CONSUMPTION IN THE SHADOW OF UNEMPLOYMENT

2014

Rodolfo G. Campos and Iliana Reggio

Documentos de Trabajo
N.º 1411

BANCO DE ESPAÑA
Eurosistema

CONSUMPTION IN THE SHADOW OF UNEMPLOYMENT

CONSUMPTION IN THE SHADOW OF UNEMPLOYMENT (*)

Rodolfo G. Campos (**)

IESE BUSINESS SCHOOL

Iliana Reggio (***)

UNIVERSIDAD CARLOS III

(*) Part of this research was carried out while Iliana Reggio was visiting the Banco de España. In writing this paper we have benefited from comments by Vicente J. Bermejo, Olympia Bover, José Manuel Campa, Juanjo Dolado, Annamaria Lusardi, Ernesto Villanueva and seminar participants at the Banco de España. We also thank an anonymous referee for many valuable comments. We gratefully acknowledge financial support by the Spanish Ministry of Economy and Competitiveness (Project ref: ECO2012-38134 and ECO2012-31358).
(**) Department of Economics, IESE Business School, Camino del Cerro del Águila 3, 28023 Madrid, Spain; email: rcampos@iese.edu.
(***) Department of Economics, Universidad Carlos III, Calle Madrid 126, 28903 Getafe (Madrid), Spain; email: ireggio@eco.uc3m.es.

Documentos de Trabajo. N.º 1411

2014

ISSN: 1579-8666 (on line)

Abstract

By how much do employed households reduce their consumption when the aggregate unemployment rate rises? In Spain during the Great Recession a 1 percentage point increase in the unemployment rate was related to a strong drop in household consumption of more than 0.7% per equivalent adult. This reduction is the response of forward-looking agents 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 indeed a reduction in quantities, and not a switch to cheaper alternatives.

Keywords: consumption, unemployment, income, life-cycle models, Spain, Great Recession.

JEL classification: E21, E24, D12, D84.

Resumen

¿En cuánto se reduce el consumo de los hogares, aun de aquellos que continúan empleados, cuando aumenta la tasa de paro? En España, durante la Gran Recesión, un punto adicional de aumento en la tasa de paro estuvo asociado a una caída de 0,7 % de caída del consumo por adulto equivalente. Esta reducción del consumo refleja el deterioro de las expectativas de ingreso de los hogares. Utilizando datos de panel que incluyen información sobre cantidades físicas, encontramos que la caída en el gasto de consumo se debe a una reducción en las cantidades, y no en los precios.

Palabras clave: consumo, desempleo, ingresos, modelos de ciclo de vida, España, Gran Recesión.

Códigos JEL: E21, E24, D12, D84.

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?

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, it 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 percent to 21.6 percent.

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 percent 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, aggregate unemployment casts a shadow on future income expectations. Households update their expectations of future income in a way that is negatively related to the unemployment rate and reduce their consumption, as predicted by forward-looking behavior.

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 this restricted subset narrows down what kind of information economic agents are likely to extract from the unemployment rate. The information contained in the unemployment rate is qualitatively different for the unemployed and the employed: because the unemployment rate provides a gauge on labor market conditions in general it is simultaneously an indication of the probability of losing a job and the probability of finding one. Whereas the unemployed will take into account both probabilities, employed workers will focus on the probability of keeping their jobs, removing the need to disentangle between the two probabilities.

The second advantage of excluding unemployed households is that it allows 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 third 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.

Fourth, the restriction to households who remain employed removes potential pitfalls in the estimation due to unobservable missing variables. For example, Carroll, Dynan, and Krane (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 to use more comprehensive definitions of consumption.[1] 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.

[1]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 zero, fifty, and a hundred percent.

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 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.[2] 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.

Finally, although beyond the scope of this investigation, our results implicitly reject the risk-sharing—or full consumption insurance—hypothesis introduced by Cochrane (1991), Mace (1991), and Townsend (1994). Their work, and subsequent work by others, tests whether, contrary to what full insurance implies, household consumption responds to idiosyncratic shocks affecting the household, such as changes in income, or variables related to it such as layoffs, sickness, etc. Full risk-sharing tends to be rejected in this literature (e.g., Nelson, 1994).[3] Notably, previous research on Spain by Bentolila and Ichino (2008) failed to reject full risk-sharing. Their study, which used individual unemployment transitions, was conducted using a different survey for a previous period (1985–1996) and used only expenditure on food items.

[2] Jappelli and Pistaferri (2010) provide a useful survey of this literature in which they distinguish between anticipated and unanticipated income shocks.

[3] However, Schulhofer-Wohl (2011) and Mazzocco and Saini (2012) have recently argued that heterogeneity in risk-aversion may produce false rejections of the risk-sharing hypothesis. On the other hand, in the presence of heterogeneous risk-aversion, less risk-sharing may be welfare improving in some environments (e.g., Campos, 2013).

In our case, the response of consumption to the unemployment rate can be interpreted as a particularly strong rejection of consumption insurance. Because the unemployment rate is external to the household, it is not subject to the moral hazard argument that is usually put forward to argue why households are unable to insure their level of consumption against individual unemployment. Thus, failure to insure points towards stronger shortcomings in the menu of financial contracts available to households. In a sense, our finding is akin to that of Attanasio and Davis (1996) who found that consumption responded to variables external to the household, in their case to movements in cohort-specific average wage rates in the US.

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) and Aguiar and Hurst (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, and find that the response to the unemployment rate was not a fall in prices.

2 Economic theory 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 ΔU_t, has a discernible effect on a household's consumption growth rate, which we denote with Δc_t. In other words, we estimate a relationship of the form

$$\Delta c_t = f(\Delta U_t) \tag{1}$$

using a consumption survey with rich data at the household level, including detailed demographic variables for household members.

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.

2.1 Economic theory

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 response to predictable income changes but should respond 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, giving rise to an extensive literature on the "excess smoothness" of consumption—which says that consumption is not sufficiently responsive to income innovations. Campbell and Deaton (1989) obtain a closed form solution at the cost of assuming a quadratic utility function. However, 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. Let $\Delta U_t \equiv U_t - U_{t-1}$ denote 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 $\Delta y_t \equiv \log Y_t - \log Y_{t-1}$ stand for the logarithmic difference of consumption and income.[4] Because these last two expressions are logarithmic differences, they can be interpreted as growth rates.

Expressing the discount factor as $\lambda = \frac{1}{1+r}$, 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

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

where E_t and E_{t-1} denote the conditional expectations taken with information available at time t and $t-1$.[5] The consumption growth rate Δc_t equals the discounted sum of innovations at time t to the expectations of the entire stream of future income growth rates Δy_{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

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

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 ΔU_t can affect consumption growth through its effect on future periods. In response to a rise in ΔU_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.

[4] We follow the convention that an uppercase letter indicates a variable in levels and a lowercase letter the logarithm of that variable.

[5] 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.

In theory, ΔU_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 ΔU_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 ΔU_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; we include unemployment of other household members as an additional control. 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, we include measures of current income growth, at the household and group level, as additional regressors when estimating the relationship between ΔU_t and Δc_t.

2.2 Empirical strategy

2.2.1 Specification and estimation

The consumption equation derived from a LC/PIH model in Section 2.1 implies a linear relationship between the consumption growth rate and innovations to income growth. The effects on consumption growth of current and future income growth rates are additively separable. If the unemployment rate U_t is used to predict expected future income growth rates $E_t \Delta y_{t+j}$ (and if the association is linear), then its time-change ΔU_t can be used as a stand-in for changes in the expectation of future income growth rates $(E_t - E_{t-1})\Delta y_{t+j}$. Moreover, according to (3) it should enter the specification linearly.[6]

[6] Although intuitive, at the moment it is just an assumption that U_t is useful in predicting Δy_{t+j}. In Section 3 we verify that the group unemployment rate has explanatory power for one period ahead income growth expectations of Spanish households using an alternative data set.

Letting i index households, g groups, and t stand for time, our specification is of the form:

$$\Delta c_{igt} = \gamma \Delta U_{gt} + \beta' X_{igt} + \alpha_t + \varepsilon_{igt}. \tag{4}$$

The covariates in X vary at the household level whereas ΔU is measured only at the group level; α is a time dummy, and ε represents the estimation error. The partial effect of the unemployment rate of the group is captured by the coefficient γ. An increase by one percentage point in the unemployment rate implies a change in γ log-percentage points in household consumption. We use the population weights available in the consumption survey to estimate equation (4) by weighted least squares.

The fact that ΔU 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, Gelbach, and Miller (2008), we apply finite sample bias correction and compute the p-values for the null hypotheses that the coefficients on ΔU are zero using a t-distribution, instead of a standard normal. 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.

2.2.2 Consumption and unemployment

Households are interviewed twice, in two consecutive years. Our variable of interest is yearly household consumption growth $\Delta c_t \equiv \log C_t - \log C_{t-1}$. 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.

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 four age groups less than 30, 30–44, 45–54, 55–64.[7] The change from one year to the other of the group unemployment rate is $\Delta U_t \equiv U_t - U_{t-1}$.

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.[8]

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.

[7]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.

[8]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.

Our sample consists of heads 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.[9]

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 (4).

The inclusion of income growth on the right hand side of (4) 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 ΔU_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 ΔU_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

[9]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).

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. The issue that remains is how to account in the estimation for the presence of the innovation to current income $(E_t - E_{t-1})\Delta y_t$ in the theoretical relationship derived in (3). To some extent, the inclusion of current income growth Δy_t may be sufficient, although this depends on what is assumed on the formation of expectations.

To illustrate the assumed expectations formation process, notice that this first term in (3) can be decomposed and written as the difference between current income growth Δy_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:

$$(E_t - E_{t-1})\Delta y_t \equiv E_t\Delta y_t - E_{t-1}\Delta y_t = \Delta y_t - E_{t-1}\Delta y_t. \tag{5}$$

If the expectation of income growth is common across households, then the influence of $E_{t-1}\Delta y_t$ will be captured by time dummies and including Δy_t will be sufficient to controls for innovations to current income (and therefore tighten the link between ΔU_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 ΔU_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: $\Delta y_t - \overline{\Delta y}_t$.

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 ΔU_t and current income growth. However, the stability of the coefficient in the face of different specifications for income growth go a long way in diffusing concerns about exclusively equating changes in ΔU_t with changes in expected future income growth. In particular, because $\overline{\Delta y}_t$ is calculated conditioning on the same variables as

ΔU_t, using it to proxy for expectations seems to be a good way of preventing the attribution of spurious explanatory power to ΔU_t. For this reason, although we tried alternative specifications, we chose to use $\overline{\Delta y_t}$ in our baseline specification.[10]

2.3 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 Figure 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 a 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

[10] In addition to the specifications of $E_{t-1}\Delta y_t$ already discussed, as a robustness check, we estimated permanent income and used it as an alternative measure of expected income growth (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, 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 current income growth innovations growth will roughly coincide with Δy_t, the actual income growth rate.

Figure 1: The unemployment rate across education and age groups

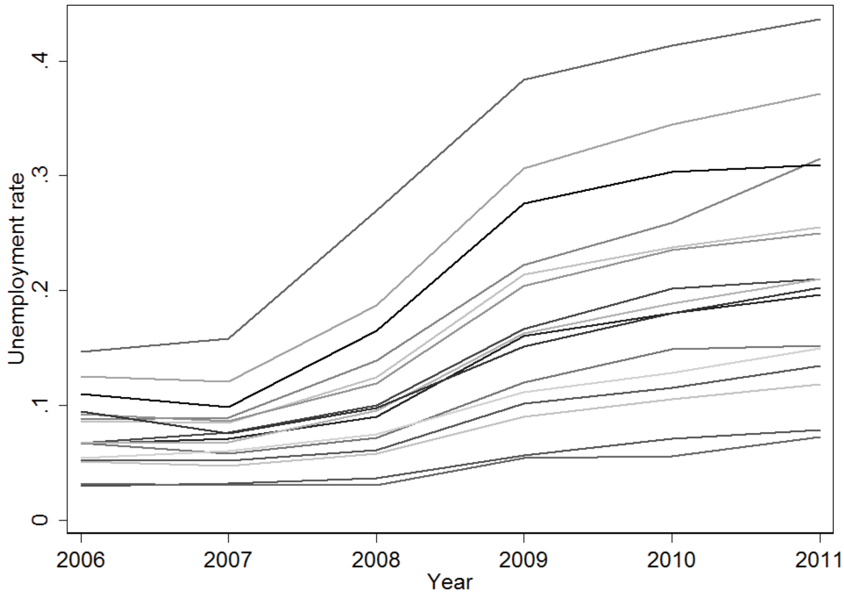

Data used to construct this figure are in the Appendix (Table 10).

series in the US (Slesnick, 1992; Garner, Janini, Passero, Paszkiewicz, and Vendemia, 2006; Heathcote, Perri, and Violante, 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).[11]

Consumption measured in the CEX is also less pro-cyclical than aggregate consumption. Campos, Reggio, and García-Píriz (2013) show that consumption measured from the CEX underestimates the cyclical correlation of aggregate consumption (PCE) with GDP by 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 the 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 equivalization 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-

[11] Pou and Alegre (2002) calculated that consumption measured in the ECPF-85 over the whole the 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.

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 6,727 households). We drop households in which the primary earner is self-employed (dropping 4,944 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 1,101 households). Missing values of some controls made us drop 9 additional households. Our final sample consists of 17,182 households.

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 percent in household consumption per equivalent adult. This result is robust to the inclusion of different income measures.

In Table 1 we show results of estimating the equation specified in (4) for different variables measuring income. We start by excluding income changes (Col. 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 Δy, the log-change in current household income, is added, the effect is only slightly lower, with a point estimate of -0.739 (Col. 2). The last two columns, which include the average change in income $\overline{\Delta y}$ (Col. 3) and the average change together with the idiosyncratic deviation from this average $\Delta y - \overline{\Delta y}$ (Col. 4) do not alter the conclusion. In both cases, the point estimate is -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.

This is consistent with the hypothesis that a rise in the unemployment rate is related to downward revisions of future income expectations. The assumption that households

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
ΔU	-0.753***	-0.739***	-0.737***	-0.737***
	(0.234)	(0.229)	(0.235)	(0.233)
Δy		0.104***		
		(0.014)		
$\overline{\Delta y}$			0.076	0.112
			(0.105)	(0.098)
$\Delta y - \overline{\Delta y}$				0.104***
				(0.015)
Observations	17,182	17,182	17,182	17,182
R-squared	0.027	0.032	0.027	0.032

*Regression of consumption expenditure growth on ΔU, various income growth variables, and controls described in Section 2.2. Standard errors (in parentheses) are clustered by group. We compute the p-values using a t-distribution. *** indicates significance at the 1% level.*

adjust income expectations in response labor market conditions seems sensible, and has been made before. For example, Stephens (2004) and Benito (2006) make a similar connection with 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, 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 consumption specification. We report pooled OLS results in Column 1 odds ratios from logit regressions in Columns 2 and 3.

The three regressions yield the same conclusion: rises in the group unemployment rate

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

	OLS	Pooled Logit	Panel Logit
U	0.608^{*}	1.074^{**}	1.055^{**}
	(0.314)	(0.037)	(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. ** and * indicate significance at the 5% and 10% level.*

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 household were responding to worsening income expectations. The quantitatively large effect on consumption, and therefore on domestic demand, operates through expectation changes.

Food consumption

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, Pistaferri, and Preston, 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 compare our results to this previous literature we repeat our regressions for food items.

The first column in Table 3 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.

Table 3: Food consumption expenditure.

VARIABLES	(1) Nondurables and Services	(2) Food	(3) Food at Home	(4) Food	(5) Food at Home
ΔU	-0.737***	-1.103***	-1.117***	-1.088***	-1.136**
	(0.233)	(0.242)	(0.372)	(0.249)	(0.386)
$\overline{\Delta y}$	0.112			0.111	-0.068
	(0.098)			(0.180)	(0.189)
$\Delta y - \overline{\Delta y}$	0.104***			0.106***	0.069**
	(0.015)			(0.022)	(0.023)
Observations	17,182	17,182	17,182	17,182	17,182
R-squared	0.032	0.011	0.009	0.015	0.010

*Regression of food expenditure growth on ΔU, income growth variables, and controls described in Section 2.2. Standard errors (in parentheses) are clustered by group. We compute the p-values using a t-distribution. *** indicates significance at the 1% level and ** indicates significance at the 5% level.*

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 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. If anything, the effect of the unemployment rate is stronger for food than for nondurables and services.

3.1 The shadow of unemployment

The main message from the results reported in Tables 1 and 3 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 ΔU. In light of the theoretical framework laid out in Section 2.1 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 equation (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 is not affected by them. In the specification closest to ours (Stephens, 2004, Table A2, cols. 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.[12] 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 of whether a household finds whether it has been laid off as the variable explaining on consumption fluctuations, what Stephens (2004) is effectively doing is identifying subjective probabilities with $(E_t - E_{t-1})\Delta y_t$, the first term in (3).[13] 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 $(E_t - E_{t-1})\Delta y_{t+j}$.

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.

[12] 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, Col. 1).

[13] This is clear from equation 3 and the surrounding discussion on p. 260–261 of Stephens (2004).

3.1.1 Consumption insurance and the permanent income hypothesis

Any regression in the form of (4) is implicitly a test of risk-sharing or full consumption insurance. This literature, which builds on work by Cochrane (1991), Mace (1991), and Townsend (1994), has tended to reject full risk-sharing. Interestingly, in the case of Spain, research by Bentolila and Ichino (2008), although for a previous period, failed to reject full risk-sharing. Bentolila and Ichino (2008) studied the period 1985–1996 and failed to find a significant effect of idiosyncratic unemployment spells on consumption defined as expenditures on food. They do not report results for items other than food. However, our results in Table 3 confirm a rejection of full risk-sharing even in the case in which consumption is restricted to food items.

The response of consumption to ΔU can be interpreted as a particularly strong rejection of consumption insurance. In this case, the absence of an insurance mechanism cannot be justified by the usual moral hazard argument. Unlike unemployment within the household, because ΔU is an average over a whole demographic group, it is not under the direct control of any household member. Thus, failure to insure against group averages suggests a more extreme case of market incompleteness.

In this line of research, Attanasio and Davis (1996) tested whether consumption responded to movements in cohort-specific variables in the US—in their case, average wage rates. They found that consumption did, in effect, respond to the change in wage rates. A consumption response could be detected over longer periods (5 years changes and over) but not for shorter periods. In comparison, our estimation tests whether consumption responds to changes in the unemployment rate over a one-year frame, implying a departure of full risk-sharing at shorter periods than those found by Attanasio and Davis (1996).

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.

3.1.2 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 ΔU 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 4. Column 1 replicates the baseline specification for comparability, and Column 2 allows for different effects of ΔU 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 public sector or the public sector. This information is not available for all households and reduces our sample to 17,174 households. Results are shown in Table 5. In the baseline

Table 4: Consumption response depending on the age of the primary earner.

VARIABLES	(1) Baseline	(2) By Age
ΔU	-0.737***	
	(0.233)	
$\Delta U \times PE25-44$		-0.783***
		(0.257)
$\Delta U \times PE45-54$		-0.581**
		(0.251)
$\Delta U \times PE55-64$		-0.229
		(0.728)
$\overline{\Delta y}$	0.112	0.138
	(0.098)	(0.106)
$\Delta y - \overline{\Delta y}$	0.104***	0.104***
	(0.015)	(0.015)
Observations	17,182	17,182
R-squared	0.032	0.033

*The first column is the baseline regression of expenditure growth, on ΔU, income growth variables, and controls described in Section 2.2. In the second column ΔU 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 1% level and ** indicates significance at the 5% level.*

specification, in Column 1, the effect of the unemployment rate is not allowed to depend on measures of job stability. We report the results from the baseline specification for the sub-sample of households for which there is information on the sector.

Table 5: Consumption response by sector.

VARIABLES	(1) Baseline	(2) by Sector
ΔU	-0.740***	
	(0.235)	
$\Delta U \times$ Private		-0.767***
		(0.258)
$\Delta U \times$ Public		-0.520
		(0.298)
$\overline{\Delta y}$	0.109	0.102
	(0.100)	(0.099)
$\Delta y - \overline{\Delta y}$	0.102***	0.102***
	(0.015)	(0.015)
Observations	17,174	17,174
R-squared	0.032	0.032

*The first column is the baseline regression of expenditure growth, on ΔU, income growth variables, and controls described in Section 2.2. In the second column ΔU is interacted with private sector and public sector dummy variables. Standard errors (in parentheses) are clustered by group. We compute the p-values using a t-distribution. *** indicates significance at the 1% level and * indicates significance at the 10% level.*

Column 2 reports results on whether working in the public of private sector makes a difference in the response to changes in the unemployment rate. The change in the unemployment rate ΔU 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 -0.52%, and not significantly different from zero.

Just as the rejection of risk-sharing could be interpreted as a validation of the self-insurance assumption underlying the LC/PIH framework, the results in this section, that younger 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.

3.2 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. Quantity reductions can be plausibly related to welfare reductions.

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 payed 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 related to a drop in prices can be addressed using EPF data. Households are asked to report, not only how much they have spent, but also the quantities they have bought of a number of goods.[14] Following the methodology in Aguiar and Hurst (2007), we define an index that measures if a household is paying more or less than the average. First, we compute p_{ijt}, the unit value paid by household i at time t for good j (where $j = 1, \ldots, J$). Using these household-specific unit values p_{ijt} we compute a weighted average price $\overline{p_{jt}}$ paid

[14] 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.

for good j at time t using quantities q_{ijt} as weights. The price index measuring how much a household overpaid or underpaid for its consumption basket is then obtained by dividing true expenditure by the cost of the bundle valued at average prices:

$$\pi_{it} = \frac{\sum_{j=1}^{J} p_{ijt} q_{ijt}}{\sum_{j=1}^{J} \overline{p_{jt}} q_{ijt}}. \tag{6}$$

Aguiar and Hurst (2007) normalize this price index to have mean one in every year. This normalization implies that $\log \pi_{it}$ measures log-deviations from the average across households. We do not normalize our price index; because our regressions of the logarithm of this index include time dummies, whether we do this normalization or not does not affect our results. To calculate the price index we use all goods with quantity data, with the exception of alcohol and tobacco. We lose 20 households for which there are no quantity data.

In Column 1 of Table 6 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 the price index $\log \pi_{it}$ on the unemployment rate, income variables, and all the covariates we used when estimating (4).[15] The results from this equation are shown in Column 2 of Table 6. The point estimate of the coefficient on ΔU 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.

Households for whom the unemployment rate increased do not pay prices that are lower than average. But do 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 $\overline{p_{jt}}$ entering the computation of the price index are different for each year. This implies that first-differencing the price index introduces a household-specific component. Households are differently affected by the time-variation in average prices $\overline{p_{jt}}$ because, from (6), this average price enters their price index according to the household-specific quantities purchased of each good q_{ijt}.

[15] 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).

Table 6: Expenditure vs. Prices.

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

*Regression of expenditure growth, price indices, and price changes on ΔU, income growth variables, and controls described in Section 2.2. Standard errors (in parentheses) are clustered by group. We compute the p-values using a t-distribution. ***, **, and * indicate significance at the 1%, 5%, and 10% level.*

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. Specifically, we substitute average prices in 2006 $\overline{p}_{j2006}$ into the denominator of (6) to obtain

$$\pi_{it}^{2006} = \frac{\sum_{j=1}^{J} p_{ijt} q_{ijt}}{\sum_{j=1}^{J} \overline{p}_{j2006} q_{ijt}}. \tag{7}$$

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. As before, we do not normalize this index in any way. The use of time dummies would entirely account for any normalization and its time change. Our measure of whether a household increased or decreased prices over the last year is the first difference of the logarithm of the index: $\Delta \log \pi_{it}^{2006}$.

We first repeat our regression using the logarithm of the new price index $\log \pi_{it}^{2006}$ as the dependent variable and regress it on the unemployment rate, income variables, and all the covariates we used when estimating (4). 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 $\Delta \log \pi_{it}^{2006}$ as the dependent variable. The coefficient on ΔU in Column 4 of Table 6 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.

As a final point, how does the result that households in our sample do not pay lower prices square with the literature? Using data from time use surveys, Aguiar and Hurst (2005, 2007), and Luengo-Prado and Sevilla (2013), documented 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

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 percent 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.

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 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.

A Appendix

Equivalence scales

Results are robust to alternative equivalence of 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 two 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.

Table 7: Robustness to alternative equivalence scales.

VARIABLES	(1) ES1	(2) ES2	(3) per capita
ΔU	-0.737***	-0.733***	-0.741***
	(0.233)	(0.232)	(0.232)
$\overline{\Delta y}$	0.112	0.114	0.109
	(0.098)	(0.098)	(0.099)
$\Delta y - \overline{\Delta y}$	0.104***	0.105***	0.102***
	(0.015)	(0.015)	(0.015)
Observations	17,182	17,182	17,182
R-squared	0.032	0.020	0.063

*The first column estimates the baseline specification of expenditure growth, on ΔU, income growth variables, and controls described in Section 2.2. 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.*

Robustness to alternative definitions of ΔU

To obtain ΔU 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 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, 55–64).

In this section we conduct a robustness exercise on the classifications used to construct groups. 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.

Table 8: Robustness to different group definitions

VARIABLES	(1) Benchmark	(2) Education 5-year age groups	(3) Education 5-year age groups skilled/unskilled
ΔU	-0.737***	-0.631**	-0.713**
	(0.233)	(0.227)	(0.320)
$\overline{\Delta y}$	0.112	0.124	0.160
	(0.098)	(0.097)	(0.098)
$\Delta y - \overline{\Delta y}$	0.104***	0.104***	0.103***
	(0.015)	(0.015)	(0.015)
Observations	17,182	17,182	17,182
R-squared	0.032	0.032	0.033

*The first column replicates the benchmark specification and regresses expenditure growth, on ΔU, income growth variables, and controls described in Section 2.2. In the second column ΔU is constructed by using a finer 5-year classification of age. In the third column ΔU 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 1% level and ** indicates significance at the 5% level.*

Additional robustness checks

In this last section we conduct a number of additional robustness checks. 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 ΔU 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 ΔU is similar to our baseline results. Finally, in Column 3 we add an interaction term between ΔU and the number of employed household members. This interaction term is not significantly different from zero and the coefficient on ΔU is, if anything, larger.

Table 9: Additional robustness checks

VARIABLES	(1) Median	(2) Males	(3) by Num. Empl.
ΔU	-0.712***	-0.695**	-1.027***
	(0.189)	(0.255)	(0.347)
$\overline{\Delta y}$	0.099	0.120	0.109
	(0.133)	(0.098)	(0.098)
$\Delta y - \overline{\Delta y}$	0.096***	0.086***	0.103***
	(0.011)	(0.018)	(0.015)
$\Delta U \times Num.\ Empl.$			0.185
			(0.120)
Observations	17,182	13,329	17,182
R-squared		0.036	0.033

*The first column estimates the baseline specification by regressing expenditure growth on ΔU, income growth variables, and controls described in Section 2.2 using quantile regression. In the second column the sample is restricted to only households in which the primary earner is male. In the third column ΔU 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 1% level and ** indicates significance at the 5% level.*

Group unemployment rates

Table 10 contains the data on group unemployment rates. Figure 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.

Table 10: Group unemployment rates by year

Education	Age	2006	2007	2008	2009	2010	2011
$<$ 1st 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.

References

Aguiar, M., and E. Hurst (2005): "Consumption versus Expenditure," *Journal of Political Economy*, 113(5), 919–948.

——— (2007): "Life-Cycle Prices and Production," *American Economic Review*, 97(5), 1533–1559.

Altonji, J. G., and A. Siow (1987): "Testing the Response of Consumption to Income Changes with (Noisy) Panel Data," *The Quarterly Journal of Economics*, 102(2), pp. 293–328.

Attanasio, O. P., and S. J. Davis (1996): "Relative Wage Movements and the Distribution of Consumption," *Journal of Political Economy*, 104(6), 1227–1262.

Benito, A. (2006): "Does job insecurity affect household consumption?," *Oxford Economic Papers*, 58(1), 157–181.

Bentolila, S., and A. Ichino (2008): "Unemployment and consumption near and far away from the Mediterranean," *Journal of Population Economics*, 21(2), 255–280.

Bewley, T. (1977): "The permanent income hypothesis: A theoretical formulation," *Journal of Economic Theory*, 16(2), 252–292.

Blundell, R., L. Pistaferri, and I. Preston (2008): "Consumption Inequality and Partial Insurance," *American Economic Review*, 98(5), 1887–1921.

Browning, M., and T. F. Crossley (2001): "Unemployment insurance benefit levels and consumption changes.," *Journal of Public Economics*, 80, 1–23.

——— (2009): "Shocks, Stocks, and Socks: Smoothing Consumption Over a Temporary Income Loss," *Journal of the European Economic Association*, 7(6), 1169–1192.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008): "Bootstrap-Based Improvements for Inference with Clustered Errors," *The Review of Economics and Statistics*, 90(3), 414–427.

Campbell, J. Y., and A. Deaton (1989): "Why Is Consumption So Smooth?," *Review of Economic Studies*, 56(3), 357–73.

Campos, R. G. (2013): "Risk-sharing and crises. Global games of regime change with endogenous wealth," *Journal of Economic Theory*, 148(4), 1624–1658.

Campos, R. G., and I. Reggio (2014): "Measurement error in imputation procedures," *Economics Letters*, 122(2), 197–202.

Campos, R. G., I. Reggio, and D. García-Píriz (2013): "Micro versus macro consumption data: the cyclical properties of the consumer expenditure survey," *Applied Economics*, 45(26), 3778–3785.

Carroll, C. D., K. E. Dynan, and S. D. Krane (2003): "Unemployment Risk and Precautionary Wealth: Evidence from Households' Balance Sheets," *The Review of Economics and Statistics*, 85(3), 586–604.

Cochrane, J. H. (1991): "A Simple Test of Consumption Insurance," *Journal of Political Economy*, 99(5), 957–76.

Flavin, M. A. (1981): "The Adjustment of Consumption to Changing Expectations about Future Income," *Journal of Political Economy*, 89(5), 974–1009.

Friedman, M. (1957): *A Theory of the Consumption Function*. Princeton University Press.

Garner, T. I., G. Janini, W. Passero, L. Paszkiewicz, and M. Vendemia (2006): "The CE and the PCE: a comparison," *Monthly Labor Review*, pp. 20–46.

Gervais, M., and P. Klein (2010): "Measuring consumption smoothing in CEX data," *Journal of Monetary Economics*, 57(8), 988–999.

Gruber, J. (1997): "The Consumption Smoothing Benefits of Unemployment Insurance," *American Economic Review*, 87(1), 192–205.

Heathcote, J., F. Perri, and L. G. Violante (2010): "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967-2006," *Review of Economic Dynamics*, 13, 15–51.

Jappelli, T., and L. Pistaferri (2010): "The Consumption Response to Income Changes," *Annual Review of Economics*, 2, 479–506.

Luengo-Prado, M. J., and A. Sevilla (2013): "Time to Cook: Expenditure at Retirement in Spain," *Economic Journal*, 123, 764–789.

Mace, B. J. (1991): "Full Insurance in the Presence of Aggregate Uncertainty," *Journal of Political Economy*, 99(5), 928–56.

Mazzocco, M., and S. Saini (2012): "Testing Efficient Risk Sharing with Heterogeneous Risk Preferences," *American Economic Review*, 102(1), 428–68.

Moulton, B. R. (1990): "An Illustration of a Pitfall in Estimating the Effects of Aggregate Variables on Micro Units," *The Review of Economics and Statistics*, 72(2), 334–38.

Nalewaik, J. J. (2006): "Current consumption and future income growth: Synthetic panel evidence," *Journal of Monetary Economics*, 53(8), 2239 – 2266.

Nelson, J. A. (1994): "On Testing for Full Insurance Using Consumer Expenditure Survey Data: Comment," *Journal of Political Economy*, 102(2), 384–94.

Pou, L., and J. Alegre (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.

Schulhofer-Wohl, S. (2011): "Heterogeneity and Tests of Risk Sharing," *Journal of Political Economy*, 119(5), 925 – 958.

Skinner, J. (1987): "A superior measure of consumption from the Panel Study of Income Dynamics," *Economics Letters*, 23, 213–216.

Slesnick, D. T. (1992): "Aggregate Consumption and Saving in the Postwar United States," *The Review of Economics and Statistics*, 74(4), 585–97.

Stephens, M. (2001): "The Long-Run Consumption Effects Of Earnings Shocks," *The Review of Economics and Statistics*, 83(1), 28–36.

——— (2004): "Job loss expectations, realizations, and household consumption behavior.," *The Review of Economics and Statistics*, 86(1), 253–269.

Townsend, R. M. (1994): "Risk and Insurance in Village India," *Econometrica*, 62(3), 539–91.

Wooldridge, J. M. (2003): "Cluster-Sample Methods in Applied Econometrics," *American Economic Review*, 93(2), 133–138.

BANCO DE ESPAÑA PUBLICATIONS

WORKING PAPERS

1220 ENRIQUE ALBEROLA, LUIS MOLINA and PEDRO DEL RÍO: Boom-bust cycles, imbalances and discipline in Europe.

1221 CARLOS GONZÁLEZ-AGUADO and ENRIQUE MORAL-BENITO: Determinants of corporate default: a BMA approach.

1222 GALO NUÑO and CARLOS THOMAS: Bank leverage cycles.

1223 YUNUS AKSOY and HENRIQUE S. BASSO: Liquidity, term spreads and monetary policy.

1224 FRANCISCO DE CASTRO and DANIEL GARROTE: The effects of fiscal shocks on the exchange rate in the EMU and differences with the US.

1225 STÉPHANE BONHOMME and LAURA HOSPIDO: The cycle of earnings inequality: evidence from Spanish social security data.

1226 CARMEN BROTO: The effectiveness of forex interventions in four Latin American countries.

1227 LORENZO RICCI and DAVID VEREDAS: TailCoR.

1228 YVES DOMINICY, SIEGFRIED HÖRMANN, HIROAKI OGATA and DAVID VEREDAS: Marginal quantiles for stationary processes.

1229 MATTEO BARIGOZZI, ROXANA HALBLEIB and DAVID VEREDAS: Which model to match?

1230 MATTEO LUCIANI and DAVID VEREDAS: A model for vast panels of volatilities.

1231 AITOR ERCE: Does the IMF's official support affect sovereign bond maturities?

1232 JAVIER MENCÍA and ENRIQUE SENTANA: Valuation of VIX derivatives.

1233 ROSSANA MEROLA and JAVIER J. PÉREZ: Fiscal forecast errors: governments vs independent agencies?

1234 MIGUEL GARCÍA-POSADA and JUAN S. MORA-SANGUINETTI: Why do Spanish firms rarely use the bankruptcy system? The role of the mortgage institution.

1235 MAXIMO CAMACHO, YULIYA LOVCHA and GABRIEL PEREZ-QUIROS: Can we use seasonally adjusted indicators in dynamic factor models?

1236 JENS HAGENDORFF, MARÍA J. NIETO and LARRY D. WALL: The safety and soundness effects of bank M&As in the EU: Does prudential regulation have any impact?

1237 SOFÍA GALÁN and SERGIO PUENTE: Minimum wages: do they really hurt young people?

1238 CRISTIANO CANTORE, FILIPPO FERRONI and MIGUEL A. LEÓN-LEDESMA: The dynamics of hours worked and technology.

1239 ALFREDO MARTÍN-OLIVER, SONIA RUANO and VICENTE SALAS-FUMÁS: Why did high productivity growth of banks precede the financial crisis?

1240 MARIA DOLORES GADEA RIVAS and GABRIEL PEREZ-QUIROS: The failure to predict the Great Recession. The failure of academic economics? A view focusing on the role of credit.

1241 MATTEO CICCARELLI, EVA ORTEGA and MARIA TERESA VALDERRAMA: Heterogeneity and cross-country spillovers in macroeconomic-financial linkages.

1242 GIANCARLO CORSETTI, LUCA DEDOLA and FRANCESCA VIANI: Traded and nontraded goods prices, and international risk sharing: an empirical investigation.

1243 ENRIQUE MORAL-BENITO: Growth empirics in panel data under model uncertainty and weak exogeneity.

1301 JAMES COSTAIN and ANTON NAKOV: Logit price dynamics.

1302 MIGUEL GARCÍA-POSADA: Insolvency institutions and efficiency: the Spanish case.

1303 MIGUEL GARCÍA-POSADA and JUAN S. MORA-SANGUINETTI: Firm size and judicial efficacy: evidence for the new civil procedures in Spain.

1304 MAXIMO CAMACHO and GABRIEL PEREZ-QUIROS: Commodity prices and the business cycle in Latin America: living and dying by commodities?

1305 CARLOS PÉREZ MONTES: Estimation of regulatory credit risk models.

1306 FERNANDO LÓPEZ VICENTE: The effect of foreclosure regulation: evidence for the US mortgage market at state level.

1307 ENRIQUE MORAL-BENITO and LUIS SERVEN: Testing weak exogeneity in cointegrated panels.

1308 EMMA BERENGUER, RICARDO GIMENO and JUAN M. NAVE: Term structure estimation, liquidity-induced heteroskedasticity and the price of liquidity risk.

1309 PABLO HERNÁNDEZ DE COS and ENRIQUE MORAL-BENITO: Fiscal multipliers in turbulent times: the case of Spain.

1310 SAMUEL HURTADO: DSGE models and the Lucas critique.

1311 HENRIQUE S. BASSO and JAMES COSTAIN: Fiscal delegation in a monetary union with decentralized public spending.

1312 MAITE BLÁZQUEZ CUESTA and SANTIAGO BUDRÍA: Does income deprivation affect people's mental well-being?

1313 ENRIQUE ALBEROLA, ÁNGEL ESTRADA and DANIEL SANTABÁRBARA: Growth beyond imbalances. Sustainable growth rates and output gap reassessment.

1314 CARMEN BROTO and GABRIEL PEREZ-QUIROS: Disentangling contagion among sovereign CDS spreads during the European debt crisis.

1315 MIGUEL GARCÍA-POSADA and JUAN S. MORA-SANGUINETTI: Are there alternatives to bankruptcy? A study of small business distress in Spain.

1316 ROBERTO RAMOS and ENRIQUE MORAL-BENITO: Agglomeration matters for trade.

1317 LAURA HOSPIDO and GEMA ZAMARRO: Retirement patterns of couples in Europe.

1318 MAXIMO CAMACHO, GABRIEL PEREZ-QUIROS and PILAR PONCELA: Short-term forecasting for empirical economists. A survey of the recently proposed algorithms.

1319 CARLOS PÉREZ MONTES: The impact of interbank and public debt markets on the competition for bank deposits.

1320 OLYMPIA BOVER, JOSE MARIA CASADO, SONIA COSTA, PHILIP DU CAJU, YVONNE MCCARTHY, EVA SIERMINSKA, PANAGIOTA TZAMOURANI, ERNESTO VILLANUEVA and TIBOR ZAVADIL: The distribution of debt across euro area countries: the role of Individual characteristics, institutions and credit conditions.

1321 BRINDUSA ANGHEL, SARA DE LA RICA and AITOR LACUESTA: Employment polarisation in Spain over the course of the 1997-2012 cycle.

1322 RODOLFO G. CAMPOS and ILIANA REGGIO: Measurement error in imputation procedures.

1323 PABLO BURRIEL and MARÍA ISABEL GARCÍA-BELMONTE: Meeting our D€STINY. A Disaggregated €uro area Short Term Indicator model to forecast GDP (Y) growth.

1401 TERESA SASTRE and FRANCESCA VIANI: Countries' safety and competitiveness, and the estimation of current account misalignments.

1402 FERNANDO BRONER, ALBERTO MARTIN, AITOR ERCE and JAUME VENTURA: Sovereign debt markets in turbulent times: creditor discrimination and crowding-out effects.

1403 JAVIER J. PÉREZ and ROCÍO PRIETO: The structure of sub-national public debt: liquidity vs credit risks.

1404 BING XU, ADRIAN VAN RIXTEL and MICHIEL VAN LEUVENSTEIJN: Measuring bank competition in China: a comparison of new versus conventional approaches applied to loan markets.

1405 MIGUEL GARCÍA-POSADA and JUAN S. MORA-SANGUINETTI: Entrepreneurship and enforcement institutions: disaggregated evidence for Spain.

1406 MARIYA HAKE, FERNANDO LÓPEZ-VICENTE and LUIS MOLINA: Do the drivers of loan dollarisation differ between CESEE and Latin America? A meta-analysis.

1407 JOSÉ MANUEL MONTERO and ALBERTO URTASUN: Price-cost mark-ups in the Spanish economy: a microeconomic perspective.

1408 FRANCISCO DE CASTRO, FRANCISCO MARTÍ, ANTONIO MONTESINOS, JAVIER J. PÉREZ and A. JESÚS SÁNCHEZ-FUENTES: Fiscal policies in Spain: main stylised facts revisited.

1409 MARÍA J. NIETO: Third-country relations in the Directive establishing a framework for the recovery and resolution of credit institutions.

1410 ÓSCAR ARCE and SERGIO MAYORDOMO: Short-sale constraints and financial stability: evidence from the Spanish market.

1411 RODOLFO G. CAMPOS and ILIANA REGGIO: Consumption in the shadow of unemployment.

BANCO DE ESPAÑA
Eurosistema

Unidad de Servicios Auxiliares
Alcalá, 48 - 28014 Madrid
E-mail: publicaciones@bde.es
www.bde.es