

Decomposing the Wealth Effect on Consumption^{*}

Monica Paiella[†]

Luigi Pistaferri[‡]

June 1, 2015

Abstract

We decompose the wealth effect on consumption into its two components. First, we distinguish between exogenous and endogenous wealth changes (due to changes in prices or portfolio choice). Second, we distinguish between anticipated and unanticipated exogenous changes. We estimate the impact on consumption of the various components using microeconomic data on consumption, wealth, and subjective asset price expectations available from the 2008-10 panel of the Italian Survey of Household Income and Wealth. We estimate an overall wealth effect of about 3 cents per (unexpected) euro increase in wealth. This effect is driven primarily from a positive consumption response to house prices. The consumption response to anticipated changes in wealth is also large and significant, of the same magnitude as the response to unanticipated changes, and similarly driven by changes in housing wealth.

^{*}This paper was written while Pistaferri was visiting EIEF, whose hospitality is gratefully acknowledged.

Thanks to Tullio Jappelli for comments. The usual disclaimers apply.

[†]University of Naples "Parthenope".

[‡]Stanford University.

1 Introduction

Whether and how much changes in wealth affect households consumption is crucial for understanding how asset prices impact the economy and to evaluate the role of monetary policy. The basic ideas and key theoretical links between wealth and consumption are typically described using the life-cycle permanent income model. According to this model, consumers accumulate and deplete their wealth in order to keep the marginal utility of consumption smoothed over time. In one version of the theory, interest rates are non-stochastic and income is the only source of uncertainty. It follows that changes in wealth reflect unexpected changes in earnings. In models with stochastic interest rates, however, households may experience an unexpected change in wealth even with constant income, due for example to asset price shocks, which will induce revisions in their optimal consumption plan. This is what is typically termed the "wealth effect".

There have been several attempts of estimating the wealth effect on consumption, using aggregate data (e.g., Lettau and Ludvigson, 2001 and 2004; Sousa, 2008) or household-level data (e.g. Dynan and Maki, 2001; Paiella, 2003; Juster et al., 2006). Cross-country comparative studies include Case et al. (2005 and 2011), Bertaut (2002) and Ludvig and Slok (2004). This research has partly being stimulated by the wide variability in asset prices of the last decades, in particular the stock market boom of the second half of the 1990s and its subsequent decline, as well as the house price boom and bust that culminated with the Great Recession of 2007-09.

Despite their explicit reference to the life-cycle permanent income model, most studies in the literature do not consider the distinction between anticipated and unanticipated changes in wealth. Another issue that is sometimes neglected in the empirical literature is the distinction between exogenous changes in wealth (due to asset price shocks) and endogenous changes (due to portfolio choice). In this paper we attempt to address both issues. To do so, we combine subjective asset price expectations from the 2008-10 Italian Survey of Household

Income and Wealth (SHIW) with ex-post price realizations to identify asset price shocks, which we then merge with data on beginning-of-period wealth to separate the unanticipated from the anticipated wealth variation. Italy is a particularly useful case to study, as household wealth is pretty high by international standards (the average wealth/income ratio is 8, compared to 6 in Germany and 5 in the US), real assets represent about 2/3 of total wealth, and debt (including mortgage debt) is low (about 80 percent of disposable income).

We argue that the "pure" wealth effect that is of interest in most of the literature is captured by the response of consumption to unanticipated wealth changes. In contrast, the response to expected wealth changes captures intertemporal substitution, not wealth effects. Since changes in asset prices reflect changes in the relative price of present vs. future consumption, the latter responds to both anticipated and unexpected wealth changes. This is unlike the distinction between anticipated and unanticipated income effects on consumption, where unanticipated income changes shift consumption but anticipated ones do not. We also isolate exogenous changes in wealth (due only to asset price shocks).

We report two main results. First, in our sample the overall wealth effect is around 1-3 cents per (unexpected) euro increase in wealth. This effect is driven primarily by a positive consumption response to house prices. In contrast, the effect of a variation in stock prices is statistically insignificant. Second, we find that the consumption response to anticipated changes in wealth is also large and significant, of the same magnitude as the response to unanticipated changes, and similarly driven by changes in housing wealth.

Our study is not the first to find evidence of a housing wealth effect exceeding the stock market wealth effect. Other studies finding similar results include Case et al. (2005, 2011), Bostic et al. (2009), Benjamin et al. (2004), and Campbell and Cocco (2007). Campbell and Cocco (2007) also distinguish between predictable and unpredictable changes and find that consumption responds to both.¹ They interpret the positive and significant effect

¹Campbell and Cocco (2007) differs from our paper because they do not have access to subjective expectations data on house prices. To estimate the effect of predictable wealth changes on consumption they

of predictable wealth changes as an indication that house prices affect consumption by relaxing borrowing constraints, along the lines of the literature on the excess sensitivity of consumption to income changes. As argued above, however, a different interpretation is that a consumption response to anticipated changes in asset prices merely reflects intertemporal substitution. Another paper related to ours is Contreras and Nichols (2010), who distinguish between permanent and transitory shocks to housing returns and find that consumption responds to both, although the effect of permanent shocks is larger.

The rest of the paper is organized as follows. In section 2, we derive an estimation framework that allows us to distinguish between responses to anticipated and unanticipated changes wealth. In section 3 we describe the data and present our empirical strategy, while Section 4 reports and discussed the results. Section 5 concludes.

2 Wealth effects and intertemporal substitution

Wealth effects on consumption are typically estimated by regressing consumption growth (or changes in consumption) on changes in wealth:

$$\Delta C_{it+1} = \alpha + \beta \Delta W_{it+1} + X'_{it+1} \gamma + \varepsilon_{it+1} \quad (1)$$

Differencing takes care of issues arising from omission of unobservable variables such as risk aversion or discount factor, which might vary systematically across the wealth distribution and contaminate estimation of the true relationship between consumption and wealth. There are several studies that take an equation like (1) as a starting point for a wealth effect analysis with micro data, such as Poterba (2000), Dynan and Maki (2001), Juster et al. (2006) and Christelis et al. (2011).

regress changes in consumption growth on house price growth and instrument house price growth with lagged values. To estimate the effect of unpredictable changes they regress consumption growth on the residual of their first-stage IV regression. Correctly separating anticipated from unanticipated wealth changes depends on the (strong) assumption that the econometrician conditions on the same information set as the individual.

Nevertheless, there are a number of problems with this regression. First, it is not clear that a regression of the change in consumption on the change in wealth measures the "wealth effect". In fact, changes in wealth arise from two different types of variation: (a) changes in the price of assets, for given portfolio composition, and (b) changes in portfolio composition, for given asset prices. To see this, note that in the presence of multiple assets, the consumer's budget constraint is defined by:

$$\begin{aligned} W_{it} &= \sum_j W_{it}^j = \sum_j p_t^j A_{it}^j \\ \sum_j W_{it+1}^j &= \sum_j R_{t+1}^j W_{it}^j + Y_{it+1} - C_{it+1} \end{aligned}$$

where W is end-of-period total wealth, Y and C are income and consumption, A^j are end-of-period shares of asset j with price p^j and gross return $R_{t+1}^j = \frac{p_{t+1}^j}{p_t^j}$, and W^j is wealth held in asset j . If there is a single asset, of course we have the usual constraint $W_{it+1} = R_{t+1}W_t + Y_{t+1} - C_{t+1}$.

We can decompose the change in wealth across two time periods as follows:

$$\begin{aligned} \Delta W_{it+1} &= \sum_j W_{it+1}^j - \sum_j W_{it}^j \\ &= \sum_j p_{t+1}^j (A_{it+1}^j - A_{it}^j) + \sum_j (p_{t+1}^j - p_t^j) A_{it}^j \\ &= \Delta W_{it+1}^E + \Delta W_{it+1}^X. \end{aligned} \tag{2}$$

The second equality comes from adding and subtracting $\sum_j p_{t+1}^j A_{it}^j$. ΔW_{it+1}^E is the change in wealth that results from portfolio shifts (and hence it is potentially endogenous), while

$$\begin{aligned} \Delta W_{it+1}^X &= \sum_j (p_{t+1}^j - p_t^j) A_{it}^j \\ &= \sum_j r_{t+1}^j W_{it}^j \end{aligned}$$

is the change in wealth that results from asset price changes (which is exogenous and not manipulable), and $r = R - 1$ is the net return.

What is commonly known as wealth effect is the response of consumption to exogenous changes in wealth (i.e., capital gains in housing or stocks). Hence, for the purpose of identifying the wealth effect, we rewrite (1) as:²

$$\begin{aligned}\Delta C_{it+1} &= \alpha + \beta \Delta W_{it+1}^X + X'_{it+1} \gamma + \varepsilon_{it+1} \\ &= \alpha + \beta \sum_j r_{t+1}^j W_{it}^j + X'_{it+1} \gamma + \varepsilon_{it+1}\end{aligned}\quad (3)$$

It is now worth noting that β in (3) captures two different effects. One is intertemporal substitution. If asset prices are expected to increase, consumers will reduce their current consumption and save more. The other effect is the actual wealth effect, i.e., the fact that unanticipated changes in asset prices induce households to modify their consumption. We can decompose the exogenous wealth increase to capture these two effects as:

$$\begin{aligned}\Delta W_{it+1}^X &= \sum_j (p_{t+1}^j - p_t^j) A_{it}^j \\ &= \sum_j (E_t p_{t+1}^j - p_t^j) A_{it}^j + \sum_j (p_{t+1}^j - E_t p_{t+1}^j) A_{it}^j \\ &= \sum_j E_{t-1} r_{t+1}^j W_{it}^j + \sum_j (r_{t+1}^j - E_t r_{t+1}^j) W_{it}^j \\ &= \Delta W_{it+1}^{XA} + \Delta W_{it+1}^{XU}\end{aligned}\quad (4)$$

where the second equality comes from adding and subtracting $E_t p_{t+1}^j$. Here ΔW_{it+1}^{XA} and ΔW_{it+1}^{XU} denote the anticipated and the unanticipated change in wealth, respectively. We can then rewrite equation (3) as:

$$\Delta C_{it+1} = \alpha + \beta_A \Delta W_{it+1}^{XA} + \beta_U \Delta W_{it+1}^{XU} + X'_{it+1} \gamma + \varepsilon_{it+1}\quad (5)$$

which allows for potentially different responses to anticipated and unanticipated wealth

²In some studies, researchers study the wealth effect associated to different types of assets, i.e., estimate:

$$\Delta C_{it+1} = \alpha + \sum_j \beta_j r_{t+1}^j W_{it}^j + X'_{it+1} \gamma + \varepsilon_{it+1}$$

where β_j measures the wealth effect associated to asset type j (housing, stocks, etc.).

changes. In this framework β_U captures the "pure" wealth effect on consumption. Regressions (1) and (3) may be unable to recover this parameter.

Unlike the distinction between anticipated and unanticipated income effects on consumption, where unanticipated income changes shift consumption but anticipated ones do not, in the wealth case both anticipated and unanticipated changes affect consumption. This can be seen clearly in an Euler equation framework:

$$\Delta C_{it+1} = \frac{1}{\gamma} (E_t r_{t+1} - \delta) + \xi_{it+1}.$$

Consumption responds both to expected changes in asset prices ($E_t r_{t+1}$), which determine the relative price of present and future consumption (the first term of 4), and to shocks to wealth induced by changes in prices (the second term of 4), which are included in the innovation term ξ_{it+1} . The parameter β_A in (5) is related to the effect of $E_t r_{t+1}$ on ΔC_{it+1} , while β_U is related to the effect of ξ_{it+1} on ΔC_{it+1} . As a consequence, estimation of a regression like (3) will yield a biased estimate of the wealth effect of consumption, with the sign of the bias depending on the magnitude of the wealth effect relative to the size of the elasticity of substitution.

3 Data

We use data from the Survey of Household Income and Wealth (SHIW), a representative survey of the Italian population. The SHIW is run bi-annually, and about half of the households are re-interviewed in the following survey. The survey collects detailed data on household consumption, income, wealth and portfolio composition, as well as demographic characteristics. We use the 2008 and 2010 surveys which include subjective expectation data on asset returns. Specifically, the survey collects individual expected returns for three broad asset classes: (a) safe assets; (b) stocks; and (c) housing. The survey technique that is used to obtain these expectations is similar to that discussed in Manski (2004), and consists of

eliciting information about two points of the subjective cumulative density function. For example, in the safe asset case household heads are first asked to report the chances that in a year's time the interest rate will be higher than today's, or $\Pr(r_{t+1}^f > r_t^f | I_{it})$ (where I_{it} is the respondent's information set at time t). Next, they are asked to report the chances that the rate will exceed today's rate by more than 1 percentage point (i.e., $\Pr(r_{t+1}^f > r_t^f + 0.01 | I_{it})$). In the stocks case, the two questions are $\Pr(r_{t+1}^s > 0 | I_{it})$ and $\Pr(r_{t+1}^s > 0.1 | I_{it})$. In the house price case, the question was asked only in 2010 and formulated slightly differently, as follows: $\Pr(r_{t+1}^H < 0 | I_{it})$ and $\Pr(r_{t+1}^H < -0.1 | I_{it})$.³ The answers to these questions allow us to characterize the distribution of expectations of future asset returns at the individual level.

The subjective expectations questions were asked to the entire sample in 2008 and to a randomly selected subsample (about half of the overall sample) in 2010. On average, around 45% of household heads answer the first of the two questions. The rest reported a "do not know" answer. While the non-response rate is high, it is comparable to the response rate obtained in other parts of the survey when asking questions involving a subjective judgement (such as lottery questions designed to measure risk aversion or intertemporal discounting). The high rate of non-response may be due to the complexity of the question. Non-responses may also reflect the fact that the subjective expectation questions were asked without preparing the respondents with a set of "warm up" questions. Finally, non-response may also reflect extreme uncertainty. Below, we present two sets of results: (a) we exclude the sub-sample answering "do not know", and (b) we impute expected returns using a model of expectation formation (as described in Section 3.1.3).

Table 1 reports the distributions of subjective expectations of asset returns, excluding cases where individuals responses imply a declining c.d.f., i.e., individuals who report

³The exact wording of the three questions is in the Appendix. Note that it is only in the safe asset case that people are asked to report expectations about future interest rates. In the two other cases, people are asked to report expectations about prices (of stocks and housing, respectively). We convert expectations about prices into expectations about returns using $R_{t+1}^j = \frac{p_{t+1}^j}{p_t^j}$.

$\Pr(r_{t+1}^f > r_t^f + 0.01 | I_{it}) > \Pr(r_{t+1}^f > r_t^f | I_{it})$ (15 percent of the total). For stocks and housing, we drop 6 percent and 10 percent of the sample, respectively. In Panel A of Table 1, we report the distribution of $\Pr(r_{t+1}^f > r_t^f | I_{it})$ (first column) and of $\Pr(r_{t+1}^f > r_t^f + 0.01 | I_{it})$ (second column). Note that in the first column we report the unconditional distribution, while in the second column we report the conditional distribution, as the follow up question was only asked to those who answered the first question and did not report $\Pr(r_{t+1}^f > r_t^f | I_{it}) = 0$. When asked about the chances of an increase in interest rates, 25% of households assigned a positive chance. Of these, 12% gave a zero chance to the event of an interest rate increase of one percentage point or more. Panel B and C repeat the same analysis for stock market returns and house prices. When asked about a stock market gain, 28% of households assigned a positive chance to that event. When asked about housing, 31% of households expected a drop in prices.

Studies of probabilistic expectations have pointed out that responses to such questions exhibit rounding to focal values, such as 5%, 10% and 25%. In addition, there is commonly heaping in responses at values of 0%, 50%, and 100%. We observe a similar phenomenon in our data (see Figure 1, where we plot the response distribution to the question on a positive stock market return), even though it seems less severe than in other surveys.⁴

3.1 Empirical strategy

While the survey we use include subjective expectations of asset returns which are rarely collected in survey data, the data have also some limitations. First, since we observe only two points of the cumulative density function, we need to impose distributional assumptions in order to recover the expected value of asset returns from the data; second, data are bi-annual; third, there is a timing discrepancy between the reported value of the stock of assets (which refers to the end of calendar years t and $t + 2$) and expected returns (which are collected at the time of the interview, typically in the middle of calendar years $t + 1$ and

⁴Response distributions for the other two asset classes look qualitatively similar.

$t + 3$); finally, as remarked above, there is non-negligible non response on the subjective expectation questions.

We now discuss how we tackle these four issues. Whenever possible, we test for our assumptions or conduct robustness checks.

3.1.1 Distributional assumptions

The responses to the probabilistic expectations questions can be used to fit individual specific subjective distributions. To compute the first two moments of these distributions, we need to make assumptions about the underlying density. We assume that household's i expectations for the return on asset j are normally distributed with mean $E_t r_{t+1}^j$ and variance $var_t r_{t+1}^j$ (where $E_t x = E(x|I_{it})$ and $var_t x = var(x|I_{it})$). In practice, each household head in the sample is asked to report:

$$\begin{aligned}\Pr(r_{t+1}^j > \alpha^j | I_{it}) &= \Phi\left(\frac{E_t r_{t+1}^j - \alpha^j}{\sqrt{var_t r_{t+1}^j}}\right), \\ \Pr(r_{t+1}^j > \beta^j | I_{it}) &= \Phi\left(\frac{E_t r_{t+1}^j - \beta^j}{\sqrt{var_t r_{t+1}^j}}\right)\end{aligned}$$

where r^j denotes the return on financial asset j ($j = f, s$), and $\Phi(\cdot)$ denotes the c.d.f. of the standard normal distribution. In the safe asset case, $\alpha^f = r_t^f$ and $\beta^f = r_t^f + 0.01$; in the stocks case, $\alpha^s = 0$ and $\beta^s = 0.1$. In the house price case, people are asked:

$$\begin{aligned}\Pr(r_{t+1}^H < \alpha^H | I_{it}) &= 1 - \Phi\left(\frac{\alpha^H - E_t r_{t+1}^H}{\sqrt{var_t r_{t+1}^H}}\right), \\ \Pr(r_{t+1}^H < \beta^H | I_{it}) &= 1 - \Phi\left(\frac{\beta^H - E_t r_{t+1}^H}{\sqrt{var_t r_{t+1}^H}}\right)\end{aligned}$$

and $\alpha^H = 0$ and $\beta^H = -0.1$.

We observe the probabilities on the left hand side from subjective reports, and α^j and β^j are either constant or depend on r_t^f which we set equal to the actual value observed in the

year of the interview. This hence becomes a system of two equations in two unknowns that can be solved for $E_t r_{t+1}^j$ and $var_t r_{t+1}^j$. Note that, in order to estimate $(E_t r_{t+1}^j, var_t r_{t+1}^j)$, we can only use respondents who answer both questions on the expected return on asset j . If more than two questions were available, one could improve the precision of the estimates or fit more flexible distributions. Moreover, the system would be over-identified. One important question is whether the assumption of normally distributed returns is appropriate. This assumption is clearly strong, but as the actual distribution of the Italian FTSE MIB returns shown in Figure 2 suggests, it is not unreasonable.⁵

In the safe asset case, the identification of the reference return r^f is somewhat complex, as the survey question makes reference to no specific safe asset (it just refers generically to the "interest rate"). We assume that the reference return is the one that investors would earn on a basket composed of bank deposits and government bills and bonds, whose returns have moved in parallel until the end of 2010. We use the average before-tax return on deposits at the end of 2008 (1.7%) and the end of 2008 return on a basket of government bonds of different maturity (4.4%). For stocks and housing, no knowledge of returns is required as households are asked the probability of a gain (a loss for housing), and the probability that the gain (loss) is 10 percent or more.

3.1.2 Bi-annual data

The regression equation (5) assumes access to annual data. However, the SHIW data are collected every other year (2008 and 2010 in our specific case). Hence we observe consumption and wealth data for 2008 and 2010 ($C_{i,08}$, $C_{i,10}$, $W_{i,08}$ and $W_{i,10}$), and one-year ahead expected returns $E_{08} r_{09}$. We adapt our estimation framework to the timing of data collection. To see how we get the equivalent of equation (5) in the bi-annual data case, start by rewriting equation (3) for a single asset in terms of the frequency of our data (omitting controls for

⁵Dominitz and Manski (2011) also make a normality assumption.

brevity):

$$\begin{aligned}\Delta C_{i,10} &= \alpha + \beta r_{10} W_{i,09} + \varepsilon_{i,10} \\ \Delta C_{i,09} &= \alpha + \beta r_{09} W_{i,08} + \varepsilon_{i,09}\end{aligned}$$

Summing up the two equations (and assuming that asset holdings in 2009 are approximately equal to those in 2008, as we do not have any information about asset holdings in 2009),⁶ we obtain:

$$\begin{aligned}C_{i,10} - C_{i,08} &= \tilde{\alpha} + \beta (r_{10} W_{i,09} + r_{09} W_{i,08}) + \varepsilon_{i,10} + \varepsilon_{i,09} \\ &= \tilde{\alpha} + \beta (p_{10} - p_{08}) A_{i,08} + \varepsilon_{i,10} + \varepsilon_{i,09} \\ &= \tilde{\alpha} + \beta ((1 + r_{09})(1 + r_{10}) - 1) W_{i,08} + \varepsilon_{i,10} + \varepsilon_{i,09}\end{aligned}$$

which is the equivalent of (3).

We next distinguish between anticipated and unanticipated wealth effects and write:

$$\begin{aligned}C_{i,10} - C_{i,08} &= \tilde{\alpha} + \beta_U [((1 + r_{09})(1 + r_{10}) - 1) - E_{08}((1 + r_{09})(1 + r_{10}) - 1)] W_{i,08} \\ &\quad + \beta_A E_{08}((1 + r_{09})(1 + r_{10}) - 1) W_{i,08} + \varepsilon_{i,10} + \varepsilon_{i,09}\end{aligned}$$

Note that we do not observe $E_{08}r_{10}$, the two-year-ahead price or return expectation. Assume that individuals know that annual returns follow an AR(1) process, i.e.,

$$r_t = \rho r_{t-1} + \xi_t$$

We can estimate ρ from data, and use the law of iterated expectations to write:

$$E_{08}r_{10} = \rho E_{08}r_{09}$$

so that:

$$E_{08}((1 + r_{09})(1 + r_{10}) - 1) \approx (1 + \rho) E_{08}r_{09}, \quad (6)$$

⁶This is an assumption that may be acceptable for housing, business wealth, and for other financial assets in the presence of inertia or adjustment costs.

if the term $r_{09}r_{10}$ is negligible. Since ρ is pre-estimated, we bootstrap the standard errors. Hence our estimating equation becomes:

$$\begin{aligned} C_{i,10} - C_{i,08} = & \tilde{\alpha} + \beta_U [(r_{09} + r_{10}) - (1 + \rho) E_{08}r_{09}] W_{i,08} \\ & + \beta_A (1 + \rho) E_{08}r_{09} W_{i,08} + \varepsilon_{i,10} + \varepsilon_{i,09} \end{aligned} \quad (7)$$

which is the equivalent of (5) adapted to the bi-annual data case.

3.1.3 Timing discrepancy

Interviews for the SHIW are typically conducted between January and October, while consumption and wealth refer to the previous calendar year. At the time of the interview, households report their expectations about asset returns over a one-year horizon. This means that while the ideal expectation of the return would be $E_{08:12}r_{09:12}$ (i.e., the expected 1-year return elicited at the end of 2008), we have instead $E_{09:m}r_{10:m}$, where m is the month of the interview. Expectations provided in the middle of year 2009 may contain new information (e.g., monetary policy intervention) released between the end of the previous calendar year 2008 and the time of the interview. This timing discrepancy may therefore induce a spurious correlation with the error term of (7). To address this issue we model expectation formation (as we illustrate below) and correct for the timing discrepancy. Our expectation formation model also allows us to impute expected returns to those who do not answer the survey questions.

Let $E_{09:m}r_{10:m}$ denote household i expectation of one year return r , with m denoting the month of the interview. We assume that subjective expectations of returns are a function of a set of demographic controls that are constant or evolve deterministically over time and of past actual returns, as follows:

$$E_{09:m}r_{10:m} = \gamma_0 + \sum_{\tau=1}^T \gamma_{\tau} r_{09:m-\tau} + \gamma_x X_i + \nu_i \quad (8)$$

We set $T = 6$. Predicted subjective expectations of annual returns at the end of 2008 are obtained using:

$$\widehat{E}_{08:12}r_{09:12} = \widehat{\gamma}_0 + \sum_{\tau=1}^T \widehat{\gamma}_{\tau}r_{08:12-\tau} + \widehat{\gamma}_x X_i, \quad (9)$$

where $r_{08:12-\tau}$ denotes the return in month 2008:12- τ . Clearly, the richer X_i , the greater the variability of predicted values.

In practice, we estimate the expectation model in (8) using subjective expectations of returns on stocks, and subjective expectations of returns on deposits and on bonds, available from the 2008 survey. The survey does not ask expectations of house prices. Hence, we retrieve this information from the 2010 survey to fit the expectation model and then predict expectations as of end of 2008. We assume that house price expectations depend on past prices in the province where the household lives, which we compute averaging self-reported house values from the SHIW. Predictions based on equation (9) are then used to compute the anticipated change in wealth. The difference between predictions as of end of 2008 and realizations in 2009 is used to compute the unanticipated change.

The estimating equation (7) now becomes:

$$\begin{aligned} C_{i,10} - C_{i,08} &= \tilde{\alpha} + \sum_k \beta_U^k \left[(r_{09}^k + r_{10}^k) - (1 + \rho^k) \widehat{E}_{08}r_{09}^k \right] W_{i,08}^k \\ &\quad + \sum_k \beta_A^k (1 + \rho^k) \widehat{E}_{08}r_{09}^k W_{i,08}^k + \varepsilon_{i,10} + \varepsilon_{i,09} \end{aligned} \quad (10)$$

where we have also allowed for the fact that we estimate the wealth effect for k different asset, where k = deposits and bonds, stocks, and real assets.⁷

3.1.4 Non-response

We approach the problem of non-response to the subjective expectation questions in two different ways. First, we analyze behavior of a reduced sample of households who respond to the subjective questions (the "Respondents" sample). Second, we impute expected returns

⁷As return on small firm shares, which are included in the real assets, we take the return on stocks. In fact, between 1995 and 2010, the return of small firm shares, based on the SHIW, tracked closely the return on the FTSE MIB, which is the stock market index of the main stock exchange in Italy.

to non-respondents using the estimates of expectation formation model discussed above (the "Whole sample"). Since non-response may be non-random, we correct our estimates for sample selection in the estimation of the expectation model in (8).

4 Empirical results

4.1 Heterogeneity of individual expectations

Table 2 reports the percentiles of the distributions of the estimated means and standard deviations of one-year ahead expectations of the returns on bank deposits, government bonds and stocks, from the 2008 survey, and housing, from the 2010 survey. When examining the findings it is helpful to have a sense of actual returns in the year preceding and following the elicitation of the expected returns. Hence, in the last two rows of the table, we report also ex-post return realizations in 2008 and 2009.

Estimates exhibit a high degree of heterogeneity. For bank deposits, the interquartile range of the expected return is about 200 basis points. Realized returns on deposits were 2.17% in 2008, and 1% in 2009, on average. Since the average expected return is 0.93%, most respondents expected returns to fall, relative to 2008. Their forecasts were indeed correct, and very close to the actual 1% return. Similar considerations apply to expectations of returns on government bonds. The average expected return was 3.6%, which is very close to the realized 3.54%. For stocks, the median expected return is -4.86%. In 2008, the Italian stock market experienced a dramatic loss, equal to almost 50 percent of its value. The decline continued until March 2009. The table suggests that most respondents expected losses on stocks also in the year ahead, but expected such losses to be much smaller than those of the previous year. Expectations, however, fell short of realizations. In fact, by the end of 2009, the stock market index was up 16.52% with respect to the beginning of the year. Finally, according to SHIW-based estimates, the average return on housing was 1.6% per year in the 2007-2008 period and 1.03% in the 2009-2010 period. In contrast, the median expected

return on housing in 2011 was much higher, at 4.31%.

The comparison between expectations and realizations shows that individual expectations were significantly close to realized returns for assets whose returns were relatively easier to predict, such as deposits or government bonds. For stocks and housing, however, the expectations were not matched by realizations. What matters, of course, is not that people formed imperfect expectations when choosing their consumption and portfolio composition, but that they acted upon such expectations (however imperfect they were).

4.2 Expectations model

Table 3 reports summary statistics of predicted individual expectations based on the estimation of the expectations model in (8). Table A1 in the Appendix reports two-step estimates of the model, where we try to correct for any sample selection due to non-random non-response to the subjective expectation questions. To control for selection, we use three variables based on information provided by the interviewers regarding the general level of understanding of the survey questions, the reliability of the answers on household income, and the general atmosphere in which the interview took place. Besides these three variables, we add a dummy for answering other subjective expectation questions in the survey. These four variables are jointly strongly significant in the probit for the probability of answering the expected returns questions (p-value < 1 percent). The Mills ratio based on this probit regression has a positive and significant coefficient in the expectations model for stocks, bank deposits, and government bonds, which suggests that self-selection is likely to be indeed an issue and lack of control may bias the estimates. Also, the positive coefficient implies that respondents tend to expect higher returns than the average household in the survey. The Mills ratio is not significant in the regression for housing.

Panel A of Table 3 reports predictions for expected returns as of end of 2008 based on equation (9). Panel B displays the ex-post returns realizations for 2009. Finally, Panel C reports expectation errors computed as the difference between return realizations in 2009 and

predicted expectations for 2009. The error is largest for stocks, around 30 percentage points. It is large also for housing, around 10 percentage points. This confirms the descriptive results presented above.

4.3 Wealth effect estimates

Our wealth effect estimates are based on a sample which is selected as follows. First, since we need to observe changes in consumption, we restrict the sample to the panel households, about 60 percent of the 2008 sample. Then, we drop households headed by individuals aged less than 18 or more than 80 (7 percent of sample). To reduce the influence of outliers, we drop households whose consumption halved or doubled between 2008 and 2010 (1 percent), those whose annual saving amounted to more than 10 times their total wealth (2 percent), and those with zero assets (including housing) (2 percent)⁸. Finally, we drop observations with anomalous reports on the subjective expected returns questions (1 percent of our sample). In our regressions, consumption consists of household expenditure on non-durable goods. Total assets are the sum of financial assets, which include end-of-year holdings of bank deposits, government and corporate bonds, and stocks, and real assets, which include end-of-year holdings of real estate (land and buildings) and shares of private businesses.

Tables 4 report the results of the estimation of the wealth effect regression (10), using the approximation in (6) for expectations of returns two years ahead. Estimates of the AR(1) process for the annual returns for such approximation are in Table A2 of the Appendix. The estimated AR(1) coefficient, ρ , ranges from 0.46 for stocks to 0.74 and 0.75 for bank deposits and bonds. For housing, information on past prices is limited and fitting the AR(1) model is not feasible. Since house prices exhibit a high degree of persistency, we set $\rho = 1$. All

⁸While in principle one would need to correct estimates for the self-selection arising from using only asset holders, in our sample this is unlikely to be an issue as only 2 percent of households have non-positive total assets.

regressions include a set of socio-demographic variables, listed in the note to the table, which allows us to control for differences in factors that may vary across the wealth distribution and contaminate the true relationship between changes in wealth and changes in spending.

Estimates in the first six column of the table are based on the "Respondent" sample. Estimates in the last three columns are based on a larger sample that includes non-respondents to the subjective expectations questions. In the first three columns, anticipated and unanticipated gains on financial assets are based on individual expectations as computed directly from the answers to the expectations questions, i.e. disregarding any time discrepancy between date of interview and date to which wealth stocks refer to. Gains on housing are always predicted as of end of 2008 using the expectation model estimated on data from the 2010 survey. We estimate that the overall wealth effect is significant and around 3 cents per unexpected euro increase in total assets. The response of consumption to anticipated changes in wealth is also significant and around 3.5 cents per euro variation. As shown in column (2) and (3), both effects are driven by real assets. In fact, when we distinguish between financial assets (as an aggregate or disaggregated into deposits and bonds, and stocks) and real assets, the effects on consumption of expected and unexpected gains on the former are both statistically insignificant, whereas the effects of gains on the latter are significant and of the same order of magnitude of the effects of gains on total assets. Consumption does not appear to respond to expected gains nor to unexpected gains on stocks even if we restrict the sample to stockholders (regressions available upon request).

Our 3% real wealth effect estimate is in line with the findings of Engelhardt (1996) for the US, regarding consumption response to gains on housing, and the estimates obtained by Disney, Henley and Jevons (2003) for the UK. However, unlike our framework, both papers look at realized gains without distinguishing between anticipated and unanticipated ones. Furthermore, our finding that consumption responses are smaller, if not negligible, for financial assets than for non-financial assets is in line with the evidence from other studies, including Case et al. (2005) and Guiso et al. (2006).

In the rest of the table, anticipated and unanticipated wealth gains are determined using predicted expectations as of end of 2008 for all assets involved, using the strategy discussed in section 3.1.3. Estimates on the "Respondents" sample are very similar to those based on actual expectations. Our results are robust to changes in the specification of the expectation model for housing returns intended to increase the degree of heterogeneity in predicted expectations.

When we extend the analysis to the "Whole sample", we obtain smaller effects: the consumption response to either anticipated or unanticipated wealth changes is about 1.3 cents per euro, but the effects are still entirely driven by real assets. The smaller response could reflect measurement error in the expectation variables. If respondents and non-respondents differ in their expectation formation mechanism and our expectation model does not fully capture such heterogeneity, expectations imputed using respondents data may not be a good description of non-respondent beliefs, leading to some attenuation bias.

5 Conclusions

In this paper, we decompose the wealth effects on consumption into its various component using subjective expectations data. Individual expectations are important determinant of choice and most economic models assign a central role to expectations regarding asset prices, future income and individual mortality. Nevertheless, the collection of expectations data is rare. Given the lack of data, researchers have typically adopted *ad hoc* assumptions and measured individual expectations using past realizations.

In our work, we combine subjective expectations data of asset returns with ex-post return realizations to distinguish between anticipated and unanticipated changes in wealth and investigate the separate consumption response to expected and unexpected changes in asset prices. Moreover, we stress that another important distinction (which is often neglected in the empirical literature on the wealth effect) is that changes in wealth are partly exogenous,

i.e. related to variations in asset prices, and partly endogenous, i.e. related to portfolio shifts. The availability of expectation data on returns allows us also to focus on exogenous changes in wealth.

We find that the consumption response to unexpected exogenous changes in wealth, i.e. the "pure" wealth effect, amounts to around 3 percent. Also, consumption responds to expected changes in asset prices, which we argue reflects intertemporal substitution. Both effects are driven by a positive consumption response to changes in house prices. Why is consumption unaffected by exogenous shocks to stock market returns? A possibility is that the extreme uncertainty surrounding the Italian stock market during our sample period may have induced households to "wait and see" before monetizing gains (or losses), a form of precautionary behavior response. The housing market was also volatile, but local factors induced much more heterogeneity. Moreover, the continuing credit market liberalization process may have resulted in better opportunities to borrow against the (modest) housing wealth gains experienced during our sample period.

A Appendix: The subjective expectation questions

The 2008 and 2010 Italian Surveys of Household Income and Wealth have a section designed to elicit individual expectations of future asset returns. Each participant in the survey is asked a set of probabilistic questions tightly worded along the lines set by Manski in several studies (e.g., Manski, 1990 and 2004). Specifically, the 2008 survey includes the following questions:

1) *On a scale from 0 to 100, what is the likelihood that in a year's time interest rates will be higher than today?*

2) *(If you gave a figure for Question [1]) What is the likelihood that they will be more than 1 percentage point higher?*

3) *On a scale from 0 to 100, what is the likelihood that if you invest in the Italian stock market today you will obtain a profit in a year's time?*

4) *(If you gave a figure for Question [3]) What is the likelihood that your investment will earn more than 10%?*

Respondents can either give a probability or answer "do not know".

Besides these same questions, the 2010 survey includes also the following questions:

5) *On a scale from 0 to 100, what is the likelihood that in a year's time house prices will be lower than today?*

6) *(If you gave a figure for Question [5]) What is the likelihood that they will fall more than 10%?*

References

- [1] Benjamin, John D., Chinloy, Peter and Jud, G. Donald, 2004. Real Estate Versus Financial Wealth in Consumption. *Journal of Real Estate Finance and Economics*, 29(3), pp. 341-54.
- [2] Bertaut, Carol, 2002. Equity Prices, Household Wealth, and Consumption Growth in Foreign Industrial Countries: Wealth Effects in the 1990s. Federal Reserve Board, International Finance Discussion Paper No. 724.
- [3] Bostic, Raphael, Gabriel, Stuart and Painter, Gary, 2009. Housing Wealth, Financial Wealth, and Consumption: New Evidence from Micro Data. *Regional Science and Urban Economics*, 39(1), pp 79-89.
- [4] Campbell, John Y. and Cocco, Joao F., 2007. How do house prices affect consumption? Evidence from micro data. *Journal of Monetary Economics*, 54(3), pp. 591-621.
- [5] Case, Karl E., Quigley, John M. and Shiller, Robert J. , 2005. Comparing Wealth Effects: the Stock Market versus the Housing Market. *Advances in Microeconomics*, 5(1), pp. 1-32
- [6] Case, Karl E., Quigley, John M. and Shiller, Robert J., 2011. Wealth Effects Revisited 1978-2009. NBER Working Papers, no. 16848.
- [7] Contreras, Juan and Nichols, Joseph, 2010. Consumption responses to permanent and transitory shocks to house appreciation. Finance and Economics Discussion Series 2010-32, Board of Governors of the Federal Reserve System (U.S.).
- [8] Christelis, Dimitrios, Georgarakos, Dimitris and Jappelli, Tullio, 2011. Wealth Shocks, Unemployment Shocks and Consumption in the Wake of the Great Recession. CSEF Working Papers 279, Centre for Studies in Economics and Finance (CSEF), University of Naples, Italy.

- [9] Dominitz, Jeff and Charles Manski, 2011. Measuring and Interpreting Expectations of Equity Returns. *Journal of Applied Econometrics*, 26: pp. 352-370.
- [10] Dynan, Karen E. and Maki, Dean M., 2001. Does Stock Market Wealth Matter for Consumption?. Finance and Economics Discussion Series 2001-23, Board of Governors of the Federal Reserve System (U.S.).
- [11] Guiso, Luigi, Paiella, Monica and Visco, Ignazio, 2006. Do Capital Gains Affect Consumption? Estimates of Wealth Effects from Italian Household Behavior. In: L. Klein (ed.), *Long Run Growth and Short Run Stabilization: Essays in Memory of Albert Ando (1929-2002)*, Edward Elgar Publishing.
- [12] Juster, F. Thomas, Lupton, Joseph P., Smith, James P., and Stafford, Frank, 2006. The Decline in Household Saving and the Wealth Effect. *The Review of Economics and Statistics*, MIT Press, 88(1), pp. 20-27.
- [13] Lettau, Martin and Ludvigson, Sydney C., 2001. Consumption, Aggregate Wealth and Stock Returns. *Journal of Finance*, 56, pp. 815-849
- [14] Lettau, Martin and Ludvigson, Sydney C., 2004. Understanding Trend and Cycle in Asset Values: Reevaluating the Wealth Effect on Consumption. *American Economic Review*, 94, pp. 276-299.
- [15] Ludwig, Alexander and Sløk, Torsten, 2004. The Relationship Between Stock Prices, House Prices and Consumption in OECD Countries. *Topics in Macroeconomics*, 4, Article 4.
- [16] Manski, Charles, 1990. The Use of Intentions Data to Predict Behavior: A Best Case Analysis. *Journal of the American Statistical Association*, 85, pp. 934-940.
- [17] Manski, Charles, 2004. Measuring Expectations. *Econometrica*, 72(5), pp. 1329-1376.

- [18] Poterba, James M., 2000. Stock Market Wealth and Consumption. *Journal of Economic Perspectives*, American Economic Association, 14(2), pp. 99-118.
- [19] Sousa, Ricardo M., 2008. Financial Wealth, Housing Wealth, and Consumption. *International Research Journal of Finance and Economics*, 19, pp. 167-191