The Effect of Movements in House Prices on Employment
Spain, 1995-2019

Lorenzo Germinetti
Student ID: 503363

Abstract

I study the effect of movements in house prices on employment. I conduct the analysis exploiting geographical variation in Spain, studying a dataset with information relative to 47 Spanish provinces from 1995 to 2019. In order to identify the elasticity of employment with respect to house prices, I use the *sensitivity instrument* approach proposed by Guren et al. (forthcoming), relying on systematic differences in the sensitivity of house prices at the province-level to cycles in house prices at the national level. I find that movements in house prices have an economically significant impact on employment. The results provide evidence of the fact that the response of households' consumption to movements in house prices plays a significant role in driving the relationship of interest. I find that a 10% larger decline in house prices in a province relative to another leads to a 1.65% to 2.02% larger decline in non-tradable employment. Part of this effect is driven by workers moving to provinces facing a less steep decline in house prices: the same 10% relative decline in house prices implies a relative decrease in non-tradable employment *per-capita* of 1.15% to 1.25%. Lastly, I obtain weak evidence of the fact that the interaction of high levels of debt with sharp declines in house prices implies a larger response of non-tradable employment to changes in the value of housing.

Master Thesis, MSc Policy Economics
Supervisor: Bas Jacobs
Second Assessor: ...

Date: August 2020
Erasmus School of Economics
Erasmus University Rotterdam

Introduction

The Spanish 2008-2013 recession was characterized by significant economic and social costs. Between the first quarter of 2008 and the last quarter of 2013 employment at the national level dropped by 17%. The burst of the housing bubble was surely key in the recession: in the same period of time, house prices declined by 30% (figure 1). The magnitude of the housing bubble was large even in comparison to similar experiences occurring in the 2000s in other countries. Taking the often-studied US case as a comparison and using annual data from the Macrohistory Database (Jordà, Schularick and Taylor, 2017) it is noticeable how house prices in the US declined by 16% from their peak in 2007 to 2011. Similarly, house prices in Spain were 160% higher in 2007 than they were in 2000, while in the same period in the US house prices rose by 60%. Figure 2 makes also clear that the increase in mortgage debt in the run up to the recession in Spain far exceeded the one witnessed in other countries such as the US, the Netherlands and Italy.

Figure 1: Employment and house prices in Spain, 1995-2019 (indexes, 2008q1=100)

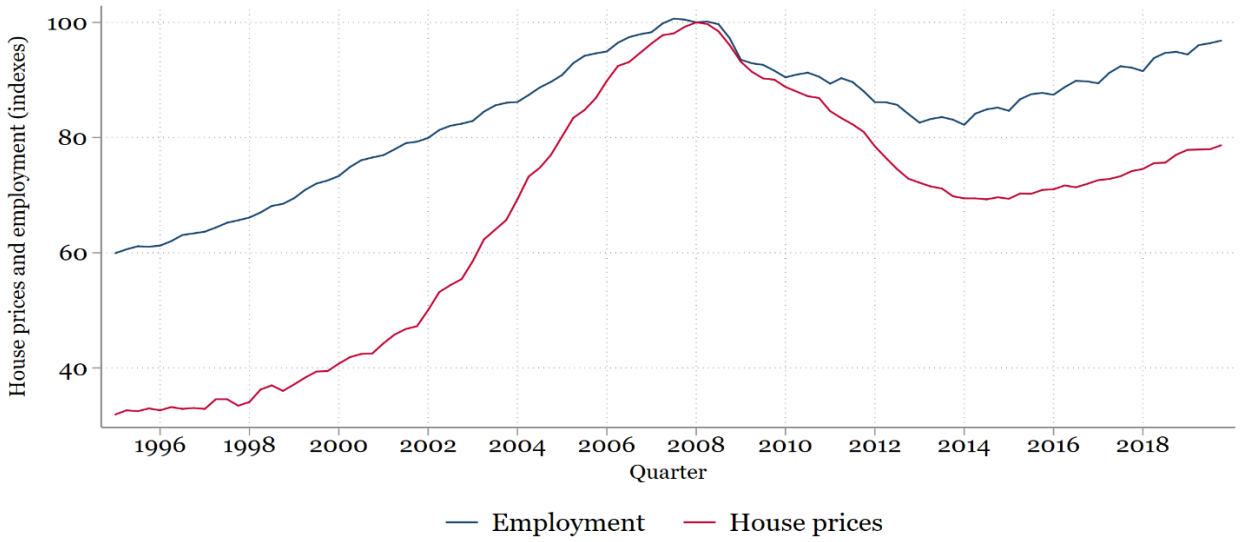

Note: Both indexes take on the value of 100 in 2008q1 and develop according to quarterly growth rates in the other periods. The data are described in more detail in section 2, the ones on employment are from the Spanish National Statistical Institute (INE), while the ones on house prices are from the Spanish Ministry of Transport, Mobility and Urban Agenda (Ministerio de Transportes, Movilidad y Agenda Urbana).

The crisis was a powerful reminder of the economic and social costs of disruptions in housing and financial markets. The events witnessed in numerous countries highlighted the relevant roles played by the balance sheets of borrowers (of households and financial intermediaries in particular) in causing and driving the recession (Gertler and Gilchrist, 2018). In particular, the economic developments focused attention on the consequences of sharp movements in house prices for the economy. Two questions became of interest: To what extent, and through which channels, do movements in house

prices lead to movements in employment? Does the interaction of movements in house prices with the state of households' balance sheets affect this relationship?

Figure 2: The real estate boom and bust: 2000-2015, selected countries

(2.a) House prices growth rates

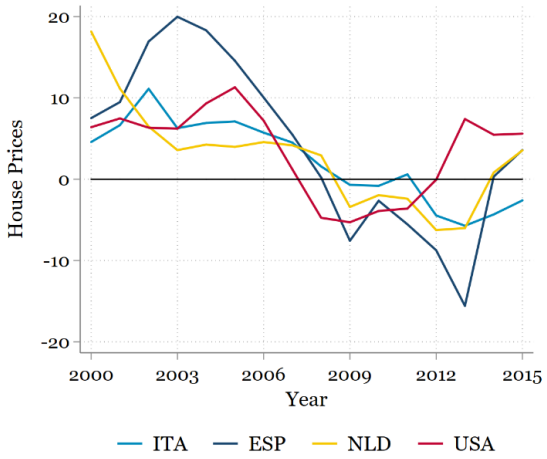

(2.b) Mortgage loans to the private sector to GDP ratio

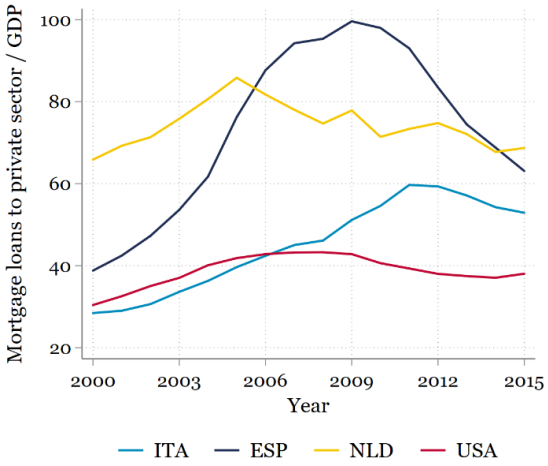

Note: The growth rates displayed in panel (2.a) are yearly growth rates in nominal house prices. The ratios displayed in panel (2.b) are between the nominal value of mortgage loans to the private sector and the nominal value of GDP. Data are from the Macrohistory Database by Jordà, Schularick and Taylor (2017).

Addressing these questions, empirical research has focused on the U.S. experience, either investigating the relationship between house prices and employment (Mian and Sufi, 2014; Guren, McKay, Nakamura and Steinsson, forthcoming; and Adelino, Schoar and Severino, 2015) or the one between housing values and households' consumption (Mian, Rao and Sufi, 2013); Kaplan, Mitman and Violante, 2017; and Aladangady, 2017). These studies generally agree on the economic significance of the effect of house prices on economic activity, while not reaching a clear consensus on: (i) what mechanism, in particular whether an effect on households' consumption or a shift in credit supply for firms, drives the relationship; and (ii) how this relationship interacts with the state of households' balance sheets (e.g. with the level of leverage or with the ratio of loan to income).

I contribute to this literature by presenting new estimates of the elasticity of employment to house prices, obtained exploiting heterogeneous developments in housing values across Spanish provinces between 1995 and 2019. In particular, I use the *sensitivity instrument* introduced by Guren et al. (forthcoming), exploiting variation in house prices developments across provinces in Spain, to retrieve said elasticity. I regard the compiled dataset and the estimation of a measure of the sensitivity of local house prices to national housing cycles which is independent from the sensitivity of local employment to aggregate shock (it is possible to refer to this measure as the *sensitivity parameter*) as valuable contributions of the present study.

Focusing on the Spanish experience, the analysis speaks to the issue of external validity with respect to the estimates obtained by other contributions to the literature. The Spanish institutional

setting is likely to differ from the US one in dimensions that are surely relevant in determining the extent to which house prices affect employment. These dimensions comprise (i) the characteristics of the mortgage market: for instance, differences in credit constraints set at mortgage origination, in reliance on mortgage refinancing, and in foreclosure laws can lead to differences in the response of households' consumption to changes in house prices; (ii) nominal rigidity in wages and prices; (iii) labor market institutions such as the generosity of unemployment insurance (iv) labor mobility across regions and economic sectors; (v) the tendency of firms to rely on housing as collateral to obtain credit; and (vi) the reliance of local governments' on property taxes and their related capacity to smooth spending over time. It is therefore of interest to compare the estimates obtained in the Spanish context to the ones presented by previous studies focusing on US data.

Furthermore, I contribute to the literature by investigating some of the questions over which a clear consensus has not been reached. In particular, I study the relative importance of the responses of households and firms to movements in house prices in driving the relationship between housing values and employment. I also exploit variation in mortgage borrowing across localities, in particular in the lead up to the 2008-2013 recession, in order to study how the interaction of movements in house prices and debt, both during good and bad times and in the context of a housing bust, influences the relationship between housing values and employment. Lastly, I ask to what extent labor mobility across regions is able to account for the heterogeneous responses of employment to different movements in house prices.

In order to conduct the analysis, I compile a dataset with information on 47 Spanish provinces between 1995 and 2019. The main challenge to identification of the elasticity of employment with respect to house prices lies in the fact that house prices and economic activity are jointly determined. In the empirical analysis, I overcome this challenge by adopting the *sensitivity instrument* approach proposed by Guren et al. (forthcoming). Such an instrument exploits variation in sensitivity of province-level house prices to national house prices movements. When house prices at the national level change, house prices in some provinces tend to respond systematically more relative to those in other provinces. Regressing the growth rate of local house prices on the growth rate of house prices at the national level and on an appropriate set of controls allows to obtain an estimate of the province level *sensitivity parameter*: the sensitivity of local house prices to national house prices cycles which is independent from the sensitivity of local employment with respect to aggregate economic shocks. The *sensitivity instrument* is then constructed by interacting this province-level parameter with the growth rate in house prices at the national level in each period. The main identifying assumption underlying this approach is that conditional on controls, there is no unobserved factor that is both correlated with house prices in the time series and that differentially affects employment in the same localities where house prices are estimated to be historically more sensitive to national housing cycles in the cross section.

The results of the analysis show that movements in house prices have an economically significant effect on employment. Estimating the elasticity of employment in (proxies for) non-tradable and tradable sectors, I obtain evidence supporting the idea that the response of households' consumption plays a significant role in driving the relationship between the value of housing and

employment. I find that a 10% larger decline in house prices in a province relative to another leads to a 1.65% to 2.01% larger decline in non-tradable employment. The fact that workers move to provinces facing a less steep decline in house prices is able to explain only part of this effect: the same 10% relative decline in house prices implies a relative decrease in non-tradable employment *per-capita* of 1.15% to 1.25%. Lastly, I obtain weak evidence of the fact that the interaction of high levels of debt with sharp declines in house prices leads to a larger response of non-tradable employment to changes in house prices.

The analysis proceeds as follows. In section 1, I present a theoretical framework that contextualizes the empirical analysis in a setting where movements in house prices that result to be heterogeneous across localities can lead to different outcomes in terms of employment. Section 2 clarifies how previous contributions to the literature have attempted to test the implications of the outlined theoretical framework, discussing their results. Section 3 describes the dataset, while section 4 exploits the data in order to present a descriptive account of the Spanish 2008-2013 recession. In section 5 I explain the identification strategy, discussing the challenges to identification as well as the choices made in the estimation of the *sensitivity parameter* to address them and presenting evidence in support of the *meaningful first stage* assumption for the *sensitivity instrument.* In section 6 I present the key results of the present analysis. In particular, section 6.1 focuses on estimates of the elasticity of total employment with respect to house prices; section 6.2 focuses on non-construction employment and on measures of non-tradable and tradable employment; section 6.3 presents result obtained when employment *per-capita* is used as dependent variable; and section 6.4 investigates the extent to which the relationship between house prices and employment in non-tradable sectors varies with varying levels of debt when house prices decline sharply. Section 7 concludes by discussing how the results obtained relate to previous contributions and proposing avenues for future research.

1 – Theory

The investigation is motivated by an understanding of the working of the economy where heterogeneous movements in house prices across localities can be drivers of fluctuations in economic activity. I briefly delve into a discussion of the economic mechanisms that can explain this relationship in sections 1.1 and 1.2.

1.1 – The relationship between house prices and economic activity

Housing is a relevant component of households' and firms' balance sheets. Changes in its value can be thought to impact economic activity, and therefore employment, through multiple mechanisms that concern the economic decision-making of households and firms.

A key channel through which changes in house prices can impact employment is households' and firms' access to credit (I refer to this channel as *the collateral lending channel* respectively for households and firms). As argued by Gertler and Gilchrist (2018), this mechanism is relevant when credit frictions, such as asymmetric information, limit borrowing and lending. As households and firms,

as borrowers, are better informed than lenders, they can gain at the expense of creditors by acting dishonestly. In such a setting, lenders will impose credit constraints which link the amount of money available to the borrower in the form of a loan to the value of collateral pledged.

The role of the value of collateral in explaining shifts in economic activity has been investigated theoretically starting with the seminal contributions of Bernanke and Gertler (1989), Bernanke, Gertler and Gilchrist (1999), Kiyotaki and Moore (1997) and Iacoviello (2005). Whereas the first three works focused solely on constraints for non-financial firms, Iacoviello (2005) introduced constraints both for entrepreneurs and for a subset of households (who are relatively impatient and therefore borrow from more patient households). Credit constraints affecting households have later been incorporated in models built in the aftermath of the financial crisis that started in 2007 (for instance Eggertsson and Krugman, 2012 and Guerrieri and Lorenzoni, 2011).

Focusing further on the response of household consumption to movements in house prices (I refer to this mechanism as the *household demand channel* in the analysis), it is possible to disentangle other mechanisms at play. Berger, Guerrieri, Lorenzoni and Vavra (2017), studying a partial equilibrium model of housing and non-housing consumption choice featuring income uncertainty, rental markets and collateralized borrowing[1] argue that it is possible to decompose the effect of consumption responses to an house prices increase in: (i) the *collateral lending channel* mentioned above, leading to more credit being available for a given housing choice; (ii) an *income effect*, whereby the increase in house prices leads households to be poorer following the increase in implicit rental costs; (iii) a *substitution effect* that makes households consume relatively more non-housing goods; and (iv) an *endowment income effect*, as households are richer relative to their initial housing holdings.

The authors arrive at two main conclusions that have important implications for the analysis presented below. First, such a realistic model of consumption predicts an arguably large consumption response to house prices movements (the authors focus attention on an aggregate elasticity of consumption to house prices of 0.23 for working age households). Second, the analysis derives a sufficient statistics formula which characterizes the consumption response of interest as depending solely on the product between the households' marginal propensity to consume and the ratio of the value of housing held over consumption. Importantly, this result implies that the elasticity of interest varies with the state of households' balance sheets: households that hold more debt relative to their assets not only tend to have a larger marginal propensity to consume out of the endowment income effect as a result a precautionary saving motive, but are also more likely to face borrowing constraints. Moreover, heterogeneity in consumption responses is predicted by the model also across tenancy regimes (renters do not respond to movements in house prices, while owners do), as well as along the income and age distributions.

As a last step in the discussion of the relationship between house prices and consumption, it appears relevant to highlight that recent theoretical contributions have argued that such a relationship is asymmetric across periods, depending on the level of house prices (Guerrieri and Iacoviello, 2017) as

[1] It is noted that a model with similar characteristics is studied also by Chen, Michaux and Roussanov (2020), who conclude that such a framework is able to account for salient features of the evolution of households' consumption and balance sheets both over time and in the cross-section of households during the US housing cycle of the early 2000s.

well as on the level of leverage of households and firms leverage (Jensen, Petrella, Ravn and Santoro, 2020). Such an asymmetry arises in models with representative borrower households and firms who result to be constrained in the amount that can be borrowed only in periods when house prices are relatively low or leverage is high, i.e. when the amount of their debt outstanding compared with the one of their assets brings them close to the borrowing constraint. When house prices are relatively high and leverage is low, loose borrowing constraints lead to the *collateral lending channel* not acting as a driver of the relationship between house prices and economic activity. In contrast, when house prices decline and leverage is high, economic agents find themselves constrained in the amount they can borrow and engage in deleveraging to restore the health of their balance sheets.

In contrast with this view, Guren, Mckay, Nakamura and Steinsson (forthcoming) argue that a model of the kind studied by Berger et al. (2017), incorporating heterogeneity across households over an entire distribution (and not across two representative households – respectively borrower and lender), rental markets and long-term mortgages, is consistent with the effect of house prices on economic activity being insensitive to changes in households' leverage. This conclusion is reached by noticing that such a model predicts a large response to house prices for households that are not close to the credit constraint. On the one hand, the possibility of becoming credit constrained in the future is taken into account by households. On the other, the response is also driven by the *endowment income effect* as well as by the *substitution effect* described above. In addition, as mortgages are long-term contracts, households are not forced to delever to satisfy a change in credit constraint during a bust, while homeowners that are pushed underwater (i.e. homeowners for whom the level of debt exceeds the value of housing held) become insensitive to changes in house prices, because they cannot access increases in housing wealth by refinancing their mortgages.

This discussion implies that it is of interest, in an empirical analysis of the kind conducted in the present study, to: (i) evaluate the relative importance of the responses respectively of households and firms to changes in house prices (section 6.2); (ii) to investigate whether the theoretical prediction of a large consumption response to such changes is confirmed by the data; and (iii) to question whether the employment response to movements in housing values is asymmetric in the interaction between leverage and low (or declining) house prices, i.e. if declines in house prices occurring when households are highly leveraged imply a relatively larger response of economic activity to house prices (section 6.4).

It is worth mentioning two additional potential mechanisms driving the relationship of between house prices and employment. First, Sole'-Olle' and Viladecans-Marsal (2019) report that Spanish municipalities obtain on average around 22% of their revenues from property taxes[2], while spending at this geographical level accounts for around 15% of public spending in the country. Changes in housing values are likely to drive changes in revenues and spending at the municipality level, and this will in turn affect employment. Second, to the extent that banking markets are local, the relationship between households' and firms' balance sheets and banks' balance sheets can also drive the fact that

[2] This figure is obtained considering that the authors report that municipalities' own revenues account for two-thirds of municipalities' current revenues; in turn, taxes account for two-thirds of municipalities' own revenues and property taxes account for 50% of tax revenues. Thus: (2/3) * (2/3) * (1/2) = 0.222, or 22.2%.

heterogeneous developments in house prices lead to different movements in employment across localities. I anticipate here that, due to data availability, I am not able to investigate the extent to which these channels drive the relationship between house prices and employment.

1.2 – Geographically heterogeneous economic developments and adjustment mechanisms

As mentioned in the introduction, the empirical analysis conducted here exploits geographical variation in house prices and employment. It is therefore warranted to discuss through which mechanisms a shift in local house prices can be thought to impact local economic activity. Focusing in particular on the effect of house prices on consumption, it appears necessary to clarify the way in which adjustment mechanisms could absorb asymmetric shocks and lead to the absence of heterogeneous outcomes across localities. The first mechanism considered relates to the differential impact of local movements in house prices on the non-tradable and tradable sectors, the corresponding adjustment in wages and the possibility of workers to move across sectors. The second concerns instead labor mobility, i.e. the movement of workers across provinces in response to geographically heterogeneous movements in house prices.

As discussed by Mian and Sufi (2014) a local decline (increase) in employment and demand might be associated with a corresponding drop (increase) in wages. Referring to a decline in house prices for ease of exposition, such a drop causes local demand to fall, driving a decline in non-tradable employment (which depends directly on local demand). If such shift would result in a decline in wages and if workers could move freely across sectors, the drop in wages would occur also in the tradable sectors, causing a corresponding increase in labor demand in these industries. Mian and Sufi (2014) present a model where the absence of nominal rigidities and labor mobility across sectors imply that a relatively larger shock to house prices in some localities does not translate into a larger drop in employment. Intuitively, the decline in non-tradable employment in these localities is compensated by the increase in tradable employment. In fact, the whole economy remains at full employment thanks to the operation of these adjustment mechanisms.

I investigate whether the employment response to movements in house prices is heterogeneous across tradable and non-tradable sectors in section 6.2. On the one hand, this investigation provides indirect evidence on the extent to which nominal rigidities and labor market frictions with respect to the movement of workers across sectors are at play. On the other, it offers an insight into the relative importance of the impact of house prices movements respectively on households' demand and on credit availability for firms. While employment in non-tradable sectors can depend on local demand, employment in tradable sectors is not affected by locality specific changes in households' spending. Furthermore, there is no apparent reason why shifts in credit availability should impact tradable and non-tradable sectors differently.

An additional labor market adjustment mechanism in response to geographically heterogeneous developments in house prices and employment is to be found in labor mobility across localities. In response to differential shifts in house prices and their corresponding changes in labor demand from

firms, driven either by movements in local consumption or in firms' credit availability, workers could decide to move to localities witnessing more favorable economic developments. Such a response would also lead to an estimation of a positive elasticity of employment with respect to house prices, although with peculiar implications. On the one hand, workers' mobility would allow individuals to insure against negative shocks incurred as a result of residing in a specific province. On the other, the extent to which workers move across provinces is related to the degree to which the estimates presented above are driven by the ability of employment in provinces that are less affected by such a negative shock to compensate for the decline in employment in more affected provinces. The impact of housing cycles on national employment would be muted in the presence of smooth labor mobility across provinces.

Following Guren et al. (forthcoming) I study the degree to which the movement of workers across provinces drives the employment response to house prices by presenting results relative to the elasticity of both non-per-capita and per-capita employment in section 6.3.

2 – Related Empirical Literature

The theoretical predictions concerning the effects of movements in house prices on employment, consumption and firms' credit availability can be tested and investigated empirically. The present analysis is most closely related to the studies by Mian and Sufi (2014); Adelino, Schoar and Severino (2015); and Guren, McKay, Nakamura and Steinsson (forthcoming). These authors study the effect of movements in house prices on employment exploiting geographical variation in house prices and employment movements in the US.

The study by Mian and Sufi (2014) has been particularly influential in establishing an interpretation of the 2007-2009 US financial crisis as driven by the deterioration in households' balance sheets (together with, notably, Mian, Rao and Sufi, 2013). The authors focus on the elasticity of non-tradable employment to housing net worth, defined as the product between house prices and the ratio of the value of the stock of housing held by households over their net worth. The focus on this variable on the right-hand-side is driven by the idea that households in different localities are differently exposed to a drop in house prices when there is heterogeneity in their balance sheets. With respect to the variable used on the left-hand-side, instead, the authors focus on two different measures of non-tradable employment: retail and restaurants employment and employment in geographically dispersed industries.

The study uses a dataset at the US county level with data relative to the period 2007-2009, thus focusing on an equation exploiting only variation in the cross-section and not over time. The analysis uses both OLS with controls and an instrumental variable approach in order to identify the causal effect of a decline in housing net worth on non-tradable employment through what they term the housing net worth channel (i.e. the effect of a shock to households' balance sheets driven by a change in house prices on households' consumption). A measure of housing supply elasticity constructed by Saiz (2010), combining geographical and regulatory restrictions at the level of Metropolitan Statistical Areas (MSAs), is used as instrument for the drop in housing values.

The authors conclude that the role played by the housing net worth channel in the decline in employment occurred in the US between 2007 and 2009 was significant. Furthermore, the study finds that the collateral lending channel for firms did not play a significant role in the US recession and that there is little evidence of labor market adjustments. They draw these last conclusions showing that (i) small businesses, theoretically more likely to incur in borrowing constraints and to use housing as collateral, were not characterized by a larger drop in employment; and (ii) employment in tradable sectors did not respond to movements in house prices during the bust. Lastly, the authors provide evidence in support of the idea that the relationship between non-tradable employment and the decline in housing net worth is not driven by changes in local governments' finances (they state that the coefficients of interest are unchanged when controlling for state spending cuts) nor by credit restrictions by banks (the relationship between employment and house prices remains unvaried for counties that are primarily served by national banks).

Adelino, Schoar and Severino (2015) investigate the causal effect of movements in house prices on employment through the collateral lending channel for firms. As the present analysis, their study also focuses on the elasticity of employment with respect to house prices. The authors compile a dataset at the US county level with information relative to the period 2002-2007. As identification strategy, their analysis also adopts an instrumental variable approach, using the Saiz (2010) measure of housing supply elasticity as instrument. As Mian and Sufi (2014), the authors exploit cross-sectional variation across counties in a given time period (corresponding to the housing boom: 2002-2007).

The authors find that small businesses (and in particular businesses in industries where start-up capital plays an important role in the creation of new firms) experienced a larger increase in employment as a consequence of the increase house values between 2002 and 2007 in the US. The authors conclude that the collateral lending channel for firms is a relevant mechanism to explain the effect of movements in house prices on employment. In particular, they argue that higher house prices allow for the creation of new small firms, as individuals can pledge collateral and are able to start a new business or become self-employed. In order to rationalize the contrasting nature of their and Mian and Sufi's (2014) results, the authors propose an asymmetric effect of housing value on employment through the collateral lending channel for firms, with this channel being operative when house prices increase and not when they decrease.

The present analysis is also greatly indebted to the study by Guren, McKay, Nakamura and Steinsson (forthcoming). Their analysis, on the one hand, develops the identification strategy and the empirical approach adopted to conduct the present empirical investigation (discussed below). On the other hand, the authors focus on the estimation of the elasticity of employment in the retail sector (as a proxy of non-tradable employment) with respect to house prices over the long term in the US, in order to investigate whether economic activity was more sensitive to movements in house prices during the early 2000s (characterized by the household boom and bust) relative to previous periods.

The study concludes that housing wealth elasticities were slightly smaller in the 2000s relative to earlier periods. In addition, no evidence is found of a boom and bust asymmetry in the elasticity of retail employment to house prices. These two findings are interpreted by the authors as pointing at the fact that the interaction of households' leverage and (negative) movements in house prices does

not lead to asymmetries in the elasticity of interest. Moreover, it is worth noticing that the authors stress that the effect of house prices on employment is driven by their effect on household consumption (they link the elasticity of retail employment with respect to house prices 1:1 with the elasticity of consumption with respect to house prices).

Another strand of the literature has instead focused directly on the relationship between house prices and consumption. Relevant contributions are the ones by Mian, Rao and Sufi (2013), Kaplan, Mitman and Violante (2017) and Aladangady (2017).

It is relevant to note that Kaplan, Mitman and Violante (2017) characterize their study as a replication of Mian, Rao and Sufi (2013) with more easily accessible data. These two analyses exploit geographic variation at the county (both studies) and ZIP code (Mian, Rao and Sufi, 2013) level in order to obtain the elasticity of consumption with respect to housing net worth (as in Mian and Sufi, 2014) as well as to house prices. They find similar and economically and statistically significant effects of house prices on consumption. The main difference between the two studies is to be found in the fact that, whereas Mian, Rao and Sufi (2013) find that the elasticity of consumption is heterogeneous along the leverage and income distribution, Kaplan, Mitman and Violante (forthcoming) find what they describe as "weak" and "more nuanced" evidence with respect to the impact of leverage. While the latter group of authors concede that the differences might arise because of the different level of geographical aggregation used, they also argue that the evidence they find is consistent with wealth effects driving the relationship between house prices and consumption (i.e. with leverage and credit constraints playing no role).

Lastly, Aladangady (2017) uses geographically linked US microdata to exploit regional heterogeneity in housing markets to identify the causal effect of house price fluctuations on consumer spending. The analysis adopts an instrumental variable approach, instrumenting changes in house prices with the interaction between the real 10-year US Treasury rate and two different measures of local housing supply elasticity: the proportion of unavailable land in a given locality, namely a Metropolitan Statistical Area (MSA) and the Wharton Land Use Regulation Index, also at the MSA level.

The disaggregated nature of the data allows the author to investigate differences in the marginal propensity to consume out of housing wealth across subpopulations of individuals. The analysis of differences between owners and renters leads to finding a negligible effect of movements in house prices on spending for renters and a positive and significant effect for homeowners. Importantly, the study further finds larger consumption responses among households with larger debt service ratios and with higher loan-to-value ratios, while no significant differences are found across the income distribution. The author concludes that collateral effects play an important role in driving the consumption response to house prices.

3 – Data

I compile a panel dataset with information on 47 Spanish provinces (I exclude the islands and the autonomous cities of Ceuta and Melilla). To this dataset, I add information on the time series

dimension of national level variables. The dataset spans the years 1995-2019 at quarterly frequency. The variables included are summarized in table 1. In what follows I discuss each variable in greater detail, explaining the data transformations carried out when preparing the final dataset.

Data on house prices in each quarter between 1995 and 2019 are obtained from the Spanish Ministry of Transport, Mobility and Urban Agenda (Ministerio de Transportes, Movilidad y Agenda Urbana). The prices refer to the average price per square meter of appraised, non-subsidized and not subject to land limitations or maximum prices (a.k.a. free) dwellings. I use the GDP deflator to obtain real house prices. The main causal variable of interest consists in the difference in logs of this variable in a given quarter compared to the same quarter in the previous year.

Data on total employment and non-construction employment are also available quarterly from 1995 to 2019. I obtain these data from the Spanish National Statistical Institute (Instituto Nacional de Estadistica, INE). I first construct the variable *non-construction employment* by subtracting construction employment from total employment. I focus on this variable in order to isolate the impact that movements in house prices can have on employment through the mechanisms explained above, disregarding the more direct effect that such fluctuations have on employment in the construction sector. Total employment and non-construction employment are used as dependent variables in the empirical analysis. Also in this case, I use the log-difference of the variable in a given quarter compared to the same quarter in the previous year. This has the perk of allowing results not to be driven by seasonal variation in employment or house prices (figure 5.b, below, points at the fact that such seasonal variation in employment might be particularly relevant).

Furthermore, I collect data on employment by six (more detailed) economic sectors in each province from the Spanish National Statistical Institute (the data are made available as part of the Spanish Regional Accounts). The variables are available yearly from 2000 to 2017. In line with the literature, focusing on measures of employment in specific economic sectors I restrict attention on: (i) employment in the manufacturing sector, considered as employment in a tradable sector; and (ii) employment in the sector comprising 'wholesale and retail trade, repair of motor vehicles, transport and storage, hospitality industry and information and communication'. Given the nature of many of the latter activities, notably retail and the hospitality industry (see Mian and Sufi, 2014 and Guren et al., forthcoming), I deem this sector to be a good proxy for employment in non-tradable sectors. With a slight abuse in terminology I refer to employment in this sector as 'non-tradable' employment in the remainder of the analysis. I then obtain *per-capita* values of retail employment by dividing the relevant variable by total province population. Once again, I use the yearly log-difference of these measures of employment as dependent variables in the analysis.

Not only are (some) of the measures of employment by *detailed* economic sector used as dependent variables, but I also use this information to construct measures of the employment shares in each sector, which are in turn used as control variables in interactions with time fixed effects. The sectors into which employment is decomposed are listed in table (1.b). I first obtain yearly measures of employment shares in each sector by dividing employment in the relevant sector by total employment. Then, given that the employment data refer to a calendar year, I assume these shares to remain constant throughout each year between 2000 and 2017. Further, I construct measures of

employment shares in the periods 1996-1999 and 2018-2019 by assuming that these variables followed, in these two time spans, the same linear time trend that they followed in the 2000-2017 period and by constructing the time series for each employment share so that it is continuous. This assumption reflects the idea that movements in employment shares in a given province are driven by long-term trends in the structure of the economy.

The Spanish National Statistical Institute is also the source of the data on yearly output (GDP) in a given province for the period 2000-2017 as well as of the data on the amount of mortgages *constituted* in a given province in a given quarter for the period 2003-2019. I combine these two variables in order to construct measures of the ratio of originated mortgage debt to output. Specifically, for each quarter between 2003 and 2019 I construct the ratio of mortgages constituted over yearly output. This variable is used in turn to construct the interaction terms in some of the regressions presented below (section 6.4). Specifically, the interaction term uses the average of this variable over the 2003-2007 period. The key idea behind this measure is that it captures part of the mortgage loan to income ratio, focusing on recently originated mortgages. Importantly, as the interaction terms in the regressions below use past values of such recently originated mortgages to output ratio, they allow to capture some degree of heterogeneity in the past reliance of provinces on mortgage debt.

Lastly, at the national level I obtain data on total employment as well as on population from the Spanish National Statistical Institute (INE). I also download data on house prices from the Spanish Ministry of Transport, Mobility and Urban Agenda and on the interest rate on mortgages with a maturity of more than three years from the Bank of Spain. When running regressions I make use of the yearly log-difference of employment and of house prices, while I include the difference from the previous year in the case of the interest rate on mortgages.

Table 1: Variables

Panel (1.a): Causal and dependent variables

Variable description	**Time Period**	**Frequency**	**Geographical Level**	**Source**
House prices (euros/m2)	1995-2019	Quarterly	Province	Spanish Ministry of TM&UA
Total employment and non-construction employment	1995-2019	Quarterly	Province	INE 1995-2001 2002-2008 2009-2019
Employment in wholesale and retail trade, repair of motor vehicles, transport and storage, hospitality industry, information and communication	2000-2017	Yearly	Province	INE
Employment in the manufacturing industry	2000-2017	Yearly	Province	INE

Panel (1.b): Variables at the province level

Variable description	Time Period	Frequency	Geographical Level	Source	Use
Employment by detailed economic sector Sectors (the letters refer to the CNAE 2009 classification): A Agriculture B-E Extractive industries, supply of electrical energy, gas, steam and air conditioning, water supply, waste management, sewerage activities and decontamination F Construction G-J Wholesale and retail trade, repair of motor vehicles, transport and storage, hospitality industry, information and communication K-N Financial and insurance activities, real estate activities, professional, scientific and technical activities, administrative and support services O-U Defense, public administration, social security, education and health, arts, recreation and entertainment, domestic activities, extraterritorial organizations and other services	2000-2017	Yearly	Province	INE	Construct the "sectoral exposure" variable, i.e. the employment share in each sector. This is used in the main empirical specification (MES)
Output	2000-2017	Yearly	Province	INE	Construct the ratio of mortgages constituted in a given quarter over yearly output
Mortgages constituted for dwellings	2003-2019	Monthly	Province	INE	Construct the ratio of mortgages constituted in a given quarter over yearly output
Population by five-year age group	1995-2019	Yearly	Province	INE	Construct employment per-capita as well as the young-adult population share in turn used in the MES and in the estimation of the sensitivity parameter (ESP)

Panel (1.c): Variables at the national level

Variable description	Time Period	Frequency	Geographical Level	Source	Use
Total employment	1995-2019	Quarterly	National	INE	MES & ESP
House prices (euros/m2)	1995-2019	Quarterly	National	Spanish Ministry of TM&UA	ESP
Interest rate on mortgages of more than 3 years	1995-2019	Monthly	National	Bank of Spain	ESP
GDP deflator	1995-2019	Quarterly	National	Bank of Spain	Construct real house prices
Total population	1995-2019	Yearly	National	INE	Construct employment per capita at the national level

Sources abbreviations: Spanish Ministry of TM&UA: Spanish Ministry of Transport, Mobility and Urban Agenda (Ministerio de Transportes, Movilidad y Agenda Urbana). INE: Spanish National Statistical Institute (Instituto Nacional de Estadistica). *Specifications abbreviations:* MES: main empirical specification; ESP: estimation of the sensitivity parameter.

Table 2: Summary Statistics Across Provinces

Variable	Mean	Std. Dev.	10th pct.	90th pct.
Real house price decline 2008q1-2013q4 (%)	-30.10	7.44	-39.17	-20.62
Employment decline 2008q1-2013q4 (%)	-17.07	4.63	-22.69	-11.51
Non-construction employment decline 2008q1-2013q4 (%)	-9.98	5.78	-16.67	-2.74
Employment per capita decline 2008q1-2013q4 (%)	-17.62	4.89	-23.42	-11.01
Non-construction employment per capita decline 2008q1-2013q4 (%)	-10.60	5.58	-16.87	-2.62
'Non-tradable' employment decline 2008-2013 (%)	-13.95	4.4	-19.26	-9.99
'Non-tradable' employment per capita decline 2008-2013 (%)	-14.91	5.04	-21.02	-7.54
House price (euros/square meter)	1175.1	369.3	774.89	1926.82
Population (thousands)	874	1122	171	1689
Employment per-capita (%)	38.5	3.8	33.0	43.7
Real house price growth rate (%)	1.33	0.98	0.11	2.22
Employment growth rate (%)	1.5	0.9	0.4	2.5
$(m/q)_{t-1}$ (%)	1.9	0.7	1.3	2.9
$(m/q)_{03-07}$(%)	3.5	1.5	2.0	5.8
Agriculture share (%)	8.5	5.1	2.8	15.5
Industry share (%)	15.5	5.5	9.2	23.6
Manufacturing share (%)	14.3	5.5	7.8	22.3
Construction share (%)	10.1	1.5	8.3	12.3
G-J sectors share ('non-tradable') (%)	27.9	3.0	24.2	32.2
K-N sectors share (%)	10.0	2.2	7.9	12.8
O-U sectors share (%)	28.0	2.9	24.3	32.3
Share of young adults (15-39 y.o.) (%)	34.0	2.6	30.4	37.4

Note: The table presents summary statistics relative to the distribution of variables *across provinces*. More precisely, the average of each variable is taken over the time series for each province. Statistics relative to the resulting distribution of means are presented. The variables $(m/q)_{t-1}$ and $(m/q)_{03-07}$ are average measures of the ratio of constituted mortgages over output respectively in the previous year and during the pre-bust period. These are used as interaction terms in regressions in section 6.4

4 – A descriptive account of the housing boom and bust in Spain

I start the analysis investigating the cross sectional patterns of employment and house prices across provinces, as well as their relationship. A key question is whether Spanish provinces experienced heterogeneous patterns in house prices and employment. Relatedly, it is of interest to ask whether, during the crisis, employment declined more in provinces that experienced larger drops in house prices.

Figures 3 and 4 explore the cross sectional and geographical distributions of the declines in house prices and employment during the recession. It is clear that the crisis, while severe across the board, had considerably different outcomes across provinces. In addition, the figures make also apparent that provinces experiencing the largest drops in house prices were *not* systematically those that

experienced the largest declines in employment. As one might expect, and as will become clear, these correlations are not sufficient to draw conclusions on the effect of house prices on employment. Aside from the fact that provinces experiencing larger house prices declines might have also been characterized by economic developments that differed from the ones occurring in other provinces, it is also possible that historical trends in employment differed systematically across provinces along the distribution of the drops in house prices.

Figure 3: House prices declines across provinces
2008q1-2013q4, percentage points

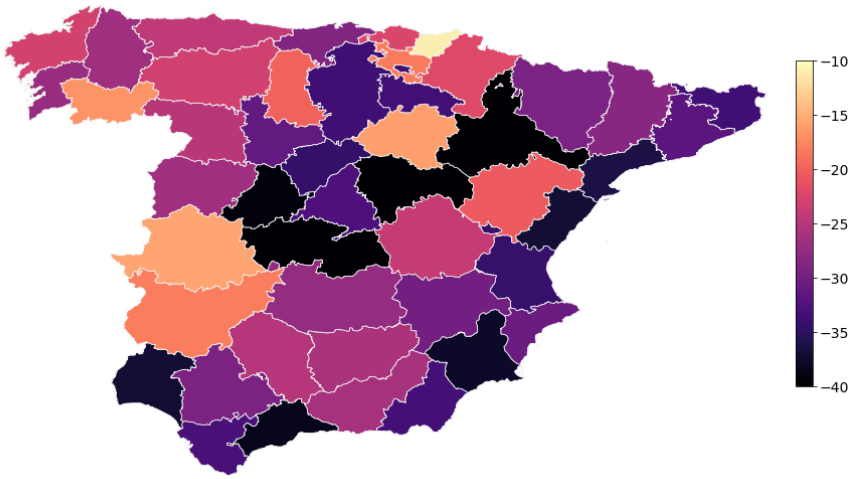

Figure 4: Employment declines across provinces
2008q1-2013q4, percentage points

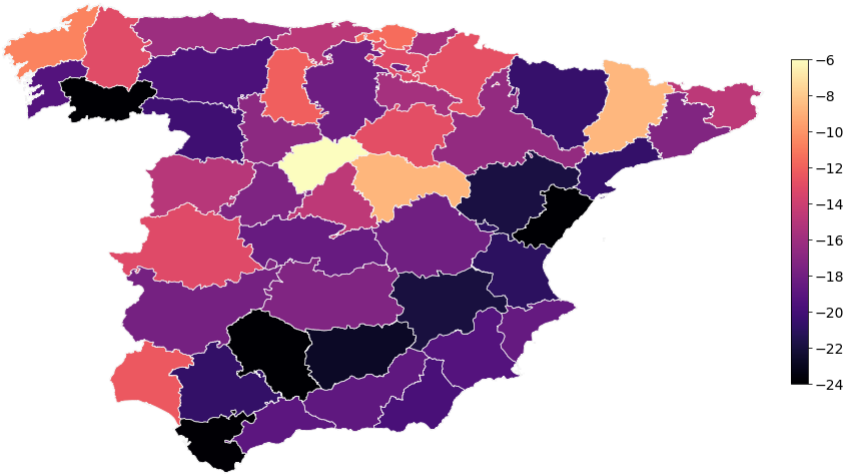

Figure 5 provides further evidence on the development of house prices and employment both in the cross-section and over time. It plots the development of house prices and employment in three different groups of provinces, experiencing respectively small, medium-sized and large house price declines between 2008q1 and 2013q4. First, it is noticeable how heterogeneous the housing cycle was across provinces. Localities experiencing a decline in house prices in the top quartile of the distribution

(in absolute value) witnessed a 40% drop in the value of dwellings, while in provinces in the bottom quartile of the distribution homes were worth 20% less after the recession compared to the 2008q1 peak. Such regional heterogeneity was also present in the boom phase of the cycle. Confirming evidence from the US experience (see, for instance, Gertler and Gilchrist, 2018), the severity of the bust was mirrored in the relative magnitude of the increases in house prices during the 2000-2008 boom. Spanish provinces experienced housing bubbles of different magnitudes.

Turning to the paths of employment, panel (5.b) presents what I regard as interesting evidence. At a first look, it appears perhaps surprising that the magnitude of the decline in employment in provinces experiencing large drops in house prices was not significantly different from the ones occurring in provinces characterized by medium-sized or even small decreases in the prices of dwellings. A more attentive look, however, reveals that provinces experiencing the largest house prices decline were on average those experiencing the highest employment growth in the period leading up to the crisis. The immediate consequence of this observation is that the *deviation from trend* in employment occurring during the crisis in provinces that experienced the largest house price decline was larger than the one witnessed in other provinces.

The evidence discussed here highlights, on the one hand, how exploiting heterogeneity in house prices movements across provinces in Spain can be a fruitful empirical strategy to investigate the consequences that such movements have on employment. On the other hand, it underlines how focusing on simple cross sectional heterogeneity, in a given period of time, in the development of house prices and employment can be misleading. Historical trends need to be taken into account.

While the patterns uncovered graphically in panel (b) of figure 5 allow to draw some tentative conclusion on the relationship between developments in house prices and movements in employment when taking into consideration heterogeneous historical trends across localities, they are nonetheless clearly insufficient to gauge the *causal* relationship between these two variables. I attempt to tackle this issue in the following sections.

Figure 5: Development of house prices and employment across provinces, with provinces grouped by severity of the house price decline between 2008q1 and 2013q4 (index: 2008q1=100)

Panel (5.a): House prices

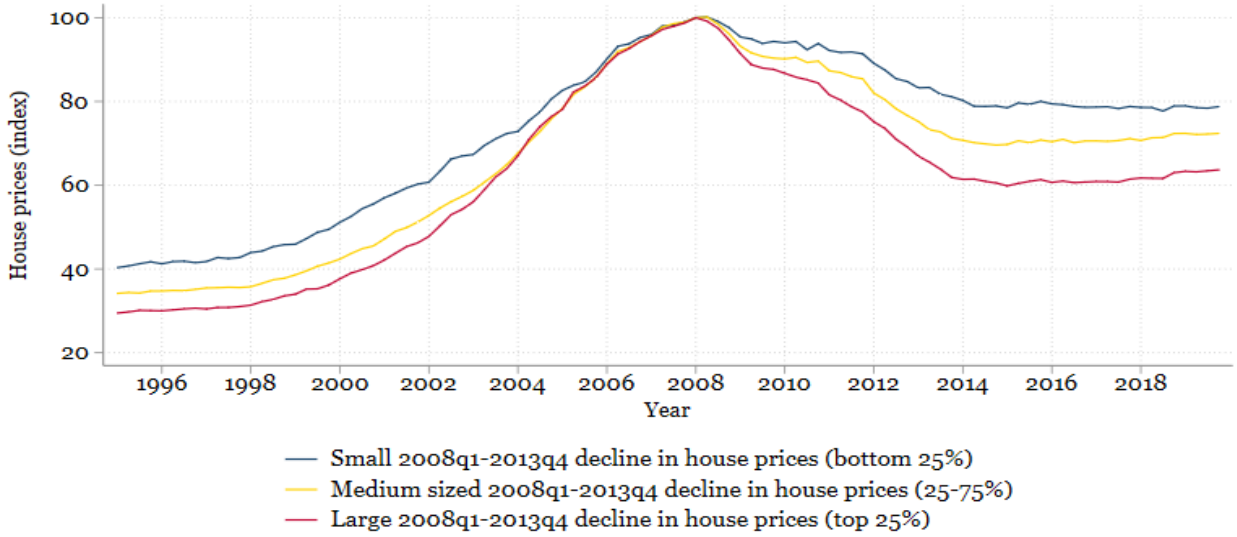

Panel (5.b): Employment

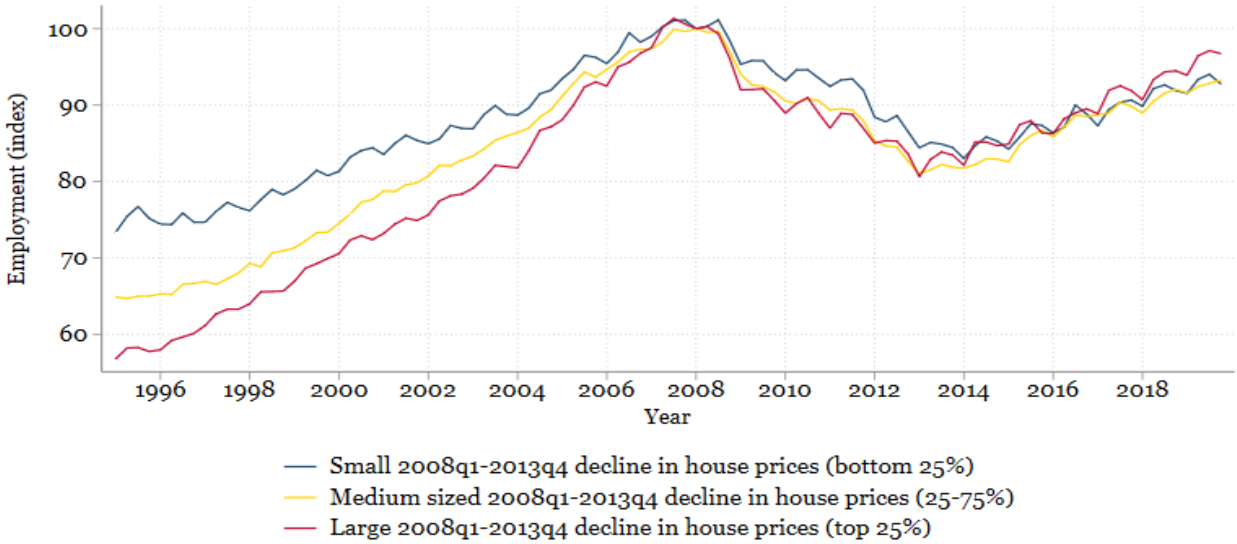

Note: Provinces are grouped by the severity of the house price decline between 2008q1 and 2013q4. The graphs plot the unweighted group averages of the house prices and employment indexes.

5 – Empirical methodology: retrieving the elasticity of employment with respect to house prices

5.1 – Empirical specification and threats to identification

In order to retrieve the elasticity of employment with respect to house prices, I use an empirical specification which exploits the panel dimension of the compiled dataset, allowing to run a regression of the log change in house prices in a given year in province l (for locality) on the log change in employment in the same period and in the same province, adding province fixed effects, time fixed effects and a vector of controls:

$$\Delta y_{l,t} = \phi_l + \rho_t + \beta \Delta p_{l,t} + \Gamma X_{l,t} + \varepsilon_{l,t} \tag{1}$$

where $\Delta y_{l,t}$ denotes the annual log-change in employment[3] in province l and at time t, $\Delta p_{l,t}$ denotes the contemporaneous annual log-change in real house prices in the same province, ϕ_l denotes a set of province fixed effects, ρ_t denotes a set of time fixed effects and $X_{l,t}$ denotes a set of controls.

The above discussion has presented evidence supporting the idea that taking into account historical trends in employment is relevant to retrieve the causal effect of interest. Panel (b) of figure 5, in particular, points at the fact that during the recession house prices dropped more in provinces that were characterized by higher historical increases in employment. The inclusion of province-level fixed effects allows to capture such heterogeneous trends in the dependent variable.

The inclusion of time fixed effects characterizes the main estimate of interest (β) as a cross-sectional estimate. The key variation exploited in the estimation of the parameter is the one across provinces in the same time period.

The main challenges for the identification of the effect of house prices on employment are the result of two facts: (i) house prices and economic activity are jointly determined, leading to the problem of reverse causality; and (ii) changes in house prices and employment can be driven by the same underlying economic shock, leading to omitted variable bias. It is possible to present this idea of the joint determination of employment and house prices growth with the following empirical model (comprising two structural equations)[4]:

$$\Delta y_{l,t} = \phi_l + \rho_t + \beta \Delta p_{l,t} + \alpha_l \Xi_t + \varepsilon_{l,t} \tag{2}$$

$$\Delta p_{l,t} = \psi_l + \theta_t + \delta \Delta y_{l,t} + \gamma_l \Psi_t + v_{l,t} \tag{3}$$

$\Delta y_{l,t}$ is the change in the log of employment in province l over year t, while $\Delta p_{l,t}$ is the change in the log of real house prices in the same province and year. Both equations include province fixed effects (ϕ_l and ψ_l, capturing heterogeneous and persistent growth rates) and time fixed effects (ρ_t and θ_t). In addition, Ξ_t and $\varepsilon_{l,t}$ constitute the error term in equation (2) and capture respectively national and idiosyncratic shocks that affect local employment. Similarly, Ψ_t and $\nu_{l,t}$ constitute the error term in

[3] In the analysis presented below, I use different measures of employment (total, non-construction, non-tradable and tradable as well as non per-capita and per-capita). To facilitate the exposition, I refer here generally to *employment*.

[4] The empirical model and the related explanation of the identification strategy follow, in part, Guren et al. (forthcoming, appendix C.1). I add here, in particular, a detailed discussion of what the potential sources of heterogeneous sensitivity of local house prices to national housing cycles are, together with their implications for identification.

equation (3) and capture respectively national and idiosyncratic shocks that affect house prices. These shocks are viewed as vectors of primitive shocks.

With regards to the issue of reverse causality, it is worth restating, following the theoretical discussion presented above, that house prices affect local employment through their effects on local households' consumption as well as through their impact on credit availability for firms pledging housing as collateral. On the other hand, a drop in local employment reduces the incomes of households (either because members of the household become unemployed or because wage growth declines), thus reducing demand for housing and house prices. Further, an increase in unemployment and economic uncertainty is likely to depress demand for housing because housing itself, being an illiquid asset, is an unattractive form of precautionary saving.

Turning to the presence of common shocks moving house prices and employment simultaneously, the literature has highlighted the relevance of various sources of demand-driven movements in house prices, such as shifts in the availability of amenities and in demographics, changes in borrowing costs and credit availability, movements in households' income driven by movements in productivity, and shifts in migration patterns (Saiz, 2010; Landvoigt, Piazzesi, Schneider, 2015). At the same time, Saiz (2010) identifies changes in input prices as the key factor for supply-driven movements in house prices. It is worth noticing here that studies focusing on understanding the causes of the US housing boom and bust in the early 2000s, such as Kaplan, Mitman and Violante (forthcoming) and Favilukis, Ludvigson and Van Nieuwerburgh (2016), have stressed that *demand-factors* are key drivers of movements in house prices. Importantly, many of these economic shocks, notably those to credit availability, borrowing costs, incomes, and migration, affect employment directly.

Given the issue of reverse causality and the likely omitted variable bias, a simple estimation of equation (2) to retrieve the elasticity of employment with respect to house prices would lead to a biased estimate of the parameter of interest, $\widehat{\beta}$.

5.2 – The sensitivity instrument

In order to overcome these challenges, I use an instrumental variable approach. In particular, I use a *sensitivity instrument* (proposed by Guren et al., forthcoming). Such an instrument exploits variation in the sensitivity of local house prices to cyclical movements in house prices at a more aggregate geographical level (in the present study: at the national level). The intuition behind the meaningfulness of the first-stage of such instrument is as follows: when house prices at the national level increase, house prices in some provinces tend to respond systematically more relative to those in other provinces. As will become clear, the key challenge is that of identifying the sensitivity of local house prices to national cycles in house prices which is independent from the sensitivity of local employment with respect to aggregate economic shocks. After having estimated the sensitivity of house prices to national cycles for a given province, it is possible to use the interaction of this sensitivity parameter (call it $\widehat{\gamma}_\iota$) together with the change in house prices at the national level (ΔP_t) as an instrument (z_t) for house prices growth at the local level ($\Delta p_{\iota,t}$). Heterogeneity in sensitivity to national cycles across localities will result in variation in the instrument predicting variation in the causal variable of interest.

In addition to the assumption of a *meaningful first stage*, a valid instrument should be as good as randomly assigned (*independence assumption*) and should affect the dependent variable of interest only through its effect on the causal variable (*exclusion restriction*). These two additional assumptions are summarized by the condition that the instrument is independent of the error term (from equation

(2): $Cov(z_t, \alpha_l \Xi_t + \varepsilon_{l,t}) = 0$). It is this condition that makes it necessary to be able to identify the sensitivity of local house prices to national cycles in house prices (driving variation in $z_t$in the cross-section) which is independent from the sensitivity of local employment with respect to aggregate economic shocks (α_l). Would this condition not be satisfied, the instrument would be invalid.

Taking a step back, it is worth discussing what the sources of heterogeneous sensitivity of local house prices to national housing cycles are likely to be. Disentangling these different sources is relevant in order to obtain a valid instrument. It is possible to think about multiple drivers of such heterogeneity: housing supply elasticity, housing demand elasticity or systematically heterogeneous propagation of national level shocks. Figure 6 presents stylized representations of these three cases. In the first case (6.a), it is shown how the same shock to housing demand leads to a heterogeneous movement in house prices in two localities with different housing supply schedules but with the same housing demand curve. The response in house prices is larger where housing supply is less elastic (supply curve in red). Measures of heterogeneous housing supply across localities have been used explicitly as instruments for house prices by the literature surveyed above. These studies tend to make use of the housing supply elasticity measures obtained for US Metropolitan Statistical Areas by Saiz (2010), who shows that measures of regulatory and geographical constraints to housing supply are a relevant predictor of the response of prices to shocks in housing demand.

Panel (6.b) presents instead the case where a shock to housing supply leads to a larger increase in prices in a locality where housing demand is relatively inelastic. I think about such a heterogeneity in housing demand as driven primarily by the composition of potential homebuyers and movers. The demographic and economic characteristics of the population in a given locality, in these regards, are likely to play a key role.

Lastly, panel (6.c) illustrates how the fact that some shocks at the aggregate level might propagate differently (in a systematic fashion) across localities can lead to a heterogeneous sensitivity of local house prices to national housing cycles. For instance, some provinces might be exposed to systematically more cyclical sectors. Variations in wages in these sectors will lead to national responses in house prices being more prominent in these particular localities as a result of increased housing demand. Similarly, differences in the banking sectors serving localities (or in the degree of competition among banks) could lead some provinces to be systematically more sensitive to credit shocks.

Figure 6: Potential drivers of heterogeneous sensitivity of local house prices to national housing cycles

(6.a) Heterogeneous housing supply elasticity

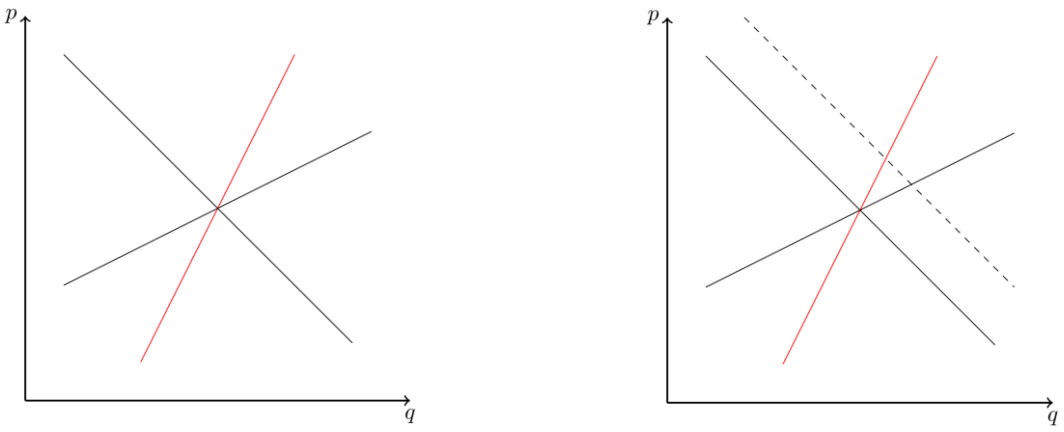

(6.b) Heterogeneous housing demand elasticity

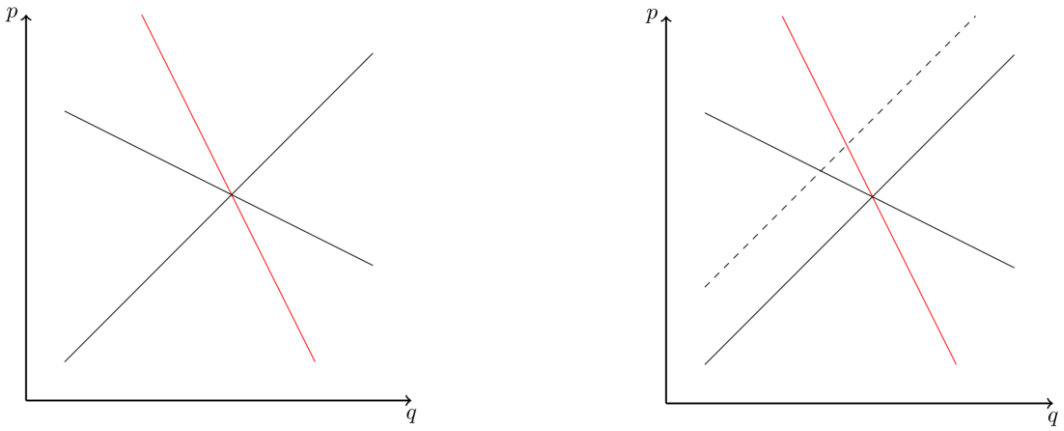

(6.c) Systematically heterogeneous propagation of national level shocks

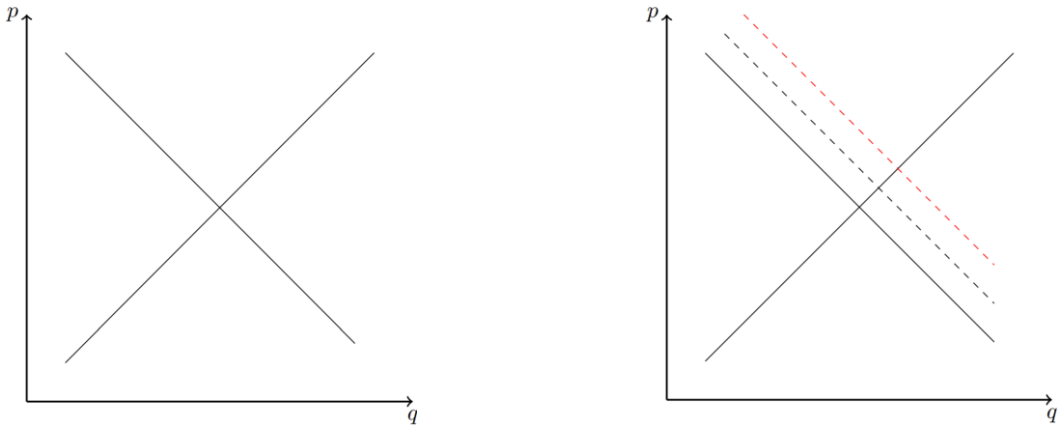

Note: Upward and downward sloping curves sketch respectively housing supply and housing demand schedules. Figures on the left in each panel represent equilibria in the housing markets of two localities in three settings: in panel (a) the two localities are assumed to have the same housing demand schedule, while one of the two localities is characterized by less elastic housing supply, in red; in panel (b) the two localities are instead assumed to have the same housing supply schedule, with heterogeneity in their housing demand (the relatively more inelastic one is again in red); lastly, panel (c) presents the case of two localities with the same housing demand and supply. The figures on the right hand side sketch the developments occurring in the three cases in the housing markets in the two localities respectively in the case of: a positive shock to housing demand that impacts the two localities in exactly the same manner (panel a); a negative supply shock that impacts the two localities in the same manner (panel b); and the case of a national positive demand shock that impacts one locality (in red) relatively more than another. The figures show three possible key determinants of a systemically higher sensitivity of local house prices to national housing cycles. Prices respond more: in localities with a less elastic housing supply following a demand shock; in localities with a less elastic housing demand following a supply shock; in localities where aggregate shocks result to be systematically more pronounced.

The *sensitivity instrument* can be constructed by estimating the *sensitivity parameter* using historical data on house prices at the national and provincial level. The first possibility for retrieving the sensitivity of house prices in different provinces to national level house price movements consists in running the regression:

$$\Delta p_{l,t} = \phi_l + \gamma_l \Delta P_t + v_{l,t} \tag{4}$$

Where ΔP_t denotes the period change in (log) real house prices at the national level and γ_l is a locality-specific coefficient. The issue with this simple strategy is that the estimate of γ_l is biased if (i) the sensitivity of employment at the local level (α_l, in equation (2)) is heterogeneous across localities; and

(ii) if the determinants of the estimated *sensitivity parameter* are correlated with provinces' heterogeneous sensitivity in employment.

To see the justification for point (i), assume that the sensitivity of localities to national house prices cycles (γ_l) is actually homogeneous across provinces. Further assuming that the sensitivity of employment (α_l) to national shocks (Ξ_t) is heterogeneous across localities[5], national shocks producing variation in employment growth across provinces and correlated with national house prices would nonetheless lead to the estimation of heterogeneous local sensitivities (γ_l) through $\delta\Delta y_{l,t}$ in equation (3). This outcome would result from the omission of changes in employment as control in equation (4)

Addressing instead point (ii), it is worth reflecting on the possibility that the characteristics making a particular housing market more sensitive to national housing cycles are correlated with the sensitivity of local employment to national shocks. As mentioned, this would result in correlation between γ_l and α_l, making the instrument invalid.

When it comes to housing supply elasticity (the first driver of heterogeneity in the sensitivity local house prices identified above), I believe it appears credible that geographical and regulatory constraints to housing supply are not correlated with the sensitivity of local employment to aggregate economic shocks. One potential issue, highlighted notably by Davidoff (2016) in the US context, is that localities with relatively inelastic housing supply differ systematically from other localities as they are more likely to be affluent places situated in coastal areas (and hence with less land available for development). Such localities will tend to be characterized by a long-term positive trend in both house prices and employment, potentially leading to bias in the estimates of interest. Although this concern is valid also in the case of Spain, in the context of the present analysis it is mitigated by the inclusion of locality fixed effects, which capture persistent trends both in employment and in house prices, as pointed out by Guren et al. (forthcoming).

Turning to heterogeneity in housing demand (panel 6.b), the possibility that local sensitivity of house prices to national shocks is correlated with that of local employment becomes relatively more credible. As shocks to housing supply impact differently localities with different populations of potential homebuyers, it is also possible that simultaneous labor market shocks (again perhaps driven by the same shock to the supply side of the economy - think for instance about a shift in prices of raw materials) will affect differently the labor market outcomes across localities, depending on the demographic and economic characteristics of the population. While I believe that this concern should be taken seriously (I account for it taking the appropriate steps in the estimation of the *sensitivity parameter,* as explained below), I would also note again that the literature points at shifts in demand for housing as primary drivers of house prices.

Lastly, I regard as the most poignant concern with respect to identification the possibility that the sensitivity of local house prices as determined by heterogeneous propagation of aggregate shocks across localities is correlated with the sensitivity of local employment to the same national shocks. As discussed above, key determinants of house prices, such as shifts in borrowing costs and credit conditions, in households' incomes and in migrations, are also likely drivers of employment. It is further possible that such primitive shocks propagate differently across localities.

These observations imply that it is possible to improve on the simple estimation strategy described by equation (4) addressing the concerns discussed here by revising the estimation of the *sensitivity parameter*, including controls in the regression of changes in house prices at the national level on movements in local house prices.

[5] Such a heterogeneity could be driven for instance by exposure to more cyclical sectors, systematically different fiscal stabilization policies conducted by local governments or different shares of workers lacking social insurance.

The instrument can be further refined by estimating γ_l controlling for local and regional changes in employment (included with locality-specific coefficients in the regression). These controls are meant to address the overall concern that the sensitivity of local house prices to national housing cycles is either driven or correlated with the sensitivity of local employment to aggregate shocks. In practice, they ensure that the variation in local house prices used for the estimation of the *sensitivity parameter* is orthogonal to local movements in employment and to the heterogeneous effect of national movements in employment across provinces. In general, the main objective of adding these controls is precisely that of allowing to capture variation in the sensitivity of local house prices to cycles in house prices at the national level which is independent from variation in the sensitivity of local employment to aggregate economic shocks.

Furthermore, specific controls can be added to address concerns relative to the potential sources of correlation between the sensitivity of house prices and the one of employment to aggregate shocks. I follow Guren et al. (forthcoming) in constructing a control meant to capture heterogeneity in local sensitivity to changes in borrowing costs. In particular, I estimate the sensitivity of local employment to changes in mortgage interest rates at the national level. I then include in the regression for the estimation of the *sensitivity parameter* the interaction of this estimated province-level sensitivity of employment to movements in mortgage interest rates with the change in mortgage interest rate itself. This control (the mortgage interest rate prediction control) is meant to ensure that correlation between the sensitivity of local house prices and the one of local employment to changes in mortgage interest rates does not drive the estimates of the *sensitivity parameter*.

Furthermore, I control for the share of young adults (15 to 39 years old) in the province population, interacting this variable with time dummies. The inclusion of the share of young adults as control captures, on the one hand, the fact that both housing and labor markets can tend to be more cyclical where young adults represent a larger share of the population. Young adults are likely to be those transitioning into home ownership, thus leading to a more "active" housing market. At the same time, younger workers are also more likely to be employed in fixed term contracts or to be transitioning into employment, thus increasing the sensitivity of the labor market to economic shocks. On the other hand, the inclusion of this variable interacted with time dummies ensures that heterogeneity in the sensitivity of local house prices to national cycles is not driven by heterogeneity in *housing demand elasticities* across localities (driven by different demographics). Intuitively, this control is meant to capture the heterogeneous effect of the shock to housing supply as described by figure (6.b), if one thinks about the heterogeneity in demand elasticities as driven by demographics.

The above discussion implies that adding additional controls in the regression for the estimation of the *sensitivity parameter* could further address concerns of bias in the obtained estimates and aid identification. For instance, one could generate prediction controls similar to the one constructed for the mortgage interest rate also for other national level variables capturing other shifts in credit supply (for instance shifts in commonly applied constraints at loan origination or measures of the health of credit markets). Furthermore, one could proxy for housing demand elasticity not only through the share of young adults, but also exploiting other dimensions of heterogeneity across localities' populations, for instance income. I do not include these variables simply because of data availability. For example, I find some of the data not to be available for a sufficiently long time period: a measure of the health of credit markets constructed by Gilchrist and Mojon (2016) is available only for the period 1999-2013; while data on incomes at the disaggregated geographical level has started being distributed by the Spanish National Statistical Institute only recently.

It is now possible to state explicitly the conditions under which the use of the described sensitivity instrument allows the identification of the elasticity of employment with respect to house

prices. The main identifying assumption is that, conditional on controls, there is no unobserved factor that is both correlated with house prices in the time series and that differentially affects employment in the same localities where house prices are estimated to be historically more sensitive to national housing cycles in the cross section (Guren et al., forthcoming, p.3).

As a last step in the identification strategy, throughout the analysis I present estimates of the elasticity of employment with respect to house prices both without additional controls and controlling for employment shares in the six sectors described in table 1.b (first row) interacted with time fixed effects. This allows to address concerns that variation in the constructed *sensitivity instrument* is driven by provinces' exposure to more cyclical sectors.

5.3 – Estimation of the sensitivity parameter(s) and the first stage

In practice, I obtain two sets of estimates of the sensitivity parameters $\widehat{\gamma_l}$. I obtain both estimates through the following empirical model:

$$\Delta p_{l,t} = \psi_l + \gamma_l \Delta P_t + \delta_l \Delta y_{l,t} + \rho_l \Delta Y_t + \Gamma X_{l,t} + v_{l,t} \tag{5}$$

The two sets of estimates differ in the sets of controls included. In the first case, with which I construct the instrument Z_1, I include only the growth rate in local and national employment. For the second set of sensitivity parameters, used to construct instrument Z_2, I also include the share of young adults (15 to 39 years old) in the province population interacted with time fixed effects and the mortgage interest rate prediction control, described above.

I use as employment variable on the right hand side the measure of total employment at the local and national level, available over the whole period 1995-2019. As the analysis presented below investigates the effect of house prices on different components and transformations of employment (either employment in specific sectors or non-per-capita and per-capita), it is worth mentioning that the results presented here with respect to the employment variables available over the whole sample are robust to the construction of various sets of instruments, each obtained by including the same employment variable on the right hand side in the estimation of the sensitivity parameter and on the left hand side in the estimation of the elasticity of employment with respect to house prices.

Figure 7 presents graphically the time series of the two instruments across the 47 provinces. Aside from clarifying the working of the empirical strategy itself, the figure is also able to present the degree of cross sectional variation in the obtained instrument. Table 3 further provides an overview of the distributions of the two sets of *sensitivity parameters*. The average *sensitivity parameter* remains consistently between 0.76 and 0.8, with minimum values of 0.1 and 0.19 and maximum values of 1.43 and 1.46. Such parameters support the idea that local house prices respond more than proportionally to national cycles in some provinces, while the response tends to be muted in other localities. In addition, it is worth noticing that the correlation between the two sets of *sensitivity parameters* γ_1 and γ_2 is 0.9884, which implies that the two estimation methods used yield similar results.

Figure 7: Instruments

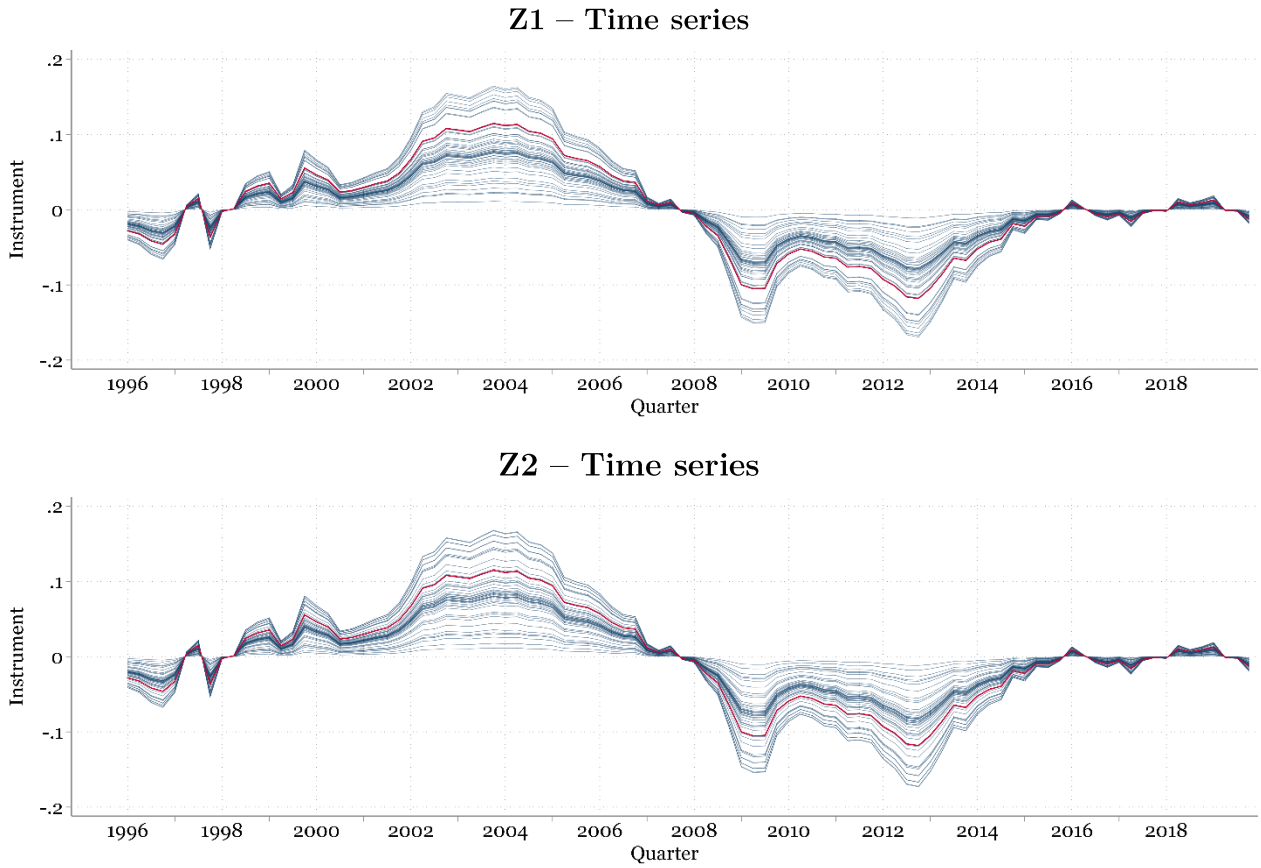

Note: The graph plots the instrument relative to total employment. The *sensitivity parameter* is obtained estimating equation (6) with employment shares in six economic sectors, share of secondary homes and share of young adults in the population as variables in the vector of controls $X_{l,t}$. These variables are interacted with time dummies. The log-change in real house prices at the national level is displayed in red. The Stata program used for this graph, *group_lines*, is courtesy of Austin L. Wright and can be found here.

Table 3: Summary Statistics of the Sensitivity Parameters

Sensitivity parameter	Mean	Std. Dev.	Min	Max
γ_1	0.76	0.34	0.10	1.43
γ_2	0.80	0.36	0.19	1.46

Note: the table presents summary statistics on the cross-sectional distribution of sensitivity parameters across 47 Spanish provinces. It is worth re-stating that sensitivity parameters vary only in the cross-section and not over time. The subscript *1* refers to the sensitivity parameter obtained when controlling only for while subscript *2* implies that also the share of young adults and the mortgage interest rate prediction control are included in the relevant regression.

Table 4 and figure 8 speak to the assumption of a *meaningful first stage*, presenting the first stage regression outputs and depicting the first stage relationships for the full quarterly 1995-2019 sample. In the analysis below, I tend to present two sets of estimates for each combination of dependent variable and instrument (Z_1 and Z_2), one without controls and one controlling for employment shares in the different economic sectors interacted with time dummies. Table 4 and figure 8 present therefore the

four first stages corresponding to each instrument-controls pair. It is possible to observe that the first stage is strong, limiting concerns that bias in the estimates can be driven by a weak instrument. The coefficients on the sensitivity instruments are positive (as expected), comprised between 0.625 and 0.656, and consistently statistically significant at the 0.1% significance level, while the F-statistics range between 89.78 and 111.04, values that result to be reassuringly well above the value of 10 commonly used in the literature as threshold for identifying strong enough instruments (Stock, Yogo and Wright, 2002).

Table 4: First stage regressions

Dependent variable: $\Delta p_{l,t}$	(1)	(2)	(3)	(4)
Sensitivity instrument	0.656***	0.625***	0.629***	0.601***
P-values	0.0	0.0	0.0	0.0
Standard errors	0.064	0.065	0.059	0.063
Instrument used	Z_1	Z_1	Z_2	Z_2
Controls?	No	Yes	No	Yes
F-statistics	**103.14**	**91.95**	**111.04**	**89.78**

Note: * indicates statistical significance at the 5% level, ** at the 1% level, and *** at the 0.1% level.

Figure 8: First Stages

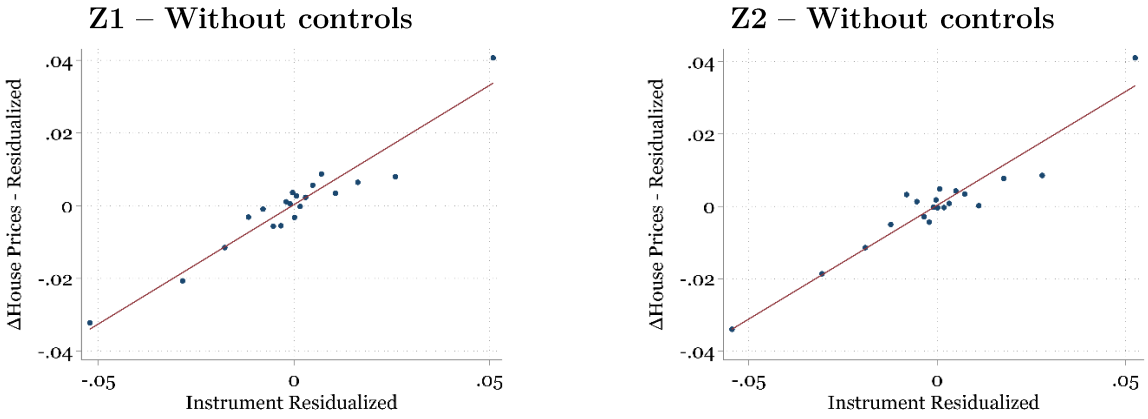

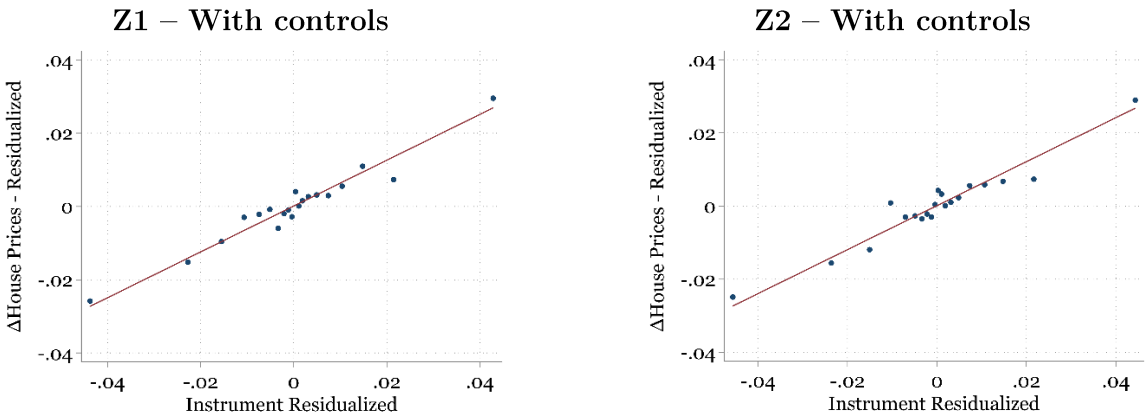

Note: The graphs depict binned scatter plots of the first stages. Both the instruments and the causal variables of interest have been residualized against province and time fixed effects and, in the cases where controls are included, against the employment shares interacted with time dummies.

6 – Estimates of the elasticity of employment with respect to house prices

In this section, I present the results of the empirical analysis. In a first step, I focus on the elasticity of total employment with respect to house prices (section 6.1). Having obtained a baseline result in this regard, I attempt to disentangle the relative importance of the *household demand channel* and of the *collateral lending channel for firms* by estimating the elasticity of employment respectively in the 'non-tradable' and in the manufacturing sectors (section 6.2).

In section 6.3 interpret the obtained estimates in light of the possible presence of labor market frictions and nominal rigidities. The estimation of the effect of movements in house prices on measures of employment *per capita* allows to characterize the extent to which the overall response of employment to housing values is driven by labor mobility across regions in response to geographically heterogeneous economic shocks.

Lastly, section 6.4 I investigate whether the interaction of high levels of debt with sharp declines in house prices implies a larger reaction of households' consumption to movements in the value of housing.

6.1 The elasticity of total employment with respect to house prices

I estimate the elasticity of employment with respect to house prices using the obtained *sensitivity instruments.* In particular, the estimate of interest is obtained through a pooled estimate of the following empirical model, already presented in equation (1):

$$\Delta y_{l,t} = \phi_l + \rho_t + \beta \Delta p_{l,t} + \Gamma X_{l,t} + \varepsilon_{l,t}$$

First, I present estimates of the elasticity of total employment with respect to house prices. The second row of table 5 reports four estimates of this elasticity: they range from 0.194 to 0.266. This implies that *a 10% decline in house prices in a province relative to another leads to a 1.94-2.66% larger decline in total employment.*

Table 5: Pooled estimates of the elasticity of total employment with respect to house prices

$\Delta y_{l,t}$: **Total employment**	(1)	(2)	(3)	(4)
$\hat{\beta}$: Δ real house prices	0.265***	0.196***	0.266***	0.194***
Standard errors	0.052	0.058	0.053	0.059
P-values	0.0	0.001	0.0	0.001
Instrument	Z_1	Z_1	Z_2	Z_2
Controlling for economic sector employment shares interacted with time dummies?	No	Yes	No	Yes
N	4290	4290	4290	4290

Note: standard errors are two-way clustered at the province and time level throughout the analysis. This allows for correlation in the time series for a given province as well as for correlation across provinces in a given time period. This choice reflects the nature of the instrument, which has a time series component which is common across provinces and a cross-sectional component which is fixed over time at the province level. * indicates statistical significance at the 5% level, ** at the 1% level, and *** at the 0.1% level.

No additional controls are used in odd-numbered columns, while estimates in even-numbered columns are retrieved controlling for employment shares in the six economic sectors defined through the *detailed* classification outlined in section 3, interacted with time dummies. As mentioned, these controls are meant to dissipate remaining concerns of bias in the obtained estimates driven by the fact that time and sector specific shocks to economic activity lead to changes both in employment and house prices. Consistently with this intuition, the estimates presented in even-numbered columns result to be lower than the ones in their odd counterparts. However, by considering the point estimates together with their respective standard errors it is also possible to conclude that excluding or including the controls does not lead to estimates that are significantly different from each other.

The first two columns present results obtained with instrument Z_1, while the results obtained in columns (3) and (4) are obtained using instrument Z_2. The results are arguably robust to the change in the way in which the instrument is obtained. Taking this consideration into account together with the observations in the previous paragraph, going forward I use the estimates obtained using instrument Z_2 and controlling for sectoral employment shares (in the case of total employment: 0.194, column 4) when presenting calculations aiming at quantifying the effects of the decline in house prices during the recession on employment as well as when comparing the results obtained here to the ones presented by previous studies.

It is perhaps necessary to restate here that the presented estimates (in table 5 as well as in the remainder of the analysis) identify the effect of movements in house prices on employment to the extent that the main identifying assumption is credible. Some factors not controlled for either in the estimation of the *sensitivity parameters* or in the empirical specifications of interest could drive a correlation between the *sensitivity instruments* and the error term in equation (1), thus making the instrument

invalid and biasing the obtained estimates. I argue nonetheless that the steps taken in the construction of the *sensitivity instrument* and the inclusion of employment shares as controls are sufficient to alleviate these concerns.

Furthermore, it is worth noticing that the magnitude these estimates appears to be substantial. For instance, considering the 2008q1-2013q4 recession and using the estimate of 0.194, it is possible to notice how moving from the 90th percentile of the distribution of real house prices decline across provinces (-20.62%) to the 10th percentile (-39.17%) leads to an additional loss in non-construction employment of 3.6%. This should be then considered in light of the fact that the (non-weighted) average loss in employment across provinces during the same period was 17.07% and that the standard deviation of the employment decline was 4.63%.

6.1 - The response of households and firms to movements in house prices

It is worth asking what drives these results. As highlighted in section 1, theoretical analyses highlight the fact that movements in house prices not only affect households' consumption, but also the capacity of firms to borrow by pledging housing as collateral. In addition, it should be noted that changes in house prices directly impact the profitability of businesses operating in the construction sector, therefore affecting employment in this industry.

The mechanisms at play within the construction sector, although of interest in their own right, are not the central interest of the present analysis. Reflecting this fact, it is possible to focus attention on the elasticity of non-construction employment with respect to house prices. Estimates relative to this elasticity are presented in table 6. In this instance, explicitly controlling for the employment share in the construction sector in each period, as done in columns 2 and 4, allows to mitigate concerns that the estimated elasticities are biased upward by the general equilibrium effects of a shock hitting the construction sector. I note that while statistically non-significant, the estimates in columns 2 and 4 result to be quite precisely estimated and close to the ones in columns 1 and 3, which in turn are significant at the 1 percent level. As expected, the point estimates of the elasticity of non-construction employment with respect to house prices result to be lower than those presented in table 5. In particular, I obtain estimates ranging from 0.098 to 0.138, implying that *a 10% larger decline in house prices in a province relative to another leads to a 0.098% to 1.38% larger decline in non-construction employment.*

Table 6: Pooled estimates of the elasticity of non-construction employment with respect to house prices

$\Delta y_{l,t}$: **Non-construction employment**	(1)	(2)	(3)	(4)
$\hat{\beta}$: Δ real house prices	0.138**	0.102	0.136**	0.098
Standard error	0.048	0.06	0.046	0.059
P-value	0.004	0.089	0.003	0.094
Instrument	Z_1	Z_1	Z_2	Z_2
Controlling for economic sector employment shares interacted with time dummies?	No	Yes	No	Yes
N	4290	4290	4290	4290

Note: standard errors are two-way clustered at the province and time level. * indicates statistical significance at the 5% level, ** at the 1% level, and *** at the 0.1% level.

I conclude that *the economic relevance of the estimates presented here adds to the evidence pointing at how movements in house prices significantly impact economic activity, in particular driving changes in the economic decision-making of households and firms.*

This evidence, however, does not provide any additional insight on which particular mechanism drives the relationship of interest. Both the theoretical discussion and the review of the empirical literature presented above have highlighted the fact that there appears not to be a clear consensus on whether the relationship between house prices and economic activity is driven by the responses of households or by those of firms. In particular, Mian and Sufi (2014) as well as Guren et al. (forthcoming) present evidence supporting the idea that the *household demand channel* drives the relationship of interest, while Adelino, Schoar and Severino (2015) argue in favor of the idea that the *collateral lending channel for firms* also plays a relevant role.

I follow these authors in investigating the relative importance of these two channels by focusing attention on measures of employment in specific sectors of the economy. As mentioned above, while changes in local demand differently impact tradable and non-tradable sectors, shifts in credit availability for firms driven by housing prices should not affect these same sectors differently. Given the available level of disaggregation in the employment variables, I am able to proxy for non-tradable employment with employment in 'wholesale and retail trade, repair of motor vehicles, transport and storage, hospitality industry and information and communication' (I refer to this sector as 'non-tradable'), while proxying for tradable employment with the measure of employment in the manufacturing sector. I estimate the elasticity of employment in these two sectors separately in order to gauge to what extent economic activity in these two sectors responds differently to movements in house prices.

Table 7: Pooled estimates of the elasticity of employment in 'non-tradable' and tradable industries with respect to house prices

Panel (7.a): Employment in 'non-tradable' industries

$\Delta y_{l,t}$: **Employment in 'non-tradable' industries**	(1)	(2)	(3)	(4)
$\hat{\beta}$: Δ real house prices	0.165***	0.194**	0.171***	0.202**
Standard error	0.042	0.066	0.037	0.077
P-value	0.0	0.003	0.0	0.009
Instrument	Z_1	Z_1	Z_2	Z_2
Controlling for economic sector employment shares interacted with time dummies?	No	Yes	No	Yes
N	762	762	762	762

Panel (7.b): Employment in the manufacturing industry

$\Delta y_{l,t}$: **Employment in manufacturing**	(1)	(2)	(3)	(4)
$\hat{\beta}$: Δ real house prices	0.111	0.1	0.11	0.111
Standard error	0.071	0.079	0.07	0.074
P-value	0.119	0.208	0.12	0.13
Instrument	Z_1	Z_1	Z_2	Z_2
Controlling for economic sector employment shares interacted with time dummies?	No	Yes	No	Yes
N	762	762	762	762

Note: standard errors are two-way clustered at the province and time level. * indicates statistical significance at the 5% level, ** at the 1% level, and *** at the 0.1% level.

Table 7 presents the obtained estimates for the elasticities of employment in the two sectors with respect to house prices. First, it is noted that the fact that the sample size shrinks drastically is the result of the fact that measures of employment at the disaggregated sectoral level are available only yearly from 2000 to 2017. Second, I observe that the point estimates of the two set of elasticities presented in tables 7.a and 7.b are respectively higher than and somewhat similar to/lower than the ones obtained for non-construction employment (from 0.098 to 0.138, table 6): while the estimates of the elasticity of 'non-tradable' employment with respect to house prices range between 0.165 and 0.202, the ones relative to employment in the manufacturing sector are consistently between 0.1 and 0.111. Despite these differences, however, I consider as particularly reassuring the fact that there is overlap between the confidence intervals of the estimates obtained in tables 6 and 7. The data on non-construction employment and employment by disaggregated economic sectors is obtained from two

different measurements by the Spanish National Statistical Institute (INE), respectively the Economically Active Population Survey and the Spanish Regional Accounts.

Lastly, and importantly, it is of interest to note that the estimates of the elasticity of employment in the manufacturing sector result to be lower than the ones of employment in 'non-tradable' sectors. While the estimates relative to employment in 'non-tradable' sectors result to be significant either at the 1% or at the 5% level, those relative to employment in the manufacturing sector are never statistically significant.

To offer an additional quantitative interpretation of the estimates relative to 'non-tradable' employment, it is possible to notice that: (i) *a 10% larger decline in a province relative to another leads to a 1.65% to 2.02% larger decline in 'non-tradable' employment*; while (ii) moving from the 90th percentile of the distribution of real house prices decline across provinces during the 2008-2013 recession (-20.62%) to the 10th percentile (-39.17%) leads to an additional loss in non-tradable employment of 3.73%, in the face of an average decline of 13.95% (and a standard deviation of 4.4%).

I conclude that *the results presented here provide some evidence of the fact that the response of households' consumption to movements in house prices plays a significant role in the relationship between house prices and employment. The household demand channel appears to be relatively more important than the collateral lending channel.*

This conclusion should be stated together with some additional observations. First, the quantitative characterization of this result, as implied by the estimates presented in table 7, is accurate to the extent that the measures of employment in 'non-tradable' industries and manufacturing are good proxies for employment respectively in non-tradable and tradable sectors. This classification follows the literature, and the evidence presented by Mian and Sufi (2014). However, misclassification of sectors in either two of the tradable/non-tradable categories would likely result in a larger difference between the actual elasticities of tradable and non-tradable employment. Second, the capacity of a comparison of the elasticity of employment in tradable and non-tradable sectors to house prices to disentangle the relative importance of the *household demand channel* and of the *collateral lending channel for firms* hinges on the assumption that firms in the two sectors do not systematically differ on their reliance on collateral to obtain credit (because, for instance, of differences in size). If firms in non-tradable sectors rely relatively more on collateral availability than those in tradable sectors, then the difference in the elasticities of employment in the two set of sectors could be driven precisely by the working of the *collateral lending channel for firms.* Unfortunately, data to investigate this possibility appear not to be available. Lastly, I believe one should be cautious in interpreting the effect of movements in house prices on 'non-tradable' employment as driven *solely* by the *household demand channel,* not only the fact that the estimates of the elasticity of manufacturing employment are positive, quite precisely estimated, and of an economically significant magnitude indicates that the *collateral lending channel for firms* could well be operative, but other mechanisms discussed above, such as the effect of movements in house prices on local governments' spending and on local banks' credit supply, could also drive, in part, the relationship.

6.3 – Adjustment mechanisms: Labor market frictions and nominal rigidities

The evidence presented above on the elasticities of total, non-construction, tradable and non-tradable employment to house prices is not compatible with a model of the economy that includes smooth wage adjustments and free movement of workers across tradable and non-tradable sectors (section 1.2). The obtained estimates point at the fact that Spanish provinces facing larger movements in house prices also witness larger changes in employment. In the event of a drop in house prices, the larger fall in employment in the non-tradable sector is not compensated by an increase in the tradable sectors. This would lead to estimates of the response of aggregate employment to house prices (tables 5 and 6) being statistically not different from zero and to negative estimates of the elasticity of manufacturing employment to house prices (table 7.b).

An additional potential adjustment mechanism to geographically heterogeneous economic developments in house prices is labor mobility (section 1.2). By estimating the elasticity of *per-capita* employment with respect to house prices, it is possible to investigate the extent to which individuals respond to heterogeneous movements in house prices across regions by migrating. Intuitively, if the response of non-per-capita employment to house prices is driven entirely by people moving towards areas that witness relatively better economic outcomes, then employment per-capita should *not* respond to movements in house prices. In contrast, an estimate of the elasticity of per-capita employment relative to house prices which is found to be positive but lower in magnitude when compared to the elasticity of non-per-capita employment indicates that while part of the decline in employment is driven by migration, there remains a share of individuals who transition either from employment into unemployment or inactivity, or vice versa.

Table 8 presents estimates of the elasticity of 'non-tradable' employment *per-capita* with respect to house prices. These estimates are smaller in magnitude when compared to their non-per-capita counterparts. Specifically the range of the estimated elasticities of 'non-tradable' employment of 0.165-0.202 presented above compares to the one of the elasticities of 'non-tradable' employment *per-capita* of 0.115-0.125. *A 10% larger decline in a province relative to another leads to a 1.15% to 1.25% larger decline in 'non-tradable' employment per-capita.*

Table 8: Pooled estimates of the elasticity of 'non-tradable' employment *per-capita* with respect to house prices

$\Delta y_{l,t}$: **Employment in 'non-tradable' industries per capita**	(1)	(2)
$\hat{\beta}$: Δ real house prices	0.115*	0.125*
Standard error	0.049	0.057
P-value	0.02	0.027
Instrument	Z_1	Z_2
Controlling for economic sector employment shares interacted with time dummies?	Yes	Yes
N	762	762

Note: standard errors are two-way clustered at the province and time level. * indicates statistical significance at the 5% level, ** at the 1% level, and *** at the 0.1% level.

This evidence suggests that, while migration is a relevant factor in explaining the response of local employment to movements in local house prices, this adjustment mechanism is not able to fully absorb heterogeneous shocks. In other words, part of the larger decline in employment in provinces facing larger house prices decline can be explained by the fact that workers migrate to provinces facing less steep drops in the value of dwellings. However, a significant share of the difference in employment outcomes is also explained by workers transitioning into unemployment or leaving the labor market.

Referring again to the 2008-2013 recession, moving from the 90th percentile of the distribution of real house prices decline across provinces (-20.62%) to the 10th percentile (-39.17%) leads to an additional loss in per-capita 'non-tradable' employment of 2.32%, in the face of an average decline of 14.91% (and a standard deviation of 5.04%).

6.4 – The role of the interaction of leverage with the housing bust

The theoretical discussion in section 1.1 has highlighted the need to question whether declines in house prices occurring when households are highly leveraged imply a relatively larger response of economic activity to house prices. Here, I again investigate these issues focusing attention on the response of 'non-tradable' employment to movements in house prices. This measure of employment is deemed to be the most appropriate to study the effect of debt on the response of households' consumption to changes in housing values as it attempts to isolate, albeit imperfectly, the *household demand channel.*

In order to speak to the question at hand, I introduce in the empirical specification an interaction term constructed as the change in real house prices multiplied by the average ratio of mortgages constituted to output over the pre-bust period (2003 to 2007) multiplied in turn by a dummy variable assuming the value of 1 during the bust period (2008 to 2013). In this case, I include the

average ratio of mortgages constituted to output over the pre-bust period (2003 to 2007) as control in the regression.

Table 9 presents the results. The estimates on the growth in real house prices result to be somewhat lower than the ones obtained when not including the interaction term. While in the case of non-per-capita 'non-tradable' employment these estimates are significant at the 5% level, results with respect to the per-capita measure are not statistically significant, although quite precisely estimated. The estimated effect of the interaction term is consistently positive and comprised between 1.284 and 1.601 across columns 1 to 4. In one of the two cases relative to non-per-capita employment, the estimate is significant at the 5% level (column 1) while the p-value relative to the result in column 2 is just above 0.05. In contrast, in the case of per-capita employment, larger standard errors lead to not be able to reject the null hypothesis of a null effect for the interaction term.

Table 9: Pooled estimates of the elasticity of 'non-tradable' employment with respect to house prices - the interaction between leverage and the housing bust

	(1)	(2)	(3)	(4)
$\hat{\beta}$: Δ real house prices	0.155*	0.166*	0.078	0.093
Standard error	0.074	0.082	0.063	0.07
P-value	0.036	0.042	0.216	0.179
Δp x $(m/q)_{03-07}$ x bust	1.601*	1.441	1.456	1.284
Standard error	0.629	0.811	1.072	1.124
P-value	0.026	0.065	0.301	0.237
Measure of employment	Non-tradable	Non-tradable	Non-tradable Per-capita	Non-tradable Per-capita
Instrument	Z_1	Z_2	Z_1	Z_2
Controlling for employment shares interacted with time dummies?	Yes	Yes	Yes	Yes
N	762	762	762	762

Note: standard errors are two-way clustered at the province and time level. * indicates statistical significance at the 5% level, ** at the 1% level, and *** at the 0.1% level. In the interaction term, m denotes mortgages and q denotes output.

To gauge the economic significance of the estimates presented here, it is possible to ask to what extent the decline in employment following a 10% decline in house prices during the housing bust for a province at the 10th percentile of the pre-bust average ratio of mortgages constituted over output (2%) differs from the one faced by a province at the 90th percentile of the same distribution (5.8%). Using the estimates in column 2, while the first province would face a decline in 'non-tradable' employment of 1.95% following a 10% decrease in house prices, the province that relied more heavily

on mortgage debt prior to the bust would experience a decline in 'non-tradable' employment of 2.5% following the same drop in housing values.

I interpret these results as providing weak evidence in support of the idea that the interaction of high levels of debt with a decline in house prices lead to a stronger relationship between house prices and economic activity.

7 – Conclusion: Discussion, relationship with previous studies and directions for future research

As argued above, a key result of the present analysis is that movements in house prices have an economically significant impact on house prices. It is worth discussing how the magnitude of the obtained estimates compares to the results presented by Guren et al (forthcoming), who present directly comparable measures of the elasticity of retail employment to house prices for the US. The first two rows of table 10 compare the baseline elasticity of 'non-tradable' employment (table 7.a, column 4) to the estimate of the elasticity of retail employment with respect to house prices obtained by Guren et al. (forthcoming) when instrumenting for changes in house prices through the sensitivity instrument (page 22, table 2, third row - obtained using only the variation in the cross-section of counties between 2007 and 2009). Rows 3 and 4 compare estimates of the elasticities of employment in the same sectors, while focusing on *per capita* measures (to allow for a better comparison, I report here the estimates obtained by the authors for the period 2000-2017, which corresponds to the one used here – page 20, table 1, third column and second row).

Table 11: Comparison of elasticity estimates with the ones found by Guren et al. (forthcoming)

Study	**Elasticity of**	**Elasticity**
Present study	'Non-tradable' employment	0.202
Guren et al. (forthcoming)	*Retail employment*	*0.116*
Present study	'Non-tradable' employment per-capita	0.125
Guren et al. (forthcoming)	*Retail employment per-capita*	*0.055*

Note: the table presents a comparison of the elasticities of various measures of employment with respect to house prices obtained in the present study with their counterparts presented by Guren et al. (forthcoming). It is worth noticing that the estimates obtained by other studies are not readily comparable: Mian and Sufi (2014) present estimates of retail and restaurant employment with respect to what they term the housing net worth shock, while Adelino, Schoar and Severino (2015) focus on the elasticity of employment with respect to house prices for specific portions of the establishment size distribution.

Arguably, the estimates obtained by the two analyses for the various measures of employment and per capita employment result to be quite different. In particular, the elasticity of non-per-capita

'non-tradable' employment found by the present study is 73% larger than the one found by Guren et al. (forthcoming), while the one relative to per-capita employment is found to be 116% larger. A first possibility that one must consider is that (part) of the difference in the two estimates could be driven by the inclusion, in the 'non-tradable' sectors considered in the present analysis, of employment in activities that result to be particularly sensitive to changes in local demand, such as restaurant employment. In addition, the difference in the two sets of estimates could also be rationalized by considering the fact that the Spanish and US institutional settings are likely to significantly differ.

The present analysis is admittedly able to speak directly only to a limited set of differences between the two countries that could drive the results. The evidence in support of the idea that higher debt and sharp house prices decline, together with the observation that the increase in debt in the run-up to the crisis and the decline in house prices during the recession were of a larger magnitude in Spain relative to the US offers a compelling explanation for the divergence of the estimates. Moreover, the comparison of the elasticities of non-per-capita and per-capita employment measures in Spain and in the US leads to conclude that part of the difference in estimates relative to non-per-capita employment could be driven by differences in labor mobility across regions. While Guren et al. (forthcoming) find that the elasticity of retail employment in the US with respect to house prices declines by 50% when moving to per-capita measures, I find that the elasticity of 'non-tradable' employment in Spain declines by 38% when considering the per-capita counterpart. Moreover,

Moving beyond the evidence presented here, it is possible to rationalize the results by focusing on differences in local governments' finances, labor market institutions, mortgage markets and composition of the firms' population. First, it has been mentioned that municipalities in Spain rely for around 22% of their revenues on property taxes. To the extent that property taxation is more relevant in Spain compared to the US for regional governments, a change in housing values is likely to drive a larger response of local governments' spending. This will in turn imply a larger reaction of local employment. Second, it is possible that a more generous unemployment insurance paired with greater rigidity in wages and prices in Spain could lead similar shifts in local consumption to elicit larger responses in local employment. Third, different institutions in mortgage markets could imply a larger response of households' consumption following movements in house prices. For instance, anecdotal evidence (Daley, 2010) on foreclosure procedures in place in Spain during the recession point to the fact that creditor protection could have led a sizeable amount of Spanish households to be hit forcefully by foreclosure. Lastly, a greater relevance of small firms (or in general of firms relying on collateral for borrowing) in the population of businesses in Spain, and in particular among firms in the 'non-tradable' sector in Spain, compared to the retail sector in the US, could also lead to higher estimates relative to Spain through the firm collateral lending channel.

Turning to a discussion of other pieces of evidence provided by the present analysis, it is possible to notice that the evidence presented above has led to conclude that the movement of workers across regions is able to account only for part of the response of employment to geographically heterogeneous developments in house prices. This finding is particularly relevant, as it implies that individuals cannot insure against negative shocks incurred as a result of residing in a specific province by relocating. As the work by Yagan (2019) has found, focusing on the US context, that the failure of migration to

provide an adjustment mechanism to locality-specific recessionary shocks is a driver of long term effects of these shocks on local employment, I regard the investigation of such labor market hysteresis (and its drivers) in other contexts as a promising avenue for future research.

Furthermore, it has been argued that a third key result is that the housing demand channel, i.e. the response of households' consumption to movements in house prices, is relatively important to explain the relationship between the value of housing and employment. This observation is drawn observing that the estimated elasticity of a measure of employment in non-tradable sectors is larger than the one relative to employment in manufacturing (a proxy for tradable employment), which consistently results to be non-statistically significant. It is worth mentioning that the results obtained with respect to employment in the manufacturing sector differ from those presented by Mian and Sufi (2014) and Guren et al. (forthcoming) for employment respectively in tradable sectors and in the manufacturing industry. While these studies also report estimates of the elasticity of manufacturing employment that are not statistically significant, the point estimates they report are closer to zero and in some cases negative. In contrast, not only the point estimates presented in table (7.b) – relative to the elasticity of manufacturing employment with respect to house prices – are consistently positive, but there is also considerable overlap between the confidence intervals of these estimates with the ones of those relative to 'non-tradable' employment. As mentioned, a possible explanation for such difference could lie in the collateral lending channel for firms being relatively more important in Spain than in the US, perhaps due to small businesses accounting for a larger share of firms.

I also find weak evidence of the fact that the interaction of high levels of debt and large declines in house prices imply a relatively larger response of non-tradable employment to movements in house prices. I deem this result to be supportive of theoretical models implying an asymmetric response of economic activity to movements in house prices over the housing cycle (Guerrieri and Iacoviello, 2017; Jensen et al., 2020). Furthermore, this finding aligns with the evidence presented by Mian, Rao and Sufi (2013) and Aladangady (2017), which implies that the consumption of households with greater leverage responds relatively more to changes in the value of dwellings. On the contrary, it contrasts with the conclusions drawn by Kaplan, Mitman and Violante (forthcoming) and Guren et al. (forthcoming) that the impact of house prices on consumption is respectively consistent with collateral effects playing no role and insensitive to changes in households' leverage. However, I would note in these regards that as noted by Kaplan, Mitman and Violante (forthcoming) and Aladangady (2017), data exploiting regional variation is not perfectly suited to study the implications of the interaction of debt with changes in house prices. Future research should strive to use microdata to provide answers around this issue.

Lastly, as mentioned above, the present analysis is silent on the capacity of potential effects of movement in house prices on local governments' spending and banks' credit supply to drive the relationship between the value of housing and employment. Given that previous research has not put particular focus on these mechanisms, I also consider their investigation as fertile ground for future research.

References

Adelino, M., Schoar, A. and Severino, F. (2015). House prices, collateral, and self-employment. *Journal of Financial Economics 117,* 288-306.

Aladangady, A. (2017). Housing Wealth and Consumption: Evidence from Geographically Linked Microdata. *American Economic Review,* 107(11), 3415-3446.

Eggertsson, G. and Krugman, P. (2012). Debt, Deleveraging, and the Liquidity Trap: a Fisher-Minsky-Koo Approach. *Quarterly Journal of Economics,* 127, 1469-1513.

Favilukis, J., Ludvigson, S.C., Van Nieuwerburgh, S. (2016). The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk Sharing in General Equilibrium. *Journal of Political Economy, 125*(1), 140-223.

Gertler, M. and Gilchrist, S. (2018). What Happened: Financial Factors in the Great Recession. *Journal of Economic Perspectives,* 32(3), 3-30.

Gilchrist, S. and Mojon, B. (2016). Credit Risk in the Euro Area. *The Economic Journal,* 128, 118-158.

Greenstone, M., Mas, A., Nguyen, H. (2020). Do Credit Market Shocks Affect the Real Economy? Quasi-experimental Evidence from the Great Recession and "Normal" Economic Times. *American Economic Journal: Economic Policy*, 12(1): 200-225.

Guren, A. M., McKay, A., Nakamura, E. and Steinsson, J. (forthcoming). Housing Wealth Effects: The Long View. *Review of Economic Studies.*

Guerrieri, L. and Iacoviello, M. (2017). Collateral Constraints and Macroeconomic Asymmetries. *Journal of Monetary Economics,* 90, 28-49.

Guerrieri, V. and Lorenzoni, G. (2017). Credit Crises, Precautionary Savings, and the Liquidity Trap. *The Quarterly Journal of Economics,* 132(3), 1427-1467.

Jensen, H., Petrella, I., Ravn, S. H., Santoro, E. (2020). Leverage and Deepening Business-Cycle Skewness. *American Economic Journal: Macroeconomics*, 12(1), 245-281.

Jordà, O., Schularick, and Taylor (2016) The Great Mortgaging: Housing Finance, Crises and Business Cycles. *Economic Policy,* 107-152.

Jordà, O., Schularick, M. and Taylor, A. M. (2017). Macrofinancial History and the New Business Cylce Facts, in NBER Macroeconomics Annual 2016, volume 31, edited by Martin Eichenbaum and Jonathan A. Parker. Chicago: University of Chicago Press.

Kaplan, G., Mitman K. and Violante, G. (forthcoming). The Housing Boom and Bust: Model Meets Evidence. *Journal of Political Economy.*

Kaplan, G., Mitman, K. and Violante, G. (forthcoming). Non-Durable Consumption and Housing Net Worth in the Great Recession: Evidence From Easily Accessible Data. *Journal of Public Economics.*

Landvoigt, T., Piazzesi, M. and Schneider, M. (2015). The Housing Market(s) of San Diego. *American Economic Review, 105*(4), 1371-1407.

Mian, A., Rao, K., Sufi, A. (2013). Household Balance Sheets, Consumption, and the Economic Slump. *American Economic Review*, 128(4), 1687-1726.

Mian, A. and Sufi, A. (2014). What Explains the 2007-2009 Drop in Employment? *Econometrica,* 82(6), 2197-2223.

Mian, Sufi and Verner (2017). Household Debt and Business Cycles Worldwide. *Quarterly Journal of Economics*, 1755-1817.

Daley, S. (2010, October 27). In Spain, Homes Are Taken but Debt Stays. *New York Times.* Retrieved from https://www.nytimes.com/2010/10/28/world/europe/28spain.html

Saiz A. (2010). The Geographic Determinants of Housing Supply. *Quarterly Journal of Economics, 125*, 1253–1296.

Sole'-Olle', A. and Viladecans-Marsal, E. (2019). Housing booms and local spending. *Journal of Urban Economics*, 113.

Stock, J., Yogo, M. and Wright, J. (2002). A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments. *Journal of Business and Economic Statistics, 20,* 518-529.

Yagan, D. (2019). Employment Hysteresis from the Great Recession. *Journal of Political Economy, 127*(5), 2505-2558.