.231)
$\Delta y$		0.104*** (0.014)		
$\overline{\Delta y}$			0.077 (0.104)	0.113 (0.098)
$\Delta y - \overline{\Delta y}$				0.104*** (0.015)
Observations	17,182	17,182	17,182	17,182
$R^2$	0.027	0.032	0.027	0.032

Regression of consumption expenditure growth on  $\Delta U$ , various income growth variables, and controls described in Section 2.2. All specifications include time dummies. Standard errors (in parentheses) are clustered by group. We compute the  $p$ -values using a  $t$ -distribution.

\*\*\* Indicates significance at the 1% level.

**Table 2**  
The group unemployment rate as a predictor of income expectations.

Variables	OLS	Pooled logit	Panel logit
$U$	0.608* (0.314)	1.074** (0.037)	1.055** (0.026)
Observations	5829	5829	5772

Regressions of a dummy variable indicating whether  $E_t \Delta y_{t+1} < 0$  on group unemployment and controls including age, education, gender, marital status, household size, home ownership, wealth, income, and time dummies. Standard errors (in parentheses) in pooled regressions are clustered by household.

\* Indicate significance at the 10% level.

\*\* Indicate significance at the 5% level.

use a median regression. We estimate our baseline regression on the sub-sample of male primary earners, and we check whether the presence of employed household members diminishes the effect of unemployment on consumption. We conclude that our result is robust to all these checks. Finally, we also consider alternative specifications allowing heterogeneous responses according to the magnitude of group unemployment rates but the evidence points towards a homogeneous response.

In Table 1 we show results of estimating the equation specified in (1) for different variables measuring income. We start by excluding income changes (Column 1) from the regressors. Without conditioning on income, consumption is related to a three-quarter of a percentage drop per point of increase in the unemployment rate. If  $\Delta y$ , the log-change in current household income, is added, the effect is only slightly lower, with a point estimate of  $-0.737$  (Column 2). The last two columns, which include the average change in income  $\overline{\Delta y}$  (Column 3) and the average change together with the idiosyncratic deviation from this average  $\Delta y - \overline{\Delta y}$  (Column 4) do not alter the conclusion. In both cases, the point estimate is very close to  $-0.737$ . Using standard errors adjusted as described in Section 2.2, all of these estimates are significantly different from zero at the 1% level.

The marginal effect of the unemployment rate on consumption is remarkably robust to the different current income measures. The results in Table 1 suggest that changes in the unemployment rate do, in fact, have a significant effect on household consumption growth that is unrelated to variation in current income.

Several surveys lack data on consumption items beyond food consumption. In some cases, a complementary data source is available, making it possible to impute consumption data (e.g., Skinner, 1987; Blundell et al., 2008; Campos and Reggio, 2014). This is not always the case, however. Therefore, much of previous research on consumption has used expenditure on food, both at home and away from home, in lieu of expenditure on nondurables and services. To make our results comparable with this previous literature we repeat our regressions for food items. As shown in the Appendix, the effect of the unemployment rate is similar.

### 3.1. The unemployment rate and future income expectations

Why do employed households reduce their consumption expenditure at times in which the unemployment rate rises? A possible explanation is that a rise in the unemployment rate is related to downward revisions of future income expectations. This association, that households adjust income expectations in response labor market conditions has been made before, e.g., by Stephens (2004) and Benito (2006) with the use of subjective job loss probabilities.

Using data from the Spanish Survey of Household Finances (EFF by its Spanish acronym) we can empirically contrast whether a rise in the group unemployment rate is related to a drop in income expectations. The EFF is a panel data set for which three waves (2002, 2005, and 2008) are available. The EFF contains a question in which households respond whether they believe that their household income over the next year will increase, stay the same, or decrease.

In Table 2 we report the results of running a regression of a dummy variable indicating whether income is expected to decrease over the next year (whether  $E_t \Delta y_{t+1} < 0$ ) on the group unemployment rate and controls similar to those in our

**Table 3**  
Consumption response depending on the age of the primary earner.

Variables	(1) Baseline	(2) By age
$\Delta U$	$-0.735^{***}$ (0.231)	
$\Delta U \times PE_{25-44}$		$-0.782^{***}$ (0.256)
$\Delta U \times PE_{45-54}$		$-0.578^{**}$ (0.251)
$\Delta U \times PE_{55-64}$		$-0.222$ (0.727)
$\overline{\Delta y}$	0.113 (0.098)	0.139 (0.105)
$\Delta y - \overline{\Delta y}$	$0.104^{***}$ (0.015)	$0.104^{**}$ (0.015)
Observations	17,182	17,182
$R^2$	0.032	0.033

The first column is the baseline regression of expenditure growth, on  $\Delta U$ , income growth variables, and controls described in Section 2.2. All specifications include time dummies. In the second column  $\Delta U$  is interacted with age dummies. Standard errors (in parentheses) are clustered by group. We compute the  $p$ -values using a  $t$ -distribution.

\*\* Indicates significance at the 5% level.

\*\*\* Indicates significance at the 1% level.

**Table 4**  
Consumption response by sector and contract type.

Variables	(1) Baseline	(2) By sector	(3) By contract type
$\Delta U$	-0.738*** (0.233)		
$\overline{\Delta y}$	0.110 (0.099)	0.103 (0.099)	0.107 (0.097)
$\Delta y - \overline{\Delta y}$	0.103*** (0.015)	0.103*** (0.015)	0.103*** (0.015)
$\Delta U \times \text{Private}$		-0.765*** (0.257)	
$\Delta U \times \text{Public}$		-0.514 (0.298)	
$\Delta U \times \text{Temporary}$			-0.844* (0.409)
$\Delta U \times \text{Indefinite}$			-0.703*** (0.216)
Observations	17,174	17,174	17,174
$R^2$	0.032	0.032	0.032

The first column is the baseline regression of expenditure growth, on  $\Delta U$ , income growth variables, and controls described in Section 2.2. All specifications include time dummies. In the second column  $\Delta U$  is interacted with private sector and public sector dummy variables. In the third column  $\Delta U$  is interacted with temporary and indefinite contract dummy variables. Standard errors (in parentheses) are clustered by group. We compute the  $p$ -values using a  $t$ -distribution.

\* Indicates significance at the 10% level.

\*\*\* Indicates significance at the 1% level.

consumption specification. We report pooled OLS results in Column 1 and odds ratios from logit regressions in Columns 2 and 3.

The three regressions yield the same conclusion: rises in the group unemployment rate are positively related to expected income drops over the next year, implying a negative relationship between  $U_t$  and  $E_t \Delta y_{t+1}$ . Expectations of future income covary negatively with group unemployment.

Summing up, during the Great Recession in Spain, household consumption expenditure exhibited a response to the unemployment rate that was quantitatively large. This drop in consumption was not due to a reduction of current income. Because the unemployment rate is related to expectations of future income growth, this suggests that households were responding to worsening income expectations and that the quantitatively large effect on consumption, and therefore on domestic demand, operated through expectation changes.

### 3.2. Consumption and income expectations in an LC/PIH setting

Why would consumption react to an increase in the unemployment rate when the household itself is not affected by unemployment? Workers do not experience an immediate drop in income if they stay employed. However, the unemployment rate, particularly that of similar households, is a signal of the evolution of their expected future labor income. According to economic theory, forward-looking agents with a consumption smoothing motive will choose to adjust their current consumption downward.

Models that belong to the life-cycle/permanent income hypothesis (LC/PIH) framework have been used extensively to study the intertemporal allocation of resources by individuals with consumption smoothing motives. The LC/PIH framework imposes structure on how consumption responds to changes in expectations of current and future income and, indirectly, to changes in the unemployment rate.

Tracing back to the work by Friedman (1957), the basic feature of LC/PIH models is that consumption is a function of the entire discounted expected future income stream of a household. In such models consumption should not respond to predictable income changes but instead to income innovations. Whenever expectations of future income change, they should immediately be reflected in a consumption adjustment.

Adapting the model by Flavin (1981), Campbell and Deaton (1989) derive an expression relating consumption innovations to income innovations. They obtain a closed form solution at the cost of assuming a quadratic utility function. Comparable solutions can be obtained for more general utility specifications by relying on approximation methods. Nalewaik (2006) derives an expression similar to that of Campbell and Deaton (1989) by using a log-linear approximation to the consumption function without assuming any particular utility function.

We will use the expression derived by Nalewaik (2006) to interpret our results. Recall that  $\Delta U_t \equiv U_t - U_{t-1}$  denotes the change of the unemployment rate between time  $t-1$  and time  $t$ , and the expressions  $\Delta c_t \equiv \log C_t - \log C_{t-1}$  and



**Table 5**  
Expenditure versus prices.

Variables	(1) Expenditure	(2) Price index	(3) Price index 2006 prices	(4) Price change 2006 prices
$\Delta U$	-1.000*** (0.273)	0.131 (0.101)	0.139 (0.116)	0.124 (0.107)
$\overline{\Delta y}$	-0.033 (0.124)	0.113* (0.059)	0.112* (0.063)	0.150** (0.064)
$\Delta y - \overline{\Delta y}$	0.078*** (0.014)	-0.001 (0.005)	-0.003 (0.005)	-0.002 (0.003)
Observations	17,162	17,162	17,162	17,161
$R^2$	0.020	0.151	0.157	0.017

Regression of expenditure growth, price indices, and price changes on  $\Delta U$ , income growth variables, and controls described in Section 2.2. All specifications include time dummies. Standard errors (in parentheses) are clustered by group. We compute the  $p$ -values using a  $t$ -distribution.

\* Indicate significance at the 10% level.

\*\* Indicate significance at the 5% level.

\*\*\* Indicate significance at the 1% level.

$\Delta y_t \equiv \log Y_t - \log Y_{t-1}$  stand for the logarithmic difference of consumption and income.<sup>10</sup> Because these last two expressions are logarithmic differences, they can be interpreted as growth rates.

Expressing the discount factor as  $\lambda = 1/(1+r)$ , where  $r$  is the interest rate, and using our timing convention, the relationship between consumption growth and expected income growth derived by Nalewaik (2006, Eq. (3)) is

$$\Delta c_t = \sum_{j=0}^{\infty} \lambda^j (E_t - E_{t-1}) \Delta y_{t+j}, \quad (2)$$

where  $E_t$  and  $E_{t-1}$  denote the conditional expectations taken with information available at time  $t$  and  $t-1$ .<sup>11</sup> The consumption growth rate  $\Delta c_t$  equals the discounted sum of innovations at time  $t$  to the expectations of the entire stream of future income growth rates  $\Delta y_{t+j}$ .

To distinguish between the present and the future it is useful to separate the first time period in the summation and write this equation as

$$\Delta c_t = (E_t - E_{t-1}) \Delta y_t + \sum_{j=1}^{\infty} \lambda^j (E_t - E_{t-1}) \Delta y_{t+j}. \quad (3)$$

The first term on the right hand side is the innovation to *current* income growth whereas, in the second term, the summation is over innovations to *future* income growth. Even if households stay employed, and are therefore likely to have a stable income in the short run, a rise in the unemployment rate  $\Delta U_t$  can affect consumption growth through its effect on future periods. In response to a rise in  $\Delta U_t$  households may revise expectations of future income growth downward. If they do so for at least some of the  $(E_t - E_{t-1}) \Delta y_{t+j}, j \geq 1$ , the whole discounted sum over future income growth rates is reduced. In short, new information on the unemployment at time  $t$  depresses consumption immediately because it conveys bad news about future income growth. Households respond to the shadow of future unemployment even though they are not (yet) affected by it.

In theory,  $\Delta U_t$  could also have an influence on  $(E_t - E_{t-1}) \Delta y_t$ , the innovation to current income growth. There are two ways in which this might be the case: either because (a) the overall unemployment rate is correlated with idiosyncratic transitions into unemployment, and through them with current income growth, or because (b) the unemployment rate and current income growth are correlated through channels other than idiosyncratic unemployment if they share a common (possibly unobservable) factor. As an example of (b), a rise in  $\Delta U_t$  may be the consequence of reduced labor demand for workers of the demographic characteristics we use to group households (age and education). If this unobservable reduction in labor demand translates into less bargaining power for workers of these characteristics it could possibly lead to lower income growth in the current period.

The link between  $\Delta U_t$  and current income growth described in (a) works through idiosyncratic unemployment at the household level. Our restriction of the sample to only households in which the primary earner stays employed already counters part of this concern. In addition, our data include the employment status of all other household members and we include unemployment of other household members as an additional control in all our regressions. The link described in (b) requires to explicitly account for current income growth at the group level in the estimation. As we describe in Section 2.2,

<sup>10</sup> This follows the convention that an uppercase letter indicates a variable in levels and a lowercase letter the natural logarithm of that variable.

<sup>11</sup> A detailed derivation of this equation is contained in Appendix A of Nalewaik (2006). This equation is the result of log-linearizing the present-value expression of the budget constraint of a household in a model with exogenous labor income and a constant interest rate coupled with the common assumption that consumption growth at time  $t+j$  is unpredictable at time  $t$ .

we include measures of current income growth, at the household and group level, as additional regressors when estimating the relationship between  $\Delta U_t$  and  $\Delta c_t$ .

How well the innovation to current income  $(E_t - E_{t-1})\Delta y_t$  in the theoretical relationship is accounted for in the estimation depends on what is assumed on the formation of expectations. Notice that this first term in (3) can be written as the difference between current income growth  $\Delta y_t$ , which is directly observed in the data, and the expectation at time  $t-1$  of current income growth  $E_{t-1}\Delta y_t$ , which is not:

$$(E_t - E_{t-1})\Delta y_t \equiv E_t\Delta y_t - E_{t-1}\Delta y_t = \Delta y_t - E_{t-1}\Delta y_t. \quad (4)$$

If the expectation of income growth is common across households, then the influence of  $E_{t-1}\Delta y_t$  will be captured by the time dummies, present in all our estimations, and including  $\Delta y_t$  will be sufficient to control for innovations to current income (and therefore tighten the link between  $\Delta U_t$  and innovations to future income growth). Of course, specifying common expectations for all households may be too restrictive. Alternatively, expectations could differ across households. If expectation formation is related to education and age, then not taking account of it may be problematic because these are the dimensions along which our main variable of interest  $\Delta U_t$  varies. We therefore consider an alternative scenario in which we assume that households can correctly forecast income growth at the group level, so that variation in  $E_{t-1}\Delta y_t$  is captured by variation in  $\overline{\Delta y}_t$ . In this case, what needs to be included as a right hand side variable is the idiosyncratic deviation from the group average:  $\Delta y_t - \overline{\Delta y}_t$ . Notice that our estimations cover these two cases.<sup>12</sup>

From the above discussion it should become clear that, unless the exact formation of expectations in the population is known, it is not possible to entirely rule out a connection between  $\Delta U_t$  and current income growth. However, the stability of the coefficient in the face of different specifications for income growth goes a long way in diffusing concerns about exclusively equating changes in  $\Delta U_t$  with changes in expected future income growth. In particular, because  $\overline{\Delta y}_t$  is calculated conditioning on the same variables as  $\Delta U_t$ , using it to proxy for expectations seems to be a good way of preventing the attribution of spurious explanatory power to  $\Delta U_t$ .

### 3.3. The shadow of unemployment

The main message from the results reported in Table 1 is that consumption was negatively and significantly affected by changes in the unemployment rate. Moreover, the inclusion of current income growth rates only slightly reduced the point estimate of the effect of the change in the unemployment rate  $\Delta U$ . In light of the results in Table 2 linking the unemployment rate with income expectations, together with the theoretical framework laid out in Section 3.2, this implies that Spanish households reduced their consumption in response to the realization of negative news on future income growth contained in the unemployment rate. During the Great Recession the rising unemployment rate cast a shadow on consumption.

Our findings complement those of Stephens (2004) for the US. He studied the effect of the subjective probability of being laid off on consumption using data from the Household Retirement Survey (HRS) and used an equation similar to our Eq. (3) for interpretation. Stephens (2004) finds that subjective probabilities are good measures of the objective occurrence of future layoffs, and therefore of future income growth. However his results are that consumption of employed workers is not affected by them. In the specification closest to ours (Stephens, 2004, Table A2, Columns 4 and 5), he does not find a significant effect of job loss probabilities, or expectation revisions of these probabilities, on log food consumption growth.<sup>13</sup> At first sight, our results would seem to contradict those of Stephens.

This is not so, however. The subjective probabilities used by Stephens (2004) are one-year ahead probabilities. Because he uses last year's probabilities and the arrival of information on whether a household has been laid off as the variable explaining consumption fluctuations, what Stephens (2004) is effectively doing in the subsample of workers who remained employed is identifying subjective probabilities with  $(E_t - E_{t-1})\Delta y_t$ , the first term in (3).<sup>14</sup> Therefore, he finds that innovations to current income growth have no effect on consumption. His results are silent on the effects on consumption of innovations to future income growth  $(E_t - E_{t-1})\Delta y_{t+j}$  for workers who do not lose their jobs.

Unfortunately, an exercise similar to that of Stephens (2004) cannot be replicated for Spain because the necessary data are lacking. Because the labor market in the US and in Spain exhibit considerable differences it is unclear to what extent what is learned for one can be transferred to the other. With this caveat in mind, if the results of Stephens (2004) and ours are jointly considered, they imply that consumption responds to news on the labor market primarily through its effect on expectations of income growth that lies at least some periods into the future.

<sup>12</sup> As a robustness check, we also estimated permanent income and used it as an alternative measure of expected income growth  $E_{t-1}\Delta y_t$  (we estimated permanent income by regressing log-income on occupation, region, education, age, age-squared, marital status, gender, home ownership, having a mortgage, the number of children in the household, and the number of earners and time dummies). All our results were essentially unaffected—this is not entirely surprising. Because permanent income is usually constructed from demographic characteristics of the household that tend to remain constant from one year to the other (with the exception of age), the time change of permanent income is roughly constant across households and idiosyncratic income deviations from permanent income will roughly coincide with  $\Delta y_t$ , the actual income growth rate.

<sup>13</sup> Findings are similar for the UK; food consumption does not respond to the subjective job loss probability in a statistically significant way (Benito, 2006, Table 3, Column 1).

<sup>14</sup> This is clear from Eq. (3) and the discussion on nondisplaced workers on Stephens (2004, p. 261).

Finally, the LC/PIH framework assumes that agents have access to only self-insurance. Agents live in the environment of the so-called Bewley model—after [Bewley \(1977\)](#)—in which they are able to transfer resources across periods via saving in an asset, but are otherwise in autarchy. Thus, although testing for risk-sharing goes beyond the purpose of this investigation, the absence of consumption insurance across households is consistent with the self-insurance assumption behind the LC/PIH model, the model we used to interpret our results.

### 3.3.1. Age and job stability

What is behind the shadow of unemployment? We have interpreted that Spanish households adjusted their consumption in response to negative news about their future signaled by the rising unemployment rate. If the unemployment rate is truly operating as a signal of future labor market conditions, then its effect should be weaker for older workers, who are closer to retirement, and therefore less affected by the labor market. At the same time, workers in the public sector, whose jobs are more stable should be less affected by their particular group unemployment rate. In Spain, as in much of Europe, public sector jobs are perceived to be more stable. Indeed, employment in the public sector has proven to be more resilient during the Great Recession; aggregate employment in the public sector increased throughout 2011 while private employment bore the brunt of the quantity adjustment in the labor market. As shown in this section, we find that the effect of the unemployment rate varies by age and sector in the way expected.

To establish whether workers close to retirement do, in effect, exhibit a muted response to the unemployment rate, we interact the change in group unemployment  $\Delta U$  with age dummies, according to the age of the primary earner. To achieve greater precision in the estimation we combine the two youngest groups, which are still far from retirement age, into a single group encompassing ages 25–44. Results are shown in [Table 3](#). Column 1 replicates the baseline specification for comparability, and Column 2 allows for different effects of  $\Delta U$  according to age groups.

The change in the unemployment rate has an effect on consumption that is strongest for younger workers, and weakest for households with primary earners aged 55–64, who are closest to retirement. In fact, the estimated effect for this group is not significantly different from zero. In contrast, households with primary earners younger than 55 exhibit negative and significant responses of consumption to the unemployment rate. The point estimate of the effect is estimated at  $-0.58\%$  for ages 45–54, and  $-0.78\%$  for ages 25–44. Households with primary earners farther away from retirement age, who will be part of the labor force for a longer period, and for whom the unemployment rate applies to a larger fraction of their income, exhibit the strongest response to the unemployment rate.

Turning to job stability, we classify households according to whether the job is in the private or the public sector and by type of contract. This information is not available for all households and reduces our sample to 17,174 households. We report the results from the baseline specification for this sub-sample in Column 1 of [Table 4](#).

Column 2 reports results on whether working in the public or private sector makes a difference in the response to changes in the unemployment rate. The change in the unemployment rate  $\Delta U$  enters the regression interacted with dummy variables indicating whether the primary earner works in the public sector or the private sector. For primary earners working in the private sector, a one percentage point in the unemployment rate is associated to a 0.77% drop in household consumption, which is close to the baseline result. For primary earners working in the public sector the point estimate is smaller, at  $-0.51\%$ , and not significantly different from zero.<sup>15</sup>

There is no large differential effect by type of contract. An increase of one percentage point in the unemployment rate is associated with a drop in consumption of 0.84% for households with primary earners with temporary contracts and of 0.70% for households with primary earners with indefinite contracts. In both cases there is a negative and significant effect of the unemployment rate.

Although the drop associated to temporary contracts is slightly larger, statistical tests do not reject equality of the coefficients. [Barceló and Villanueva \(2010\)](#) find that in Spain workers with temporary contracts accumulate more financial wealth to build a buffer stock of savings. Their finding can explain why the response by temporary contract workers is not as strong, because they can draw down their stock of savings, partially offsetting the effect of larger job instability.

The results in this section, that older workers and public sector employees are less responsive to the unemployment rate, provides an additional check on the interpretation of what households see in the unemployment rate. The evidence suggests that increases in the unemployment rate were, in effect, related to worsening labor market conditions, and were perceived as such by Spanish households during the Great Recession. At the same time, Spanish households reacted to this increase in ways consistent with their particular demographic and sectoral characteristics.

### 3.4. Prices versus quantities

Whether a fall in expenditure is an actual fall of consumption, or merely a reduction in the cost of the basket purchased by a household, has quite different implications for welfare. Arguably, a household's well-being will be less affected by drops in expenditure that are due to lower prices than if the reduction is in quantities. On the other hand, quantity reductions can be plausibly related to welfare reductions.

<sup>15</sup> Results are similar if only workers with permanent contracts are used in the regression.

Aguiar and Hurst (2005) distinguish between consumption and consumption expenditure and find that, in the US, the unemployed, and also retirees, exhibit drops in expenditure that are not necessarily tied to drops in actual consumption. The reason is that households have more time to shop and to cook, and therefore spend less on food. Luengo-Prado and Sevilla (2013) confirm this finding for retirees in Spain: over the period 1985–2004 households paid lower prices for food after retirement, partly explaining away the steep consumption expenditure drop at retirement.

Whether the expenditure drop due to the rising unemployment rate is explained by a drop in prices is a question that we address using EPF data, by adapting the methodology in Aguiar and Hurst (2007). This methodology allows us to construct a price index that measures how much a household overpaid or underpaid for its consumption basket. Households report, not only how much they have spent, but also the quantities purchased of a number of goods.<sup>16</sup>

In Column 1 of Table 5 we report the results of re-estimating our baseline equation restricting expenditure to only the subset of items for which there is quantity information. To analyze whether drops in expenditure are due to drops in prices we run a regression of the logarithm of a price index à la Aguiar and Hurst (2007) on the unemployment rate, income variables, and all the covariates we used when estimating (1).<sup>17</sup> Results are shown in Column 2 of Table 5. The point estimate of the coefficient on  $\Delta U$  in Column 2 is not significantly different from zero; if anything, it is slightly positive. A rise in the unemployment rate is not associated with households paying prices that are lower than average. Whereas the unemployment rate lowers expenditure, it does not lower the price index.

An alternative question could be if households that face a higher unemployment rate pay lower prices than the year before. The price index of Aguiar and Hurst (2007) is not particularly suited to answer this question because it uses a time-varying benchmark. Average prices entering the computation of the price index are different for each year. To resolve the potential problem posed by time-variation in average prices, we modify the price index so that it values consumption bundles relative to average prices in a benchmark year. The price index now measures how much a household overpaid or underpaid relative to average prices in 2006, the year that is also used as base year in Spain's consumer price index.

We first repeat our regression using the logarithm of the new price index as the dependent variable and regress it on the unemployment rate, income variables, and all the covariates we used when estimating (1). Results are presented in Column 3, and are similar to those in Column 2. The change in the benchmark year does not overturn the result obtained for the price index of Aguiar and Hurst (2007). Next, we re-estimate the equation using the log-change in the price index as the dependent variable. The coefficient on  $\Delta U$  in Column 4 of Table 5 is not statistically different from zero. A rise in the unemployment rate from one year is not associated with drops in the price paid by households from one year to the other.

In conclusion, the drop in expenditure is not explained by lower prices. Both in the cross-sectional comparison, and along the time dimension, the reduction in expenditure due to the unemployment rate was unrelated to price drops. With our data, we cannot rule out an increase in home production. However, the absence of a fall in prices indicates a reduction in actual quantities purchased, and therefore a plausible effect on welfare. The finding in the previous literature suggests that households who pay less for their consumption are also using more time for shopping and home production (because they are unemployed or because they are retired). This provides an intuitive explanation for why we do not observe price reductions. Households who remain employed do not have any extra time that frees up for these alternative uses.

#### 4. Conclusion and discussion

During the Great Recession, Spanish households in which the primary earner was not afflicted by an idiosyncratic unemployment experience reduced their consumption at the rate of 0.7% per percentage point increase in the unemployment rate. Because this response was unrelated to a drop in contemporaneous income, it suggests forward-looking behavior. An increase in the unemployment rate contains new information on future income streams. As predicted by economic theory, in a way consistent with the permanent income hypothesis, households adjusted their consumption downward accordingly.

As an alternative explanation for this drop in consumption we explored whether the fall of expenditure was driven by a drop in prices rather than in quantities. We found that the drop in consumption expenditure was due to a reduction in quantities purchased, not lower prices. Therefore, the additional margin of increased shopping time, which has been found to explain expenditure drops at retirement or during unemployment spells, cannot account for the steep reduction in consumption expenditure by employed households in response to the unemployment rate.

We have organized the interpretation of our results around the LC/PIH framework using the log-linear approximations derived by Nalewaik (2006). As is well known, by disregarding higher order terms, first-order log-linear approximations implicitly rule out precautionary motivations. Therefore, a potentially fruitful avenue for further research is to analyze whether the rise in the unemployment rate conveys information not only on the expectation of future income but also on its variance, and—if this is the case—whether it has an effect on the consumption of employed households.

Because the unemployment rate was rising throughout the period considered, results are somewhat one-sided. Spanish data from the Great Recession do address the evolution of household consumption in a context of a rapidly rising unemployment rate, but are silent on the effect of a decreasing unemployment rate. In this respect, and using the classification

<sup>16</sup> This is done for a variety of food items and for certain utilities. Given the kind of consumption goods, it seems reasonable to interpret drops in quantities as reductions of welfare.

<sup>17</sup> In particular, our controls include age and age-squared, and therefore account for the hump-shaped life-cycle behavior of prices that was documented by Aguiar and Hurst (2007).

by Jappelli and Pistaferri (2010), our results add to the literature studying the effects of negative shocks on consumption. Moreover, the downward revision of household consumption expenditure due to the shadow of unemployment is quantitatively large. Because it affects households not themselves directly impacted by unemployment, it highlights the existence of an important channel through which a rising unemployment rate has a deleterious impact on domestic demand.

## Acknowledgments

In writing this paper we have benefited from comments by Vicente J. Bermejo, Moshe Buchinsky, Olympia Bover, José Manuel Campa, Juanjo Dolado, Annamaria Lusardi, Ernesto Villanueva, seminar participants at Banco de España, and the EEA-ESEM 2014 Congress. We also thank three anonymous referees and an editor for their valuable comments. We gratefully acknowledge financial support by the Spanish Ministry of Economy and Competitiveness (Project refs. ECO2012-38134 and ECO2012-31358). Some of our data are from INE, the Spanish statistical agency. INE is not responsible for any of the results obtained in this paper.

## Appendix A

### A.1. Food consumption expenditure

As is the case with expenditure on consumption, expenditure on food is also negatively affected by  $\Delta U$ . The first column in Table 6 replicates our benchmark specification and is included for comparability purposes. We report results for two measures of expenditure on food. The usual definition of food is the sum of food consumed either at home or away from home. We also use a measure of food that excludes food consumed away from home. The results for these alternative definitions of food appear in Columns 2 and 3 without income controls and in Columns 4 and 5 with income controls. Our results indicate that changes in the unemployment rate affect food expenditure whether measured at home or as the sum of

**Table 6**  
Food consumption expenditure.

Variables	(1) Nondurables and services	(2) Food	(3) Food at home	(4) Food	(5) Food at home
$\Delta U$	−0.735*** (0.231)	−1.105*** (0.241)	−1.122*** (0.372)	−1.090*** (0.248)	−1.142*** (0.386)
$\overline{\Delta y}$	0.113 (0.098)			0.110 (0.178)	−0.072 (0.192)
$\Delta y - \overline{\Delta y}$	0.104*** (0.015)			0.106*** (0.023)	0.069** (0.023)
Observations	17,182	17,182	17,182	17,182	17,182
$R^2$	0.032	0.011	0.009	0.015	0.010

Regression of food expenditure growth on  $\Delta U$ , income growth variables, and controls described in Section 2.2. All specifications include time dummies. Standard errors (in parentheses) are clustered by group. We compute the  $p$ -values using a  $t$ -distribution.

\*\* Indicates significance at the 5% level.

\*\*\* Indicates significance at the 1% level.

**Table 7**  
Robustness to alternative equivalence scales.

Variables	(1) ES1	(2) ES2	(3) Per capita
$\Delta U$	−0.735*** (0.231)	−0.731*** (0.231)	−0.739*** (0.231)
$\overline{\Delta y}$	0.113 (0.098)	0.115 (0.098)	0.111 (0.099)
$\Delta y - \overline{\Delta y}$	0.104*** (0.015)	0.105*** (0.015)	0.102*** (0.015)
Observations	17,182	17,182	17,182
$R^2$	0.032	0.020	0.063

The first column estimates the baseline specification of expenditure growth, on  $\Delta U$ , income growth variables, and controls described in Section 2.2. All specifications include time dummies. In the second column the dependent variable is constructed using the ES2 equivalence scale. In the third column the dependent variable is the log-change of per-capita consumption. Standard errors (in parentheses) are clustered by group. We compute the  $p$ -values using a  $t$ -distribution.

\*\*\* Indicates significance at the 1% level.

at home and outside. Point estimates in either case indicate that a one-percentage point increase in the unemployment rate is related to a drop of roughly 1.1% in food consumption. These coefficients are significantly different from zero but not statistically different from the point estimate for nondurables and services.

### A.2. Robustness to alternative equivalence scales

Results are robust to alternative equivalence scales. The first column of Table 7 uses the standard OECD equivalence scale that we use in the main text, in which the first adult in the household is weighted by 1, successive adults are weighted by 0.7, and dependents are weighted by 0.5. Column 2 uses the modified equivalence scale, in which the first adult is weighted by 1, successive adults by 0.5, and dependents by 0.3. These two equivalence scales are taken directly from the EPF dataset. Finally, the third column shows results for per-capita data, in which all household members are weighted equally.

### A.3. Robustness to alternative definitions of $\Delta U$

To obtain  $\Delta U$  we group households according to the education level and age of the primary earner (in the second interview) and calculate a group-specific unemployment rate as the ratio of the unemployed to the labor force in each one of these groups. In the main text we followed the standard classification used by INE when cross-tabulating data from its labor

**Table 8**  
Robustness to different group definitions.

Variables	(1) Benchmark	(2) Education 5-year age groups	(3) Education 5-year age groups skilled/unskilled
$\Delta U$	-0.735*** (0.231)	-0.631** (0.226)	-0.711** (0.319)
$\overline{\Delta y}$	0.113 (0.098)	0.125 (0.097)	0.161 (0.098)
$\Delta y - \overline{\Delta y}$	0.104*** (0.015)	0.104*** (0.015)	0.104*** (0.015)
Observations	17,182	17,182	17,182
$R^2$	0.032	0.032	0.033

The first column replicates the benchmark specification and regresses expenditure growth, on  $\Delta U$ , income growth variables, and controls described in Section 2.2. All specifications include time dummies. In the second column  $\Delta U$  is constructed by using a finer 5-year classification of age. In the third column  $\Delta U$  is constructed using the finer age classification and conditioning on skilled/unskilled occupations. Standard errors (in parentheses) are clustered by group. We compute the  $p$ -values using a  $t$ -distribution.

\*\* Indicates significance at the 5% level.

\*\*\* Indicates significance at the 1% level.

**Table 9**  
Outliers, sample, and other household members.

Variables	(1) Median	(2) Males	(3) By num. empl.
$\Delta U$	-0.712*** (0.189)	-0.695** (0.255)	-1.027*** (0.347)
$\overline{\Delta y}$	0.099 (0.133)	0.120 (0.098)	0.109 (0.098)
$\Delta y - \overline{\Delta y}$	0.096*** (0.011)	0.086*** (0.018)	0.103*** (0.015)
$\Delta U \times \text{num. empl.}$			0.185 (0.120)
Observations	17,182	13,329	17,182
$R^2$		0.036	0.033

The first column estimates the baseline specification by regressing expenditure growth on  $\Delta U$ , income growth variables, and controls described in Section 2.2 using quantile regression. All specifications include time dummies. In the second column the sample is restricted to only households in which the primary earner is male. In the third column  $\Delta U$  is interacted with the number of employed household members. Standard errors (in parentheses) are clustered by group. We compute the  $p$ -values using a  $t$ -distribution.

\*\* Indicates significance at the 5% level.

\*\*\* Indicates significance at the 1% level.

force survey and used four education groups (less than 1st cycle in high school, completed 1st cycle in high school, completed 2nd cycle in high school, higher education) and four age groups (less than 30, 30–44, 45–54, and 55–64).

We conduct a robustness exercise on the classifications used to construct groups and present the results in Table 8. Column 2 uses finer 5-year age bins and Column 3 adds whether the primary earner's job is skilled or unskilled as a criterion. In all regressions standard errors are clustered by group, using the appropriate group in each case. Results are overall similar to those obtained in our main analysis.

#### A.4. Outliers, sample, and other household members

To ascertain whether results are affected by outliers we use quantile regression methods to run a median regression in Column 1 of Table 9. The coefficient on  $\Delta U$  is similar to the one obtained in our baseline regression, indicating that the result is not driven by outliers. In Column 2 we restrict the sample to only male primary earners, who are frequently used as the sample of interest because they are likely to have a less elastic labor supply. Again, the coefficient on  $\Delta U$  is similar to our

**Table 10**  
Robustness to different specifications.

Variables	(1) Benchmark	(2) Interaction with $U$	(3) Below or above the median
$\Delta U$	−0.735*** (0.231)	−0.793 (0.574)	−0.735*** (0.231)
$\Delta U \times U$		0.003 (0.022)	
$\Delta U \times U_{low}$			0.018 (0.576)
$\bar{\Delta y}$	0.113 (0.098)	0.122 (0.117)	0.113 (0.097)
$\Delta y - \bar{\Delta y}$	0.104*** (0.015)	0.104*** (0.015)	0.104*** (0.015)
Observations	17,182	17,182	17,182
$R^2$	0.032	0.032	0.032

The first column replicates the benchmark specification and regresses expenditure growth, on  $\Delta U$ , income growth variables, and controls described in Section 2.2. All specifications include time dummies. The second column adds an interaction between  $\Delta U$  and  $U$ . The third column adds an interaction between  $\Delta U$  and a dummy variable for facing an unemployment rate below the median value of 10% ( $U_{low}$ ). Standard errors (in parentheses) are clustered by group. We compute the  $p$ -values using a  $t$ -distribution.

\*\*\* Indicates significance at the 1% level.

**Table 11**  
Group unemployment rates by year.

Education	Age	2006 (%)	2007 (%)	2008 (%)	2009 (%)	2010 (%)	2011 (%)
< 1 st cycle high school	< 30	14.6	15.8	26.9	38.4	41.3	43.6
	30–44	12.5	12.1	18.7	30.6	34.4	37.2
	45–54	8.8	8.9	13.9	22.2	25.9	31.5
	55–64	6.7	7.6	10.0	16.6	20.2	21.0
1st cycle high school	< 30	11.0	9.9	16.5	27.6	30.3	31.0
	30–44	8.6	8.4	12.5	21.4	23.8	25.6
	45–54	6.6	7.1	9.0	16.0	18.0	19.6
	55–64	6.7	5.7	7.2	12.0	14.9	15.2
2nd cycle high school	< 30	9.2	8.6	11.9	20.4	23.5	25.0
	30–44	6.7	6.8	9.6	16.3	18.9	21.0
	45–54	5.4	6.0	7.5	11.2	12.8	14.9
	55–64	5.2	5.2	6.1	10.1	11.5	13.4
Higher education	< 30	9.5	7.5	9.8	15.1	18.0	20.2
	30–44	5.1	4.7	5.8	9.0	10.5	11.8
	45–54	2.9	3.2	3.7	5.6	7.1	7.9
	55–64	3.2	3.0	3.0	5.4	5.6	7.2

Group unemployment rates calculated from the Spanish labor force survey (EPA) as the ratio of the unemployed to the active population in each group. Averages of quarterly data.

baseline results. Finally, in Column 3 we add an interaction term between  $\Delta U$  and the number of employed household members. This interaction term is not significantly different from zero and the coefficient on  $\Delta U$  is, if anything, larger.

#### A.5. Robustness to alternative specifications

The specification used in all our tables has percentage point difference on the right hand side. There may be reasons to suspect that the response to unemployment rate is heterogeneous across groups. For example, the same point increase in the current unemployment rate may have different implications for life-time earnings depending on the groups' long-term unemployment rates. To investigate this possibility, we conduct several robustness checks on the specification. In Table 10 we report the results from two alternative specifications. The specification in Column 2 includes as an additional control the interaction between  $\Delta U$  and  $U$ . The estimated coefficient on this interaction is not significant, suggesting that the impact of  $\Delta U$  does not change with the unemployment rate the group faces. In Column 3 the control added is the interaction between  $\Delta U$  and a dummy variable for facing an unemployment rate below the median value of 10% ( $U_{low}$ ).

Both alternatives suggest that the response to unemployment rate is not heterogeneous across groups.

#### A.6. Group unemployment rates by year

Table 11 contains the data on group unemployment rates. Fig. 1 plots the evolution of each of these 16 group unemployment rates over time. Because data in the Spanish labor force survey (EPA) are quarterly, we averaged group unemployment rates over all four quarters in a year.

### Appendix B. Supplementary data

Supplementary data associated with this article can be found in the online version at <http://dx.doi.org/10.1016/j.eurocorev.2015.04.006>.

### References

- Aguiar, M., Hurst, E., 2005. Consumption versus expenditure. *J. Polit. Econ.* 113 (5), 919–948.
- Aguiar, M., Hurst, E., 2007. Life-cycle prices and production. *Am. Econ. Rev.* 97 (5), 1533–1559.
- Altonji, J.G., Siow, A., 1987. Testing the response of consumption to income changes with (noisy) panel data. *Q. J. Econ.* 102 (2), 293–328.
- Barceló, C., Villanueva, E., 2010. The Response of Household Wealth to the Risk of Losing the Job: Evidence from Differences in Firing Costs. Banco de España Working Papers 1002, Banco de España.
- Benito, A., 2006. Does job insecurity affect household consumption? *Oxf. Econ. Pap.* 58 (1), 157–181.
- Bentolila, S., Cahuc, P., Dolado, J.J., Barbanchon, T.L., 2012. Two-tier labour markets in the great recession: France versus Spain. *Econ. J.* 122 (562), F155–F187.
- Bewley, T., 1977. The permanent income hypothesis: a theoretical formulation. *J. Econ. Theory* 16 (2), 252–292.
- Blundell, R., Pistaferri, L., Preston, I., 2008. Consumption inequality and partial insurance. *Am. Econ. Rev.* 98 (5), 1887–1921.
- Browning, M., Crossley, T.F., 2001. Unemployment insurance benefit levels and consumption changes. *J. Public Econ.* 80, 1–23.
- Browning, M., Crossley, T.F., 2009. Shocks, stocks, and socks: smoothing consumption over a temporary income loss. *J. Eur. Econ. Assoc.* 7 (6), 1169–1192.
- Cameron, A.C., Gelbach, J.B., Miller, D.L., 2008. Bootstrap-based improvements for inference with clustered errors. *Rev. Econ. Stat.* 90 (3), 414–427.
- Campbell, J.Y., Deaton, A., 1989. Why is consumption so smooth? *Rev. Econ. Stud.* 56 (3), 357–373.
- Campos, R.G., Reggio, I., 2014. Measurement error in imputation procedures. *Econ. Lett.* 122 (2), 197–202.
- Campos, R.G., Reggio, I., García-Piriz, D., 2013. Micro versus macro consumption data: the cyclical properties of the consumer expenditure survey. *Appl. Econ.* 45 (26), 3778–3785.
- Carroll, C.D., Dynan, K.E., Krane, S.D., 2003. Unemployment risk and precautionary wealth: evidence from households' balance sheets. *Rev. Econ. Stat.* 85 (3), 586–604.
- Costain, J., Jimeno, J.F., Thomas, C., 2010. Employment Fluctuations in a Dual Labor Market. Banco de España Working Papers 1013, Banco de España.
- Flavin, M.A., 1981. The adjustment of consumption to changing expectations about future income. *J. Polit. Econ.* 89 (5), 974–1009.
- Friedman, M., 1957. *A Theory of the Consumption Function*. Princeton University Press, Princeton.
- Garner, T.I., Janini, G., Passero, W., Paszkiewicz, L., Vendemia, M., 2006. The CE and the PCE: a comparison. *Mon. Labor Rev.*, 20–46.
- Gervais, M., Klein, P., 2010. Measuring consumption smoothing in CEX data. *J. Monet. Econ.* 57 (8), 988–999.
- Gruber, J., 1997. The consumption smoothing benefits of unemployment insurance. *Am. Econ. Rev.* 87 (1), 192–205.
- Heathcote, J., Perri, F., Violante, L.G., 2010. Unequal we stand: an empirical analysis of economic inequality in the United States, 1967–2006. *Rev. Econ. Dyn.* 13, 15–51.
- Jappelli, T., Pistaferri, L., 2010. The consumption response to income changes. *Ann. Rev. Econ.* 2, 479–506.
- Luengo-Prado, M.J., Sevilla, A., 2013. Time to cook: expenditure at retirement in Spain. *Econ. J.* 123, 764–789.
- Moulton, B.R., 1990. An illustration of a pitfall in estimating the effects of aggregate variables on micro units. *Rev. Econ. Stat.* 72 (2), 334–338.
- Nalewaik, J.J., 2006. Current consumption and future income growth: synthetic panel evidence. *J. Monet. Econ.* 53 (8), 2239–2266.
- Pou, L., Alegre, J., 2002. La encuesta continua de presupuestos familiares (1985–1996): descripción, representatividad y propuestas de metodología para la explotación de la información de los ingresos y el gasto. Fundación de las cajas de ahorros confederadas para la investigación económica y social.
- Skinner, J., 1987. A superior measure of consumption from the panel study of income dynamics. *Econ. Lett.* 23, 213–216.
- Slesnick, D.T., 1992. Aggregate consumption and saving in the postwar United States. *Rev. Econ. Stat.* 74 (4), 585–597.
- Stephens, M., 2001. The long-run consumption effects of earnings shocks. *Rev. Econ. Stat.* 83 (1), 28–36.
- Stephens, M., 2004. Job loss expectations, realizations, and household consumption behavior. *Rev. Econ. Stat.* 86 (1), 253–269.
- Wooldridge, J.M., 2003. Cluster-sample methods in applied econometrics. *Am. Econ. Rev.* 93 (2), 133–138.