



LIETUVOS BANKAS
EUROSISTEMA

Advance Information and Consumption Insurance: Evidence and Structural Estimation

Working Paper Series

No 108 / 2022

Advance Information and Consumption Insurance: Evidence and Structural Estimation*

Marcelo Pedroni[†]

Swapnil Singh[‡]

Christian A. Stoltenberg[§]

* We are thankful to Adrien Auclert, Francesco Caselli, Thomas Douenne, Wouter Den Haan, Christian Haefke, Albert Jan Hummel, Albert Marcet, Jaanika Meriküll, Benjamin Moll, Franc Klaassen, Frank Kleibergen, Bérengère Patault, Arne Uhlendorff, Hakki Yazici, and Dimitrios Tsomocos for valuable feedback and comments that helped to improve the paper. Views expressed in the paper are those of authors, and do not necessarily reflect the view of their institutions.

[†] University of Amsterdam, m.pedroni@uva.nl.

[‡] Bank of Lithuania and Kaunas University of Technology. Email: ssingh@lb.lt

[§] University of Amsterdam and Tinbergen Institute, c.a.stoltenberg@uva.nl

© Lietuvos bankas, 2022

Reproduction for educational and non-commercial purposes is permitted provided that the source is acknowledged.

Gedimino pr. 6, LT-01103 Vilnius

www.lb.lt

Working papers describe research in progress by the author(s) and are published to stimulate discussion and critical comments.

The series is managed by the Applied Macroeconomic Research Division of the Economics Department and the Center for Excellence in Finance and Economic Research.

All papers in the series are refereed by internal and external experts.

The views expressed are those of the author(s) and do not necessarily represent those of the Bank of Lithuania.

ABSTRACT

We show that advance information on future income can be identified from the correlation between consumption growth and future income growth conditional on current income growth. Employing PSID data, we find that this conditional correlation is positive and significant. We use this evidence to structurally estimate a standard incomplete markets model and discover that US households possess enough advance information to reduce their income forecast errors by 15%. This significantly affects the measurement of consumption insurance. With advance information, 25% more income shocks pass through to consumption on average, and more than twice as much for the 5% asset poorest.

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

JEL Classification: C23, D12, D31, D52, D81, E21, G52.

1. INTRODUCTION

How well can households insure their consumption against unexpected income changes? The answer to this question is crucial for the design of any policy that aims at providing additional insurance to households. To measure consumption insurance, it is necessary to determine 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 any private advance information about their future income, in which case their foreknowledge coincides with that of an econometrician. On the one hand, this assumption is consistent with the results in [Blundell, Pistaferri, and Preston \(2008\)](#) who—using data from the Panel Study of Income Dynamics (PSID)—find that the *unconditional* correlation between consumption growth and future income growth is statistically insignificant. On the other hand, there is mounting direct evidence that individual expectations are predictive of subsequent income realizations even after information available to an econometrician is taken into account.¹

In this paper, we reconcile these findings by making the following contributions. First, in a stylized income-fluctuation model, we show that it is the correlation of consumption growth with future income growth *conditional* on current income growth that is most useful to identify advance information. Second, we use the same dataset as [Blundell et al. \(2008\)](#) and find that this conditional correlation is positive and significantly different from zero. Third, we use this evidence to inform a structural model and estimate a degree of advance information consistent with the one obtained from direct evidence. Finally, we show that the standard practice of ignoring advance information leads to a substantial overestimation of consumption insurance. Taken together, these contributions amount to a new method for measuring advance information and consumption insurance even if data on individual income expectations is absent.

We consider first an income-fluctuation model in which households possess private advance information on their future income innovations. To facilitate the insurance of permanent income shocks, households have access to an informal private insurance scheme. We show that households revise their consumption in response 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, con-

¹The predictive power of subjective expectations for future realizations has been demonstrated, for example, for earnings ([Dominitz and Manski, 1997](#); [Dominitz, 1998](#); [Kaufmann and Pistaferri, 2009](#)), for future job losses ([Hendren, 2017](#); [Campbell, Carruth, Dickerson, and Green, 2007](#)), and for the probability of leaving unemployment ([Mueller, Spinnewijn, and Topa, 2021](#)).

sumption reacts to future income changes only in proportion to the amount of advance information that households possess. Conditioning on current income growth, a positive response of consumption to future income changes is indicative of the presence of advance information. Considering the amount of advance information identified by this response, the consumption response to current income changes can then be used to identify consumption insurance.²

We use the insights from the stylized model to investigate whether US households possess advance information. More specifically, we use the PSID dataset 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 unconditional correlation is not significantly different from zero with a somewhat nonintuitive negative sign. The picture changes, however, when we control for current income growth. Consumption growth is then significantly positively correlated with future income growth (and with current income growth), suggesting that US households indeed possess advance information about 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, which explains the somewhat nonintuitive negative sign. 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 nonfood categories of nondurable consumption expenditures in the PSID before 1999 from the Consumer Expenditure Survey (CEX). [Attanasio and Pistaferri \(2014\)](#) propose imputing 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 greater detail. With this alternative imputation method, the same picture emerges; the conditional correlation of current consumption growth with future income growth is significantly positive, while the unconditional correlation is negative.

²This implies that, in principle, it is possible to 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 \(2014\)](#).

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. With advance information and 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 an economic model with more quantitatively relevant preferences, a more realistic income process, and 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 \(2020\)](#). Second, we supplement households' self-insurance with an informal partial insurance scheme, similar to the one in [Guvenen and Smith \(2014\)](#). We use indirect inference to structurally estimate this more quantitative model. In the estimation procedure, the simple regression described above serves as an auxiliary model. As in the stylized economy, we find that the correlation of consumption growth with future income growth identifies the degree of advance information.

With the estimated economy at hand, we can 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? We find that advance information reduces the mean forecast error of US households about the income one year ahead by 15%. Moreover, we find that 27% of households' unexpected income changes pass through to consumption. Ignoring advance information biases this pass through downward by 25%, leading to an overestimation of the degree of consumption insurance. Advance information and insurance both reduce the response of consumption to current income changes. Thus, to match the response in the data, a model without advance information exclusively attributes it to insurance.

The overestimation of insurance is magnified at the bottom of the wealth distribution. For the five-percent asset poorest households, 58% of all unexpected income changes result in consumption changes with advance information. Without advance information, only 29% of all income shocks pass through to consumption. While the degree of informal partial insurance is not very precisely estimated, both advance information and consumption insurance as a whole are.

Similar to the test proposed by [Blundell et al. \(2008\)](#), the evidence we provide here is indirect since it does not rely on individual income expectations data directly but detects advance information from the relationship of current consumption with future income changes. Our empirical finding that US households do possess advance information echoes growing direct evidence on the predictive power of individual expectations. With US data, [Dominitz \(1998\)](#) estimates that individual income expectations reduce income forecast errors between 12 and 21%. Our estimate of 15% is consistent with these findings and, in that sense, we reconcile the direct and indirect evidence on advance information.

Standard macroeconomic consumption-savings models typically assume that households have no advance information. Noteworthy exceptions are [Guvenen and Smith \(2014\)](#), [Kaplan and Violante \(2010\)](#) and [Stoltenberg and Uhlenhorff \(2022\)](#). [Guvenen and Smith \(2014\)](#) estimate a life-cycle model for the US economy with households that have advance information about 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. \(2008\)](#), and the insurance that emerges in a life-cycle standard incomplete markets model with a standard calibration. Using Italian household data, [Stoltenberg and Uhlenhorff \(2022\)](#) document that households' consumption choices depend on their earnings expectations and show that the relationship can be rationalized by households with advance information about their future earnings. Our paper complements this literature by structurally estimating the degree of advance information possessed by US households and exploring its implications for consumption insurance.

Our work is also related to recent work that seeks to measure income uncertainty ([Guvenen, Karahan, Ozkan, and Song, 2021](#); [Arellano, Blundell, Bonhomme, and Light, 2021](#); [Arellano, Blundell, and Bonhomme, 2017](#)). Using administrative data, [Guvenen et al. \(2021\)](#) find that the distribution of earnings changes shows non-Gaussian features such as negative skewness. [Arellano et al. \(2017\)](#) provide arguments that earnings shocks are characterized by nonlinearities in persistence. [Arellano et al. \(2021\)](#) build on the earlier work by [Arellano et al. \(2017\)](#) to investigate heterogeneous consumption responses to income shocks with nonlinear persistence. Advance information also directly impacts households' income uncertainty. Since the aforementioned papers abstract from advance information, we view our results as complementary to theirs.

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. In Section 5, we lay out our estimation method, and present in Section 6 our estimates of advance information and consumption insurance. We conclude in Section 7.

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.³ In this economy, we analytically show how the consumption response to current and future income changes is shaped by advance information and partial insurance.

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, 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, $u(c) = c - bc^2$.

Households' income is subject to permanent income shocks. That is, the initial income level, y_0 , is given, and households' future income levels follow a random walk $y_t = y_{t-1} + \eta_t$, with innovations $\eta_t \sim \mathcal{N}(0, \sigma^2)$. Households observe past and current income levels and have advance information only about the next period's income innovation, so that

$$\mathbb{E}_t[y_{t+1}] = y_t + \kappa\eta_{t+1} = (1-\kappa)y_t + \kappa y_{t+1}. \quad (2.1)$$

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's income with certainty, that is, $\mathbb{E}_t[y_{t+1}] = y_{t+1}$. Households have only advance information about the next period's income innovation which implies $\mathbb{E}_t[y_{t+1+j}] = (1-\kappa)y_t + \kappa y_{t+1}$,

³Our formulation of advance information and partial insurance follows exactly the one in [Güvenen and Smith \(2014\)](#), who consider a two-period version of the model.

for all $j \geq 0$.⁴

With self-insurance, households cannot insure permanent income shocks. To allow for insurance, we therefore consider an informal partial insurance scheme controlled by parameter $0 \leq \theta \leq 1$. Informal insurance 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*

$$c_t - c_{t-1} = (1 - \kappa) \left(1 - \frac{\theta r}{1 + r} \right) (y_t - y_{t-1}) + \frac{\kappa}{1 + r} (y_{t+1} - y_t), \quad \text{for all } t \geq 1. \quad (2.2)$$

Proof. See Appendix A. □

It follows from Proposition 2.1 that, in principle, it is 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 innovation, $y_{t+1} - y_t = \eta_{t+1}$, only affects current consumption, $c_t - c_{t-1}$, in proportion to how much advance information the household has about the innovation, 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 .⁵

Following Mace (1991), consumption insurance is traditionally measured by regressing consumption on income growth.⁶ Equation (2.2) 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

⁴An equivalent representation of advance information are households with a private signal s_t that reveals in period t a fraction κ of their income innovation in the next period η_{t+1} . With a random walk, the conditional income expectation $\mathbb{E}[y_{t+1}|y_t, s_t]$ is then given by Equation (2.1).

⁵To obtain equation (2.2), the model must have at least three periods, which is why the two-period version of the model in Guvenen and Smith (2014) cannot deliver it.

⁶The 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 (2.2) would be replaced by log deviations from the steady state, so that overtime differences become growth rates.

insurance. Note that equation (2.2) can also be written as follows:

$$c_t - c_{t-1} = \left(1 - \frac{\theta r}{1+r}\right) (y_t - \mathbb{E}_{t-1}[y_t]) + \frac{\kappa}{1+r} (y_{t+1} - y_t), \quad (2.3)$$

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 (2.3) 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 expectation data. In the first step, advance information is identified from the response of current consumption to future income growth. Then, 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.⁷ The PSID is one of the longest-running and largest income panel datasets in the world but does not include information on households' income expectations, which is why we adopt the second approach in the following sections.

Identifying advance information requires an additional step when income does not follow a random walk. Then, income changes do not solely capture innovations to income, and current and future income changes can be correlated.⁸ We therefore proceed in two steps. First, in the next section, we use the insights from Equation (2.2) to have a fresh look at the data and establish that current consumption growth is correlated with future income growth if current income growth is controlled for. Second, in the following sections, we use this evidence as an auxiliary model to structurally estimate a standard incomplete markets model with advance information and informal partial insurance by indirect inference.

⁷For 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 (2.3). In the more realistic model presented in Section 4, households also insure consumption with precautionary savings.

⁸Whether 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.

3. EMPIRICAL RESULTS

Using panel data on consumption and income, the objective of this section is to empirically detect whether consumption growth responds to future income growth. For our baseline results, we use the PSID dataset constructed by [Blundell et al. \(2008\)](#) and provide evidence that the response, controlling for current income growth, is positive and significant. As robustness exercises, we consider an alternative consumption imputation method proposed by [Attanasio and Pistaferri \(2014\)](#) and apply different sample selection criteria. Overall, the evidence in this section suggests that US households possess advance information on their future income.

In what follows, the key variables of interest are after-tax household income net of asset income and nondurable household consumption expenditures. Both variables are residual variables that follow after controlling for a vector of observable household characteristics.⁹

3.1 Baseline estimation results

One take-away from the previous section is that advance information impacts the correlation of current consumption and future income changes. Throughout this section, we 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,

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

with $\Delta x_{it} \equiv \log(x_{it}) - \log(x_{it-1})$ denoting the growth rate of variable x . We refer to the coefficient $\tilde{\beta}_{\Delta y_{t+1}}$ 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. The corresponding regression equation is

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

Here, a conditional coefficient $\beta_{\Delta y_{t+1}}$ that is significantly different from zero is suggestive of advance

⁹A 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

	(1)	(2)
Δy_{t+1}	-0.013 (0.018)	0.045** (0.019)
Δy_t		0.184*** (0.021)
Observations	10502	10471

Source: Panel Study of Income Dynamics 1978–1992.

Notes: Baseline specification. The table reports the result of regressing current consumption growth on future income growth, including or excluding current income growth. Year-fixed effects. Standard errors are clustered at the household level.

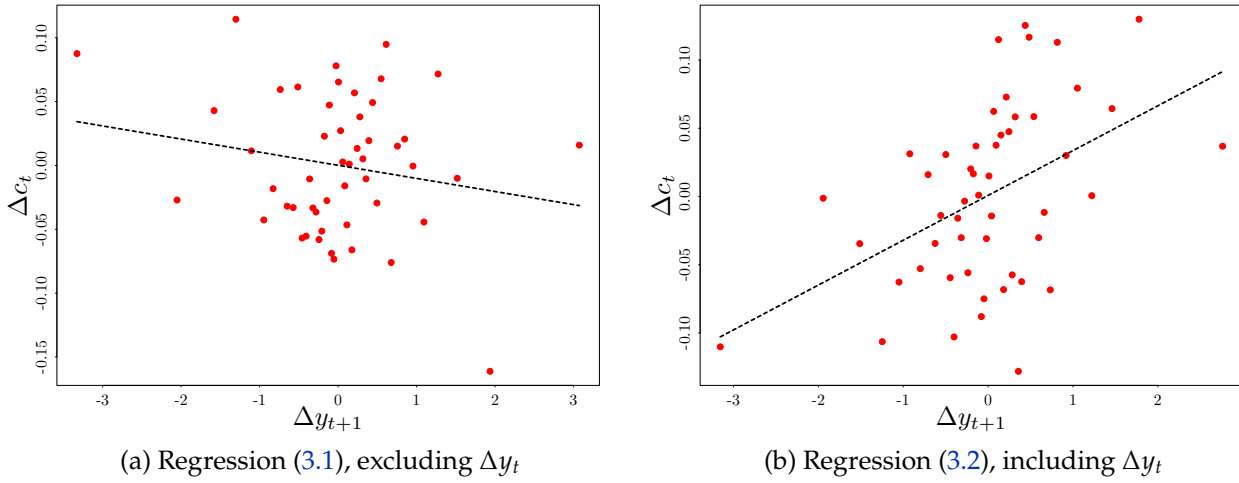
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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 first column displays the correlation for the regression equation (3.1), while the second column shows the results for the regression equation (3.2). We take year fixed effects and cluster standard errors at the household level. The main messages are as follows. The first regression equation yields a (nonsignificant) negative correlation between current consumption and future income growth. The second regression, however, indicates a significant positive correlation. The coefficient $\beta_{\Delta y_{t+1}}$ is precisely estimated with p -values of 0.014.

In Appendix C, we provide the outcomes of several robustness exercises. In particular, we find that our baseline estimation results are robust to including asset income or employing innovations in income growth as regressors (see Tables 6 and 7). In Table 8, we further entertain the same test to detect advance information as Blundell et al. (2008), who investigate whether the unconditional covariances satisfy $\text{cov}(\Delta c_{it}, \Delta y_{it+j}) = 0$, for all years t , and $j \geq 1$. In Column 1, we reproduce their test results for the unconditional covariances and cannot reject the null hypothesis with p -values of at least 25%. However, once we control for previous income changes, we safely reject the null hypothesis (see Columns 2 and 3), confirming our baseline estimation results in Table 1.

Figure 1: Consumption growth and future income growth



Source: Panel Study of Income Dynamics 1978–1992.

Notes: Baseline specification. The figure plots the nonparametric 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.

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

$$\tilde{\beta}_{\Delta y_{t+1}} = \beta_{\Delta y_{t+1}} + \delta \beta_{\Delta y_t}, \quad (3.3)$$

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, 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 $\tilde{\beta}_{\Delta y_{t+1}} < \beta_{\Delta y_{t+1}}$. Thus, the correlation between current consumption and future income growth in the first regression tends to be downward biased by $\delta \beta_{\Delta y_t}$ as a result of

Figure 2: Current income growth and future income growth



Source: Panel Study of Income Dynamics 1978–1992.

Notes: Baseline specification. The figure plots the nonparametric 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.

omitting current income growth as a regressor.

The traditional approach to measuring consumption insurance is to regress current consumption growth on current income growth,

$$\Delta c_{it} = \hat{\beta}_0 + \hat{\beta}_{\Delta y_t} \Delta y_{it} + \hat{\eta}_{it}. \quad (3.4)$$

The take-away from the estimations in Table 1 is that this measure also tends to be downward biased when future income growth is omitted in equation (3.4).¹⁰ In particular, it follows that

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

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 $\hat{\beta}_{\Delta y_t} = 0.151 < 0.184 = \beta_{\Delta y_t}$, resulting in a tendency to overestimate consumption insurance.

As an intermediate summary, we find that, controlling for current income growth, consumption is positively correlated with future income growth, which suggests that US households possess advance information about their future income. In what follows, we investigate whether this finding is robust

¹⁰The estimation results of (3.2) and (3.4) resemble the test for strict exogeneity proposed by Wooldridge (2010) on Page 324. A coefficient $\beta_{\Delta y_{t+1}}$ that is significantly different from zero signals a violation of strict exogeneity.

Table 2: Standardized consumption growth regressions

Sample: Age group:	SEO sample excluded				SEO sample included			
	30-65		20-65		30-65		20-65	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Panel A: CEX Imputation (BPP, 2008)</i>								
Δy_{t+1}	-0.010 (0.014)	0.034** (0.014)	-0.008 (0.013)	0.034*** (0.013)	-0.029** (0.012)	0.024** (0.011)	-0.023** (0.010)	0.029*** (0.010)
Δy_t		0.138*** (0.016)		0.133*** (0.014)		0.155*** (0.013)		0.152*** (0.011)
Observations	10502	10471	13030	12994	16012	15979	20347	20308
<i>Panel B: PSID Imputation (Attanasio and Pistaferri, 2014)</i>								
Δy_{t+1}	-0.030** (0.013)	0.022* (0.013)	-0.030** (0.012)	0.025** (0.012)	-0.034*** (0.010)	0.024** (0.010)	-0.032*** (0.010)	0.026*** (0.009)
Δy_t		0.151*** (0.014)		0.159*** (0.013)		0.160*** (0.011)		0.162*** (0.010)
Observations	10073	10058	11911	11895	16335	16316	19594	19573

Source: Panel Study of Income Dynamics 1978–1992.

Notes: The table reports the result of regressing current consumption growth on future income growth, including or excluding current income growth as a 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 **bold**. 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

to employing an alternative method for imputing consumption and different sample selection criteria.

3.3 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 as follows. 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.¹¹ Further, 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 households 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. The empirical evidence in this section suggests that households do have advance information about their future income.

In the following, we use the evidence from this section to inform a theoretical model. More specifically, we employ indirect inference with the regression (3.2) as an auxiliary model to estimate how much advance information US households possess in the context of a standard incomplete markets model. Given an estimated amount of advance information, consumption insurance can then be quantified by the consumption response to unexpected income changes, $y_{it} - \mathbb{E}_{it-1}[y_{it}]$, with expectations that take into account households' advance information.¹²

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 $\mathbb{E}_0 [\sum_t \beta^t u(c_t)]$. 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 the next period's income. The household (y, s) pair follows a Markov process specified below. Households can only accumulate a risk-free asset, a , and face borrowing

¹¹The exception is the PSID imputation in the lower panel of Column 2, for which the coefficient is positive at the 10% level (with a p -value of 0.067).

¹²This logic resembles insights from [Jappelli and Pistaferri \(2010\)](#), who also advocate for estimating the correlation of consumption changes with unexpected income changes.

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

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

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

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

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

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

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's income is informative with probability κ , so that

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

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,¹³ it must be that

$$\Pr(s = y'|y) = \Pr(y'|y), \quad (4.2)$$

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' .¹⁴

5. QUANTITATIVE EXERCISE

First, we provide functional forms for the income process and household preferences and describe how we choose the structural parameters. One part of the parameter vector is preset, while the other part of the parameter vector—including the precision of signals and informal partial insurance—is structurally estimated by indirect inference. Eventually, we define atheoretical measures of advance information and consumption insurance.

5.1 Preset parameters and the income process

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 $\underline{a} = 0$.¹⁵ The interest rate is set to $r = 3.9\%$, the annual real risk-free rate for the US post-1980 as documented by [Jordà, Knoll, Kuvshinov, Schularick, and](#)

¹³In particular, the condition in equation (4.2) guarantees that $\sum_{s \in Y} \Pr(s|y) \Pr(y'|y, s) = \Pr(y'|y)$.

¹⁴Below, 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 $\Pr((s_P, s_T) = (y'_P, y'_T)|y_P) = \Pr(y'_P|y_P) \Pr(y'_T)$.

¹⁵In Appendix E, we conduct robustness exercises with respect to these choices.

Table 3: Preset parameters

	Parameter	Value
σ	Relative risk aversion	2.000
\underline{a}	Borrowing constraint	0.000
r	Interest rate	3.9%

Taylor (2019).

Following the standard practice in the literature, we allow household income to have a persistent and a transitory component; that is,¹⁶

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

with

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

5.2 Parameters estimated with indirect inference

We use indirect inference to estimate the remaining structural parameters collected in the vector $\Theta = (\sigma_P, \sigma_T, \rho, \kappa, \theta, \beta)$. The parameters $(\sigma_P, \sigma_T, \rho)$ control the dynamics of the income process. Given a particular income process, the remaining parameters shape households' consumption-savings decisions and the wealth distribution. Indirect inference relies on a parsimonious auxiliary model to connect the consumption-savings decisions of households in the data to the decisions in the economic model described in Section 4. The methodology follows [Guvenen and Smith \(2014\)](#).

The auxiliary model comprises a consumption equation, and an equation capturing income dynamics. The first equation is the regression equation (3.2), which captures how current consumption growth depends on current and future income growth:

$$\Delta c_{it} = \beta_0 + \beta_{\Delta y_t} \Delta y_{it} + \beta_{\Delta y_{t+1}} \Delta y_{it+1} + \epsilon_{it}. \tag{5.1}$$

¹⁶[Ejrnaes and Browning \(2014\)](#) shows that this specification is equivalent to an ARMA(1,1)-process with a single innovation term.

The estimates in the survey data are $\beta_{\Delta y_t} = 0.184$ (s.e. 0.021), and $\beta_{\Delta y_{t+1}} = 0.045$ (s.e. 0.019).¹⁷ The second equation represents the dynamics of after-tax income

$$\log(y_{it}) = \beta_{y,0} + \beta_{y_{t-1}} \log(y_{it-1}) + \epsilon_{y,it}, \quad \epsilon_{y,it} \sim (0, \sigma_y^2). \quad (5.2)$$

In the survey data, we estimate $\beta_{y_{t-1}} = 0.777$ (s.e. 0.009), a residual standard deviation, $\sigma_y = 0.285$ (s.e. 0.006), and the regression coefficient of current income growth on future income growth, $\delta = -0.325$ (s.e. 0.011).¹⁸

With indirect inference, the parameter vector Θ is chosen to minimize the distance between the coefficients of the auxiliary model estimated on the survey data and the ones estimated on the simulated data stemming from the economic model. The coefficients of the auxiliary model are given by the consumption responses to current and future income growth, $\beta_{\Delta y_t}$ and $\beta_{\Delta y_{t+1}}$ from Equation (5.1), and by $\beta_{y_{t-1}}$ and σ_y from the income equation (5.2).¹⁹ The empirical section shows that the autocorrelation of income growth has important implications for the covariance of current consumption with future income growth, which is why we also target the regression coefficient of current on future income growth, δ . To capture basic features of the wealth distribution, we include the average wealth-to-income ratio in the set of auxiliary model parameters, which equals 2.90 (s.e. 0.055) in the sample.²⁰

In total, the auxiliary model comprises six coefficients, which equals the length of the parameter vector Θ . Thus, we consider a case of exact identification and estimate the parameters in steps. In the first step, we analytically solve for $(\sigma_P, \sigma_T, \rho)$ such that $(\beta_{y_{t-1}}(\Theta), \sigma_y(\Theta), \delta(\Theta))$ is equal to its data counterpart, $(\hat{\beta}_{y_{t-1}}, \hat{\sigma}_y, \hat{\delta})$.²¹ This yields $\sigma_P = 0.192$, $\sigma_T = 0.162$ and $\rho = 0.891$. Given $(\sigma_P, \sigma_T, \rho)$, for each candidate pair of signal precision κ and informal partial insurance θ , we choose the discount factor β to match the average wealth-to-income ratio in the sample. In the next step, we estimate κ and θ by minimizing the distance between the model estimates $(\beta_{\Delta y_t}(\Theta), \beta_{\Delta y_{t+1}}(\Theta))$ and their data

¹⁷These standard errors are computed from a nonparametric bootstrap with 500 repetitions.

¹⁸Notice that the auxiliary model need not be correctly specified. Instead, its main purpose is to identify the structural parameters in Θ .

¹⁹In our baseline estimation, we do not include the residual standard deviation in Equation (5.1), σ_ϵ , as a coefficient in the auxiliary model, which is estimated equal to 0.399 (s.e. 0.007) in the survey data. As a robustness exercise, we extend Θ with the standard deviation of a classical measurement error in consumption expenditures. Such a measurement error, however, does not affect the estimates of the key coefficients $\beta_{\Delta y_t}$ and $\beta_{\Delta y_{t+1}}$, which is why we decided not to target the residual standard deviation in our baseline estimation.

²⁰We 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 home value minus the outstanding mortgage. The income used to compute the ratio is the total family income.

²¹In Appendix D, we derive the analytical expressions.

counterparts ($\hat{\beta}_{\Delta y_t}, \hat{\beta}_{\Delta y_{t+1}}$). The identification and the estimation results for κ, θ are discussed in the next section.

5.3 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 household income forecast error defined as

$$\tilde{\kappa}(\kappa) = \frac{\text{MSFE}_y - \text{MSFE}_{y,s}(\kappa)}{\text{MSFE}_y}, \quad 0 \leq \tilde{\kappa}(\kappa) \leq 1 \quad (5.3)$$

where

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

$\pi(y)$ 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.²² In addition to its model independence, the measure $\tilde{\kappa}$ has the advantage that we can readily compare our estimates of advance information in Section 6 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 income-consumption pass-through coefficient β_{INS} of the following regression

$$\Delta c_{it} = \text{constant} + \beta_{INS} \{\log(y_{it}) - \log(\mathbb{E}[y_{it} | y_{it-1}, s_{it-1}])\} + \epsilon_{it}, \quad (5.4)$$

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

²²For our specification of signals in Section 4.2, one can show with tedious but straightforward algebra that $\tilde{\kappa} = \kappa^2$.

6. ESTIMATION RESULTS

We begin with a discussion on identification. Afterward, we present the parameter estimates. With these estimates, we quantify the amount of advance information US households possess and how well they can insure against unexpected income shocks. Eventually, we study counterfactual scenarios without advance information to demonstrate that ignoring advance information results in a quantitatively important overestimation of consumption insurance, particularly for households in the left tail of the wealth distribution.

6.1 Identification of signal precision and informal partial insurance


In the simple model from Section 2, we show that κ and θ 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 also extends to the quantitatively more realistic model when we employ the two coefficients as part of our auxiliary model in Equation (5.1).

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 the identification from κ stems from 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, confirming our auxiliary model choice.

6.2 Parameter estimates

In Table 4, we summarize the parameter estimates of the last two estimation steps described in Section 5.2 and the implications of these estimates of 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 error of 0.023. We estimate an informal partial insurance parameter of $\theta = 0.158$, which is less precisely estimated than the other parameters with a standard error of 0.111.

Figure 3: Identification of signal precision and informal partial insurance



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

The parameter estimates yield that compared to an econometrician, households' advance information on future income shocks reduces their forecast error by $\tilde{\kappa} = 15.3\%$, which is also precisely estimated. This value is consistent with the direct evidence provided in [Dominitz \(1998\)](#), who reports reductions in mean-squared forecast errors when conditioning on income expectations between 12 and 21%. 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 precisely estimated, the income-consumption pass-through as a measure of total consumption insurance is precisely estimated with a standard error of 0.012.

Our estimates of 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 E](#) for details on the robustness exercises).²³

²³In our baseline estimation, we do not include the residual standard deviation in Equation (5.1), σ_ϵ , as a coefficient in the auxiliary model. As a robustness exercise, we include the standard deviation as an additional coefficient of the auxiliary model and use it to estimate the standard deviation of a classical measurement error in log consumption expenditures as an additional structural parameter. Without affecting our other parameter estimates, we estimate a standard deviation of the error of $\sigma_c = 0.278$, which is similar to the value of 0.36 found in [Güvener and Smith \(2014\)](#).

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.

6.3 Does advance information matter for consumption insurance?

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 consumption insurance, in particular for households in the left tail of the wealth distribution. In Table 5 and Figure 4, we summarize our key findings on economies with and without advance information.

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 $\hat{\beta}_{\Delta y_t} = 0.151$ 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 $\hat{\beta}_{\Delta y_t} = 0.151 < 0.184 = \beta_{\Delta y_t}$, 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 .²⁴ 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 approximately 25%.

Scenario 2: no advance information but targeting the coefficients $\beta_{\Delta y_t}, \beta_{\Delta y_{t+1}}$. 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

²⁴The reason why ignoring advance information leads to an overestimation of insurance can be understood in the simple model of Section 2. First, note that κ and θ reduce the first coefficient in equation (2.2). Thus, given a target for this coefficient, restricting $\kappa = 0$ leads to a higher estimate for θ , which, in turn, reduces the pass through in equation (2.3).

Table 5: Counterfactual scenarios

	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


compared to ours is that their model does not feature advance information on future income shocks.²⁵ Unlike in Scenario 1, their auxiliary model does, however, allow consumption to respond to future income changes. Following the spirit of their analysis, we continue to assume that households do not possess advance information about their future income shocks. Unlike in Scenario 1, however, we estimate θ by targeting the two 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.²⁶ 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, although less than in Scenario 1, the model still predicts too much insurance with a pass through of 24%, which is more than 10% lower than our baseline estimate of 27% obtained in the model with advance information.

Insurance across the wealth distribution with and without advance information. Figure 4 shows how consumption insurance varies across the wealth distribution under the different scenarios considered above. As one would expect, in every scenario households with higher levels of wealth are better able to insure their consumption (leading to lower β_{INS} coefficients). The degree to which wealthier households are better insured is, however, higher in the benchmark. 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

²⁵Their 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.

²⁶The parameter θ is also more precisely estimated than in the baseline with a standard error of 0.077.

Figure 4: Consumption insurance conditional on wealth



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

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.

7. CONCLUSION

In this paper, we have investigated whether US households possess advance information on their future income. In a simple income-fluctuation model, we first show that when households possess advance information, they change consumption in response to current and future income changes. Based on these insights, we take a fresh look at the data and find—in contrast to earlier findings—that consumption growth is positively correlated with future income changes if current income changes are controlled for. Without controlling for current income growth, the relationship between current consumption growth and future income growth is negative and statistically insignificant, confirming the results from the earlier literature.

To explain the positive correlation, we have estimated a structural model with indirect inference and find that US households possess advance information that reduces their income forecast errors

by 15%. Our estimates imply that the conventional approach to abstract from advance information results in an overestimation of consumption insurance. While on average 27% of all unexpected income changes are passed through to consumption with advance information, 25% fewer income shocks pass through without advance information. 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.
- ARELLANO, M., R. BLUNDELL, AND S. BONHOMME (2017): "Earnings and consumption dynamics: a nonlinear panel data framework," *Econometrica*, 85, 693–734.
- ARELLANO, M., R. BLUNDELL, S. BONHOMME, AND J. LIGHT (2021): "Heterogeneity of Consumption Responses to Income Shocks in the Presence of Nonlinear Persistence," Mimeo, University of Chicago.
- 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., F. KARAHAN, S. OZKAN, AND J. SONG (2021): "What do data on millions of U.S. workers reveal about lifecycle earnings dynamics?" *Econometrica*, 89, 2303–2339.
- 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): "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households." *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.
- 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.
- STOLTENBERG, C. A. AND A. UHLENDORFF (2022): "Consumption choices and earnings expectations: empirical evidence and structural estimation," Mimeo, University of Amsterdam.

WOOLDRIDGE, J. M. (2010): *Econometric Analysis of Cross Section and Panel Data*, vol. 1 of *MIT Press Books*,
The MIT Press, 2nd edition ed.

APPENDIX: FOR ONLINE PUBLICATION

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 obtain

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

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

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

Then, using the Euler equations, we obtain

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

and it follows that

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

Next, using the fact that

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

and

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

we obtain

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

□

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.²⁷ 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, a 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 the socioeconomic characteristics of households. The retrospective nature of the survey implies that information collected in a year refers to the previous year. For the purpose of this study, this information includes detailed food expenditures (since 1968) and other nondurable expenditures (in more detail since 1999), wages and income of household members.²⁸ Since 1968, the PSID has consistently collected information on detailed food expenditure within the household.²⁹ The information on food expenditure includes food at home, food away from home and, if used, the value of food stamps. In 1999, the PSID started to collect information on other nondurable consumption components. This measure covers approximately 70 percent of nondurable spending from national accounts and matches well with aggregates from National Income and Product Accounts (NIPA) (Blundell, Pistaferri, and Saporta-Eksten, 2016). Specifically, during the period 1999 to 2005, the PSID collected consistent information on car maintenance expenditures, health expenditures, rent and utility expenditures, gasoline and transportation expenditures, and childcare and education expenditures. Other expenditure categories, such as clothing and entertainment, were

²⁷For a detailed description of the PSID, see Hill (1993); McGonagle, Schoeni, Sastry, and Freedman (2012).

²⁸The retrospective nature of the survey is not consistent for all categories. While income questions capture the income position of the household in the previous year, the same cannot be said for food expenditure questions. For detailed discussion and alternative views, see Hall and Mishkin (1982) and Altonji and Siow (1987).

²⁹Food expenditure information was not collected in 1973, 1988 and 1989.

added after 2005. For the sake of consistency, we only use those nondurable categories that have been 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 males, who experience no change in headship except for small changes in family compositions. Hence, we drop the Latino subsample or households that experience a change in headship. Further, households that experience dramatic changes 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 the 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 a female head or head below the age of 25 years, and drop households with a head or spouse (if present) that have an hourly wage below half the minimum wage.³⁰

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 the primary consumption basket and (ii) revision of the consumer price index. Hence, CEX collects information on a wide range of expenditure categories along with information on income and sociodemographic characteristics of the household.³¹

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 CEX data to impute consumption in the PSID data over the period 1980–1992. Specifically, they estimate a demand function for food that depends on nondurable 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 nondurable expenditure in the PSID data. In [Blundell et al. \(2008\)](#), the

³⁰[Attanasio and Pistaferri \(2014\)](#) also drop the SEO sample. However, we do not do that, to provide and compare SEO results with [Blundell et al. \(2008\)](#).

³¹For comparison between PSID consumption data collected between 1999 and 2011, and the CEX data, see [Li, Samancioglu, and Schoeni \(2014\)](#)

food expenditure is the total annual expenditure on food at home and outside. The nondurable 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 nondurable consumption net of food expenditure on socioeconomic variables, relative prices and a polynomial food expenditure.³² The net nondurable expenditure included home insurance, utility bills (electricity, heating, water, and miscellaneous), car insurance and repairs, gasoline and other transportation expenditures, expenditures on childcare and education, health-related expenditures, and rent.³³ Using estimated coefficients, nondurable expenditure over the period is computed as the total expenditure on food plus predicted net nondurable expenditure.³⁴

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 1992–1996.³⁵

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: EMPIRICAL RESULTS

In this section, we entertain three robustness exercises. First, we include asset income in our measure of household income. Second, we consider consumption-growth regressions that use innovations to income-growth rates instead of income growth rates as independent variables. Further, we conduct the covariance test provided in [Blundell et al. \(2008\)](#) to detect advance information but also control for current income growth.

Including asset income. In our main specification, we proxy income with household labor income. This measure of income is used to generate the results in Tables 1 and 2. As a robustness exercise, we consider

³²Inclusion of food expenditure in nondurable consumption at this stage will lead to bias due to correlated errors in food. Hence, the net nondurable consumption is used.

³³As clothing and entertainment are added only since 2005, these items are excluded from the nondurable consumption definition.

³⁴All nominal values are deflated using the consumer price index.

³⁵The imputation is done by NBER TAXSIM version 32.

Table 6: Consumption growth regressions: asset income included

	(1)	(2)
Δy_{t+1}	-0.027 (0.020)	0.042** (0.020)
Δy_t		0.206*** (0.023)
Observations	10522	10506

Source: Panel Study of Income Dynamics 1978–1992.

Notes: Asset income included. The table reports the result of regressing current consumption growth on future income growth, including or excluding current income growth. Year-fixed effects. Standard errors are clustered at the household level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: Consumption growth regressions: income-growth innovations

	(1)	(2)
e_{t+1}	0.043** (0.019)	0.047** (0.021)
e_t		0.205*** (0.023)
Observations	10391	8911

Source: Panel Study of Income Dynamics 1978–1992.

Notes: The table reports the result of regressing current consumption growth on the innovations to income growth, including or excluding current income growth innovations. Year-fixed effects. Standard errors are clustered at the household level.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

the total income of households, that is, labor income plus asset income. Table 6 presents the corresponding consumption-growth regression results, which are very similar to the baseline results from Table 1.

Consumption growth and innovations. In the simplified model of Section 2, income follows a random walk and income growth captures the innovations to income in the absence of advance information. The survey data, on the other hand, reveal that income is characterized by mean reversion, and income growth is not equal to innovations in income. To capture innovations, we consider innovations to income growth as regressors for consumption growth. The corresponding estimation results are displayed in Table 7, with innovations stemming from an AR(1)-process for the growth rates, $\Delta y_{it} = \gamma_0 + \gamma_1 \Delta y_{it-1} + e_{it}$. As in the baseline, the correlation of consumption growth with future income growth innovations is significantly different from zero. Thereby, the regression coefficients for future innovations in Columns (1) and (2) are very similar, indicating a lower correlation of the regressors than in the baseline.

Covariance test to detect advance information. In Table 8, we implement the test for advance information entertained in [Blundell et al. \(2008\)](#). It tests the joint hypothesis that the covariance of consumption growth with future income growth is equal to zero for all sample years. Practically, 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 significantly different from zero. [Blundell et al. \(2008\)](#) find that there is no evidence of advance information. We replicate their result in Column (1) of Table 8.

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 approximately 4%, indicating the presence of advance information about the household's income two years ahead.

Table 8: 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%

Notes: 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 $\Delta y_{t+i}, 0 \leq i \leq j - 1$.

D. IDENTIFICATION OF INCOME PROCESS PARAMETERS

We target the following moments from the data: $\hat{\beta}_{y_{t-1}}$, $\hat{\sigma}_y$ from the income equation and the regression coefficient of current on future income growth $\hat{\delta}$.

Using that the persistent shock $\log(y_t^P)$ evolves according to $\log(y_t^P) = \rho \log(y_{P,t-1}) + \epsilon_{P,t}$ and for transitory shock $\epsilon_{T,t}$, we receive the first restriction to pin down the parameters of the transitory-persistent income process

$$\hat{\beta}_{y_{t-1}} = \frac{\text{cov}[\log(y_t), \log(y_{t-1})]}{\text{var}(y)} = \frac{\rho}{1 - \rho^2} \frac{\sigma_P^2}{\text{var}(y)}, \quad (\text{D.1})$$

where we used that the variance of log income is $\text{var}(y) = \hat{\sigma}_y^2 / (1 - \hat{\beta}_{y_{t-1}}^2)$, $\text{var}[\log(y_t^P)] = \sigma_P^2 / (1 - \rho^2)$, and $\sigma_T^2 = \text{var}(y) - \text{var}[\log(y_t^P)]$. The regression coefficient of current on future income growth is given by

$$\hat{\delta} = \frac{\text{cov}(\Delta_{y_t}, \Delta_{y_{t+1}})}{\text{var}(\Delta_{y_{t+1}})}. \quad (\text{D.2})$$

The variance in the denominator is

$$\begin{aligned} \text{var}(\Delta_{y_{t+1}}) &= \text{var} \left[(\rho - 1) \log(y_t^P) + \epsilon_{P,t+1} + \epsilon_{T,t+1} - \epsilon_{T,t} \right] \\ &= (\rho - 1)^2 \text{var} \left[\log(y_t^P) \right] + \sigma_P^2 + 2\sigma_T^2 \\ &= \sigma_P^2 \frac{(\rho - 1)^2 + 1 - \rho^2}{1 - \rho^2} + 2 \left[\text{var}(y) - \frac{\sigma_P^2}{1 - \rho^2} \right] \\ &= 2 \left[\text{var}(y) - \frac{\rho}{1 - \rho^2} \sigma_P^2 \right]. \end{aligned}$$

Using equation (D.1), the expression can be simplified to yield the following expression:

$$\text{var}(\Delta_{y_{t+1}}) = 2 \text{var}(y) (1 - \hat{\beta}_{y_{t-1}}).$$

The covariance in the numerator of $\hat{\delta}$ is

$$\begin{aligned}
\text{cov}(\Delta_{y_{t+1}}, \Delta_{y_t}) &= \text{cov} \left[\log(y_{t+1}^P) - \log(y_t^P) + \epsilon_{T,t+1} - \epsilon_{T,t}, \log(y_t^P) - \log(y_{t-1}^P) + \epsilon_{T,t} - \epsilon_{T,t-1} \right] \\
&= \text{cov} \left[\rho \log(y_t^P) + \epsilon_{P,t+1} - \log(y_t^P) + \epsilon_{T,t+1} - \epsilon_{T,t}, \rho \log(y_{t-1}^P) + \epsilon_{P,t} - \log(y_{t-1}^P) + \epsilon_{T,t} - \epsilon_{T,t-1} \right] \\
&= \text{cov} \left\{ (\rho - 1) \left[\rho \log(y_{t-1}^P) \right] + \epsilon_{P,t} + \epsilon_{P,t+1} + \epsilon_{T,t+1} - \epsilon_{T,t}, (\rho - 1) \log(y_{t-1}^P) + \epsilon_{P,t} + \epsilon_{T,t} - \epsilon_{T,t-1} \right\} \\
&= (\rho - 1)^2 \rho \text{var} \left[\log(y^P) \right] + (\rho - 1) \sigma_P^2 - \sigma_T^2 \\
&= \sigma_P^2 \left[\frac{\rho(\rho^2 - 2\rho + 1) + (\rho - 1)(1 - \rho^2)}{1 - \rho^2} \right] - \text{var}(y) + \frac{\sigma_P^2}{1 - \rho^2} \\
&= \sigma_P^2 \left[\frac{\rho(\rho^2 - 2\rho + 1) + (\rho - 1)(1 - \rho^2) + 1}{1 - \rho^2} \right] - \text{var}(y) \\
&= \sigma_P^2 \frac{\rho(2 - \rho)}{1 - \rho^2} - \text{var}(y).
\end{aligned}$$

Using again equation (D.1), we obtain

$$\text{cov}(\Delta_{y_{t+1}}, \Delta_{y_t}) = \text{var}(y)(2 - \rho)\hat{\beta}_{y_{t-1}} - \text{var}(y) = \text{var}(y)(2\hat{\beta}_{y_{t-1}} - \rho\hat{\beta}_{y_{t-1}} - 1),$$

and we can solve for ρ

$$\rho = \frac{2\hat{\beta}_{y_{t-1}} - 1 - 2\hat{\delta}(1 - \hat{\beta}_{y_{t-1}})}{\hat{\beta}_{y_{t-1}}}, \quad (\text{D.3})$$

and for σ_P using equation (D.1)

$$\sigma_P = \sqrt{\frac{1 - \rho^2}{\rho} \frac{\hat{\sigma}_y^2}{1 - \hat{\beta}_{y_{t-1}}^2} \hat{\beta}_{y_{t-1}}}. \quad (\text{D.4})$$

Given ρ, σ_P^2 , we receive for σ_T

$$\sigma_T = \sqrt{\frac{\hat{\sigma}_y^2}{1 - \hat{\beta}_{y_{t-1}}^2} - \frac{\sigma_P^2}{1 - \rho^2}}. \quad (\text{D.5})$$

E. 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. We adjust the discount factor, β , so that the mean wealth-to-income ratio is maintained at 2.9, and adjust the signal precision, κ , and the informal partial insurance parameter, θ , 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 a negative net worth of -0.201 .³⁶

³⁶We compute this ratio using the 1984 and 1989 PSID wealth modules and take the average.

Table 9: 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 9 summarizes the results. Notice that, while the estimates of informal partial insurance, θ , vary significantly in response to the changes to σ and \underline{a} , this variation is consistent with the relatively large standard error reported in Table 4. On the other hand, the estimates of κ and our model-independent measures of advance information and consumption insurance, $\tilde{\kappa}$ and β_{INS} , are robust to these variations.