Labour Supply Responses to Financial Wealth Shocks: Evidence from Italy*

Renata Bottazzi (University of Bologna and Institute for Fiscal Studies, London)

Serena Trucchi (Cardiff Business School)

Matthew Wakefield (University of Bologna and Institute for Fiscal Studies, London)

Abstract

We look at how strongly shocks to wealth affect labour supply, using Italian data. We use asset price shocks to provide a measure of wealth changes that is exogenous to the household's saving and labour supply. Results point to significant effects of wealth on: hours of work; whether agents leave their jobs; and, labour earnings. The magnitude of these effects can be substantial, for example for individuals who suffered larger wealth losses during the financial crisis. Responses are similar for men and women on average, but older working-age individuals have relatively strong responses that drive the population results. Short-run effects are somewhat persistent.

Keywords: Labour Supply; Financial wealth shocks; Wealth effects.

JEL codes: D15, J22

---

* The name order of authors is alphabetical.

Contact details: Bottazzi: renata.bottazzi (at) unibo.it, Department of Statistical Sciences - University of Bologna, via Belle Arti 41, Bologna, 40126 (BO), Italy;
Trucchi: trucchis@cardiff.ac.uk, Cardiff Business School, University of Cardiff;
Wakefield: matthew.wakefield (at) unibo.it, Department of Economics, University of Bologna.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Sklodowska-Curie grant agreement no. 655770 (Trucchi). Much of the work for this paper was completed while Trucchi was lecturer at Ca' Foscari University of Venice.
We thank Orazio Attanasio, Richard Blundell, Carl Emmerson, Giovanni Gallipoli, Søren Leth-Petersen, Michał Myck and Ian Preston, for helpful discussions, and James Cloyne and two anonymous referees for comments that greatly improved the paper. We also thank participants at: the 6th HFC workshop at the Banque Centrale du Luxembourg, June 2018; the Ce^2 Workshop, Warsaw, July 2018; the 33rd Annual Congress of the EEA, Cologne, August 2018; the XXXIII annual conference of the Italian Association of Labour Economists (AIEL), Ancona, September 2018; a Bank of Italy seminar, Rome, November 2018, and, at the Netspar International Pensions Workshop, Leiden, January 2019.
The usual disclaimer applies.

I. Introduction

When faced with wealth shocks, do individuals adjust their labour supply? How strong are any adjustments? In a world in which increasing longevity, and modifications to pension systems, are encouraging individuals to accumulate private wealth, such questions are increasingly pertinent. Answers to these questions will be important for designing pensions and other policies that shape asset returns and incentives to accumulate wealth. The aggregate nature of asset price movements also means that these labour supply responses are important for trend and business cycle movements in economic output.

It is well-understood that forward-looking models suggest that when faced with unexpected changes (or "shocks") to lifetime resources, households should adjust their consumption and saving behaviour. Further, such models of "smoothing" suggest that households should adjust on other margins too: when consumption and leisure are both 'normal' goods, a negative (positive) shock to resources should lead to a decrease (increase) in purchases of both. There is a very extensive literature on how consumption responds to resource shocks. The last ten years, for example, has seen a plethora of papers looking at the links between housing wealth (and mortgage debt), and consumption,[1] and more broadly at the effects of resource shocks on consumption.[2] These papers in turn form part of a long established literature looking at wealth effects in consumption (Poterba (2000) and Paiella (2000) provide excellent surveys of the earlier literature).

Given the vast literature on consumption, it is perhaps surprising that the topic of resource shocks and leisure (or labour supply) has received much less attention (though see the discussion later in this section). Part of the reason for this difference in attention may be the nature of constraints on individuals' labour-supply choices. While agents can often make very fine adjustments to consumption expenditures, and so respond quickly and smoothly to shocks, the labour-supply choice is often between working full-time, or working part-time, or not working. In addition, while consumption spending is limited by a budget constraint, it may be possible to choose to adjust labour supply only if labour market conditions (labour

[1] Important empirical contributions include Mian, Rao and Sufi (2013), Cloyne et al (2019), Ganong and Noel (forthcoming) and Kaplan, Mitman and Violante (forthcoming), while Berger et al (2018) and Violante, Kaplan and Mitman (forthcoming) put more focus on taking models to the data.

[2] See, for example, Carroll, Otsuka and Slacalek (2011), Banks et al., (2012), Christelis, Georgarakos and Jappelli (2015), Carroll et al. (2017), Paiella and Pistaferri (2017) and Bottazzi, Trucchi and Wakefield (2020).

demand) allow. We aim to provide evidence on the responsiveness of labour supply to resource shocks, using data on Italy in the mid and late 2000s, a period that includes a considerable shock to financial asset values. The size of the shock is helpful as a source of potentially coarse adjustments in labour supply, and our analysis points to relatively large adjustments in hours being an important source of the effects that we find (see Section IV. 2.). In order to convince readers that the effects that we find are driven by labour supply choices, we control for various potential confounders that might lead to correlations between wealth shocks and movements in labour demand (see Section III and Section IV.1.).

Another challenge for studies aiming to identify the wealth (or "income") effect that we are interested in, is to find a source of variation in resources that is independent of household labour-supply choices, and that will not also be associated with a change in behaviour due to a "substitution" (or wage) effect. To address this, we import from the consumption literature the idea that a shock to asset values can provide a source of variation in wealth that is exogenous to households' behaviour. The shock comes from the 2006 – 10 period that we choose for our study. Italy's FTSE-MIB fell by more than 60% between May 2007 and March 2009, with a large part of this fall in the middle part of 2008.[3] Households that held wealth in stocks thus suffered a sudden, potentially large and mostly unanticipated shock to their financial wealth.[4] Using this shock to asset values as a source of variation and exogeneity, we present estimates for the responsiveness of labour supply to changes in resources based on an Instrumental Variables (IV) estimator and its associated reduced form. The IV estimator is related to that developed by Banks et al (2012) to estimate the propensity to consume out of a wealth shock, and we have used a similar estimator in our own study of consumption responses in Italy (Bottazzi et al., 2013 and 2020).[5] The details of our estimators are discussed in Section III.

Our data come from the Bank of Italy's Survey on Household Income and Wealth (SHIW). These data contain detailed information on respondents' labour-market outcomes, and our dependent variables use information on whether individuals work, on the number of hours

[3] The evolution of stock prices in Italy is documented in more detail in Appendix Figure A1.
[4] Our emphasis on the effects of financial wealth is related to the fact that, unlike in the US and UK, house prices in Italy did not drop dramatically in the early part of the Great Recession (Agenzia delle Entrate and Associazione Bancaria Italiana, 2019).
[5] The IV estimator has also been used by Crawford (2013) to look at the effect of wealth changes on the retirement plans of older people in England.

that they work, and on the level of their labour income. The data also contain rich information on households' demographic and economic characteristics including, crucially for our purposes, on their asset ownership and the values of asset holdings. The survey sample is representative of the Italian resident population and there is a rotating panel component. This impressive set of characteristics of the dataset, coupled with exogeneity from the wealth shock, make our study of broad interest for those wishing to understand labour supply responses to resource shocks.

While the literature on resource shocks and labour-supply is not as extensive as that on wealth and consumption, we are not the first researchers to look at labour-supply responses to wealth and wealth shocks. To meet the need for exogenous, non-wage related shocks to financial resources, some recent papers have looked at how households respond to lottery wins.[6] These interesting papers have provided broadly consistent results. Based on Swedish lottery data, Cesarini et al. (2017) estimate that winners of substantial lottery prizes adjust their earnings to offset about 1% of the lottery prize in each of the first ten years after the lottery win; they also find that this response is mostly due to a reduction of hours (rather than shifting to lower paid employment), and structural estimation suggests lifetime marginal propensities to earn in the 15 – 17 % range. Instead for Dutch lottery winners, Picchio et al. (2017) find a marginal propensity to earn of around (minus) 5%, with this response spread across 3 years after the win when large prize winners are included in the sample, or seeming more immediate if large prizes are excluded. Earnings responses to lottery wins were also found in the seminal work of Imbens et al. (2001), who had data on Massachusetts lottery players. Prizes in the Massachusetts lotteries were spread over a 20 year period rather than being paid in a single lump sum, and in their favoured specifications the authors found earnings responses equivalent to around 11% of the annualised prize in each of the first six years after a lottery win. These different papers also found results that were consistent in other dimensions, for example pointing to weak (Picchio et al., 2017), or little or no (Cesarini et al., 2017; Imbens et al., 2001), differences in responses for men and women. Our contribution to this literature comes from exploiting a different shock to wealth, and so

[6] Other papers have looked at how labour supply responds to inheritances, though few have been able to make the distinction between expected and surprise receipts (the importance of the distinction is clear in the analysis of retirement of Brown et al., 2010). Hurst and Lusardi (2004) and Disney and Gathergood (2009) have considered the relationship between wealth and entrepreneurship, but both papers conclude there is little evidence that inheritances or house price shocks encourage self-employment by unbinding liquidity constraints.

estimating effects for a different sample of the population (those who hold risky financial wealth, rather than those who play lotteries) which seems interesting given the potential cyclical importance of asset price shocks. In addition, in contrast to a lottery win, the wealth shock we exploit is negative and it is possible that responses to negative shocks differ from responses to positive news.[7]

Other papers have focussed, more closely than we do, on this cyclical element of asset price shocks. Coile and Levine (2011) find evidence that households in the US of around retirement age responded in their labour supply to the recent stock market crash, but this effect did not fully offset the effects of unemployment on these older workers. Using a related methodology, Disney, Radcliffe and Smith (2015) find little evidence of wealth effects on labour supply in the UK. Using different empirical variation, Disney and Gathergood (2018) find significant effects of house price shocks on the labour supply of younger individuals and older men in Britain. Using a similar method to isolate house price shocks, Milosch (2014) finds effects on the labour supply of younger women (particularly those with children, and in response to positive shocks) and older married men (responding to negative shocks) in the US. Our data and method allow us to focus on quantifying the labour-supply response to shocks to financial wealth.

A preview of our main results is as follows. Our baseline point estimates indicate that a reduction in risky financial wealth of 1000 euro would lead to 2.4 to 3.1 hour increase in annual labour supply, and a reduction in the likelihood of leaving work of between 0.055 and 0.09 percentage points. When combined with the (large) mean losses in risky wealth among holders of such wealth in our sample, these estimates suggest average increases in labour supply of between one part-time working week and one full-time working week, or a decrease of between 0.5 and 1 percentage point (or 10 and 20 percent) in the likelihood of leaving work. In financial terms, we find that for every thousand euros of loss in wealth, labour earnings were increased by between 55 and 68 euro in the year in which effects are measured. We find that our results are mainly driven by older workers (between 50 and retirement age),[8] and that responses are relatively similar for men and women. We also find

[7] This kind of asymmetry has been found for consumption responses to house-price changes (Disney, Henley and Gathergood, 2010).

[8] An effect inducing older workers to stay in the labour market may in itself have important aggregate consequences given recent evidence (Boeri, Garibaldi and Moen, 2016, and Bertoni and Brunello, 2017) that the

tentative evidence that the responses persist for at least one survey wave after the period of the wealth shock.

The paper is organised as follows. Section II describes the data that we exploit and section III details our empirical method. Section IV, V and VI present and discuss results. In Section IV, we present results for our baseline sample for changes in hours of work and in the likelihood of leaving employment, and then for changes in labour earnings. In section V we check robustness to restricting the sample to holders of risky assets and explore heterogeneity in results across populations defined by age and sex. The results in Section VI look at whether effects persist for a sample wave after the period of the wealth shock. Section VII concludes.

II. Data

The Italian Survey on Household Income and Wealth (SHIW) is a representative sample of the resident population. From 1987 onward the survey is conducted every other year (with a two-year gap between 1995 and 1998) and covers about 24,000 individuals and 8,000 households[9]. The panel component of the survey sample involves approximately 50% of households being re-interviewed in the following wave.

The survey contains a rich set of household and person characteristics as well as information on incomes and savings, and on household wealth and labour supply. Wealth data is rich, containing both participation and value for a range of financial assets, housing wealth, and businesses. For the purpose of our analysis, we use data for the years 2004-2010. In this way we are able to observe changes in labour supply and wealth between 2006 and 2008, and between 2008 and 2010. The information from 2004 (and 2006 and 2008) is used as required to construct lagged variables. The variation provided by the period of the large adjustment to financial asset values in 2007-08, is helpful for our empirical method.

We now describe the SHIW variables that we exploit, beginning with the labour-supply

retention of workers nearing retirement age has restricted employment opportunities for younger individuals in Italy, particularly in times of recession. These papers exploit increases in pension ages as a source of exogenous variation in the employment of older workers. The reforms that they exploit should not confound our study since individual financial wealth shocks will be unrelated to changes in pension age and because the biggest variation in pension ages (and the "Monti" Government that enacted it) occurred after the period of our data.

[9] A household is a group of individuals related by blood, marriage or adoption and sharing the same dwelling.

and wealth variables that are particularly important for our study.

1. SHIW labour-supply variables.

The SHIW dataset provides detailed information on labour supply, including regarding whether agents work, and about hours of work, potentially across multiple jobs. There is also information on sector and industry of employment, and on whether individuals are self-employed or work as employees. Our main dependent variables use information on whether agents have work, and on hours of work.

Basic descriptive statistics for these variables in our sample, and for households with and without risky financial assets, are provided in Table 1. In the table, hours of work are annual hours worked by an individual over the reference year,[10] and change in hours of work are the difference in annual hours worked between the current survey year and the previous one. Being in work is defined as having any paid job in the survey year, and those recorded as leaving work are individuals who are not in work in the current survey year but were in work in the previous one.

Figures 1 and 2 provide a fuller picture of labour supply in our data, as they respectively display the distribution (CDF) of hours in our sample, and the distribution of changes in hours in our sample (as a histogram). Both graphs display these distributions for our full (regression) sample (panel (a)), and for those in the sample that hold risky financial wealth (panel (b)), and separate plots for the different years in our data are overlaid. In the full sample, the distributions look rather stable across the different years. Among those with risky assets, we see a (rightward) shift in the CDFs towards working more hours (Figure 1 (b)), at the time of the major wealth shock in our data (between 2006 and 2008), and this shift results in a higher proportion working full-time (around 2000 hours or more) from 2008. In the change in hours data (Figure 2 (b)), the 2006-2008 period sees a greater number of increases in hours in the range 600 to 1000 hours in the year, compared to the later (2008-10) period. These

[10] We compute annual hours of work from survey responses regarding average weekly hours while in work during the calendar year. Note that SHIW surveys are named after the year to which data refer, and interviews are conducted at the beginning of the following calendar year (so SHIW 2008 contains information on 2008, based on interviews conducted in early 2009). For hours of work, we include hours in all jobs excluding occasional ones and respondents are specifically asked to include overtime hours. To make sure hours properly reflect time worked during the year, we also exploit survey information on the number of weeks or months during which the individual worked.

differences in the distributions among the sample with risky wealth are likely to be important for our results in Section IV.

TABLE 1
Descriptive statistics: labour supply

	Hours of Work		Change in Hours of Work		
	Mean	**Median**	**Mean**	**Median**	**N**
Full regression sample	973.4	960	-65.59	0	7140
2008	970.9	864	-52.51	0	3526
2010	975.8	960	-78.36	0	3614
HH with risky assets	1137.2	1440	-87.53	0	1206
2008	1107.7	1440	-82.39	0	602
2010	1166.6	1440	-92.65	0	604
	Work		**Leave Work**		
	Mean	**Median**	**Mean**	**Median**	**N**
Full regression sample	0.5462	1	0.0549	0	7140
2008	0.5383	1	0.0573	0	3526
2010	0.5540	1	0.0526	0	3614
HH with risky assets	0.6360	1	0.0498	0	1206
2008	0.6262	1	0.0482	0	602
2010	0.6457	1	0.0513	0	604

Source: authors' analysis of SHIW data. Sample as in baseline regressions.

FIGURE 1
CDFs of the distribution of hours worked

(a) Full Regression Sample

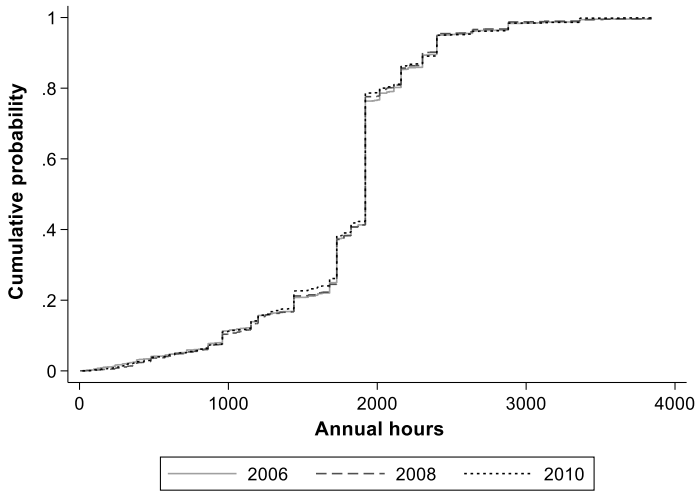

(b) Risky-asset holders

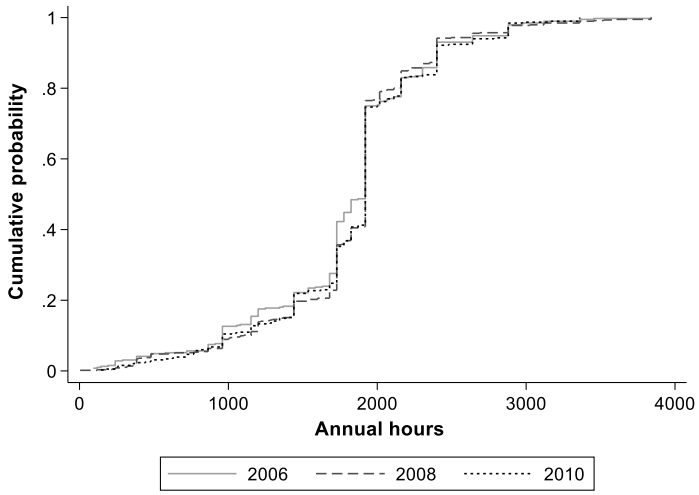

Note: Excludes those who never work.

FIGURE 2
Histogram of the change in hours worked

(a) Full Regression Sample

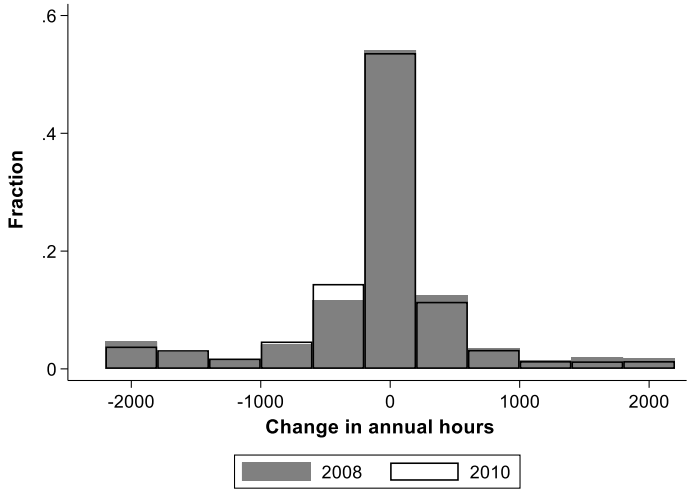

(b) Risky-asset holders

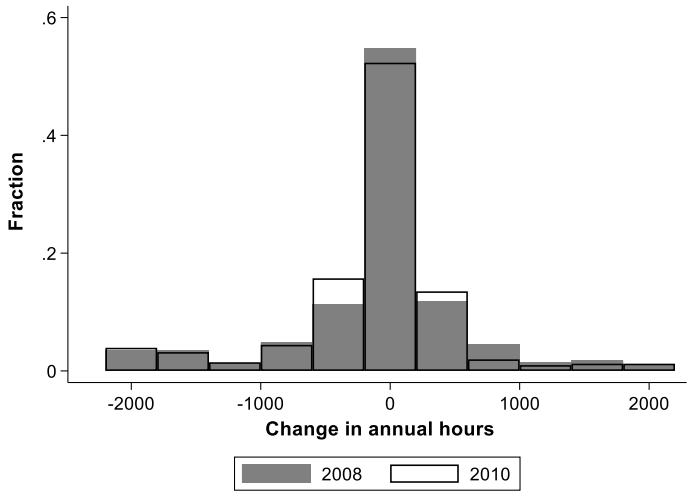

Note: Excludes those who never work.

2. SHIW financial wealth variables.

The SHIW dataset collects detailed information on household portfolios. Respondents are asked about ownership of, and about amounts of wealth held in, each of many types of asset. Assets are grouped in broad categories: cash (bank accounts and saving certificates); Italian government bonds (with different durations); domestic bonds and investment funds; Italian shares; foreign bonds and shares; and, other minor categories. Within each broad category individuals are asked about a detailed set of assets. SHIW also provides information on

household wealth in several types of mutual funds, and these funds can be categorised according to the extent to which they expose the holder to stock-market risk.

If survey respondents report that they hold an asset, they are then asked about how much wealth they held in that asset at the 31st of December in the year after which the survey wave is named (i.e. December 31st 2008 for the "2008 SHIW").[11] Respondents are first asked to indicate in to which of several bands of value their asset fell and then to report a point amount for this value. Failure to report a point amount results in the household being asked whether the value of their holding is nearer to the bottom, middle or top of the band. Since not all individuals give a point amount we use some imputed values for wealth. In imputation we use band and bottom/middle/top information to allocate values by asset.[12]

Since our main regressions are in first differences (see Section III) we have to be careful about the fact that imputation could considerably increase noise to signal ratio, especially where individuals report holdings in the relatively broad top bands of asset values. For this reason, in our sample selection we exclude from the sample households who do not provide a point amount and ever report being in the top bands (imputed wealth in a single asset above 150 000 euros with no upper limit).[13] Our sample selection also requires information on the variables included in our regressions and panel information (for a subset of variables) for three consecutive waves (to have a difference and lagged information), and we select individuals between ages 25 and 69. In households with more than one member, we keep the household head and his or her spouse.[14] We end up with a sample of around 7000 person-year observations.

We have described the labour-supply "outcome" variables for our sample, and the wealth variables that are contained in our data. Appendix Table A1 provides descriptive statistics for

[11] Having end of year wealth means we have data on households at close to the top of the stock market (at the end of 2006) and at close to the bottom of the crash (at the end of 2008).
[12] To have a homogeneous measure of asset values we do not use imputed values provided by the Bank of Italy, since they are not available for the 2004 wave. We need to rely on imputation by the Bank of Italy for (the sum of) three types of deposit in 2006, since information on the band they belong to is not available.
[13] When it comes to our analysis of wealth effects there could still be a concern regarding different awareness of wealth shocks from respondents who do not report the point amount of their asset value. To look at whether this could affect our results, we estimate our baseline specifications (Table 3) on samples with stricter requirements on non-imputation of wealth. Results are robust and are available on request.
[14] We also perform our analysis including other adult household members. Details are available on request.

our regression sample in terms of a range of characteristics, many of which become control variables in our regressions. Households with risky financial wealth will be important for our analyses, and the appendix table lists statistics for our full regression sample, and also for the set of individuals from households that have at least some risky financial wealth.[15] Those with risky wealth are similar to the full sample in terms of characteristics such as age and gender balance, but differ somewhat from the full sample in other dimensions. They are (almost by construction) relatively wealthy, and they are also relatively likely to be in work and to have permanent contracts, and they have comparatively higher incomes. They also tend to be more educated, more likely to be from the north of Italy, and more likely to be drawn from certain industrial sectors (such as finance or education and other public services), compared to the full sample. In Sections III and IV.1 we discuss how we attempt to ensure that these differences are unlikely to be problematic for our analyses.

III. Methodology

It is familiar that forward-looking models suggest that when faced with unexpected changes (or "shocks") to lifetime resources, households should adjust their consumption and saving behaviour. Further, such models of "smoothing" would suggest that households should adjust on other margins too, including through their choices over leisure and labour-supply (see, for instance, Low, 2005). We aim to understand whether, and how strongly, wealth shocks affect labour supply. To investigate this, we relate changes (first-differences) in labour supply choices, to changes in the value of (financial) wealth:

$$\Delta l_{ht} = \alpha + \omega \, \Delta w_{ht} + \varepsilon_{ht} \qquad (1)$$

Where: h and t indicate household and time period; l is a labour supply choice; w is the relevant measure of wealth; α and ω are coefficients and ε is an error term; and, Δ indicates "first difference" so that $\Delta l_{ht} = l_{ht} - l_{h(t-1)}$, with differences of other variables defined analogously.

Simply implementing equation (1) empirically by relating changes in labour supply to contemporaneous changes in wealth, is unlikely to provide a value of ω that can be

[15] More precisely, the subsample is agents that had risky financial wealth when we first observed them in our panel sample. These individuals are approximately 17% of our full regression sample.

interpreted as an unbiased estimate of the effect of wealth changes on labour supply. The complication is one of endogeneity since one way to increase wealth is to work and earn more. That is, an agent who works (and earns) more will increase wealth more than an otherwise similar agent who works less. This will generate a positive correlation between wealth changes and labour supply changes, even if (or when) the wealth changes are not causing adjustments to labour supply. Failure to account for this would thus lead to an upwards bias in the estimated coefficient.

We apply a method of dealing with this endogeneity that has been used in the consumption literature, and regress the change in labour supply on the "passive" part of the change in wealth.[16] The "passive" part of the change in wealth is the part that comes from capital gains and changes in asset values, rather than the part that is generated by choices about how much to earn, spend and save. [17]

To arrive at a value for the passive part of the change in wealth, we take a fixed wealth portfolio for each household, and calculate how the value of this portfolio would have changed due to changes in asset values and in the absence of any active saving (or dissaving) by the household. More concretely, consider calculating the change in the value of this fixed portfolio (hereafter "the calculated change in wealth") for an individual whose change in labour supply and wealth are observed for the period 2006 to 2008. A candidate fixed portfolio is the amounts of assets held in 2006. The household might (for example) have a certain amount of cash deposits, domestically held shares, and domestically held bonds.[18] Real values for these holdings by the end of 2008 can be calculated by applying the relevant real interest rate to the cash deposits, and the real change in the relevant price index for stocks and bonds, to up- (or down-) rate the values of the initial holdings. This will give a final value of the portfolio, and the calculated change in wealth is this final value less the initial value of the portfolio.

[16] The method we apply analyses changes in variables and so we work in first-differences. Differencing also brings the standard empirical advantages of conditioning out a fixed effect and potentially reducing the informational burden of the estimator (as wealth changes are likely to be easier to measure than wealth levels).
[17] The idea that this passive change can be used to deal with this endogeneity, dates back at least to Dynan and Maki (2001).
[18] The list of assets classes used in our empirical application, and the price indices and interest rates that we apply to them, are described in Appendix A of Bottazzi, Trucchi and Wakefield (2020).

In the previous paragraph, we described a calculation of the passive change in wealth for *(t-1)* to *(t)*, as based on the fixed portfolio from *(t-1)*. In fact in our empirical work we use portfolio information from *(t-2)*. That is, when we are dealing with changes in wealth (and labour supply) between 2008 and 2010, the portfolio information comes from 2006; portfolio information from 2004 is used when dealing with changes between 2006 and 2008. Taking a second lag ensures that the portfolio measure is not affected by measurement error from a survey period used in constructing differences of wealth (and labour supply) outcomes. In particular, in this way the portfolio measure will not be contaminated by the same measurement error that affects our measure of changes in observed wealth.[19] We use Δw_t^{fp} to denote our calculated value of the passive part of the change in wealth, and this is the calculated change in the value of the fixed portfolio from *(t-2)*.

A key part of our empirical strategy is to replace Δw_{ht} with Δw_{ht}^{fp} when estimating the relationship between labour supply outcomes and wealth changes described in equation (1). We also always exploit the richness of our data and include an "$\boldsymbol{X}$" vector of covariates that can help with precision and (as discussed in detail below) identification. Thus, the equation to be estimated becomes:

$$\Delta l_{ht} = \alpha + \omega \, \Delta w_{ht}^{fp} + \boldsymbol{X'_{ht}\beta} + \varepsilon_{ht} \qquad (2)$$

note that the labels on some coefficients, and for the error term, are the same in equations (1) and (2): this is for convenience and should not be taken as implying that estimating of the two equations would yield identical results.

For equation (2) to accurately measure the relationship of interest, we would need that the change in the value of the fixed portfolio accurately captures the "passive" part of the change in wealth. It is possible that the measure is not entirely accurate: our observations come at two-year intervals and in the period between observations households might take actions that adjust their exposure to asset price changes. If this means that the "passive" effect of changes in asset values on wealth is actually smaller than the values we calculate, then estimation of (2) would yield an underestimate of the size of the effect of wealth

[19] Being free of this contamination is particularly important when implementing the IV estimator described in the next paragraphs. The method of using lags is relatively standard for dealing with endogeneity in differenced panel data models, and is familiar from the literature on estimating log linear approximations to Euler equations (see the discussion of Attanasio and Weber, 1993, p.634, or Banks, Blundell and Tanner, 1998, especially footnote 8).

changes on labour supply. Even if our calculated variable does not capture "passive" changes in wealth entirely accurately, it can be expected to be correlated with actual changes in wealth and is unaffected by active saving decisions and thus unaffected by the influence of labour supply on wealth that we described above. Thus, the calculated change in wealth is the ideal "excluded variable" to construct an instrument for actual changes in wealth. This leads us to the following instrumental variables (IV) estimator:

$$\Delta l_{ht} = \alpha^{IV} + \omega^{IV} \widehat{\Delta w_{ht}} + \boldsymbol{X'_{ht}\beta^{IV}} + \varepsilon_{ht}^{IV} \qquad (3)$$

where: $\widehat{\Delta w_{ht}}$ is the predicted change in the relevant measure of wealth based on the following first-stage equation for the observed (reported) change in the value of a household's financial wealth (Δw_{ht}):

$$\Delta w_{ht} = \gamma + \varphi \, \Delta w_{ht}^{fp} + \boldsymbol{X'_{ht}\delta} + \mu_{ht} \qquad (4)$$

In our empirical results in the next section, we present estimates of ω from both equations (2) and (3); equation (2) is the reduced form of the IV estimator described by equations (3) and (4).

These specifications effectively compare labour supply outcomes for those with larger changes in (shocks to) wealth, to outcomes for those with smaller changes. Therefore, our estimates should not pick up the effects of aggregate changes that (equally) affected the labour supply of all agents. To identify effects of wealth shocks we need that, in the absence of the shocks, changes in labour supply outcomes would, *conditional on other regressors*, have been similar across those that do or do not suffer shocks.

In the light of this identifying assumption, the inclusion of a rich set of covariates can be important for convincing identification, as well as for precision. The exact set of regressors that we include will be detailed when we discuss empirical results in Section IV, and here we sketch some points relating to identification. All of our regressions include controls for demographic characteristics, but there might be additional potential "confounders" for the measurement of wealth effects. For example, one might think that portfolio shares, and therefore wealth shocks, are unlikely to be randomly distributed across households in different regions or working in different industries, and we already presented descriptive evidence that those with risky financial wealth are unusually likely to be drawn from the north of Italy and from certain industries (see the last paragraph of Section II, and Appendix Table

A1). If different regions or industries also faced different changes in labour market conditions, then correlations between wealth shocks and changes in labour market outcomes might not reflect only labour supply choices (as we would wish). We aim to control for such confounding factors by introducing, alongside the regressor for the local unemployment rate, region-year and industry-year effects, in a similar fashion to the method elaborated by Cloyne et al. (2019) in their innovative paper on measuring the effect of house prices on household borrowing. In addition to introducing fixed effects, in Section IV we also discuss how our estimates are affected when we analyse subsamples that are more homogenous in terms of region or industry, and so hopefully also in terms of labour demand shocks, than our baseline sample.

Given the nature of the crucial "Δw_{ht}^{fp}" variable, the key exogenous variation in wealth that we are exploiting is that generated by asset price changes. One way to justify that such changes come as shocks would be to note that asset price movements are highly persistent (permanent), so that the best guess of future prices are current prices and deviations from this are surprises. [20] Furthermore, in our case the biggest source of variation in asset prices comes from the 2007-2008 stock-market crash and it seems reasonable to suppose that price falls in this period were largely unanticipated (especially by individuals who remained in the stock market). Thus, the large change in asset prices in 2007-2008 is important for providing us with variation that is both substantial and exogenous.[21] The idea of using asset price changes as a source of plausibly exogenous variation in wealth has been exploited by researchers in other contexts. To investigate the effect of wealth on consumption, Banks et al. (2012) propose an IV strategy similar to that described above and apply it for a sample of older English households; in Bottazzi, Trucchi and Wakefield (2013 and 2020) we follow a similar approach with representative data for Italy. Banks et al. (2012) also look at other outcomes, notably expectational outcomes. Crawford (2013) finds little effect of wealth shocks on the retirement plans of older people in the England. Schwandt (2017) also exploits variation from asset price changes and finds some effects of wealth shocks on the health

[20] Data available in the 2008 and 2010 waves of data indicate that individuals also did not expect wealth shocks to reverse rapidly: around 60% of individuals report that the probability of making gains on the Italian stock market within a year is less than 15%. The expected persistence of the loss in wealth may help to explain why behavioural responses can be detected.

[21] In principle, individuals could also be exposed to wealth shocks through pension holdings. However, in Italy private pension schemes have not been widely held, and they are only slightly correlated to ownership of risky assets in our sample. Moreover, only a small minority of those with pension funds report owning stock funds.

among a sample of older Americans.

Descriptive statistics for the constructed change in wealth variable are provided in Table 2. More details on the construction of the variable, and on the comparison between actual (reported) and calculated changes in wealth are provided in Bottazzi, Trucchi and Wakefield (2020; see particularly the final part of Section III, and online Appendix A).

TABLE 2
Descriptive statistics: change in calculated risky financial wealth

	Statistics of distribution of change in calculated risky financial wealth						
	Mean	**P10**	**P25**	**P50**	**P75**	**P90**	**N**
Full regression sample:							
All	-1291	0	0	0	0	0	7140
2008	-2720	-7088	0	0	0	0	3526
2010	103	0	0	0	0	135	3614
Owners of risky wealth only:							
All	-7644	-20473	-8061	1	220	614	1206
2008	-15930	-35811	-18298	-8061	-4725	-2419	602
2010	613	30	65	220	541	1133	604

Source: authors' analysis of SHIW data. Sample as in baseline regressions.

IV. Empirical results for our baseline sample

1. Results for hours and participation

We present results for the estimators described in the previous section, for outcome variables that capture total changes in hours worked and in the decision of whether to exit employment.[22] Table 3 presents results for the change in hours worked and Table 4 presents results with an indicator of whether the agent left work during the two-year period between SHIW surveys as the outcome variable.

[22] We also considered entering work as an outcome variable but did not get significant results. Output available on request.

TABLE 3
Baseline results for change in hours of work

Dependent Variable: Change in Hours of Work				
	RF	**IV**	**RF**	**IV**
Δ risky financial wealth	**-2.396****	**-3.129***	**-2.390***	**-3.121***
	(1.169)	(1.688)	(1.222)	(1.748)
Δ housing wealth	0.072	0.116*	0.068	0.113*
	(0.053)	(0.061)	(0.053)	(0.062)
Couple	8.647	7.340	-1.577	-3.298
	(17.42)	(17.66)	(18.11)	(18.39)
Δ no. of people in HH	15.59	17.37	19.26	20.98
	(17.93)	(18.05)	(18.05)	(18.13)
Male	-91.93***	-92.18***	-15.37	-16.01
	(12.89)	(12.97)	(16.28)	(16.39)
High-school education	12.99	11.65	15.32	13.95
	(14.90)	(14.99)	(16.11)	(16.19)
Post-school education	56.01**	54.29**	46.55*	43.80*
	(23.00)	(23.45)	(24.51)	(25.00)
Regional unemployment rate	7.908*	9.491**	2.687	5.585
	(4.679)	(4.797)	(7.846)	(8.075)
Year 2010	-32.93*	-31.75*		
	(18.68)	(18.81)		
Central Italy	-4.067	0.476		
	(18.02)	(18.37)		
Southern Italy	-53.77	-61.87*		
	(36.53)	(37.13)		
Public sector employee	-24.61	-23.91	9.92	21.20
	(19.66)	(19.85)	(52.06)	(33.64)
Self employed	-46.53	-47.08	4.51	4.53
	(32.05)	(32.19)	(34.59)	(34.65)
5-year age dummies	Yes	Yes	Yes	Yes
Initial total wealth dummies	Yes	Yes	Yes	Yes
Years of contributions	No	No	Yes	Yes
Sector-year indicators	No	No	Yes	Yes
Region-year indicators	No	No	Yes	Yes
Constant	-129.8***	-135.9***	-2.964	-71.06
	(44.34)	(45.96)	(109.92)	(108.96)
# Observations	7140	7140	6891	6891

Note: *: significant at 10% level; **: significant at 5% level; ***: significant at 1% level.
Controls include "initial" *total* wealth dummies, with *total* wealth measured in 2004 for the 2006-08 sample and in 2006 for the 2008-10 sample: a dummy for zero or negative wealth and decile dummies for positive wealth.
Sectors for the Sector-year indicators are: Agriculture, Utilities, Construction, Retail, Transport, Finance, Real Estate, Domestic Services, Education and other Public Services, Extra Territorial, and Entertainment.
Regions for the Region-year indicators are: North, Centre, South.
Details of the first stage for the IV regressions in columns 2 and 4 of this table, are reported in Appendix Table A2. *F-test* of first stage is 12.58 and 12.11 for column 2 and 4 respectively.

TABLE 4
Baseline results for leave work

Dependent Variable: Leave Work				
	RF	**IV**	**RF**	**IV**
Δ risky financial wealth	**0.000678****	**0.000886****	**0.000547****	**0.000715***
	(0.000266)	(0.000436)	(0.000279)	(0.000424)
Δ housing wealth	-0.0000191	-0.0000317	-0.0000142	-0.0000244
	(0.0000154)	(0.0000183)	(0.0000157)	(0.0000181)
Couple	0.000534	0.000904	0.00433	0.00473
	(0.00753)	(0.00754)	(0.00769)	(0.00768)
Δ no. of people in HH	0.00169	0.00118	-0.000217	-0.000610
	(0.00751)	(0.00752)	(0.00757)	(0.00756)
Male	0.0185***	0.0185***	-0.0145**	-0.0144**
	(0.00538)	(0.00540)	(0.00649)	(0.00650)
High-school education	-0.0131**	-0.0127**	-0.00779	-0.00747
	(0.00611)	(0.00613)	(0.00616)	(0.00617)
Post-school education	-0.0285***	-0.0280***	-0.0133	-0.0127
	(0.00922)	(0.00934)	(0.00994)	(0.0100)
Regional unemployment rate	-0.00169	-0.00213	0.00230	0.00164
	(0.00208)	(0.00211)	(0.00280)	(0.00284)
Year 2010	-0.00356	-0.00389		
	(0.00674)	(0.00677)		
Central Italy	0.000813	-0.000204		
	(0.00723)	(0.00729)		
Southern Italy	0.0204	0.0227		
	(0.0168)	(0.0170)		
Public sector employee	0.0481***	0.0479***	0.0386**	0.0384***
	(0.00916)	(0.00918)	(0.0155)	(0.0113)
Self employed	0.0304***	0.0305***	0.0100	0.0100
	(0.00980)	(0.00989)	(0.0103)	(0.0103)
5-year age dummies	Yes	Yes	Yes	Yes
Initial total wealth dummies	Yes	Yes	Yes	Yes
Years of contributions	No	No	Yes	Yes
Sector-year indicators	No	No	Yes	Yes
Region-year indicators	No	No	Yes	Yes
Constant	0.0427**	0.0444**	-0.0448	-0.0173
	(0.0171)	(0.0175)	(0.0382)	(0.0385)
# Observations	7140	7140	6891	6891

Note: *: significant at 10% level; **: significant at 5% level; ***: significant at 1% level.
Controls include "initial" *total* wealth dummies, with *total* wealth measured in 2004 for the 2006-08 sample and in 2006 for the 2008-10 sample: a dummy for zero or negative wealth and decile dummies for positive wealth.
Sectors for the Sector-year indicators are: Agriculture, Utilities, Construction, Retail, Transport, Finance, Real Estate, Domestic Services, Education and other Public Services, Extra Territorial, and Entertainment.
Regions for the Region-year indicators are: North, Centre, South.
Details of the first stage for the IV regressions in columns 2 and 4 of this table, are reported in Appendix Table A2. *F-test* of first stage is 12.58 and 12.11 for column 2 and 4 respectively.

In these tables, the parameter on the main coefficient of interest (the change in wealth) is displayed in bold in the first row. The wealth variable that we use is the change in the value

of risky financial wealth, which is wealth that has some exposure to stock-market risk (either because the wealth is directly held in stocks, or because it is wealth held in a wrapper product such as a mutual fund, that includes some exposure to the stock market). This is wealth that was particularly exposed to the stock market fluctuations of 2007-2008 that provide us with a key source of variation. The second reported coefficient in all of the regressions is that on the reported change in the value of housing wealth, and this provides a useful comparison to the main coefficient.

Both Tables 3 and 4 present results from four regressions that are in reduced form (RF) and IV (second stage) pairs. For the IV regressions, the instrument is significantly correlated with the endogenous regressor and the F-test shows that we do not have a problem of weak instruments; details of the first stage regressions for Tables 3 and 4 are reported in Appendix Table A2. The difference between the two pairs of regressions is in the set of regressors. All of the regressions include regressors for demographic characteristics and some controls to capture (local) economic conditions and for the sector of employment (public or private), as well as controls for "initial" (*t-1*) total financial wealth.[23] In line with the idea of controlling for potential confounders of labour supply effects that is outlined in Section III, the expanded specifications in the right-hand columns of the tables more fully control for region and sector effects that may change over time, by replacing region and year indicators with a full set of interacted region-year indicators, and by including a set of indicators of interactions between (more detailed) sector of employment and year.,. Alongside these variables, the expanded specifications also include a flexible set of indicator variables to capture effects on labour supply behaviour of years of contributions to Italy's public pension (or social security) system. It is potentially important to include such indicators since labour supply decisions are likely to be affected by public sector pension accrual, at least near to retirement age.[24]

In Table 3, the coefficient of interest is remarkably stable at around -2.4 for the reduced-form specifications, and approximately -3.1 in the IVs, and these estimates are always

[23] The full set of regressors is detailed in the tables. Where necessary, the note to the table clarifies exactly what the variables measure.

[24] The indicators that we include are designed to capture flexibly the difference between individuals who have few years of contributions and so are far from pension eligibility, and individuals who have larger numbers of years of contributions and so are close to receiving a generous pension. In particular, we use indicators for groups of years of contributions that are particularly narrow once years of contributions are 30 or more, and we also interact these with an indicator for being a public-sector employee since public and private sectors have, at times, been treated differently.

significantly different from 0 at least at the 10% level. Given that the variable is measured in thousands of (2010) euros, this can be interpreted as indicating that, for every 1000 euro of increase in wealth, annual hours decrease by, on average, 2.4 or 3.1 hours per year. By contrast, the coefficient on the change in housing wealth (which is also measured in thousands of euros) is around 30 times smaller in size than the main coefficient of interest and it has the opposite sign. The change in housing wealth having comparatively little impact on our outcomes is consistent throughout the results reported in the paper, and we shall not comment on it further. A fuller assessment of whether the effect from the main variable of interest is substantial, is postponed until the subsection IV. 2.

Effects of the wealth shock on hours of work may reflect some workers making modest adjustments to their hours, but may also be partly due to some workers making "extensive margin" decisions to participate rather than to quit their jobs or otherwise stay out of work. To look at whether there is an effect from wealth shocks to labour market exits, we turn to the results for "leave work" in Table 4. The coefficients in the first row of the table are significant at the 5% level, or, in the case of fuller IV specification, at the 10% level. The point estimates suggest that a 1000 euro increase in wealth is associated with an increase in the probability of leaving work of between 0.055 percentage points and 0.089 percentage points. Again, a discussion of whether these effects are economically substantial is postponed to the next subsection.

Aside from the change in wealth, a few of the other regressors reported in the tables are consistently statistically significant. In the main specifications, the correlations are unsurprising. More educated individuals show more positive changes in hours and are less likely to leave work: since our sample starts at age 25, we are observing these individuals as they progress up the career path. Being male is negatively associated with the change in hours, and increases the likelihood of leaving work, but these patterns are both consistent with men participating more and working longer hours, and so having greater scope to reduce hours and leave work. The employment-type dummies (public sector employee and self-employed), both have positive and significant coefficients in the leave work regressions, but the "public sector employee" dummy must be carefully interpreted once sector-year interactions and years of contributions indicators (which are interacted with public sector in order to reflect pension system rules) are included.

A set of potentially important regressors are the indicators for the level of household wealth. Being wealthier might, all else equal, encourage reduced labour supply (an income effect) and temper precautionary responses to shocks, and is also correlated with exposure to risky assets.[25] In order to avoid an omitted variable bias whereby the effect of the level of wealth on labour supply adjustments is attributed to the coefficient on the change in risky wealth, the results presented in the body of the paper always come from specifications that include indicators for having zero financial wealth and of financial wealth decile group (among those with positive wealth). These variables are based on the value of total financial wealth, lagged by two survey periods (i.e. four years).[26] Results with and without these wealth dummies are reported in Appendix Tables A3 and A4 and show that including the wealth indicators does affect our coefficients of interest somewhat in the change in hours regression,[27] but the indicators are less important (and significant) in the "leave work" regressions. The sensitivity of our results to the financial wealth indicators is related to the fact that our sample includes a large amount of heterogeneity in wealth. If we restrict our sample to include only agents that held the risky financial assets that were exposed to the asset-price shock, we get a sample that is much more homogeneous in terms of wealth levels, and the size and significance of our estimates are much less sensitive to controlling for the wealth level. A fuller discussion of results for the risky-asset owner subsample is provided in Section V.1. and Appendix Table A6.

As anticipated in Section III, one potential worry for proper identification of wealth effects is that shocks to labour market conditions, for example at the region or industry level, and exposure to risky assets (and therefore to wealth shocks), might be correlated. The expanded specifications in the right-hand columns of the Tables 3 and 4 attempt to deal with this through the inclusion of a full set of region-year and sector of employment-year indicators, among the regressors. The results in these tables, and in Appendix Tables A3 and A4, show that the inclusion of the various extra controls does not substantially affect our main

[25] In our regression sample around 17% of agents have exposure to risky financial wealth, but this proportion is around 40% among the decile group with the most financial wealth. This also means that average reported losses in wealth at the time of the financial crisis are much larger in the wealthiest group than in the whole sample.
[26] Using *total* financial wealth means that the wealth groups are not based on the same wealth used to construct our main variable of interest, and lagging by two periods ensures the measure of wealth is not taken from the same survey of the differenced outcome variable for the regression.
[27] Coefficients on the wealth dummies (available on request) indicate larger increases in hours for those in lower wealth groups.

coefficients of interest. More specifically, the main results are shown to be insensitive to the inclusion of: the change in the value of housing wealth; the region-year and sector-year interactions; and, the addition of variables capturing contributions to the state pension system.

The right-most column of each of Appendix Table A3 and A4 illustrates the effect of adding individual fixed-effects to our fullest specifications reported in Table 3 and 4. Adding the fixed effect to the differenced equation means that we rely on within individual variation in wealth changes (and labour supply) to identify effects, rather than exploiting cross-sectional variation in wealth shocks and labour supply outcomes. This means that the fixed-effects equations can be a powerful means to control for potentially unobserved variables that might confound the relationship of interest at the cross-sectional level, and indeed this "triple differencing" strategy is important in Cloyne et al. (2019). Given that our data contains only two differenced observations per individual (plus an extra lag for creating the instrument), we do not have the ideal setting to apply this fixed-effects estimator. Nonetheless, we find it reassuring that our point estimates for the change in hours regression are insensitive to adding the fixed effect, even if the coefficient is no longer significant at standard levels (we would have significance at the 13% level). The result in the "leave work" regression is more fragile as the coefficient is approximately halved by the addition of the fixed effect. To anticipate subsection IV. 3., when the outcome is the change in labour income, the coefficient is robust to the addition of the fixed effect and it remains significant at the 10% level.

As an additional way to address the issue of potential confounders, we also perform our analysis on subsamples that are chosen in such a way as to reduce heterogeneity in characteristics that could be a source of correlation between exposure to risky assets and labour market outcomes. For example, the results based on a sample of those who have at least some risky financial assets, that will be discussed in subsection V.1., show that our findings are not generated by the labour supply behaviour of those who had no risky wealth and so were exposed to zero wealth shocks. Further analysis on subsamples, chosen on the basis of industrial sector or region of residence, leads us to conclude that our results are preserved. In particular, our results (point estimates and significance) are preserved if we restrict our sample to individuals attached to industrial sectors in which exposure to risky assets is seen to be higher, on average (finance, real estate, education and other public

services, utilities, retail and entertainment) or if we keep only individuals who reside in the north and centre of Italy – regions that tend to be richer and have a higher proportion of stockholders – though in this case the reduction in sample size does affect significance.[28]

We have found our results to be robust to a number of other changes to our specification. Coefficients are not substantially affected by restricting the sample to those who were in work when we first observed them, or to owners of risky assets who were in work, though the reduction in sample size does affect significance.[29] Removing small numbers of self-employed, or of individuals who moved house between survey waves, does not noticeably change the main coefficient or its significance.[30] Regarding the set of regressors, adding a measure of risk preference has almost no impact on our main results (see also Table 4.3 of Bottazzi et al. 2017). Finally, we explored adding controls for having a permanent contract, living in an "urban" area,[31] and sharing the home with young dependents, or with at least one elderly person, or with someone who is in poor health: adding these regressors separately or all together (with or without controls for region-year and sector-year and for risk aversion) did not have a substantial effect on our main coefficients of interest.[32]

2. Interpreting the results on hours and participation

We take our baseline results as providing robust evidence that the wealth shock had a significant impact on labour supply. To look at whether these effects are substantial, we consider them in the context of actual changes of wealth in our sample. Our data span the 2007 – 08 stock market crash. As described in Section III (Table 2), among those in our sample who held risky financial wealth and so were subject to the wealth shock, the average (calculated) change in such wealth is a loss of approximately 7600 euros. Our estimated coefficients of approximately -2.4 or -3.1 (see Table 3), suggest that an agent who suffered a loss in wealth of this magnitude would increase their labour supply by between 18 and 24

[28] All results mentioned in this paragraph that are not reported in the paper or the appendix, are available on request from the authors.
[29] In the change in hours regressions, point estimates for the reduced form are actually larger in the sample of workers than in our full sample, but differences in the first stage mean the IV is similar in the two samples. When we keep only workers that have risky assets, our point estimates are generally larger than in the full sample.
[30] Indeed, if anything, points estimates get larger in magnitude with these changes. Eliminating those who move also gets rid of a few apparently perverse significant results for the change in house value variable.
[31] A centre of residence with at least 40 000 residents.
[32] Results available on request.

hours per year.[33] This average loss in wealth comes across the 2006 – 2008 sample, a period of substantial stock market losses, and the 2008 – 2010 sample during which asset values were much more stable. Considering the earlier period in isolation, among stockholders even the median wealth loss is slightly more than 8000 euros and the (mean) average wealth loss is almost 16000 euros. Our estimates indicate that an individual suffering a loss in wealth of 16000 euro would increase their labour supply by between 38 and 50 hours per year. In other words, on average, the mean losses in risky wealth observed in our sample would have led to individuals increasing their labour supply by between one part-time working week and one full-time working week. Given that mean annual hours in our sample (around 970 hours for the whole sample or 1140 hours across those who own risky assets) approximately correspond to a full year working part-time, we find these changes in hours to be non-trivial.

Using the coefficients for the leave work specifications reported in Table 4, we can also calculate what wealth shocks of different sizes would mean for the likelihood of leaving work. We arrive at predicted effects of a 0.4 to 0.7 percentage point reduction in the likelihood of leaving work for an individual who suffered a wealth loss of around 7600, or of around 0.9 to 1.4 percentage points if we consider an individual who suffered the loss of 16000 euros. Given that the baseline likelihood of leaving work in our sample is around 5 percentage points (across a two-year observation period), these predicted effects amount to a 10 to 20 percent change and again seem economically important.

We can also use our results to provide predictions of what proportion of the change in hours in our sample due to wealth shocks, is generated by individuals adjusting on the "extensive" margin by not quitting. We base our predictions on the more parsimonious IV specifications reported in Tables 3 and 4 (with key coefficients -3.129 and 0.000886). The results from the "change in hours" regression allow us to calculate how much we think hours change due to the wealth shock (a combination of intensive and extensive margin adjustments), while the results from the "leave work" regression allow us to recover a prediction for the change in hours due to changes in decisions regarding exiting work.[34] We

[33] These numbers are arrived at by multiplying the mean change in wealth (7.644 thousand euros) by the smallest and largest coefficients from the first row of Table 3.

[34] The method to arrive at these predictions involves comparing predicted changes in hours based on our regression results and observed changes in wealth, to similar predictions when the change in wealth is counter-factually set to zero. Predictions based on the "leave work" regression are of the change in the likelihood of leaving work, and we multiply this by observed hours at *t–1* to convert it into a change in hours. Hours at *t-1* are

find that, among those in our sample who hold risky wealth, the average total change in hours due to the wealth shock is around 21 hours per year, and the average change due to changes in the probability of leaving work induced by wealth shocks, is approximately 7 hours per year. In other words, approximately one-third of the total change in hours seems to be due to changes in labour market exits, meaning around two-thirds of the total change is due to adjustments in hours from agents who continue to work.

The effects that we find thus seem in large part to be due to agents making adjustments to their hours on the intensive margin. We may still ask whether the effects are coming from a large number of people making small adjustments, or reflect a smaller number of people making larger adjustments.[35] This issue is of added interest given that the Italian labour market has been substantially reformed since the early 2000s, and consequences of these reforms have included the granting of greater flexibility in working-time arrangements and reduced restrictions on additional or overtime work and on stipulating flexible working arrangements (see Devicienti et al., 2018).[36]

We have undertaken different analyses to look empirically at whether larger or smaller adjustments in hours have been important. First, we truncated our sample to include only individuals with a change in hours of magnitude of 500 hours per year or less, and re-ran our change in hours specification. We found that this led to a large fall in our baseline estimates, and that results were no longer significant. This suggests that larger adjustments are important in driving our results. Second, we looked directly at overtime hours. Working paid overtime could be one way to make a small adjustment to hours. Agents in our data are asked to include overtime hours when reporting their hours of work, but we also have a separate measure of overtime hours. At a descriptive level we see very little change in average

a natural benchmark given that we are interested in how hours would have changed if agents had exited work between *t-1* and *t*.
For the interested reader, full details of the method are described in Section B of our appendix.

[35] We could also ask whether the changes in hours seem to come from individuals that change jobs, or from individuals that stay with the same employer. The data allow us to identify whether individuals had changed job between survey waves. Among current workers in our regression sample, 131 had changed job since the previous survey wave; this amounts to approximately 1.8% of the regression sample or almost 3.4% of workers in the sample. Given the small numbers of individuals changing jobs, and the fact that their hours changes are not drastically different on average from the changes for other workers, it is unlikely that the results that we find are driven mainly by individuals that change jobs. We do not observe a sufficiently large number of individuals that change jobs to investigate this further.

[36] Policy interventions also contributed to a sizeable increase in the share of temporary employees in the Italian workforce during the 2000s, and this share was close to the European average by 2010 (Cappellari et al., 2012).

overtime hours in our sample over time. We also performed our baseline regressions with change in overtime hours, and with an indicator for working a positive amount of overtime, as the dependent variable, and we found small and insignificant coefficients on the change in risky wealth variable. Adjustments on the overtime margin do not seem decisive in driving our results. Finally, at a more descriptive level, looking at our data on hours of work for the sample that hold risky assets, perhaps the most noticeable change in the distribution of hours between the period before and after the wealth shock (so, comparing 2008 and 2010 to 2006), is that we see a larger chunk of individuals working full-time after the shock (see Figure 1 (b)); in terms of changes in hours, adjustments in the period when the shock occurred (2006-08) are more concentrated in the (plus) 600 – 1000 hours per year range, compared to adjustments in the period without the large negative wealth shock (2008-10, see Figure 2 (b)). Putting together this regression and descriptive evidence, we tentatively conclude that individuals that made relatively large adjustments to their hours, such as working full-time rather than part-time, are important in driving our results. Given that adjustments of this type are likely to be more feasible than finer adjustments in many institutional settings, this would seem to add interest to our results even outside of the Italian context.

In sum, our baseline estimates indicate that the average wealth losses experienced in 2007-08 by those with risky financial wealth led to an increase in labour supply of between one part-time and one full-time working week per year. Our estimates also suggest that the fall in probability of leaving work due to these wealth losses is between 0.5 and 1 percentage point, and we calculate that the change in hours from this reduction in labour market exits would account for around a third of the total adjustment in hours in the sample.

3. Change in labour income

If adjustments in labour supply are effective as a means to offset shocks to financial resources, we would expect the behavioural responses found in the previous subsection to also be reflected in earnings. Looking at labour earnings as an outcome provides a means to investigate this directly. Table 5 presents results for our baseline sample, but with the change in (net of tax) labour income as the dependent variable.

TABLE 5
Results for labour income

	RF	IV	RF	IV
Dependent Variable: Change in Labour Income				
Δ Risky Financial Wealth	**-55.38****	**-66.87****	**-56.62****	**-68.25****
	(25.79)	(26.50)	(26.93)	(27.42)
Δ housing wealth	0.67	1.57	0.41	1.34
	(1.10)	(1.21)	(1.10)	(1.22)
Years of contributions	No	No	Yes	Yes
Region-year	No	No	Yes	Yes
Sector-year	No	No	Yes	Yes
# Observations	5938	5938	5785	5785

Note: *: significant at 10% level; **: significant at 5% level; ***: significant at 1% level. Controls include "initial" *total* wealth dummies, with *total* wealth measured in 2004 for the 2006-08 sample and in 2006 for the 2008-10 sample: a dummy for zero or negative wealth and decile dummies for positive wealth. The smaller sample size relative to Tables 3 and 4 is a result of needing data on earnings.
Sectors for the Sector-year indicators are: Agriculture, Utilities, Construction, Retail, Transport, Finance, Real Estate, Domestic Services, Education and other Public Services, Extra Territorial, and Entertainment.
Regions for the Region-year indicators are: North, Centre, South.
F-test of first stage is 10.70 and 10.31 for column 2 and 4 respectively.

The results indicate that for every thousand euros of loss in wealth, labour earnings were increased by between 55 and 68 euro per year. These results are stable across the specifications reported in the table, and Appendix Table A5 provides more details. The appendix table shows that: in line with the results for change in hours and for leave work, controls for wealth levels are important for the main coefficient of interest; the results for change in labour income remain significant at the ten per cent level also after allowing for an individual fixed effect.

The estimate of a short-run marginal propensity to earn of minus five to seven per cent[37] is similar in magnitude to the response to lottery wins found by Picchio et al. (2017), though a little larger than the short-run responses identified by Cesarini et al. (2017).[38] We can relate

[37] This marginal propensity to earn simply takes the one-year change in income, relative to the size of the total loss of wealth.
[38] The Picchio et al. (2017) study reports an immediate 5% response when the largest prizes are excluded from the estimation sample, while Cesarini et al. find an MPE of around minus 1% which is almost constant in each of the first ten years after the prize win, adding up to a lifetime response of 15 to 17% that is inferred using a structural model. The case studied by Imbens et al. (2017) is somewhat different since the lottery prizes are paid out in installments of 20 years, and they find responses to the annual payments of around minus 11% in each of the first six years after the prize win.

these earnings estimates to the results reported in Table 3. Our baseline estimates for hours correspond to an earnings response of about 36 to 47 euros per thousand euro of change in wealth per year.[39] These values are slightly lower than the point estimates in Table 5, and such a difference could be due to the wealth shock leading to workers increasing work effort (pushing for promotions, for example) as well as their hours. However, the numbers based on the change in hours results are well within the confidence bands for, and indeed less than one standard deviation away from, the estimates in Table 5, and so we feel that the two sets of estimates are reasonably consistent.

V. Heterogeneity

1. Holders of "risky" financial assets

As mentioned in Section IV.1, our baseline sample includes both households with and without exposure to risky financial wealth. Households that do not hold risky assets provide information that helps to identify coefficients on variables other than the main change in risky financial wealth variable. To check that including these households does not substantially alter our estimates of our main coefficient of interest, we ran our regressions on the subsample of households that have exposure to risky financial wealth.[40] Restricting our sample in this way also results in better balance on observables (such as region or industrial sector) than our full sample.[41] Table 6 reports the results for the change in hours and leave work specifications, for this subsample.

[39] The average hourly wage among workers in our sample is around 15 euros; multiplying this by 2.39 (the RF-estimate for the hours response) gives 36 (to the nearest euro), while 15 times the IV estimate of 3.12 gives 47.
[40] Precisely, the subsample is those who have a non-zero value for the change in the value of the fixed portfolio that is the crucial variable to identify our IV and reduced form (OLS) estimators. This means that the sample is of those who held risky assets at the relevant lag (usually, of 2 survey periods).
[41] In order to look at balance on observables in the sample of those with risky wealth, we compared those whose risky wealth (which is highly correlated with wealth shocks) is above or below the median in the sample; details of the comparison are available from the authors on request.

TABLE 6

Robustness: sample including only holders of risky financial wealth

	RF	IV	RF	IV
Dependent Variable: Change in Hours of Work				
Δ Risky Financial Wealth	**-3.411****	**-3.806***	**-3.639****	**-4.107***
	(1.614)	(2.022)	(1.688)	(2.138)
Δ housing wealth	0.078	0.135	0.088	0.150
	(0.081)	(0.092)	(0.084)	(0.098)
Dependent Variable: Leave Work				
Δ Risky Financial Wealth	**0.001005*****	**0.001122****	**0.001151*****	**0.001300****
	(0.00383)	(0.000550)	(0.000396)	(0.000598)
Δ housing wealth	0.000010	-0.000007	0.000016	-0.000003
	(0.000018)	(0.000023)	(0.000019)	(0.000027)
Years of contributions	No	No	Yes	Yes
Region-year	No	No	Yes	Yes
Sector-year	No	No	Yes	Yes
# Observations	1206	1206	1184	1184

Note: *: significant at 10% level; **: significant at 5% level; ***: significant at 1% level. Controls include "initial" *total* wealth dummies, with *total* wealth measured in 2004 for the 2006-08 sample and in 2006 for the 2008-10 sample: a dummy for zero or negative wealth and decile dummies for positive wealth.
Sectors for the Sector-year indicators are: Agriculture, Utilities, Construction, Retail, Transport, Finance, Real Estate, Domestic Services, Education and other Public Services, Extra Territorial, and Entertainment.
Regions for the Region-year indicators are: North, Centre, South.

The results in Table 6 are similar to those in Tables 3 and 4. Indeed, the effects are slightly stronger, both in terms of point estimates and significance, when we use the subsample. As with our main results, the results in Table 6 are based on regressions that include dummies for financial wealth decile group; Appendix Table A6 shows that in this subsample the coefficient of interest is not substantially affected (and significance is not changed) if the wealth level dummies are dropped. We interpret the results from analysis of this subsample as indicating that results based on our full sample certainly do not exaggerate estimates of the main coefficients of interest. In the remainder of the paper we stick to the broader sample and the more conservative estimates.

2. Older individuals

There are reasons to think that the labour supply of older individuals might be particularly responsive to the wealth shock we investigate. Older households tend to have more financial wealth and so are more likely to have been substantially exposed to the wealth

shock that is important for our estimation strategy. Workers with established employment are also those who are more likely to have some flexibility to adjust their hours of work, and who might be considering whether to leave jobs or reduce their hours as they enter or move towards retirement. Table 7 therefore reports results for the subsample of those aged 50 to 69.

The results for OLS regressions are very much in line with those reported in Tables 3 and 4, while the results for the IV specifications suggest stronger effects (at least in terms of point estimates) in the older sample. In our exactly identified system, the bigger difference between the reduced form and the IV for this older sample reflects that the correlation at the first stage is less strong (with a coefficient of around 0.5 instead of 0.8). However, this is not a reflection of a weak instrument: the F-test gives a value in excess of 16 for this older subsample (full first stage results available on request).

TABLE 7
Older subsample, ages 50 – 69

	RF	IV	RF	IV
Dependent Variable: Change in Hours of Work				
Δ Risky Financial Wealth	**-2.822****	**-6.287***	**-2.822****	**-6.501***
	(1.379)	(3.281)	(1.436)	(3.577)
Δ housing wealth	0.110*	0.241**	0.098*	0.236**
	(0.057)	(0.103)	(0.057)	(0.109)
Dependent Variable: Leave Work				
Δ Risky Financial Wealth	**0.000766****	**0.001706***	**0.000710****	**0.001635***
	(0.00350)	(0.000910)	(0.000369)	(0.000989)
Δ housing wealth	-0.000024	-0.000060	-0.000018	-0.000052
	(0.000020)	(0.000033)	(0.000020)	(0.000034)
Years of contributions	No	No	Yes	Yes
Region-year	No	No	Yes	Yes
Sector-year	No	No	Yes	Yes
# Observations	4473	4473	4343	4343

Note: *: significant at 10% level; **: significant at 5% level; ***: significant at 1% level. Controls include "initial" *total* wealth dummies, with *total* wealth measured in 2004 for the 2006-08 sample and in 2006 for the 2008-10 sample: a dummy for zero or negative wealth and decile dummies for positive wealth.
Sectors for the Sector-year indicators are: Agriculture, Utilities, Construction, Retail, Transport, Finance, Real Estate, Domestic Services, Education and other Public Services, Extra Territorial, and Entertainment.
Regions for the Region-year indicators are: North, Centre, South.

While the results for this older sample are in line with, or stronger than, the results for our baseline sample, we get much weaker patterns of coefficients, and usually insignificant results, for younger subsamples (results available on request). It is therefore clear that older working age individuals and those around retirement age, are important in driving our main results. This is indicative that part of the effects that we find reflect some workers postponing reductions in hours in the run up to retirement. In the context of a tight labour market in which the retention of older workers may have restricted opportunities for young adults to find work (cf. Boeri, Garibaldi and Moen, 2016; and, Bertoni and Brunello, 2017), the effect of the wealth shock on the labour market attachment of older workers is potentially important.

3. Men and women

Table 8 shows how our estimates vary if we split our baseline sample in to subsamples of men and women. The point estimates for both men and women are similar to our baseline results, though with reduced sample sizes significance levels are reduced. If there is a difference between the two samples it is that point estimates are slightly stronger for men when we consider the change in hours margin, though even here we cannot reject equal responses across genders. The similar responses for men and women may seem at odds with received wisdom that women's labour supply is relatively more responsive to financial incentives (see Keane, 2011), but it is worth noting that papers that exploit lotteries in order to estimate the income effect that we aim to identify also find no evidence that women respond more strongly than men (see Cesarini et al., 2017; Picchio et al., 2017; and, Imbens et al., 2001).[42,43]

[42] The Picchio et al. (2017) paper reports some specifications in which men appear to have stronger earnings responses to a win than women, but the difference vanishes after year 0 (the year of the win) and seems to be related to including a few very big lottery wins in the sample.

[43] We attempted to look at whether these effects involve men and women responding "jointly" within couples, or men responding in some households and women in others. If anything results pointed towards joint responses, but the findings are not significant (results available on request). Similarly, we did not find significant evidence on whether the responses of men and women might reflect different awareness of shocks to financial wealth. Both topics seem interesting for future analysis, but may require different data and/or methods from those used in the current paper.

TABLE 8

Labour supply for women and men

	RF	IV	RF	IV
Subsample: Females				
Dependent Variable: Change in Hours of Work				
Δ risky financial wealth	**-1.848**	**-2.338**	**-1.936**	**-2.458**
	(1.289)	*(1.732)*	*(1.412)*	*(1.858)*
Δ housing wealth	0.050	0.086	0.053	0.091
	(0.069)	*(0.073)*	*(0.067)*	*(0.072)*
Dependent Variable: Leave Work				
Δ risky financial wealth	**0.000781 ****	**0.000989**	**0.000700 ***	**0.000889**
	(0.000394)	*(0.000626)*	*(0.000427)*	*(0.000628)*
Δ housing wealth	-0.000019	-0.000034	-0.000015	-0.000029
	(0.000021)	*(0.000026)*	*(0.000021)*	*(0.000026)*
Years of contributions	No	No	Yes	Yes
Region-year	No	No	Yes	Yes
Sector-year	No	No	Yes	Yes
# Observations	3835	3835	3666	3666
Subsample: Males				
Dependent Variable: Change in Hours of Work				
Δ risky financial wealth	**-2.899**	**-3.843**	**-2.905**	**-3.882**
	(1.923)	*(2.889)*	*(1.969)*	*(2.998)*
Δ housing wealth	0.100	0.150	0.092	0.142
	(0.082)	*(0.101)*	*(0.084)*	*(0.104)*
Dependent Variable: Leave Work				
Δ risky financial wealth	**0.000696 ***	**0.000922**	**0.000605**	**0.000809**
	(0.000372)	*(0.000636)*	*(0.000379)*	*(0.000629)*
Δ housing wealth	-0.000018	-0.000030	-0.000013	-0.000023
	(0.000022)	*(0.000027)*	*(0.000024)*	*(0.000027)*
Years of contributions	No	No	Yes	Yes
Region-year	No	No	Yes	Yes
Sector-year	No	No	Yes	Yes
# Observations	3305	3305	3225	3225

Note: *: significant at 10% level; **: significant at 5% level; ***: significant at 1% level. Controls include "initial" *total* wealth dummies, with *total* wealth measured in 2004 for the 2006-08 sample and in 2006 for the 2008-10 sample: a dummy for zero or negative wealth and decile dummies for positive wealth.

Sectors for the Sector-year indicators are: Agriculture, Utilities, Construction, Retail, Transport, Finance, Real Estate, Domestic Services, Education and other Public Services, Extra Territorial, and Entertainment.

Regions for the Region-year indicators are: North, Centre, South.

VI. Persistence of effects

The results presented so far identify how labour supply responds to the change in wealth in the period immediately following the shock. It is of interest to also look, to the extent that data allow, at whether these short-run effects persist. We can investigate this for those members of our 2006-08 sample who we also observe in 2010. From our rolling panel dataset we have almost 3000 observations for which we observe long "two-wave" (2006 – 2010) differences (and our instrument).

For this sample, Table 9 presents results for our baseline hours of work and leave work specifications. The first two columns provide results for the 2006-10 differences; in the case of the leave work indicator, this long difference is an indicator of having left work in either of the intervals 2006-08 or 2008-10. The middle columns of the table are our baseline (2006-08) specifications but for the reduced sample, and the final two columns show results with 2008-10 differences as the dependent variable. Since the regressors that we use are consistent across the columns, the coefficients in the first two columns are the sum of the coefficients in the analogous columns for the 2-year differences.

The results for the change in hours indicate that the effects identified in our baseline specifications are persistent. The coefficients for the 2006-10 change are significant and only slightly smaller in magnitude than the short-run (2006-08) response to the wealth shock, due to a small but not significant reversal in 2008-10. The coefficients across the leave work regressions point to a similar story, though in this case only the coefficient for the short-run effect is significant. We thus interpret the evidence as indicating that the short-run effects identified in our baseline regressions are not immediately fully reversed. Such persistence is in line with the findings of Cesarini et al. (2017) and Picchio et al. (2017), although it is perhaps less obvious that the effects we find should persist since the responses we identify often involve workers that are approaching retirement age increasing their labour supply. On the other hand, it is important to be clear that our finding of persistence does not necessarily imply that households have higher hours for three or four years after the shock. The persistent effect indicates that those who suffered the wealth shock between 2006 and 2008 had a larger change in hours between 2006 and 2008, and also between 2006 and 2010, than those who suffered smaller shocks. This could reflect some households increasing their hours (relative to their choice in the absence of the shock) between 2006 and 2008, and not

reversing this after another two years, or it could reflect some households that suffer the wealth shock increasing hours in the 2006-2008 period, and *other* households, that also suffered the shock, increasing hours between 2008 and 2010. Even though our empirical strategy does not allow us to distinguish between these possibilities, the finding of persistence seems noteworthy.

TABLE 9
Persistence of effect

	RF	IV	RF	IV	RF	IV
	2006-10		2006-08		2008-10	
Dependent Variable: Change in Hours of Work						
Δ Risky Financial Wealth 2006-08	-3.348 ** *(1.113)*	-3.910 ** *(1.570)*	-4.048 ** *(1.332)*	-4.728 ** *(1.962)*	0.700 *(1.249)*	0.818 *(1.480)*
Δ housing wealth	0.139** *(0.0695)*	0.185** *(0.0777)*	0.199 ** *(0.0667)*	0.254 ** *(0.0801)*	-0.0594 *(0.0582)*	-0.0690 *(0.0639)*
# Observations	2870					
Dependent Variable: Leave Work						
Δ Risky Financial Wealth 2006-08	0.000584 *(0.000628)*	0.000682 *(0.000757)*	0.000686 ** *(0.000320)*	0.000802 * *(0.000437)*	-0.000102 *(0.000548)*	-0.000119 *(0.000637)*
Δ housing wealth	*-0.0000333* *(0.0000275)*	*-0.0000413* *(0.0000315)*	-0.0000234 *(0.0000201)*	-0.0000328 *(0.0000223)*	-0.00000989 *(0.0000183)*	-0.00000849 *(0.0000211)*
# Observations	2870					

Note: *: significant at 10% level; **: significant at 5% level; ***: significant at 1% level. Controls include "initial" *total* wealth dummies, with *total* wealth measured in 2004 for the 2006-08 sample and in 2006 for the 2008-10 sample: a dummy for zero or negative wealth and decile dummies for positive wealth.

5. Conclusions

We have looked at whether shocks to asset values lead to labour supply adjustments, using Italian data. We used asset price shocks to provide a measure of wealth changes that is exogenous to households' saving and labour supply behaviour.

Our results suggest that wealth losses led to some increases in hours worked, and reductions in numbers leaving jobs. The magnitude of these effects could be substantial for those suffering larger wealth shocks (although such shocks are concentrated among relatively few owners of risky assets). For example, when combined with the mean losses in risky wealth among holders of such wealth, our point estimates suggest average increases in labour supply of between one part-time working week and one full-time working week. Looking at the extensive margin for the same group, we found a decrease of between 0.5 and 1 percentage point (or 10 and 20 percent) in the likelihood of leaving work. Using labour income as our outcome variable, we find a marginal propensity to earn of between (minus) 5 and 7 percent. Our baseline findings measure short-run responses to the wealth shock, but we also show evidence of at least some persistence in these responses.

Examining population subgroups allows us to investigate heterogeneity in effects. In terms of age, we find that older subgroups are important in driving our results as the clearest responses come from those of older working-age and around retirement age. We find little evidence that the labour supply of men and women responds differently to the wealth shock.

The evidence in this paper indicates that households use labour-supply, as well as the spending and saving margin, to smooth out shocks. Since shocks can be aggregate in nature, this may have important macroeconomic implications. Increases in labour supply may smooth out the adverse impact of a negative shock. On the other hand, it is also possible that, in a tight labour market situation, older workers staying in employment may reduce job opportunities for their younger counterparts. Opportunities for younger workers have been particularly hit during the current health crisis, and it is possible that when activity can restart these individuals will be joined in the labour force by older workers who experienced negative resource shocks as the crisis took effect. The aggregate impacts of the effects we discuss remain a topic for further work.

References

Agenzia delle Entrate and Associazione Bancaria Italiana (2019), Rapporto Immobiliare 2019 – Il Settore Residenziale, Osservatorio del Mercato Immobiliare (Agenzia delle Entrate), https://www.agenziaentrate.gov.it/portale/documents/20143/263076/rapporto+immobiliare2019+ri_RI2019_Residenziale_20190523.pdf/a175f856-2363-dda7-da64-b1c0e544eb12

Attanasio, Orazio P., and Guglielmo Weber (1993), "Consumption Growth, the Interest Rate and Aggregation", *The Review of Economic Studies* 60, 631-49.

Banks, James, Richard Blundell and Sarah Tanner (1998), "Is there a Retirement-Savings Puzzle", *American Economic Review* 88(4), 769 – 788.

Banks, James, Rowena Crawford, Thomas F. Crossley and Carl Emmerson (2012), *The Effects of the Financial Crisis on Older Households in England*, IFS Working Paper W12/09.

Berger, David, Veronica Guerrieri, Guido Lorenzoni and Joseph Vavra, (2018). "House Prices and Consumer Spending," *Review of Economic Studies*, Oxford University Press, vol. 85(3), pages 1502-1542

Bertoni, Marco and Giorgio Brunello (2017), *Does Delayed Retirement Affect Youth Employment? Evidence from Italian Local Labour Markets*, IZA DP No. 10733.

Boeri, Tito, Pietro Garibaldi and Espen Moen (2016), *A Clash of the Generations? Increase in Retirement Age and Labour Demand for Youth*, INPS WorkINPS papers number 1, July 2016, ISSN 2532-8565, (https://www.inps.it/docallegatiNP/Mig/InpsComunica/WorkInps_Papers/1_WorkINPS_Papers_1luglio2016.pdf)

Bottazzi, Renata, Serena Trucchi and Matthew Wakefield, (2013), *Wealth effects and the consumption of Italian households in the Great Recession*, IFS Working Paper W13/21.

Bottazzi, Renata, Serena Trucchi and Matthew Wakefield, (2017), *Labour Supply Responses to Financial Wealth Shocks: Evidence from Italy*, IFS Working Paper W17/29.

Bottazzi, Renata, Serena Trucchi and Matthew Wakefield (2020), "Consumption responses to a large shock to financial wealth: evidence from Italy", *The Scandinavian Journal of Economics*, 122(2), 762-789.

Brown, Jeffrey R.; Courtney C. Coile and Scott J. Weisbenner (2010), “The Effect of Inheritance Receipt on Retirement”, *The Review of Economics and Statistics* 92(2), 425 – 434.

Cappellari, L., Dell’Aringa, C. and Leonardi, M. (2012), *Temporary Employment in Italy*, CESIfo DICE Report.

Carroll, C. D., Otsuka, M. and Slacalek, J. (2011), “How Large Are Housing and Financial Wealth Effects? A New Approach”, *Journal of Money, Credit and Banking* 43(1), 55–79.

Carroll, C. D., Slacalek, J., Tokuoka, K. and White, M. N. (2017), “The distribution of wealth and the marginal propensity to consume”, *Quantitative Economics* 8, 977 – 1020.

Cesarini, David; Erik Lindqvist; Matthew J. Notowidigdo and Robert Östling (2017), “The Effect of Wealth on Individual and Household Labor Supply: Evidence from Swedish Lotteries”, *The American Economic Review* 107(12), 3917-3946.

Christelis, D., Georgarakos, D. and Jappelli, T. (2015), "Wealth Shocks, Unemployment Shocks and Consumption in the Wake of the Great Recession," *Journal of Monetary Economics* 72, 21-41.

Cloyne, James, Kilian Huber, Ethan Ilzetzki, and Henrik Kleven. 2019. "The Effect of House Prices on Household Borrowing: A New Approach." *The American Economic Review*, 109 (6): 2104-36.

Coile, Courtney C., and Phillip B. Levine (2011), “The Market Crash and Mass Layoffs: How the Current Economic Crisis May Affect Retirement”, *The B.E. Journal of Economic Analysis and Policy*, Berkeley Electronic Press, vol. 11(1), 22 pages.

Crawford, Rowena (2013), “The effect of the financial crisis on the retirement plans of older workers in England”, *Economics Letters* 121, 156-159.

Devicienti, F., Grinza, E. and Vannoni, D. (2018), "The impact of part-time work on firm productivity: evidence from Italy," *Industrial and Corporate Change*, vol. 27(2), pages 321-347.

Disney, Richard, and John Gathergood (2018), “House Prices, Wealth Effects and Labour Supply”, *Economica 85*, 449-478.

Disney, Richard, and John Gathergood (2009), “Housing Wealth, Liquidity Constraints and Self-employment”, *Labour Economics 16*, 79 – 88.

Disney, Richard, John Gathergood and Andrew Henley (2010), “House Price Shocks, Negative Equity and Household Consumption in the United Kingdom”, *Journal of the European Economic Association 8(6)*, 1179 – 1207.

Disney, Richard, Anita Radcliffe and Sarah Smith (2015), “Booms, Busts and Retirement Timing”, *Economica 82*, 399 – 419.

Dynan, Karen E., and Dean M. Maki (2001), *Does Stock Market Wealth Matter for Consumption?*, Board of Governors of the Federal Reserve Systems “FEDS” (Finance and Economics Discussion Series) paper 2001-23, May 2001.

Ganong, Peter and Pascal Noel (2020), “Liquidity vs. Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession”, *The American Economic Review* 110(10), 3100-3138.

Hurst, Erik and Annamaria Lusardi (2004), “Liquidity Constraints, Household Wealth, and Entrepreneurship”, *Journal of Political Economy 112(2)*, 319 – 347.

Imbens, Guido W.; Donald B. Rubin and Bruce I. Sacerdote (2001), “Estimating the Effect of Unearned Income on Labor Earnings, Savings, and Consumption: Evidence from a Survey of Lottery Players”, *The American Economic Review* 91(4), 778 – 794.

Kaplan, Greg, Kurt Mitman and Gianluca Violante (2020), “Non-Durable Consumption and Housing Net Worth in the Great Recession: Evidence from Easily Accessible Data”, *Journal of Public Economics* 189, 104176.

Kaplan, Greg, Kurt Mitman and Gianluca Violante (2020), "The Housing Boom and Bust: Model Meets Evidence," *Journal of Political Economy* 128(9), 3285–3345.

Keane, Michael P. (2011), “Labor Supply and Taxes: A Survey”, *Journal of Economic Literature* 49(4), 961 – 1075.

Low, Hamish W. (2005), “Self-insurance in a life-cycle model of labour supply and savings”, *Review of Economic Dynamics* 8, 945 – 975.

Mian, A., Rao, K. and Sufi, A. (2013), “Household Balance Sheets, Consumption, and the Economic Slump”, *The Quarterly Journal of Economics* 128(4), 1687-1726.

Milosch, Jennifer, (2014), “House Price Shocks and Labour Supply Choices”, chapter 3 of Jennifer Milosch, *Economic Shocks and Household Decisions*, PhD dissertation, University of California, Santa Barbara, (https://www.alexandria.ucsb.edu/downloads/qr46r114t).

Paiella, M. (2009), “The Stock Market, Housing and Consumer Spending: A Survey of the Evidence on Wealth Effects”, *Journal of Economic Surveys* 23(5), 947-73.

Paiella, M. and Pistaferri, L. (2017), “The Anatomy of the Wealth Effect”, *The Review of Economics and Statistics* 99(4), 710-721.

Picchio, Matteo; Sigrid Suetens and Jan C. van Ours (2017), “Labour Supply Effects of Winning a Lottery”, *The Economic Journal* 128 (June), 1700-1729.

Poterba, J. M. (2000), “Stock Market Wealth and Consumption”, *Journal of Economic Perspectives* 14 (2), 99–118.

Schwandt, Hannes (2018), “Wealth Shocks and Health Outcomes: Evidence from Stock Market Fluctuations”, *American Economic Journal: Applied Economics*, 10(4): 349 - 377.