European Economic Review 78 (2015) 39–54

Contents lists available at ScienceDirect

European Economic Review

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

Consumption in the shadow of unemployment ☆

Rodolfo G. Campos [a,*], Iliana Reggio [b]

[a] Department of Economics, IESE Business School, Camino del Cerro del Águila 3, 28023 Madrid, Spain
[b] 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.

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?

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

* Corresponding author.
E-mail addresses: rcampos@iese.edu (R.G. Campos), ireggio@eco.uc3m.es (I. Reggio).

http://dx.doi.org/10.1016/j.euroecorev.2015.04.006
0014-2921/

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%.[1]

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.[2] 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

[1] In comparison, in the US the unemployment rate averaged 4.8% in 2007, and peaked at 9.6% in 2010.

[2] 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%.