TI 2022-032/VI
Tinbergen Institute Discussion Paper

Advance Information and Consumption Insurance: Evidence from Panel Data

Marcelo Pedroni[1]
Swapnil Singh[2]
Christian Stoltenberg[1,3]

1 University of Amsterdam
2 Bank of Lithuania
3 Tinbergen Institute

Advance Information and Consumption Insurance: Evidence from Panel Data

Marcelo Pedroni
University of Amsterdam

Swapnil Singh
Bank of Lithuania

Christian A. Stoltenberg
University of Amsterdam
Tinbergen Institute

PRELIMINARY

April 15, 2022

Abstract

We investigate whether US households possess advance information about their future income and what this means for consumption insurance. Based on insights from a theoretical model, we propose a new test to detect advance information, which requires only panel data on consumption and income. Using the Panel Study of Income Dynamics, we find—in contrast to the existing literature—strong support for the existence of advance information. We use this evidence to estimate a standard incomplete markets model and find that advance information reduces households' income forecast errors by 15%. Our estimation results imply that 27% of all unexpected income changes are passed through to consumption. Ignoring advance information leads to a significant overestimation of consumption insurance and even more so at the bottom of the wealth distribution.

Keywords: income risk, advance information, consumption insurance, panel data, incomplete markets.

JEL classifications: C23, D12, D31, D52, D81, E21, G52

1. Introduction

How well can US households can insure their consumption against unexpected income changes? The answer to this question is crucial for policies that aim at providing additional insurance to households. Pivotal for answering the question is to quantify which part of observed income changes are unforeseen and which parts are known to the household in advance. It is common to assume that households do not have advance information on their future income in which case their foreknowledge coincides with that of an econometrician. This assumption is consistent with the findings from Blundell, Pistaferri, and Preston (2008) who—using data from the Panel Study of Income Dynamics (PSID)—argue that households indeed do not possess advance information about their future income.

In this paper, we propose a novel test to detect advance information that is derived from the analytical solution of a simple consumption-savings model. Using this test and the same PSID data, we find strong support for the hypothesis that US households do possess advance information, which allows them to reduce the mean squared errors of their one-year-ahead income forecasts by roughly 15%. Given this income uncertainty, we estimate that households cannot insure 27% of unexpected income changes, which amounts to less consumption insurance than what previous studies that abstract from advance information find.

Our analytical results are derived in a stylized permanent-income economy with certainty equivalence in which households possess private advance information on their future income. As in Guvenen and Smith (2014), households also have access to an informal private insurance scheme to partially insure their consumption against unexpected income changes. We show that households' revise their consumption according to both, current and future changes in income. While the consumption response to current income depends on both advance information and the degree of insurance, consumption reacts to future income changes only in proportion to the amount of advance information that households possess.

These results have immediate implications for two relevant issues. First, they imply that one can, in principle, disentangle advance information and insurance using only panel data on consumption and income growth, an issue raised in Kaufmann and Pistaferri (2009) and, more recently, in Guvenen and Smith (2014). Second, testing for the existence of advance information in the data requires regressing consumption growth *jointly* on current and future income changes. Regressing consumption changes only on future income changes to detect advance information results in an omitted-variable bias if current and future income changes are correlated.

To revisit whether US households possess advance information, we employ the PSID dataset as compiled by Blundell et al. (2008) and run two regressions. First, we regress current consumption growth exclusively on future income growth, and reproduce their findings: the corresponding correlation is not significantly different from zero with a somewhat non-intuitive negative sign, indicating that US households do not possess advance information. The picture changes, however, when we additionally control for current income growth. Consumption growth is then significantly positively correlated with both, current and future income changes, providing evidence that US households do possess advance information on their future income. The reason for the different results is an omitted-variable problem that arises because the first regression does not control for income changes in the current period. Due to mean reversion, current and future income growth are negatively correlated. When omitting current income growth, the correlation of consumption with future income growth is therefore downward biased. These estimates are robust findings. They apply both for the nationally representative subsample, the Survey Research Center (SRC) sample, but also when we additionally include the sample that targets low-income households, the Survey of Economic Opportunity (SEO), or consider a different assumption about the age bracket of the household head.

The dataset compiled by Blundell et al. (2008) imputes the missing non-food categories of non-durable consumption expenditures in the PSID before 1999 from the Consumer Expenditure Survey (CEX). Attanasio and Pistaferri (2014) propose to impute the missing consumption expenditures in the earlier years of the PSID from later years of the same panel when information on consumption expenditures is collected in more detail. For this alternative imputation method, the same picture emerges; the correlation of current consumption growth with current and future income growth is significantly positive, and the latter correlation is negative when we do not control for current income changes.

Similarly to the test proposed by Blundell et al. (2008), our test detects advance information indirectly since it does not rely on individual income expectations data directly, but uncovers advance information from the relationship of current consumption with future income changes. Our main empirical finding that US household possess advance information echoes growing direct evidence that finds a strong correlation between individual expectations and subsequent realizations, even when other information available to an econometrician is taken into account. Dominitz and Manski (1997), Dominitz (1998) and Kaufmann and Pistaferri (2009) provide evidence for a significant relationship between expected and future income. With US data, Dominitz (1998) estimates that individual income expectations reduce income forecast errors between 12 and 21%. With our paper, we therefore reconcile the indirect with the direct evidence on advance information.[1]

The analytical results from the stylized model also clarify why advance information matters for the measurement of consumption insurance. Importantly, the traditional approach to use the correlation of current consumption growth and current income growth is only informative about consumption insurance in the absence of advance information. Otherwise, without any direct information on households' income expectations, estimating the exact amount of insurance requires a structural model. The PSID does not contain direct information on households' income expectations. For this reason, we also consider a more realistic economic model without certainty equivalence and with occasionally binding borrowing constraints. The model is an extension of the standard incomplete-markets model. The first extension is the introduction of advance information, which we model by considering households that receive signals about their future income innovations, similarly to Singh and Stoltenberg (2020). Secondly, we supplement the households self insurance with an informal partial insurance scheme. We structurally estimate this economy to address the following policy-relevant questions: How much do US households know about their future income? How well can households insure against unexpected changes to their income?

As in the stylized economy, we find that the correlation of consumption growth with future income growth identifies the degree of advance information. Advance information reduces the mean forecast error of US households by about 15%. Further, we find that 27% of households' unexpected income changes pass through to consumption changes. Ignoring that US households possess advance information biases this pass through downward by 25%, leading to an overestimation of the degree of consumption insurance. This overestimation is magnified at the bottom of the wealth distribution. While the degree of informal partial insurance cannot be precisely estimated, both advance information and consumption insurance as a whole are.

Standard macroeconomic consumption-savings models typically assume that households have no

[1]The predictive power of subjective expectations for future realizations has not only been demonstrated for earnings but also in other contexts, including the risk of job loss (Hendren, 2017; Campbell, Carruth, Dickerson, and Green, 2007), and probability of leaving unemployment (Mueller, Spinnewijn, and Topa, 2021). Manski (2017) summarizes the state of this literature.

advance information. Noteworthy exceptions are Guvenen and Smith (2014), and Kaplan and Violante (2010). Guvenen and Smith (2014) estimate a life-cycle model for the US economy with households that have advance information on their deterministic income profiles. Kaplan and Violante (2010) investigate whether advance information could bridge the gap between consumption insurance as estimated in Blundell et al. (2008) to the insurance that emerges in a life-cycle standard incomplete markets model with a standard calibration. In our paper, we ask whether households' possess advance information about their future income shocks and what this implies for consumption insurance.

In Section 2, we present a simple permanent-income economy with advance information and informal partial insurance. In Section 3, we provide our main empirical results. In Section 4, we describe a quantitative structural model which we use, in Section 5, to obtain our estimates of advance information, informal partial insurance, and overall consumption insurance. We conclude in Section 6.

2. Advance information and partial insurance: a stylized economy

In this section, we consider a stylized economy in which households have advance information about their future income and have access to an informal partial insurance scheme.[2] In this economy, we establish that advance information and partial insurance can be identified from a simple regression of consumption growth on current and future income changes.

Consider an income-fluctuation model in which infinitely-lived households solve

$$\max_{\{c_t, a_{t+1}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t [u(c_t)], \quad \text{subject to } c_t + a_{t+1} = (1+r)a_t + y_t^d, \quad \text{for all } t \geq 0,$$

and also subject to a no-Ponzi constraint, and with a_0 given. Here c_t, a_t, and y_t^d denote household consumption, asset, and disposable income in period t. The parameter β is the household discount factor, and $u(\cdot)$ is the utility function. The interest rate, r, is exogenous. For simplicity, we assume that $\beta(1+r) = 1$ and that the utility function is quadratic, $u(c) = c - bc^2$.

The initial income level, y_0, is given and future income levels follow a random walk $y_t = y_{t-1} + \eta_t$, with $\eta_t \sim \mathcal{N}(0, \sigma^2)$. Households observe past and current income levels and have advance information

[2] Our formulation of advance information and partial insurance follows exactly the one in Guvenen and Smith (2014) who consider a two-period version of the model.

about next period's income, so that

$$\mathbb{E}_t[y_{t+1+j}] = (1-\kappa)y_t + \kappa y_{t+1}, \quad \text{for all } j \geq 0.$$

The parameter $0 \leq \kappa \leq 1$ controls the degree of advance information; when $\kappa = 0$, there is no advance information and $\mathbb{E}_t[y_{t+1}] = y_t$, in line with the random-walk assumption. When $\kappa = 1$, households know their next period income with certainty, that is $\mathbb{E}_t[y_{t+1}] = y_{t+1}$.

Informal partial insurance, with parameter $0 \leq \theta \leq 1$, brings disposable income, y_t^d, closer to the income level expected by the household, so that

$$y_{t+1}^d = (1-\theta)y_{t+1} + \theta \mathbb{E}_t[y_{t+1}].$$

With $\theta = 0$ there is no informal partial insurance, whereas with $\theta = 1$ disposable income in the next period is not at all risky. With these definitions, we can establish the following proposition.

Proposition 2.1. *Optimal consumption-savings decisions by the households imply*

$$c_t - c_{t-1} = (1-\kappa)\left(1 - \frac{\theta r}{1+r}\right)(y_t - y_{t-1}) + \frac{\kappa}{1+r}(y_{t+1} - y_t), \quad \textit{for all } t \geq 1. \tag{2.1}$$

Proof. See Appendix A. □

It follows from Proposition 2.1 that it is, in principle, possible to disentangle advance information from informal partial insurance using only data on consumption and income changes. In particular, controlling for changes in current income, $y_t - y_{t-1}$, a future income surprise, $y_{t+1} - y_t$, only affects current consumption, $c_t - c_{t-1}$, in proportion to how much advance information the household has, which is intuitive. As a result, κ can be identified with that coefficient alone. Given κ, the coefficient on current income changes can then be used to identify θ.[3] Unless current and future income changes are uncorrelated, testing for the existence of advance information in the data requires regressing consumption changes jointly on current and future income changes.[4]

Following Mace (1991), consumption insurance is traditionally measured by regressing consump-

[3]To obtain equation (2.1), the model must have at least three periods, which is why the two-period version of the model in Guvenen and Smith (2014) cannot not deliver it.

[4]Whether current and future income growth are correlated depends on the stochastic process for residual income in the data. For example, if income follows a random walk as in our stylized economy, the correlation is zero, but if income is persistent mean reversion implies that the correlation is negative.

tion on income growth.[5] Equation (2.1) clarifies that this measure represents the degree of consumption insurance only in the absence of advance information, that is, if $\kappa = 0$; otherwise, the correlation between consumption and income growth is the result of both advance information and consumption insurance. Notice that equation (2.1) can also be written as follows

$$c_t - c_{t-1} = \left(1 - \frac{\theta r}{1+r}\right)(y_t - \mathbb{E}_{t-1}[y_t]) + \frac{\kappa}{1+r}(y_{t+1} - y_t), \tag{2.2}$$

where the expectation operator takes into account that households have advance information. Thus, advance information affects the pass through of unexpected income changes, $y_t - \mathbb{E}_{t-1}[y_t]$, to consumption via its effect on income expectations directly, but also since future income changes are an additional—and correlated—source of consumption changes. Equation (2.2) implies two possibilities to measure consumption insurance in the presence of advance information. One approach is to use direct evidence on household subjective income expectations. The second approach, which is the approach we propose in this paper, does not require expectations data. In the first step, advance information is identified from the correlation of current consumption with future income growth. Given advance information, households' income expectations follow from a structural model of expectation formation, which pins down the "true" income surprises from a household perspective, $y_t - \mathbb{E}_{t-1}[y_t]$, and the first coefficient then exclusively measures consumption insurance.[6] The PSID is one of the longest-running and largest income panel data sets in the world, but does not include information on households' income expectation, which is why we follow the second approach in the following sections.

In the next section, we use the theoretical insights from this section to revisit the empirical evidence on whether US households possess advance information about their future income.

3. Empirical Results

The objective of this section is to empirically detect whether US households possess advance information about their future income using panel data on consumption and income. For our baseline results

[5] The simple linear-quadratic model analyzed above can be thought of as an approximate version of a more realistic model in which the level variables that appear in equation (2.1) would be replaced by log deviations from steady state, so that over-time differences become growth rates.

[6] For the simple model in this section, with certainty equivalence, there is a one-to-one mapping between the degree of informal partial insurance, θ, and overall consumption insurance, which is captured by the term $r\theta/(1+r)$ in equation (2.2). In the more realistic model presented in Section 4, households also insure consumption with precautionary savings.

we use the PSID dataset constructed by Blundell et al. (2008). As robustness exercises, we consider an alternative consumption imputation method proposed by Attanasio and Pistaferri (2014) and apply different sample selection criteria. As our main empirical result, we find that US households possess significant amounts of advance information about their future income.

For our baseline, we choose the PSID dataset constructed by Blundell et al. (2008). The key variables of interest are after-tax household income net of asset income and non-durable household consumption expenditures. Both variables are residual variables that follow after controlling for a vector of observable household characteristics.[7]

3.1 Testing for advance information: baseline estimation results

One take-away of the previous section is that testing for advance information requires estimating the correlation of current consumption and future income changes. Throughout this section, we therefore compare the estimation results from two different regression equations. In the first equation, we follow an approach similar to the one in Blundell et al. (2008) and regress current consumption growth exclusively on future income growth,

$$\Delta\log(c_{it}) = \widetilde{\beta}_0 + \widetilde{\beta}_{\Delta y_{t+1}}\Delta\log(y_{it+1}) + \tilde{\epsilon}_{it}, \tag{3.1}$$

with $\Delta\log(x_{it}) \equiv \log(x_{it}) - \log(x_{it-1})$ denoting the growth rate of variable x. A coefficient $\widetilde{\beta}_{\Delta y_{t+1}}$ significantly different from zero is intended to detect advance information. We refer to this coefficient as the unconditional regression coefficient because it is estimated without conditioning on income growth in the current period. In the second regression, we follow the theoretical results from the previous section and regress current consumption growth on both future and current income growth,

$$\Delta\log(c_{it}) = \beta_0 + \beta_{\Delta y_t}\Delta\log(y_{it}) + \beta_{\Delta y_{t+1}}\Delta\log(y_{it+1}) + \epsilon_{it}. \tag{3.2}$$

Here, a conditional coefficient $\beta_{\Delta y_{t+1}}$ that is significantly different from zero exclusively signals advance information, while the conditional coefficient $\beta_{\Delta y_t}$ can be affected by both consumption insurance and advance information.

In Table 1 and Figure 1, we report our baseline estimates of the two regressions. For each model, the

[7]A detailed description of the data, the sample selection and the corresponding methods used for imputing consumption data are described in detail in Appendix B.

Table 1: Baseline estimates: consumption growth regressions

	Year fixed effects		Year + Household fixed effects	
	(1)	(2)	(3)	(4)
Δy_{t+1}	−0.013	0.045**	−0.020	0.046**
	(0.018)	(0.019)	(0.019)	(0.019)
Δy_t		0.184***		0.185***
		(0.021)		(0.023)
Observations	10502	10471	10423	10391

Source: Panel Study of Income Dynamics 1978–1992
Description: The table reports the result of regressing current consumption growth on future income growth, including or excluding current income growth. Baseline specification in Columns (1)-(2) takes year-fixed effects, and Columns (3)-(4) additionally household fixed effects into account.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

first column displays the correlation for the regression equation (3.1), while the second column shows the results for the equation (3.2). The main messages are the following. The first regression equation yields a (non-significant) negative correlation between current consumption and future income growth. The second regression, however, indicates a significant positive correlation. These estimations results apply with year fixed effects in Columns (1)–(2), henceforth baseline specification, but also hold when we additionally include household fixed effects in Columns (3)–(4).[8] The coefficient $\beta_{\Delta y_{t+1}}$ is precisely estimated with p-values of 0.014 (year fixed effects) and 0.017 (year and household fixed effects).[9]

3.2 Why do the estimated coefficients of consumption on future income growth differ?

To fix ideas, assume that, as in the theoretical model, consumption growth depends only on current and future income growth. Then, the relationship between the correlation coefficients of consumption

[8] In their Table 5, Blundell et al. (2008) test whether the unconditional covariances satisfy $\text{cov}(\Delta c_{it}, \Delta y_{it+j}) = 0$, for all years t, and $j \geq 1$. They cannot reject the null hypothesis with p-values of at least 25%, indicating no evidence for advance information. In the first column of Table 9 in Appendix D, we replicate their tests results. In the second and third columns, we also test $\text{cov}(\Delta c_{it}, \Delta y_{it+j}) = 0$, for all t, $j \geq 1$, but condition the covariances on previous Δy_{t+j}'s. We safely reject the null hypothesis $\text{cov}(\Delta c_{it}, \Delta y_{it+1}) = 0$, for all t. Hence, we find strong evidence for the existence of advance information and confirm our main findings as displayed in Table 1.

[9] The estimation results are robust to including asset income. In Table 8 in Appendix D, we provide the corresponding regression table.

Figure 1: Consumption growth and future income growth

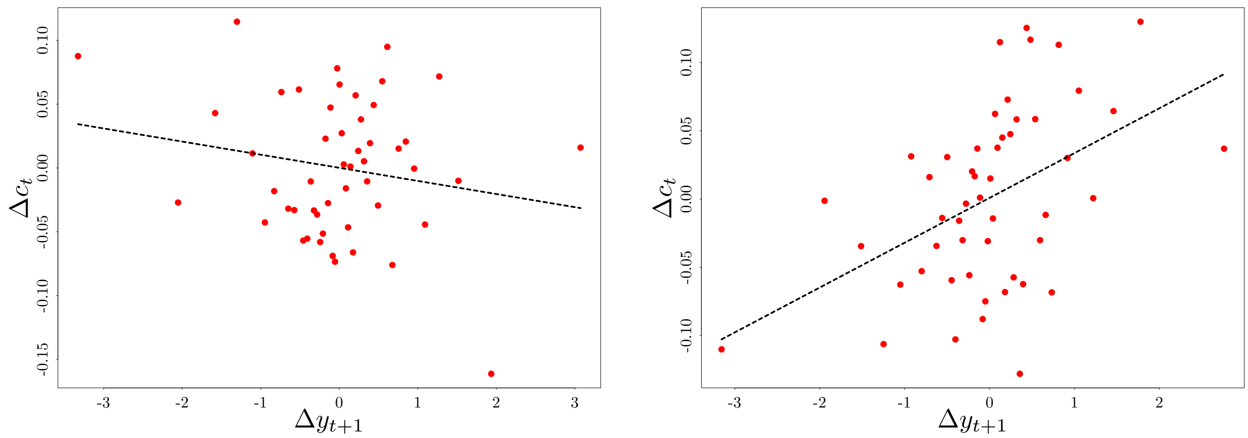

(a) Regression (3.1), excluding Δy_t (b) Regression (3.2), including Δy_t

Source: Panel Study of Income Dynamics 1978–1992
Description: Baseline specification. The figure plots the non-parametric relationship, binned scatter plots, between current consumption and future income growth when current income growth is included or excluded. In this sample, consumption is imputed using the CEX data. The left panel refers to the case when current income growth is excluded while estimating the relationship between current consumption growth and future income growth. In the right panel, current income growth is included. Each binned scatter plot is constructed using 50 equal-sized bins.

growth with future income growth in the first and second regressions is given by the following

$$\widetilde{\beta}_{\Delta y_{t+1}} = \beta_{\Delta y_{t+1}} + \delta \beta_{\Delta y_t}, \tag{3.3}$$

with δ being the coefficient that results from regressing current income growth on future income growth. Equation (3.3) states, given that $\beta_{\Delta y_t} \neq 0$, only if current and future income growth are uncorrelated ($\delta = 0$), is the estimated correlation of current consumption with future income growth identical for the two regression equations.

As illustrated in Figure 2, this is not the case in the PSID data. Instead, mean reversion implies that current and future income are significantly negatively correlated with $\delta = -0.325$, precisely estimated with a standard error of 0.011. Given the negative correlation and the fact that $\beta_{\Delta y_t}$ is significantly positive in all regressions shown in Table 1, it follows that $\widetilde{\beta}_{\Delta y_{t+1}} < \beta_{\Delta y_{t+1}}$. Thus, the correlation between current consumption and future income growth in the first regression is downward biased by $\delta \beta_{\Delta y_t}$ as a result of omitting current income growth as a regressor.

The traditional approach to measure consumption insurance is to regress current consumption

Figure 2: Current income growth and future income growth

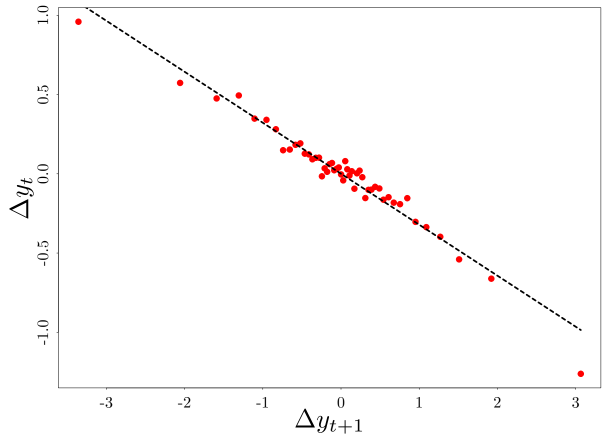

Source: Panel Study of Income Dynamics 1978–1992
Description: Baseline specification. The figure plots the non-parametric relationship, binned scatter plots, between current and future income growth. Coefficient from regressing current income growth on future income growth, $\delta = -0.325$, standard error of 0.011. The binned scatter plot is constructed using 50 equal-sized bins.

growth on current income growth,

$$\Delta \log(c_{it}) = \widehat{\beta}_0 + \widehat{\beta}_{\Delta y_t} \Delta \log(y_{it}) + \widehat{\eta}_{it}. \tag{3.4}$$

Following similar arguments as above, an omitted-variable bias also affects this measure when households possess advance information because future income growth is omitted in equation (3.4). The resulting bias is

$$\widehat{\beta}_{\Delta y_t} - \beta_{\Delta y_t} = \hat{\delta}\beta_{\Delta y_{t+1}} < 0, \tag{3.5}$$

where $\hat{\delta}$ denotes the coefficient that results from regressing future income on current income growth. As a result, we estimate a downward-biased coefficient $\widehat{\beta}_{\Delta y_t} = 0.151 < 0.184 = \beta_{\Delta y_t}$, which would lead to an overestimation of consumption insurance.

As an intermediate summary, we find that, controlling for current income growth, consumption is positively correlated with future income growth, which indicates that US household possess advance information on their future income. In what follows, we investigate whether this finding is robust to employing an alternative method for imputing consumption and different sample selection criteria.

3.3 Testing for advance information: alternative imputation and sample selection

Attanasio and Pistaferri (2014) propose an alternative procedure to impute consumption expenditures that relies solely on the information provided in the PSID. In a nutshell, the authors use the more detailed information on consumption expenditures in later years of the PSID to impute consumption expenditures in earlier years.

As in our baseline, we compare the estimation results of the two regressions (3.1)-(3.2). To compare the estimation results for both imputation procedures, we standardize the variables in our regressions so that all variables have a unit standard deviation and a mean of zero. The standardized regression coefficients are displayed in Table 2, where the upper panel contains the estimation results for the baseline and the lower panel the results for Attanasio and Pistaferri (2014)'s imputation.

The main messages are the following. For both imputation methods, the regression coefficient of consumption growth with respect to future income growth is negative in the first regression, but positive in the second regression. For the second regression, the correlation of current consumption with future income growth is significantly positive for all but one specification at least at the 5% level. For one specification, the correlation is positive at the 10% level (with a p-value of 0.067). Furthermore, all regression coefficients are very similar across the different samples and both imputation procedures.

The analytical results in Proposition 2.1 clarify why it matters whether household have advance information on their future income. Only if they do not, is the traditional approach valid to measure consumption insurance with the covariance of current consumption growth and current income growth. Nevertheless, the main take-away from this section is that households do possess advance information on their future income such some income changes are already known in advance. The theoretical model from the previous section illustrates, in equation (2.2), that consumption insurance can then be identified from the consumption response to unexpected income changes, $y_{it} - \mathbb{E}_{it-1}[y_{it}]$, with expectations that take into account households' advance information. This resembles insights from Jappelli and Pistaferri (2010) who also advocate for estimating the correlation of consumption changes with unexpected income changes. This can be done either directly by employing data on households' subjective income expectations, or indirectly using the income expectations that stem from an estimated theoretical economic model with advance information. The PSID does not include information on households' income expectation, which is why we follow the second approach in the following sections.

Table 2: Standardized consumption growth regressions

Sample:	SEO sample excluded				SEO sample included			
Age group:	30-65		20-65		30-65		20-65	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: CEX Imputation (BPP)								
Δy_{t+1}	**−0.010**	**0.034****	−0.008	0.034***	−0.029**	0.024**	−0.023**	0.029***
	(0.014)	**(0.014)**	(0.013)	(0.013)	(0.012)	(0.011)	(0.010)	(0.010)
Δy_t		**0.138*****		0.133***		0.155***		0.152***
		(0.016)		(0.014)		(0.013)		(0.011)
Observations	**10502**	**10471**	13030	12994	16012	15979	20347	20308
Panel B: PSID Imputation								
Δy_{t+1}	−0.030**	0.022*	−0.030**	0.025**	−0.034***	0.024**	−0.032***	0.026***
	(0.013)	(0.013)	(0.012)	(0.012)	(0.010)	(0.010)	(0.010)	(0.009)
Δy_t		0.151***		0.159***		0.160***		0.162***
		(0.014)		(0.013)		(0.011)		(0.010)
Observations	10073	10058	11911	11895	16335	16316	19594	19573

Source: Panel Study of Income Dynamics 1978-1992
Description: The table reports the result of regressing current consumption growth on future income growth, including or excluding current income growth as regressor. All variables are standardized. Odd numbered columns exclude and even numbered columns include current income growth. The Panel *A* refer to Blundell, Pistaferri and Preston (2008) sample, covering period 1978–1992. In this sample, consumption is imputed using the CEX data. Standardized baseline specification in **boldface**. The Panel *B* also covers the period 1978–1992, but consumption is imputed using PSID consumption data 1999–2015. Columns (1)-(4) exclude SEO samples. Columns (5)-(8) report results when both SRC and SEO samples are included in the regression. In columns (1)-(2), and (5)-(6), household head's age is restricted between 30 and 65 years. In other columns, household head's age is restricted between 20 and 65 years. All regressions include year fixed effects. Standard errors are clustered at the household level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

4. Quantitative model

In this section, we present a quantitative version of the standard incomplete markets model with advance information and informal partial insurance.

There is a continuum of households with time-separable preferences $\mathbb{E}_0\left[\sum_t \beta^t u(c_t)\right]$. In every period, each household exogenously supplies one unit of labor, and receives a labor income (income) of $y \in Y$, and a signal $s \in Y$ about next period income. The household (y, s) pair follows a Markov process specified below. Households can only accumulate a risk-free asset, a, and face borrowing constraints such that the set of possible values for a is given by $A \equiv [\underline{a}, \infty)$. Let $Z \equiv A \times Y^4$, so that households are indexed by their individual states $(a, y, s, \hat{y}, \hat{s}) \in Z$ where $\hat{y}$ and $\hat{s}$ denote the income and signal received by the household in the last period. Given an interest rate r, each household chooses policy functions $c(a, y, s, \hat{y}, \hat{s})$, and $a'(a, y, s, \hat{y}, \hat{s})$ to solve

$$v(a, y, s, \hat{y}, \hat{s}) = \max_{c,a'} u\left(c(a, y, s, \hat{y}, \hat{s})\right) + \beta \sum_{y' \in Y} \sum_{s' \in Y} v\left(a'(a, y, s, \hat{y}, \hat{s}), y', s', y, s\right) \Pr\left(y', s' \mid y, s\right),$$

subject to

$$c(a, y, s, \hat{y}, \hat{s}) + a'(a, y, s, \hat{y}, \hat{s}) = y^d(y, \hat{y}, \hat{s}) + (1 + r)a, \quad \text{and} \quad a'(a, y, s, \hat{y}, \hat{s}) \geq \underline{a},$$

where the household's disposable income is given by

$$y^d(y, \hat{y}, \hat{s}) = (1 - \theta)y + \theta \sum_{\tilde{y} \in Y} \tilde{y} \Pr\left(\tilde{y} \mid \hat{y}, \hat{s}\right).$$

4.1 Equilibrium

We consider the stationary equilibrium of this model defined as follows.

Definition 1. *Given r, a **stationary equilibrium** is a value function v, policy functions c and a', and a stationary distribution λ, such that:*

1. *The policy functions c and a' solve the household's problem, and v is the corresponding value function;*
2. *The stationary distribution λ satisfies*

$$\lambda(\mathcal{Z}) = \int_Z Q\left((a, y, s, \hat{y}, \hat{s}), \mathcal{Z}\right) d\lambda, \quad \textit{for all } \mathcal{Z} \textit{ in the Borel } \sigma\textit{-algebra of } Z,$$

where Q is the transition probability measure consistent with the Markov process for (y, s) and the policy function a'.

4.2 Markov process for income and signals

Household income follows a Markov process with transition probabilities $\Pr(y'|y)$. The signal that the household receives about their next-period income is informative with probability κ, so that

$$\Pr(y'|y,s) = \begin{cases} \kappa + (1-\kappa)\Pr(y'|y), \text{ if } s = y' \\ (1-\kappa)\Pr(y'|y), \text{ otherwise.} \end{cases} \tag{4.1}$$

To conclude the description of the signal process we need to specify the probability of receiving a particular signal. For the signal process to be consistent with the income process,[10] it must be that

$$\Pr(s = y'|y) = \Pr(y'|y), \tag{4.2}$$

that is, the probability of receiving a signal $s = y'$ given current income y must be equal to the probability of transitioning from income y to y'.[11]

5. Quantitative results

In this section, we first define atheoretical measures of advance information and consumption insurance. Afterwards, we describe how we calibrate the income process and other preset parameters of the structural model. We then argue that the logic for identifying advance information and informal partial insurance from Section 2 extends to the more quantitatively realistic model of Section 4. We then estimate the two key parameters: signal precision κ, and informal partial insurance θ. With these estimates, we finally quantify the amount of advance information US households possess and how well they can insure unexpected income shocks.

[10] In particular, the condition in equation (4.2) guaranties that $\sum_{s\in Y}\Pr(s|y)\Pr(y'|y,s) = \Pr(y'|y)$.

[11] Below we allow income to have permanent and transitory components, y_P and y_T respectively. The signal process can be easily extended to this case by combining the two components into one large Markov process. Alternatively, to keep the state space small, one can *equivalently* consider signal pairs (s_p, s_T) such that $\Pr((y'_P, y'_T)|(y_P, s_p, s_T)) = \kappa + (1-\kappa)\Pr(y'_P|y_P)\Pr(y'_T)$, if $(s_p, s_T) = (y'_P, y'_T)$, and $\Pr((y'_P, y'_T)|(y_P, s_p, s_T)) = (1-\kappa)\Pr(y'_P|y_P)\Pr(y'_T)$ otherwise, and $\Pr((s_p, s_T) = (y'_p, y'_T)|y_p) = \Pr(y'_P|y_P)\Pr(y'_T)$.

5.1 Measuring advance information and consumption insurance

Households' advance information reduces their income forecast error relative to an econometrician who predicts future income solely on the basis of current income. A model independent measure of advance information is given by the relative reduction of households income forecast error defined as

$$\tilde{\kappa}(\kappa) = \frac{\mathrm{MSFE}_y - \mathrm{MSFE}_{y,s}(\kappa)}{\mathrm{MSFE}_y}, \qquad 0 \leq \tilde{\kappa}(\kappa) \leq 1 \tag{5.1}$$

where

$$\mathrm{MSFE}_y = \sum_y \pi(y) \sum_{y'} \Pr(y'|y) \{\log(y') - \mathbb{E}[\log(y') \mid y]\}^2$$
$$\mathrm{MSFE}_{y,s}(\kappa) = \sum_{y,s} \pi(y,s) \sum_{y'} \Pr(y'|y,s) \{\log(y') - \mathbb{E}[\log(y') \mid y,s]\}^2 \leq \mathrm{MSFE}_y,$$

$\pi(y)$ is the invariant distribution of income, and $\pi(y,s)$ is the joint invariant distribution of income and signals. In what follows, we use this measure to quantify households' advance information.[12] In addition to its model independence, the measure $\tilde{\kappa}$ has the advantage that we can readily compare our estimates for advance information in Section 5.3 to the direct evidence on the predictive power of individual income expectations.

In the model, households can use self-insurance and informal partial insurance to guard their consumption against unexpected income changes. For this reason, we measure consumption insurance via the consumption-pass through coefficient β_{INS} of the following regression

$$\Delta c_{it} = \text{constant} + \beta_{INS} \{\log(y_{it}) - \log(\mathbb{E}[y_{it} \mid y_{it-1}, s_{it-1}])\} + \epsilon_{it}, \tag{5.2}$$

where a coefficient value of one indicates no insurance and a value of zero indicates full insurance against unexpected income changes.

5.2 Preset and calibrated parameters

The preset and calibrated parameters are organized in Table 3. We set the degree of relative risk aversion and the borrowing limit to $\sigma = 2$ and $\underline{a} = 0$.[13] The interest rate is set to $r = 3.9\%$, the annual real

[12] For our specification of signals in Section 4.2, one can show with tedious but straightforward algebra that $\tilde{\kappa} = \kappa^2$.

[13] In Appendix C, we conduct robustness exercises with respect to these choices.

risk-free rate for the US post 1980 as documented by Jordà, Knoll, Kuvshinov, Schularick, and Taylor (2019). Following the standard practice in the literature, we allow the household income to have a persistent and a transitory component, that is,[14]

$$\log(y_{it}) = \log\left(y_{it}^{P}\right) + \log\left(y_{it}^{T}\right),$$

with

$$\log\left(y_{it}^{P}\right) = \rho \log(y_{P,it-1}) + \epsilon_{P,it}, \quad \epsilon_{P,it} \sim \mathcal{N}(0, \sigma_P^2), \quad \text{and} \quad \log(y_{T,it}) = \epsilon_{T,it}, \quad \epsilon_{T,it} \sim \mathcal{N}(0, \sigma_T^2).$$

We, then, choose the three parameters (ρ, σ_P, and σ_T) to match three moments of income dynamics observed in our benchmark PSID sample with simulated data. The first two moments are the regression coefficient, $\beta_{y_{t-1}} = 0.777$ (s.e. 0.009),[15] and the residual standard deviation, $\sigma_y = 0.285$ (s.e. 0.006), of the following regression equation capturing the dynamics of after-tax income levels

$$\log(y_{it}) = \beta_{y,0} + \beta_{y_{t-1}} \log(y_{it-1}) + \epsilon_{yit}, \quad \epsilon_{yit} \sim (0, \sigma_y^2). \tag{5.3}$$

The empirical section shows that auto-correlation of income growth plays an important role to detect advance information. For this reason, we also target the regression coefficient of current on future income growth, $\delta = -0.325$ (s.e. 0.011). To discretize the persistent component, we use the procedure described in Tauchen and Hussey (1991). We, then, obtain model counterparts for the targeted moments from Monte Carlo simulations.

The remaining parameters—the discount factor β, signal precision κ, and the informal partial insurance parameter θ—are estimated within the model.

5.3 Estimated parameters, identification and implications

For each candidate pair of signal precision κ and informal partial insurance θ, the discount factor β is chosen to match the average wealth-to-income ratio of 2.9 in the sample.[16] In the next step, we estimate

[14] Ejrnæs and Browning (2014) shows that this specification is equivalent to an ARMA(1,1)-process with a single innovation term.

[15] These standard errors are computed from a non-parametric bootstrap with 500 repetitions.

[16] We compute the wealth-to-income ratio using the 1984 and 1989 PSID wealth module. Our measure of wealth includes the total net value of farm or business assets, savings through financial assets, net value of durable assets including housing, and value of trust funds, private annuities, and IRAs. This measure of wealth includes home equity which is defined as the

Table 3: Calibrated parameters of quantitative model

Parameter	Value	Description
ρ	0.890	auto-covariance of persistent component of log income
σ_P	0.216	standard deviation of innovation to persistent component of log income
σ_T	0.100	standard deviation of innovation to transitory component of log income
σ	2.000	degree of relative risk aversion
$\underline{a}$	0.000	borrowing constraint
r	3.9%	interest rate

κ and θ by minimizing the squared deviations of the two regression coefficients from equation (3.2) estimated with simulated data from the model, which we denote by $\beta_{\Delta y_t}(\kappa,\theta)$ and $\beta_{\Delta y_{t+1}}(\kappa,\theta)$, and their counterparts in the PSID data with our baseline specification from Column 2 of Table 1

$$f(\kappa,\theta) = \left[\beta_{\Delta y_t}(\kappa,\theta) - \beta_{\Delta y_t}\right]^2 + \left[\beta_{\Delta y_{t+1}}(\kappa,\theta) - \beta_{\Delta y_{t+1}}\right]^2. \tag{5.4}$$

To be more precise, we target $\beta_{\Delta y_t} = 0.184$ (s.e. 0.021), and $\beta_{\Delta y_{t+1}} = 0.045$ (s.e. 0.019).

Identification. In the simple model from Section 2, we show that advance information and informal partial insurance can be identified from the coefficients of regressing consumption growth on current and future income growth. Here, we investigate whether the identification logic from the simple model extends also to the quantitatively more realistic model.

In Figure 3, we plot absolute deviations of the two regression coefficients estimated on data simulated in the structural model from their survey data counterparts. The right panel of Figure 3 shows that κ is to a large extent already identified by the regression coefficient that captures the covariances of current consumption with future income growth, $\beta_{\Delta y_{t+1}}$. Conditional on κ, the second regression coefficient $\beta_{\Delta y_t}$ identifies informal partial insurance θ. Thus, the identification logic from the simple model in Section 2 also applies in the quantitative model from this section.

Estimation results. In Table 4, we summarize the estimated parameters and the implications for advance information and consumption insurance. Both regression coefficients and the mean of the net-worth-to-income ratio distribution are exactly matched by the model. The discount factor β is precisely estimated with a standard value of 0.944. Signal precision is found to be $\kappa = 0.394$, with a standard

home value minus the outstanding mortgage. The income used to compute the ratio is the total family income.

Figure 3: Identification of signal precision and informal partial insurance

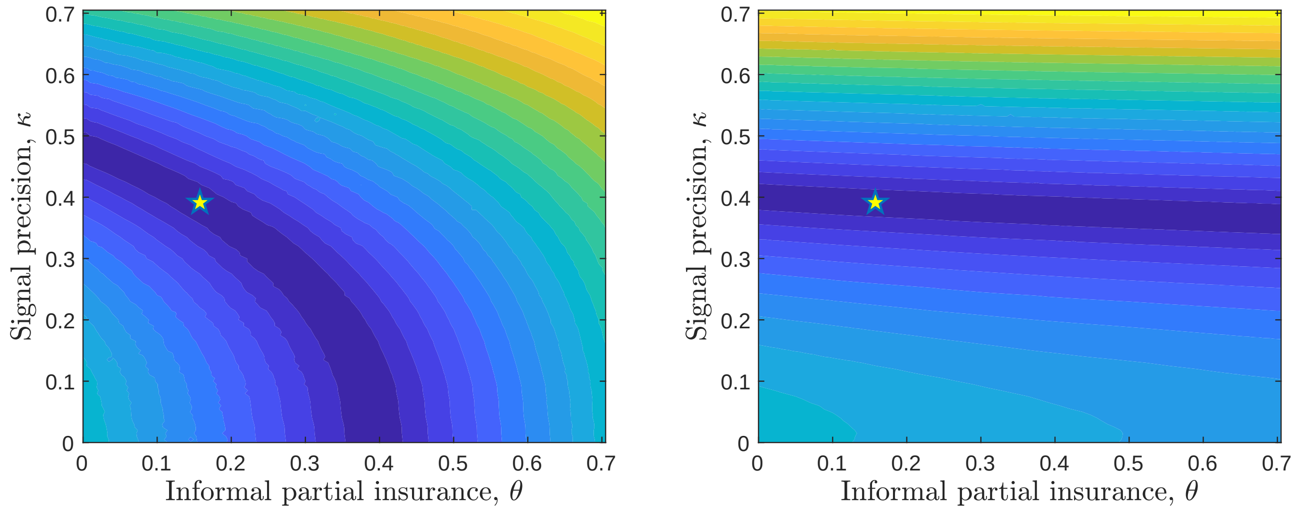

(a) Current income growth coefficient

(b) Future income growth coefficient

Notes: Coefficients resulting from regressing consumption growth on current and future income. Absolute deviations of the regression coefficients estimated in the structural model,$\beta_{\Delta y_t}(\kappa,\theta)$, $\beta_{\Delta y_{t+1}}(\kappa,\theta)$, from survey data estimates as functions of κ,θ. Dark blue (yellow) indicates the lowest (highest) deviation.

error of 0.023. Compared to an econometrician, households' advance information of income shocks reduces their forecast error by $\tilde{\kappa} = 15.3\%$, also precisely estimated. This figure is consistent with the direct evidence of Dominitz (1998) who reports reduction in mean-squared forecast errors when conditioning on income expectations between 12 and 21%. We estimate an informal partial insurance $\theta = 0.158$, though less precisely, with a standard error of 0.111. Taken together, these estimates imply that US households only partially insure shocks to their income with 26.9% of all unexpected income changes being reflected in consumption changes, $\beta_{INS} = 0.269$. While the informal partial insurance parameter θ is not very precisely estimated, the income-consumption pass through β_{INS} as a measure of total consumption insurance is precisely estimated with a standard error of 0.012. As can be seen in the first column of Table 6, consumption insurance is heterogeneous across the wealth distribution; for households in the first quintile 50% of all unexpected income changes are passed through to consumption while for the last quintile the pass-through is merely 15%.

Our estimates for advance information, κ and $\tilde{\kappa}$, as well as for consumption insurance, β_{INS}, are further robust to changing the degree of risk aversion or the borrowing limit (see Appendix C for details on the robustness exercises).

Table 4: Parameter estimates: advance information and consumption insurance

I. Parameter estimates		
Signal precision, κ	Informal partial insurance, θ	Discount factor, β
0.391 (0.023)	0.158 (0.111)	0.944 (0.001)

II. Implications	
Advance information, $\tilde{\kappa}$	Consumption insurance, β_{INS}
0.153 (0.018)	0.269 (0.012)

Notes: Standard errors computed from a parametric bootstrap with 500 repetitions in parentheses.

5.4 Counterfactual scenarios

In the following, we compute two scenarios that show that neglecting advance information results in downward-biased estimates of the income-consumption pass-through, indicating too much insurance. Furthermore, we ask whether informal partial insurance is quantitatively important to reliably estimate the amounts of advance information and consumption insurance. In Table 5, we summarize our key findings for the counterfactual scenarios.

Scenario 1: no advance information and targeting a downward biased regression coefficient. Suppose that advance information is assumed to be absent, with $\kappa = 0$, and that the downward-biased regression coefficient $\widehat{\beta}_{\Delta y_t} = 0.151$ is targeted to estimate θ. In this case, we estimate $\theta = 0.741$, a substantially higher estimate than in our baseline estimation with advance information. The reason is twofold. First, with $\widehat{\beta}_{\Delta y_t} = 0.151 < 0.184 = \beta_{\Delta y_t}$, a lower coefficient is targeted, requiring a larger amount of informal partial insurance. Second, without advance information, the response of consumption to income changes is entirely attributed to insurance, informal and self-insurance, which also tends to increase the estimate of θ. More importantly, the resulting amount of consumption insurance is overstated; instead of an income-consumption pass-through of 0.269 as in the baseline, only 0.201 of all unexpected income shocks are reflected in consumption changes, a decrease of about 25%.

Scenario 2: no advance information but targeting the correct moments. Our estimated parameter for informal partial insurance, $\theta = 0.158$, is lower than the corresponding value found in the related work of Guvenen and Smith (2014) who estimate $\theta = 0.451$. One difference in their work compared

Table 5: Counterfactuals

	Parameter estimates			Implications			
	κ	θ	β	$\tilde{\kappa}$	β_{INS}	$\beta_{\Delta y_t}$	$\beta_{\Delta y_{t+1}}$
Benchmark	0.391	0.158	0.944	0.153	0.269	0.184	0.045
Scenario 1	0.000	0.741	0.948	0.000	0.201	0.151	0.025
Scenario 2	0.000	0.414	0.946	0.000	0.242	0.182	0.023
Scenario 3	0.438	0.000	0.943	0.191	0.287	0.189	0.050

to ours is that their model does not feature advance information on future income shocks.[17] Unlike in Scenario 1, their auxiliary model does, however, allow consumption to respond to future income changes. To follow the spirit of their analysis, we continue to assume that households do no possess advance information on their future income shocks. Unlike in Scenario 1, however, we estimate θ by targeting the two unbiased regression coefficients $\beta_{\Delta y_t} = 0.184$ and $\beta_{\Delta y_{t+1}} = 0.045$, and find $\theta = 0.414$, which is very similar to Guvenen and Smith (2014)'s estimate.[18] As expected, the model without advance information has difficulty capturing the consumption response to future income changes, but reproduces the response to current income changes. Nevertheless, though less than in Scenario 1, the model still predicts too much insurance with a pass through of 24%, which is 10% lower than our baseline estimate of 27% obtained in the model with advance information.

Scenario 3: no informal partial insurance but targeting the correct moments. Given that informal partial insurance is less precisely estimated, it is natural to investigate whether its omission is relevant to reliably estimate consumption insurance and advance information. For this exercise, we restrict $\theta = 0$, and estimate κ by targeting the two regression coefficients $\beta_{\Delta y_t} = 0.184$ and $\beta_{\Delta y_{t+1}} = 0.045$. We compute $\kappa = 0.438 > 0.391$. Thus, relative to the baseline, households possess more advance information reducing their forecast errors more, by $\tilde{\kappa} = 0.191$. The income-consumption pass through is also higher, with $\beta_{INS} = 0.287$. The two regression coefficients are approximated well given that only one parameter, κ, is used to match the two targets. In particular, the model slightly over-states the consumption response to current income growth—rationalizing the higher income-consumption pass through—as well as the consumption response to future income growth, which explains the higher

[17]Their goal is to capture the life-cycle dynamics of income, which is why they consider advance information on the deterministic growth rate of income, but no advance information on future income shocks.

[18]The parameter θ is also more precisely estimated than in the baseline with a standard error of 0.077.

Table 6: Consumption insurance: conditional on wealth quintile

	Benchmark	Scenario 1	Scenario 2	Scenario 3
β_{INS}, $a \in [0, 5th]$	0.576	0.291	0.477	0.646
β_{INS}, $a \in (5th, 10th]$	0.532	0.300	0.469	0.579
β_{INS}, $a \in (10th, 20th]$	0.409	0.298	0.372	0.435
β_{INS}, $a \in (20th, 50th]$	0.282	0.229	0.259	0.299
β_{INS}, $a \in (50th, 100th]$	0.178	0.146	0.164	0.188

Notes: Consumption insurance conditional on the wealth quantiles at the beginning of the period.

estimate for $\tilde{\kappa}$.

Different insurance implications across the wealth distribution. Table 6 shows how consumption insurance varies across the wealth distribution under the difference scenarios considered above. As one would expect, in every scenario households with higher levels of wealth are able to better insure their consumption (leading to lower β_{INS} coefficients). The degree to which wealthier households are better insured is, however, higher in the scenarios with advance information (Benchmark, and Scenario 3). As a result, the overestimation of consumption insurance that follows from abstracting from advance information is significantly magnified at the bottom of the wealth distribution. For the 5-percent poorest households, 58% of all shocks are passed through to consumption in the benchmark, while this number equals 29% in Scenario 1. This is not a result of more households being borrowing constrained in the economies with advance information—in fact, the opposite is true. Instead, it is a result of the following mechanism. With advance information, households save less when they receive a positive signal about their future income and are then especially unprepared in the event the signals are incorrect. To put it another way, the worst surprise in an economy without advance information is to receive a negative income shock, whereas with advance information it is to receive a positive income signal and then a negative income shock.

6. Conclusion

In this paper, we first show, in a simple income-fluctuation model, that when households possess advance information they change consumption in response to current as well as future income changes.

We argue that, as a result of mean reversion, current and future income growth are negatively correlated and, therefore, an unbiased test of the presence of advance information requires inspecting the effect of future income growth on consumption growth *controlling for current income growth*. Using consumption and income data from the PSID, we conduct this test and find strong evidence for the presence of advance information. Without controlling for current income growth, the relationship between current consumption growth and future income growth is negative and statistically insignificant. However, controlling for current income growth leads to a positive and statistically significant relationship. We use this evidence to estimate a standard incomplete markets model and find that advance information is consistent with a reduction in the income forecast errors of US households of 15%. Our estimation results also imply that 27% of all unexpected income changes are passed through to consumption. Ignoring advance information would imply a pass-through of 20%, significantly overestimating the amount of consumption insurance. The overstating of consumption insurance is more pronounced for households at the bottom of the wealth distribution. Abstracting from advance information we find a pass-through rate of 29% for the bottom 5% of the wealth distribution, this number nearly doubles when advance information is introduced.

References

Altonji, J. G. and A. Siow (1987): "Testing the response of consumption to income changes with (noisy) panel data," *The Quarterly Journal of Economics*, 102, 293–328.

Attanasio, O. and L. Pistaferri (2014): "Consumption Inequality over the Last Half Century: Some Evidence Using the New PSID Consumption Measure," *American Economic Review*, 104, 122–26.

Blundell, R., L. Pistaferri, and I. Preston (2008): "Consumption inequality and partial insurance," *American Economic Review*, 98, 1887–1921.

Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016): "Consumption Inequality and Family Labor Supply," *American Economic Review*, 106, 387–435.

Campbell, D., A. Carruth, A. Dickerson, and F. Green (2007): "Job insecurity and wages," *The Economic Journal*, 117, 544–566.

Dominitz, J. (1998): "Earnings expectations, revisions, and realizations," *The Review of Economics and Statistics*, 80, 374–388.

Dominitz, J. and C. F. Manski (1997): "Using expectations data to study subjective income expectations," *Journal of the American Statistical Association*, 92, 855–867.

Ejrnæs, M. and M. Browning (2014): "The persistent–transitory representation for earnings processes," *Quantitative Economics*, 5, 555–581.

Guvenen, F. and A. A. Smith (2014): "Inferring labor income risk and partial insurance from economic choices," *Econometrica*, 82, 2085–2129.

Hall, R. and F. S. Mishkin (1982): "mThe Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households. n," *Econometrica*, 50, 461.

Hendren, N. (2017): "Knowledge of future job loss and implications for unemployment insurance," *American Economic Review*, 107, 1778–1823.

Hill, M. (1993): "PSID User's Guide," *The Institute for Social Research, Ann Arbor, MI*.

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

Jordà, O., K. Knoll, D. Kuvshinov, M. Schularick, and A. M. Taylor (2019): "The rate of return on everything, 1870–2015," *Quarterly Journal of Economics*, 134, 1225–1298.

Kaplan, G. and G. L. Violante (2010): "How much consumption insurance beyond self-insurance?" *American Economic Journal: Macroeconomics*, 2, 53–87.

Kaufmann, K. and L. Pistaferri (2009): "Disentangling insurance and information in intertemporal consumption choices," *American Economic Review*, 99, 387–92.

Li, G., M. Z. Samancioglu, and R. Schoeni (2014): "Estimates of Annual Consumption Expenditures and Its Major Components in the PSID in Comparison to the CE," *American Economic Review*, 104, 132–35.

Mace, B. J. (1991): "Full insurance in the presence of aggregate uncertainty," *Journal of Political Economy*, 99, 928–956.

Manski, C. F. (2017): "Survey Measurement of Probabilistic Macroeconomic Expectations: Progress and Promise," Working Paper 23418, National Bureau of Economic Research.

McGonagle, K. A., R. F. Schoeni, N. Sastry, and V. A. Freedman (2012): "The Panel Study of Income Dynamics: Overview, recent innovations, and potential for life course research," *Longitudinal and life course studies*, 3.

Mueller, A. I., J. Spinnewijn, and G. Topa (2021): "Job Seekers' Perceptions and Employment Prospects: Heterogeneity, Duration Dependence, and Bias," *American Economic Review*, 111, 324–63.

Singh, S. and C. A. Stoltenberg (2020): "Consumption insurance with advance information," *Quantitative Economics*, 11, 671–711.

Tauchen, G. and R. Hussey (1991): "Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models," *Econometrica*, 59, 371–96.

Appendix

A. Proof of Proposition 2.1

Proof. Fix some period $t \geq 0$. The Euler equations imply that consumption follows a random walk,

$$c_t = \mathbb{E}_t[c_{t+j}], \quad \text{for all } j \geq 0.$$

Iterating forward on the sequential budget constraints we get

$$\sum_{j=0}^{\infty} \frac{c_{t+j}}{(1+r)^j} = (1+r)a_t + \sum_{j=0}^{\infty} \frac{y^d_{t+j}}{(1+r)^j},$$

and applying the expectation operator $\mathbb{E}_t[\cdot]$ to both sides implies

$$\sum_{j=0}^{\infty} \frac{\mathbb{E}_t[c_{t+j}]}{(1+r)^j} = (1+r)a_t + \sum_{j=0}^{\infty} \frac{\mathbb{E}_t[y^d_{t+j}]}{(1+r)^j}.$$

Then, using the Euler equations we obtain

$$c_t = ra_t + \frac{r}{(1+r)} \sum_{j=0}^{\infty} \frac{\mathbb{E}_t[y^d_{t+j}]}{(1+r)^j},$$

and it follows that

$$c_{t+1} - c_t = c_{t+1} - \mathbb{E}_t[c_{t+1}] = \frac{r}{(1+r)} \sum_{j=0}^{\infty} \frac{\mathbb{E}_{t+1}[y^d_{t+1+j}] - \mathbb{E}_t[y^d_{t+1+j}]}{(1+r)^j}.$$

Next, using the fact that

$$\begin{aligned}\mathbb{E}_t[y^d_{t+1+j}] &= (1-\theta)\mathbb{E}_t[y_{t+1+j}] + \theta\mathbb{E}_t[\mathbb{E}_{t+j}[y_{t+1+j}]] = (1-\theta)\mathbb{E}_t[y_{t+1+j}] + \theta\mathbb{E}_t[y_{t+1+j}] \\ &= \mathbb{E}_t[y_{t+1+j}] = (1-\kappa)y_t + \kappa y_{t+1}, \quad \text{for all } j \geq 0,\end{aligned}$$

and

$$y^d_{t+1} = y_{t+1} - (1-\kappa)\theta(y_{t+1} - y_t),$$

we obtain

$$\begin{aligned}
c_{t+1} - c_t &= \frac{r}{1+r}\left\{ y^d_{t+1} - \mathbb{E}_t[y^d_{t+1}] + \sum_{j=1}^{\infty} \frac{\mathbb{E}_{t+1}[y^d_{t+1+j}] - \mathbb{E}_t[y^d_{t+1+j}]}{(1+r)^j} \right\} \\
&= \frac{r}{1+r}\left\{ y_{t+1} - (1-\kappa)\theta(y_{t+1} - y_t) - [(1-\kappa)y_t + \kappa y_{t+1}] + \sum_{j=1}^{\infty} \frac{(1-\kappa)y_{t+1} + \kappa y_{t+2} - [(1-\kappa)y_t + \kappa y_{t+1}]}{(1+r)^j} \right\} \\
&= \frac{r}{1+r}\left\{ y_{t+1} - (1-\kappa)\theta(y_{t+1} - y_t) - [(1-\kappa)y_t + \kappa y_{t+1})] + \frac{(1-\kappa)y_{t+1} + \kappa y_{t+2} - [(1-\kappa)y_t + \kappa y_{t+1})]}{r} \right\} \\
&= \frac{(1-\kappa)[1+(1-\theta)r]}{1+r}(y_{t+1} - y_t) + \frac{\kappa}{1+r}(y_{t+2} - y_{t+1}).
\end{aligned}$$

□

B. Details on data and empirical results

PSID data. The Panel Study of Income Dynamics is a nationally representative household panel survey operating since 1968. The dataset follows original households and members who moved away from these households.[19] The initial survey contained two groups of families. First, an oversample of 1,872 low-income families called Survey of Economic Opportunity (SEO) sample. Second, nationally representative sample of 2,930 families, called Survey Research Center (SRC) sample. The survey was conducted annually during the period 1968–1997, and biennially since 1999.

The objective of the PSID is to collect retrospective information on socio-economic characteristics of the household. The retrospective nature of the survey implies information collected in a year refers to the previous year. For the purpose of this study, this information include detailed food expenditures (since 1968) and other non-durable expenditure (in more detail since 1999), wages and income of households members.[20] Since 1968, the PSID has consistently collected information on detailed food expenditure within the household.[21] The information on food expenditure include food at home, food away from home and, if used, the value of food stamps. Starting 1999, the PSID started to collect information on other non-durable consumption components. This measure covers about 70 percent of non-durable spending from national accounts and matches well with aggregate from National Income and Product Accounts (NIPA) (Blundell, Pistaferri, and Saporta-Eksten, 2016). Specifically, during the period 1999 and 2005, the PSID collected consistent information on car maintenance expenditure, health expenditures, rent and utility expenditures, gasoline and transportation expenditures, and child care and education expenditures. Other expenditure categories, clothing and entertainment, were added

[19]For a detailed description of the PSID, see Hill (1993); McGonagle, Schoeni, Sastry, and Freedman (2012).

[20]The retrospective nature of the survey is not consistently true for all categories. While income questions capture income position of the household in the previous year, same cannot be said for food expenditure questions. For detailed discussion and alternative views, see Hall and Mishkin (1982) and Altonji and Siow (1987).

[21]Food expenditure information was not collected in 1973, 1988 and 1989.

after 2005. For the sake of consistency, we only use those nondurable categories which are present since 1999.

We choose the same sample as proposed by Blundell et al. (2008) or Attanasio and Pistaferri (2014) when implementing their respective imputation methods. For Blundell et al. (2008), the objective is to focus on 1968 sample households, inhabited by continuously married couples and headed by male, who experience no change in headship except for small changes in family compositions. Hence, we drop Latino subsample or households which experience a change in headship. Further, households which experience dramatic change in their family composition are also dropped. As our main focus is on consumption insurance and advance information associated with income risk, we drop households with heads who are more than 65 years old. The last sample selection is done on the basis of level and growth of income of households. Hence, households reporting income less than \$100 or an income growth above 500% or below -80% are dropped from the sample. The sample selection by Attanasio and Pistaferri (2014) is less stringent. We use the PSID sample over the period 1968–2014, drop Latino and immigrant samples, correct for food outliers, drop households with female head or head below the age of 25 years, and households with head or spouse (if present) that have an hourly wage below half the minimum wage.[22]

CEX data. While the PSID's main focus is on income, the CEX survey focuses more on consumption expenditures. The data is collected by the Bureau of Labor Statistics (BLS) and serves two purposes: (i) construction of primary consumption basket, and, (ii) revision of the consumer price index. Hence, CEX collects information on wide range of expenditure categories along with information on income and socio-demographic characteristics of the household.[23]

The data is compiled through the use of two complementary surveys: (i) the Diary survey, and, (ii) the Interview survey. Through the Diary survey, information is collected on small and frequently purchased items which includes items on food, personal care etc. The Interview survey, follows a household for a maximum of five quarters. The first quarter is used to collect information on basic sample characteristics. Detailed questions on income and expenditure are asked over the next four quarters.

Imputation procedure. As mentioned above, we use two methods to impute consumption expenditures in the PSID to eventually construct a panel dataset on consumption and income. The first method, used by Blundell et al. (2008), uses the CEX data to impute consumption in the PSID data over the period 1980–1992. Specifically, they estimate a demand function for food which depends on non-durable expenditure, relative prices, and demographic characteristics of the household. As, PSID and CEX contain information on food consumption, the inversion of the estimated demand function gives non-durable expenditure in the PSID data. In Blundell et al. (2008), the food expenditure is the total annual expenditure on food at home and outside. The non-durable

[22] Attanasio and Pistaferri (2014) also drop SEO sample. However, we do not do that, to provide and compare SEO results with Blundell et al. (2008).

[23] For comparison between PSID consumption data collected between 1999 and 2011, and the CEX data, see Li, Samancioglu, and Schoeni (2014)

expenditure is the sum of food expenditure, alcohol, tobacco, utility services, transport, gasoline, personal care, clothing and footwear.

The second method, proposed by Attanasio and Pistaferri (2014), uses the expanded categories in the PSID expenditure data over the period 1999–2014, to impute consumption expenditures in the PSID over the period 1968–1997. They regress non-durable consumption net of food expenditure on socioeconomic variables, relative prices and a polynomial food expenditure.[24] The net non-durable expenditure included home insurance, utility bills (electricity, heating, water, and miscellaneous), car insurance and repairs, gasoline and other transportation expenditures, expenditure on childcare and education, health related expenditures, and rent.[25] Using estimated coefficients, non-durable expenditure over the period is computed as the total expenditure on food plus predicted net non-durable expenditure.[26]

Income measure. For our main specification, we focus on household labor income. Hence, from the total family income we subtract federal taxes and asset income. PSID does not provide information on federal taxes after the survey year 1991. Hence, we use NBER TAXSIM to impute federal taxes for the period 1992–1996.[27]

To construct $\log(y_{it})$, we regress the logarithm of income on year of birth dummy, family size dummy, number of children within family dummy, dummy for an income earned by person other than head or spouse, dummy for children residing outside the house, and retrospective survey year dummy interacted with education level dummy, race dummy, employment or unemployment dummy, and the region dummy. The residual from this regression is $\log(y_{it})$. We analogously construct residual consumption $\log(c_{it})$ with the imputed consumption data.

C. Robustness: risk aversion and borrowing limit

In this section, we consider different values for two preset parameters: the coefficient of relative risk aversion (CRRA), σ, and the borrowing constraint, $\underline{a}$. When changing each parameter we keep interest rates and the parameters of the income process fixed at their baseline calibration values. On the other hand, we allow the discount factor, β, to adjust so that the mean wealth-to-income ratio is maintained at 2.9, and the signal precision κ and the informal partial insurance parameter θ adjust to match $\beta_{\Delta y_t} = 0.184$ and $\beta_{\Delta y_{t+1}} = 0.045$. For the CRRA, σ, we consider halving and doubling the baseline value of 2. For the borrowing constraint, $\underline{a}$, we introduce as an additional target the mean wealth-to-income ratio of households with negative net worth of -0.201.[28]

[24] Inclusion of food expenditure in non-durable consumption at this stage will lead to bias due to correlated errors in food. Hence, the net non-durable consumption is used.

[25] As clothing and entertainment is added only since 2005, hence these items are excluded from the non-durable consumption definition.

[26] All nominal values are deflated using the consumer price index.

[27] The imputation is done by NBER TAXSIM version 32.

[28] We compute this ratio using the 1984 and 1989 PSID wealth modules and take the average.

Table 7: Robustness with respect to σ and $\underline{a}$

	Varied param.		Parameter estimates			Implications	
	σ	$\underline{a}$	β	κ	θ	$\tilde{\kappa}$	β_{INS}
Benchmark	2.0	0.000	0.944	0.391	0.158	0.153	0.269
Halving CRRA	1.0	0.000	0.955	0.391	0.248	0.153	0.266
Doubling CRRA	4.0	0.000	0.909	0.406	0.000	0.165	0.274
Disciplining $\underline{a}$	2.0	−0.188	0.944	0.401	0.075	0.161	0.272

Table 7 summarizes the results. First notice that, while the estimates for informal partial insurance, θ, vary significantly in response to the changes to σ and $\underline{a}$, this variation is consistent with the standard errors reported in Table 4. On the other hand, the estimates for κ and our model-independent measures of advance information and consumption insurance, $\tilde{\kappa}$ and β_{INS}, are exceptionally robust to these variations.

D. Robustness: Empirical Results

In our main specification, we proxy income with the labor income of the household. This measure of income is used to generate the results in Tables 1 and 2. To check the robustness of our result with respect to this choice, we consider as an alternative measure the total income of household, that is labor income plus asset income. Table 8 presents the corresponding consumption-growth regression results, which are very similar to the baseline results from Table 1.

Table 8: Consumption growth regressions: asset income included

	Year fixed effects		Year + Household fixed effects	
	(1)	(2)	(3)	(4)
Δy_{t+1}	-0.027	0.042**	-0.036*	0.043**
	(0.020)	(0.020)	(0.020)	(0.021)
Δy_t		0.206***		0.209***
		(0.023)		(0.025)
Observations	10522	10506	10443	10427

Source: Panel Study of Income Dynamics 1978–1992
Description: The table reports the result of regressing current consumption growth on future income growth, including or excluding current income growth. Measure of income includes labor earnings and asset income. Columns (1)-(2) takes year-fixed effects (as baseline specification), and Columns (3)-(4) additionally household fixed effects into account.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

In Table 9, we implement another test for advance information. The test is implemented in the following way: the covariance between current income growth and future income growth is regressed on year dummies (without a constant term), and the joint $F-$statistic along with the corresponding $p-$value is computed. The intuition is that, in the presence of advance information, future income shocks are revealed to individuals. Hence, the covariance of current consumption growth with future income growth should be statistically different from zero. Blundell et al. (2008) implement this test in their Table 5 and find that there is no evidence of advance information. We replicate their result in column (1) of Table 9.

However, as discussed in the main text, to identify the presence of advance information, we have to include current income growth and future income growth simultaneously. Hence, in column (2), we first partial out the effect of current income growth on current consumption growth, future income growth, and year dummies, and then implement the test. We can see that, after partialling out the effect of current income growth, the $p-$value for the joint significance test is close to zero for the covariance between current consumption growth and one-period-ahead future income growth.

In column (3), we take this partialling out approach one step further. For instance, to construct the joint significance test of current consumption growth and two-period-ahead income growth, we partial out current income growth as well as one-period-ahead income growth. Then, even for two-period-ahead income growth, the $p-$value of the joint test is around 4%, indicating the presence of advance information about the household's income two years ahead.

Table 9: Covariance tests

	Unconditional (BPP, 2008) (1)	Conditional on Δy_t (2)	Conditional on all Δy_{t+i} (3)
$\text{cov}(\Delta c_t, \Delta y_{t+1}) = 0, \forall t$	25%	0%	0%
$\text{cov}(\Delta c_t, \Delta y_{t+2}) = 0, \forall t$	27%	10%	4%
$\text{cov}(\Delta c_t, \Delta y_{t+3}) = 0, \forall t$	74%	51%	12%
$\text{cov}(\Delta c_t, \Delta y_{t+4}) = 0, \forall t$	68%	12%	50%

Description: Baseline specification. Table contains the p-values of the joint test $\text{cov}(\Delta c_t, \Delta y_{t+j}) = 0, \forall t$, given $j \leq 1 \leq 4$. Column 1 displays the test for the unconditional covariances (as in Table 5 of Blundell et al. (2008)). Column 2 tests the covariances conditional on Δy_t and Column 3 the covariances conditional on all Δy_{t+i}, $0 \leq i \leq j-1$.