UNIVERSITY OF AMSTERDAM

UvA-DARE (Digital Academic Repository)

Advance information and consumption insurance: Evidence and structural estimation

Pedroni, M.; Singh, Swapnil; Stoltenberg, C.A.

DOI
10.1016/j.jmoneco.2025.103748

Publication date
2025

Document Version
Final published version

Published in
Journal of Monetary Economics

License
CC BY

Link to publication

Citation for published version (APA):
Pedroni, M., Singh, S., & Stoltenberg, C. A. (2025). Advance information and consumption insurance: Evidence and structural estimation. *Journal of Monetary Economics*, *151*, Article 103748. https://doi.org/10.1016/j.jmoneco.2025.103748

General rights
It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations
If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: https://uba.uva.nl/en/contact, or a letter to: Library of the University of Amsterdam, Secretariat, P.O. Box 19185, 1000 GD Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (https://dare.uva.nl)

Download date:18 Nov 2026

Journal of Monetary Economics 151 (2025) 103748

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme

Advance information and consumption insurance: Evidence and structural estimation☆

Marcelo Pedroni [a],*, Swapnil Singh [b], Christian A. Stoltenberg [a]

[a] University of Amsterdam and Tinbergen Institute, The Netherlands
[b] Bank of Lithuania and Kaunas University of Technology, Lithuania

ARTICLE INFO

Keywords:
Advance information
Consumption insurance
Structural estimation
Income risk
Panel data
Incomplete markets

ABSTRACT

We show that households' private 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. Without advance information, the marginal benefits of public insurance are underestimated by an order of magnitude for some of the poorest wealth quantiles.

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 parts 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 et al. (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 for identifying advance information. Second, we use the same dataset as Blundell et al. (2008) and find that this conditional

☆ We are thankful to Adrien Auclert, Chris Busch, Francesco Caselli, Thomas Douenne, Wouter Den Haan, Christian Haefke, Albert Jan Hummel, Dirk Krueger, 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.

* Corresponding author.
E-mail address: m.pedroni@uva.nl (M. Pedroni).

[1] The 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 et al., 2007), and for the probability of leaving unemployment (Mueller et al., 2021).

https://doi.org/10.1016/j.jmoneco.2025.103748
Received 1 December 2023; Received in revised form 17 January 2025; Accepted 17 January 2025
Available online 27 January 2025
0304-3932/© 2025 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

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 and an underestimation of the effects of additional public insurance. Taken together, these contributions amount to a new method for measuring advance information and consumption insurance, even in the absence of data on individual income expectations.

We first consider an income-fluctuation model where households possess private advance information about their future income innovations. Households want to insure against persistent income shocks and have access to an informal private insurance scheme, as proposed by Guvenen and Smith (2014). We show that households adjust their consumption in response to both current and future income innovations. The consumption response to current income innovations depends on both the extent of advance information and the degree of insurance, while consumption responds to future income innovations only in proportion to the advance information that households possess.[2] Further, we find that using the correlation between consumption growth and future income growth – instead of future income innovations – to detect advance information requires controlling for current income growth. Without this control, the correlation is downward biased because current and future income growth are negatively correlated. Even with advance information, this downward bias, combined with the loss of predictive power from omitting current income growth, can render the unconditional correlation not significantly different from zero. Thus, relying on this correlation to test for advance information may lead to the incorrect conclusion that households have no advance information.

We use the insights from the stylized model to investigate whether US households possess advance information. 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, reproducing their findings: the corresponding unconditional correlation is not significantly different from zero and has a negative sign. However, in line with the logic from the stylized model, the picture changes when we control for current income growth. In this case, consumption growth is significantly positively correlated with both future and current income growth. These findings are robust. For example, they apply to both the nationally representative subsample, the Survey Research Center sample, and when we include the Survey of Economic Opportunity sample, which targets low-income households, when we consider different assumptions about the age bracket of the household head, or when we apply the alternative imputation procedure proposed by Attanasio and Pistaferri (2014).

A positive conditional correlation between consumption growth and future income growth is not definitive evidence of advance information. This is because income growth encompasses more than just innovations in income. A portion of income growth is predictable due to income persistence and mean reversion. Households' advance information, however, pertains specifically to future income innovations. To disentangle these components in the data, we estimate various income processes and extract the corresponding income innovation residuals. We then regress consumption growth on the estimated current and future income innovation residuals. Our results show that consumption growth is significantly positively correlated with future income innovation residuals, providing convincing evidence that US households indeed possess advance information.

We also use the stylized model to show that advance information matters for measuring consumption insurance. Traditionally, the correlation between current consumption growth and current income growth is used to assess consumption insurance. However, this method is only accurate in the absence of advance information. When advance information is present and direct data on households' income expectations is unavailable, a structural model becomes necessary. Notably, the PSID and most other large consumption and income panel datasets do not include direct information on households' income expectations.

To measure advance information and consumption insurance in the PSID, we 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, modeled by considering households that receive signals about their future income realizations. Second, we supplement households' self-insurance with an informal partial insurance scheme. We use indirect inference to structurally estimate this quantitative model. In the estimation procedure, the regression of consumption growth on current and future income growth serves as an auxiliary model. We find that the coefficient of this regression on future income growth identifies the degree of advance information in the model, in line with the stylized model.

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 poorest five percent of households, 58% of all unexpected income changes result in consumption changes with advance information, while without advance information, only 29% of all income shocks pass through to consumption. This overestimation, in turn, leads to a significant underestimation of the effect of additional public insurance. Ignoring advance information results in underestimating the elasticity of consumption insurance with respect to public insurance by an order of magnitude for households in some of the poorest wealth quantiles.

[2] This 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 by Kaufmann and Pistaferri (2009) and more recently by Guvenen and Smith (2014).

Literature review

Similarly 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 between current consumption and future income changes. Our empirical finding that US households possess advance information aligns with growing direct evidence on the predictive power of individual expectations. For instance, using US data, Dominitz (1998) estimates that individual income expectations reduce income forecast errors by 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. Notable exceptions include Kaplan and Violante (2010), and Stoltenberg and Uhlendorff (2025).[3] Kaplan and Violante (2010) investigate whether advance information on permanent shocks can bridge the gap between the level of consumption insurance estimated by Blundell et al. (2008) and the insurance predicted by a standard incomplete markets life-cycle model. In contrast, we introduce advance information (as well as partial insurance) in the model, use it to structurally estimate the amount of advance information possessed by US households, and then investigate the implications for consumption insurance. Using data from Italian households, Stoltenberg and Uhlendorff (2025) document that consumption choices depend on earnings expectations, suggesting that households possess advance information about their future earnings. Their approach relies on data on subjective income expectations, while our method requires only consumption and income data. Additionally, our structural estimation quantifies the degree of advance information and its implications for consumption insurance, whereas their study focuses on the relationship between consumption expenditures and income expectations.

The signal structure we introduce for modeling advance information is equivalent to the news specification used in Barsky and Sims (2011) for continuous income processes. However, solving standard incomplete-market models typically requires discretizing the income process. For discrete income processes, the method we introduce – with discrete signals that inform about future income states – requires fewer state variables, making computations with advance information and partial insurance feasible.

In a recent paper, Commault (2022) shows that the Blundell et al. (2008) measure of the pass-through of transitory income shocks to consumption is biased downwards. This bias arises from the MA(1) structure of the transitory component of income, which allows past shocks to affect consumption growth. Commault (2022) proposes a robust estimator of the pass-through that accounts for this dependence. The bias we identify in this paper is of a different nature. It stems from the dependence of consumption growth on future income shocks due to advance information. Moreover, our structural estimation of advance information and partial insurance allows consumption growth to depend on current, past, and future shocks.

Our work is also related to recent efforts to measure income uncertainty (Guvenen et al., 2021; Arellano et al., 2024, 2017). Using administrative data, Guvenen et al. (2021) find that the distribution of earnings changes exhibits non-Gaussian features such as negative skewness. Arellano et al. (2017) argue that earnings shocks are characterized by nonlinear persistence. Arellano et al. (2024) 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. Section 4 describes a quantitative structural model. In Section 5, we lay out our estimation method and present our estimates of advance information and consumption insurance in Section 6. In Section 7, we entertain counterfactual scenarios and study policy implications of advance information for consumption insurance. We conclude in Section 8.

2. Advance information and partial insurance: a stylized economy

In this section, we consider a stylized economy where households possess advance information about their future income and have access to an informal partial insurance scheme. We provide an analytical framework to illustrate how advance information and partial insurance shape the consumption response to both current and future income changes. These analytical results help motivate and interpret our empirical findings in Section 3. In Section 4, we extend this model in several dimensions and show that the main conclusions remain valid.

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

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

and a no-Ponzi constraint, with a_0 given. Here, c_t, a_t, and y_t^d denote household consumption, assets, 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 $\beta(1+r) = 1$ and that the utility function is quadratic, $u(c) = c - bc^2$.

Household income follows a persistent stochastic process. The initial income level, y_0, is given, and future income levels follow $y_t = \rho y_{t-1} + \eta_t$, with innovations $\eta_t \sim \mathcal{N}(0, \sigma^2)$ and $\rho \in [0, 1]$ as the auto-correlation coefficient. Households observe past and current

[3] Guvenen and Smith (2014) estimate a life-cycle model for the US economy in which households learn about their deterministic income profiles. Information about income profiles is distinct from the advance information about future shocks, which is the focus of this paper.

income levels and have advance information about the next period's income innovation, so that[4]

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

The parameter $\kappa \in [0,1]$ controls the degree of advance information; when $\kappa = 0$, there is no advance information and $\mathbb{E}_t[y_{t+1}] = \rho y_t$. 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 advance information only about the next period's income innovation, which implies that $\mathbb{E}_t[y_{t+1+j}] = (1-\kappa)\rho^{j+1} y_t + \kappa \rho^j y_{t+1}$, for all $j \geq 0$.

We also allow for an informal partial insurance scheme controlled by the parameter $\theta \in [0,1]$. 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*

$$c_t - c_{t-1} = \alpha_1 (y_t - \rho y_{t-1}) + \alpha_2 (y_{t+1} - \rho y_t), \tag{2}$$

with

$$\alpha_1 \equiv \frac{r}{1+r} \frac{1+r-\theta(1+r-\rho)}{1+r-\rho} (1-\kappa) \geq 0, \quad \textit{and} \quad \alpha_2 \equiv \frac{r}{1+r} \frac{1}{1+r-\rho} \kappa \geq 0.$$

Proof. See Appendix A.1. □

Proposition 2.1 shows that it is possible to disentangle advance information from informal partial insurance using only data on consumption and income changes. Specifically, controlling for current income innovations, $\eta_t = y_t - \rho y_{t-1}$, a future income innovation, $\eta_{t+1} = y_{t+1} - \rho y_t$, affects current consumption, $c_t - c_{t-1}$, in proportion to the household's advance information about the innovation (i.e. α_2 is proportional to κ). As a result, κ can be identified from this coefficient alone. Given κ, the coefficient on current income changes identifies θ.[5]

Given a particular income process, Eq. (2) provides a clear empirical strategy to test for advance information. First, estimate the income-process parameter ρ and extract the income innovation residuals, $\hat{\eta}_t = y_t - \hat{\rho} y_{t-1}$. Second, regress consumption growth on both current and future innovation residuals. If consumption growth is correlated with future innovation residuals, this indicates the presence of advance information. We follow this strategy below to obtain convincing evidence for advance information, which we then corroborate via structural estimation.

Our structural estimation via indirect inference requires an auxiliary model—a regression equation that can be used to identify the key elements of the model. By analogy with Eq. (2), we use the following regression that is independent of specific assumptions about the income process, replacing income innovations with income changes:

$$c_t - c_{t-1} = \beta_0 + \beta_{\Delta y_t} (y_t - y_{t-1}) + \beta_{\Delta y_{t+1}} (y_{t+1} - y_t) + \epsilon_t.$$

As argued above, the key to identifying advance information is the consumption response to future income changes, $\beta_{\Delta y_{t+1}}$. However, it is crucial that this coefficient be conditional on current income changes. We show below that the unconditional correlation suffers from omitted-variable bias if income growth is serially correlated. In the following lemma, we establish that income growth is negatively serially correlated when income is persistent.

Lemma 2.1. *Current and future income growth are negatively correlated:*

$$\delta \equiv \frac{\operatorname{cov}(y_t - y_{t-1}, y_{t+1} - y_t)}{\operatorname{var}(y_{t+1} - y_t)} = \frac{\rho - 1}{2} \leq 0.$$

Proof. See Appendix A.2. □

Given the negative serial correlation of income growth, we show in the following proposition that the unconditional correlation between consumption growth and future income growth is biased downward.

[4] This is equivalent to households having signals s_t that reveal a fraction κ of their next period's income innovation η_{t+1} in period t. The conditional income expectation $\mathbb{E}[y_{t+1}|y_t, s_t]$ is given by Eq. (1). This is also equivalent to the news specification from Barsky and Sims (2011). If income follows $y_{t+1} = \rho y_t + \epsilon_t + \mu_{t+1}$, with ϵ_t as the known component and μ_{t+1} as the true innovation, the same conditional expectation holds if κ is the ratio of standard deviations of ϵ_t and η_t.

[5] Our formulation of advance information and partial insurance follows Guvenen and Smith (2014), who consider a two-period version of the model. To obtain Eq. (2), the model must have at least three periods.

Proposition 2.2. *Consider the following consumption-growth regression:*

$$c_t - c_{t-1} = \widetilde{\beta}_0 + \widetilde{\beta}_{\Delta y_{t+1}}(y_{t+1} - y_t) + \widetilde{\epsilon}_t.$$

The regression results in a downward-biased estimate of the correlation between consumption growth and future income growth. The bias is given by:

$$\widetilde{\beta}_{\Delta y_{t+1}} - \beta_{\Delta y_{t+1}} = \delta\beta_{\Delta y_t} = \frac{\rho - 1}{2}\left[(\alpha_1 + \alpha_2)\frac{1-\rho}{3-\rho} + \alpha_1\rho\right] \leq 0.$$

Proof. See Appendix A.3. □

It follows that, to accurately detect advance information in the data, it is essential to control for current income growth when estimating the correlation between consumption growth and future income growth. The unconditional correlation, measured by $\widetilde{\beta}_{\Delta y_{t+1}}$, is biased downward when current and future income growth are negatively correlated. Even with advance information, this downward bias, combined with the loss of predictive power from omitting current income growth, can render the unconditional correlation not significantly different from zero. Thus, relying on this correlation to test for advance information may lead to the incorrect conclusion that households have no advance information.

To understand the implications for consumption insurance, notice that Eq. (2) can also be written as follows:

$$c_t - c_{t-1} = \frac{\alpha_1}{1-\kappa}\left(y_t - \mathbb{E}_{t-1}[y_t]\right) + \alpha_2(y_{t+1} - \rho y_t), \tag{3}$$

where, importantly, the expectation operator accounts for households' advance information. Therefore, to directly measure consumption insurance, it is necessary to have detailed data on household income expectations – and not the expectations of an econometrician. Our approach bypasses the fact that household expectation data is often unavailable. We derive households' income expectations from an estimated structural model of expectation formation with advance information, allowing us to identify the "true" income surprises from the household's perspective, $y_t - \mathbb{E}_{t-1}[y_t]$. This ensures that the first coefficient correctly measures consumption insurance.

We proceed in two steps. First, leveraging insights from Proposition 2.1, we reexamine the data to show that current consumption growth is positively correlated with future income growth when controlling for current income growth. We also show that consumption growth correlates with future innovation residuals from various estimated income processes. Second, we use this evidence as an auxiliary model within the framework of indirect inference to structurally estimate a standard incomplete markets model. This model includes persistent and transitory income shocks, precautionary savings, advance information, and informal partial insurance. We use the model to measure advance information, partial insurance, and consumption insurance.

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, when controlling for current income growth, the response is positive and significant. Following the logic of the stylized model from the previous section, we also find that consumption growth is correlated with future income innovations from various estimated income processes.

In what follows, the key variables of interest are after-tax household income (excluding asset income) and nondurable household consumption expenditures. Both variables are *residuals*, obtained after controlling for a vector of observable household characteristics.[6]

3.1. Baseline estimation results

One key takeaway from the previous section is that advance information impacts the correlation between 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 Blundell et al. (2008) and regress current consumption growth exclusively on future income growth:[7]

$$\Delta c_{it} = \widetilde{\beta}_0 + \widetilde{\beta}_{\Delta y_{t+1}}\Delta y_{it+1} + \widetilde{\epsilon}_{it}, \tag{4}$$

where $\Delta x_{it} \equiv \log(x_{it}) - \log(x_{it-1})$ denotes the growth rate of variable x. We refer to the coefficient $\widetilde{\beta}_{\Delta y_{t+1}}$ 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:

$$\Delta c_{it} = \beta_0 + \beta_{\Delta y_t}\Delta y_{it} + \beta_{\Delta y_{t+1}}\Delta y_{it+1} + \epsilon_{it}. \tag{5}$$

[6] A detailed description of the data, sample selection, and methods used for imputing consumption data are provided in Appendix B.

[7] The stylized linear-quadratic model analyzed above can be thought of as an approximate version of a model in which the level variables that appear in Eq. (2) would be replaced by log deviations from the steady state, so that overtime differences become growth rates.

Table 1
Baseline estimates: consumption growth regressions.
Source: Panel Study of Income Dynamics 1978–1992.

	(1)	(2)
Δy_{t+1}	−0.013	0.045**
	(0.018)	(0.019)
Δy_t		0.184***
		(0.021)
Observations	10 502	10 471

Notes: Baseline specification. The table reports the results of regressing current consumption growth on future income growth, with and without including current income growth. Year fixed effects are included. Standard errors are clustered at the household level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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

In Table 1, we report our baseline estimates for the two regressions. The first column displays the correlation from regression Eq. (4), while the second column shows the results from regression Eq. (5). We include year fixed effects and cluster standard errors at the household level. The main findings 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 a p-value of 0.014.

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

The results in Table 1 directly align with Lemma 2.1 and Proposition 2.2. First, in the PSID data, current and future income are significantly negatively correlated, with $\delta = -0.325$, precisely estimated with a standard error of 0.011. Second, given this negative correlation and the significant positive values of $\beta_{\Delta y_t}$ in all regressions shown in Table 1, it follows that $\widetilde{\beta}_{\Delta y_{t+1}} < \beta_{\Delta y_{t+1}}$. Omitting current income growth as a regressor leads to a downward bias in the correlation between current consumption and future income growth in the first regression by $\delta\beta_{\Delta y_t}$.

The traditional approach to measuring consumption insurance involves regressing current consumption growth on current income growth:

$$\Delta c_{it} = \widetilde{\beta}_0 + \widetilde{\beta}_{\Delta y_t} \Delta y_{it} + \widetilde{\eta}_{it}. \tag{6}$$

Invoking the logic of Proposition 2.2 once again, the takeaway from Table 1 is that this measure also tends to be downward biased when future income growth is omitted from Eq. (6). Specifically, it follows that:

$$\widetilde{\beta}_{\Delta y_t} - \beta_{\Delta y_t} = \tilde{\delta}\beta_{\Delta y_{t+1}} < 0,$$

where $\tilde{\delta}$ is the coefficient from regressing future income on current income growth. As a result, we estimate a downward-biased coefficient $\widetilde{\beta}_{\Delta y_t} = 0.151$, which is less than the true value $\beta_{\Delta y_t} = 0.184$.

In Appendix C, we present several robustness exercises. Our baseline estimation results are robust to including asset income (Table C.1), using income growth innovations as regressors (Table C.2), and applying the alternative imputation procedure proposed by Attanasio and Pistaferri (2014) (Table C.3). In Table C.4, we replicate the test for advance information used by Blundell et al. (2008), who examine whether the unconditional covariances satisfy $\text{cov}(\Delta c_{it}, \Delta y_{it+j}) = 0$ for all years t and $j \geq 1$. We reproduce their finding that these unconditional covariances are statistically insignificant. However, once we control for previous income changes, we safely reject the null hypothesis, confirming our baseline estimation results.

3.3. Do US households possess advance information on their future income?

While the estimation results from Table 1 are suggestive, they do not conclusively establish that households have advance information about their future income. Since current and future income growth are negatively correlated, income growth does not exclusively reflect income innovations. Conditional on the information available in period t and on current income growth, future income growth can be expressed as:

$$\Delta y_{t+1} = \underbrace{\widetilde{\mathbb{E}}_t[\log(y_{t+1})] - \log(y_t)}_{\text{known in } t} + \underbrace{\eta_{t+1}}_{\text{innovation}},$$

where $\widetilde{\mathbb{E}}_t$ denotes the expectation operator for an econometrician without access to households' advance information. Future income growth has two components: the known component, $\widetilde{\mathbb{E}}_t[\log(y_{t+1})] - \log(y_t)$, and the innovation component, η_{t+1}. Advance information pertains to signals that households receive about the second component, the innovation η_{t+1}. If consumption growth is positively correlated with the known component, the coefficient $\beta_{\Delta y_{t+1}}$ could be significantly different from zero even without advance information, potentially due to transitory shocks.

Table 2
Consumption-growth regressions on current and future income innovation residuals.
Source: Panel Study of Income Dynamics 1978–1992.

	ARMA(1,1) (1)	ARMA(1,2) (2)	ARMA(2,1) (3)	ARMA(2,2) (4)
Current innovation residuals, $\hat{\eta}_t$	0.236***	0.226***	0.236***	0.222***
	(0.027)	(0.031)	(0.027)	(0.032)
Future innovation residuals, $\hat{\eta}_{t+1}$	0.060**	0.065**	0.060**	0.064**
	(0.024)	(0.028)	(0.024)	(0.028)

Notes: The table reports the result of regressing current consumption growth on current and future income innovations residuals of various estimated ARMA income processes: $\log(y_{it}) = \rho_1 \log(y_{it-1}) + \rho_2 \log(y_{it-2}) + \theta_1 \eta_{it-1} + \theta_2 \eta_{it-2} + \eta_{it}$. Year-fixed effects are included. Standard errors are clustered at the household level. There are 7569 observations for Columns (1) and (3), and 6290 for Columns (2) and (4). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

To further investigate the presence of advance information, we proceed as follows. Eq. (2) suggests that the correlation of consumption growth with future income innovation residuals $\hat{\eta}_{t+1}$ from an estimated income process indicates advance information. By regressing on the innovation residuals, we can isolate the effect of $\widetilde{\mathbb{E}}_t[\log(y_{t+1})]$ on $\beta_{\Delta y_{t+1}}$.

First, we estimate ARMA models for log income to obtain the econometrician's expected income, $\widetilde{\mathbb{E}}_t[\log(y_{t+1})]$, excluding households' advance information. Ejrnæs and Browning (2014) demonstrate that stochastic processes with two shocks – one permanent or persistent and one transitory – can be represented by an ARMA process with a single shock. For instance, the canonical persistent-transitory process corresponds to an ARMA(1,1), while the permanent-transitory process with MA(1) transitory shocks, as used by Blundell et al. (2008) and more recently by Commault (2022), can be represented as an ARMA(1,2).

Second, we regress consumption growth on both the current and future innovation residuals, $\hat{\eta}_t$ and $\hat{\eta}_{t+1}$, obtained from the first step. Specifically, we run the following regression:

$$\Delta c_t = \beta_{\hat{\eta}_t} \hat{\eta}_t + \beta_{\hat{\eta}_{t+1}} \hat{\eta}_{t+1} + v_t.$$

Regression coefficients on future income innovation residuals, $\beta_{\hat{\eta}_{t+1}}$, that are significantly different from zero indicate the presence of advance information. The corresponding estimation results are presented in Table 2. We find all coefficients $\beta_{\hat{\eta}_{t+1}}$ to be significantly different from zero, with p-values well below five percent. These results suggest that the positive correlation of consumption with future income growth is not solely due to transitory shocks.

Next, to estimate the degree of advance information households possess without direct data on their expectations and to measure the implications for consumption insurance, we need a structural model. While the estimation results for Eq. (5) do not imply advance information by themselves, we argue that they are useful for identifying advance information conditional on a given income process and a structural model. More specifically, we employ indirect inference using regression (5) as an auxiliary model to estimate how much advance information US households possess in the context of a standard incomplete markets model.[8] Given an estimated amount of advance information, consumption insurance can be quantified by the consumption response to unexpected income changes, $\log(y_{it}) - \mathbb{E}_{it-1}[\log(y_{it})]$, with expectations that account for households' advance information.[9]

4. Structural model

In this section, we present a quantitative version of the standard incomplete markets model incorporating 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 each period, households receive income $y \in Y$ and observe a signal $s \in Y$ about their next period income. The (y, s) pair follows a Markov process specified below. Households can accumulate a risk-free asset, a, and face borrowing constraints, such that the set of possible values for a is $A \equiv [\underline{a}, \infty)$. 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 in the last period, with $Z \equiv A \times Y^4$. 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

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

The definition of the stationary equilibrium for this economy can be found in Appendix D.

[8] We use regression (5) as the auxiliary model in part because it is independent of the income process specification. As a robustness exercise, we also used the results in Table 2 as the auxiliary model and obtained similar results.

[9] This is consistent with the insights from Jappelli and Pistaferri (2010), who also advocate for estimating the correlation between consumption changes and unexpected income changes.

Table 3
Preset parameters.

Parameter		Value
σ	Relative risk aversion	2.0
$\underline{a}$	Borrowing constraint	0.0
r	Interest rate	3.9%

Signal structure. Household income follows a Markov process with transition probabilities $\Pr(y'|y)$. The signal structure is modeled as in the stylized model of Section 2, in a parsimonious way that minimally extends the state space. The transition probabilities, conditional on receiving signal $s \in Y$, are given by:

$$\Pr(y'|y,s) = \begin{cases} \kappa + (1-\kappa)\Pr(y'|y), & \text{if } s = y' \\ (1-\kappa)\Pr(y'|y), & \text{otherwise.} \end{cases} \tag{7}$$

Each period, the signal observed by the household indicates which of the possible next-period income levels they will receive. With probability κ, the signal is informative and their next-period income is exactly what the signal predicted. With probability $(1-\kappa)$, the signal is uninformative, and the transition probabilities revert to the unconditional Markov process, $\Pr(y'|y)$. The parameter $\kappa \in [0,1]$ controls the degree of advance information. When $\kappa = 0$, there is no advance information and $\Pr(y'|y,s) = \Pr(y'|y)$. When $\kappa = 1$, households know their next-period income with certainty, that is, $\Pr(y'|y, s = y') = 1$ and $\Pr(y'|y, s \neq y') = 0$.

To complete the description of the signal process, we need to specify the probability of receiving a particular signal. For the signal process to be internally consistent with the income process and Bayes' rule, it must be that

$$\Pr(s = y'|y) = \Pr(y'|y), \tag{8}$$

that is, the probability of receiving a signal $s = y'$ given current income y must equal the probability of transitioning from income y to y'.[10] This ensures that the signal process is consistent with rational expectations.

5. Quantitative exercise

In this section, we first specify the functional forms for the income process and household preferences, along with the preset parameters. We then discuss the parameters that are structurally estimated using indirect inference, including the precision of signals and the extent of informal partial insurance. Finally, we define measures for advance information and consumption insurance.

5.1. Preset parameters and the income process

The preset parameters are organized in Table 3. The utility function is given by $u(c) = c^{1-\sigma}/(1-\sigma)$, with the degree of relative risk aversion set to $\sigma = 2$, and the borrowing limit is set to $\underline{a} = 0$.[11] The interest rate is set to $r = 3.9\%$, corresponding to the annual real risk-free rate for the US post-1980 as documented by Jordà et al. (2019). Household income has persistent and transitory components:

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

with

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

After estimating the parameters of these processes, as described below, we use the Tauchen and Hussey (1991) method to discretize them. This results in a Markov process for income y with transition probabilities $\Pr(y'|y)$. The corresponding income-conditional expectations, $\widetilde{\mathbb{E}}[y'|y]$, represent the expectations of an econometrician without advance information. Applying the signal structure described in the previous section, we obtain households' income expectations with advance information, $\mathbb{E}[y'|y,s]$.[12]

[10] In particular, the condition in Eq. (8) guarantees that $\sum_{s \in Y} \Pr(s|y)\Pr(y'|y,s) = \Pr(y'|y)$.

[11] In Appendix F, we conduct robustness exercises with respect to these choices.

[12] The signal process can be easily extended to this persistent/transitory income process by combining the two components into one large Markov income process. Alternatively, to keep the state space small, one can exploit the persistent/transitory structure and *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. The signal probabilities are given by $\Pr((s_P, s_T) = (y'_P, y'_T)|y_P) = \Pr(y'_P|y_P)\Pr(y'_T)$.

Table 4
First-step parameter estimates.

Parameter		Value
σ_P	Standard deviation, persistent shocks	0.192
σ_T	Standard deviation, transitory shocks	0.162
ρ	Auto-correlation coefficient	0.891

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. 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.[13] The first equation is the regression equation (5), which captures how current consumption growth depends on current and future income growth:

$$\Delta c_{it} = \beta_0 + \beta_{\Delta y_t} \Delta y_{it} + \beta_{\Delta y_{t+1}} \Delta y_{it+1} + \epsilon_{it}. \tag{9}$$

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).[14] The second equation represents the dynamics of after-tax income:

$$\log(y_{it}) = \beta_{y,0} + \beta_{y_{t-1}} \log(y_{it-1}) + \epsilon_{y,it}, \quad \epsilon_{y,it} \sim \mathcal{N}(0, \sigma_y^2). \tag{10}$$

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).

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 those estimated on the simulated data from the economic model. The moments 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 Eq. (9), and by $\beta_{y_{t-1}}$ and σ_y from the income Eq. (10).[15] 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.[16]

In total, the auxiliary model comprises six moments, which matches the length of the parameter vector Θ. Thus, we consider a case of exact identification, allowing us to estimate the parameters in steps. In the first step, we use only income data to estimate the parameters of the income process $(\sigma_P, \sigma_T, \rho)$. In the second step, we employ income and consumption data to identify κ and θ, and use the wealth-to-income ratio to pin down β.

More precisely, in the first step, we analytically solve for $(\sigma_P, \sigma_T, \rho)$ such that $(\beta_{y_{t-1}}(\Theta), \sigma_y(\Theta), \delta(\Theta))$ equals its income-data counterpart, $(\hat{\beta}_{y_{t-1}}, \hat{\sigma}_y, \hat{\delta})$.[17] The corresponding parameter estimates are displayed in Table 4.[18] 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. Then, 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 counterparts $(\hat{\beta}_{\Delta y_t}, \hat{\beta}_{\Delta y_{t+1}})$. The identification and estimation results for κ and 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 based on current income. We measure advance information by the relative reduction in household income forecast error,

[13] Note that the auxiliary model need not be correctly specified. Instead, its main purpose is to identify the structural parameters in Θ.

[14] These standard errors are computed from a nonparametric bootstrap with 500 repetitions.

[15] In our baseline estimation, we do not include the residual standard deviation in Eq. (9), σ_ϵ, as a coefficient in the auxiliary model, which is estimated at 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 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.

[16] We compute the wealth-to-income ratio using the 1984 and 1989 PSID wealth modules. Our measure of wealth includes the total net value of farm or business assets, savings through financial assets, the net value of durable assets including housing, and the 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.

[17] In Appendix E, we derive the analytical expressions.

[18] We also estimate ρ using the GMM approach described in Hryshko and Manovskii (2018). For our baseline sample, we obtain $\rho = 0.872$ with a standard error of 0.012, which is close to the value reported in Table 4.

defined as:

$$AI \equiv \frac{\text{MSFE}_y - \text{MSFE}_{y,s}}{\text{MSFE}_y} \in [0, 1], \tag{11}$$

with

$$\text{MSFE}_y \equiv \sum_y \pi(y) \sum_{y'} \Pr(y'|y) \left\{ \log(y') - \widetilde{\mathbb{E}} \left[\log(y') \mid y \right] \right\}^2,$$

$$\text{MSFE}_{y,s} \equiv \sum_{y,s} \pi(y,s) \sum_{y'} \Pr(y'|y,s) \left\{ \log(y') - \mathbb{E} \left[\log(y') \mid y, s \right] \right\}^2 \leq \text{MSFE}_y .$$

Here, $\pi(y)$ is the invariant distribution of income, and $\pi(y, s)$ is the joint invariant distribution of income and signals. We use this measure to quantify households' advance information.[19] Unlike the parameter κ that controls the precision of signals in the model, the measure AI allows for direct comparison of our estimates of advance information in Section 6 to the direct evidence on the predictive power of individual income expectations.

In the model, households use self-insurance and informal partial insurance to protect their consumption against unexpected income changes. Therefore, we measure consumption insurance via the income-consumption pass-through coefficient β_{INS} in the following regression:

$$\Delta c_{it} = \text{constant} + \beta_{INS} \left\{ \log(y_{it}) - \log \left(\mathbb{E}[y_{it} \mid y_{it-1}, s_{it-1}] \right) \right\} + \epsilon_{it}, \tag{12}$$

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

6. Second-step estimation results, advance information and insurance

We begin with a discussion on identification. Next, we present the parameter estimates. Using these estimates, we quantify the amount of advance information that US households possess and assess their ability to insure against unexpected income shocks. Finally, we study counterfactual scenarios without advance information to demonstrate that ignoring advance information results in a significant overestimation of consumption insurance and an underestimation of the effects of public 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 this identification strategy extends to the quantitatively more realistic model when we use these two coefficients as part of our auxiliary model in Eq. (9).

Fig. 1 plots the absolute deviations of the two regression coefficients estimated on data simulated in the structural model from their survey data counterparts. The right panel of Fig. 1 shows that the identification of κ stems from the regression coefficient that captures the covariance 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 to the quantitative model in this section, validating our auxiliary model choice.

6.2. Parameter estimates and implications

In Table 5, we summarize the parameter estimates from the final two steps described in Section 5.2, along with the implications of these estimates 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 at a standard value of 0.944. Signal precision is estimated at $\kappa = 0.391$ with a standard error of 0.023. The informal partial insurance parameter θ is estimated at 0.158, albeit less precisely, with a standard error of 0.111.[20]

The parameter estimates imply that, compared to an econometrician, households' advance information on future income shocks reduces their forecast error by $AI = 15.3\%$, a value that is precisely estimated. This finding aligns with the direct evidence provided in Dominitz (1998), who reports reductions in mean-squared forecast errors when conditioning on income expectations ranging from 12 to 21%.

Furthermore, the parameter estimates indicate that US households only partially insure against income shocks, 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, this income-consumption pass-through, as a measure of total consumption insurance, is precisely estimated with a standard error of 0.012. Looking at both income shocks separately, the pass-through coefficient for persistent shocks amounts to $\beta_{INS,P} = 0.395$, while transitory shocks are better insured with $\beta_{INS,T} = 0.093$.[21]

[19] For our specification of signals in Section 4, one can show with tedious but straightforward algebra that $AI = \kappa^2$.

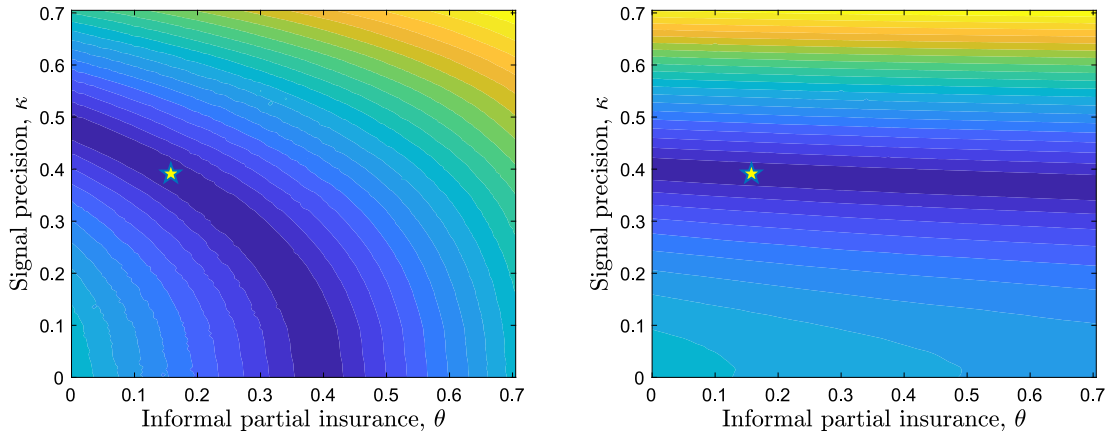

(a) Current income growth coefficient (b) Future income growth coefficient

Fig. 1. 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.

Table 5
Second-step 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, AI	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.

Robustness. Our estimates of advance information, κ and AI, as well as consumption insurance, β_{INS}, are robust to changes in the degree of risk aversion or the borrowing limit. We also obtain similar estimates of advance information and consumption insurance when we use the regression results of consumption growth with respect to income innovations from Table 2 as the auxiliary model, instead of the consumption–income growth regression. Details on these robustness exercises can be found in Appendix F.

In Section 3.3, we discuss why a coefficient $\beta_{\Delta y_{t+1}}$ significantly different from zero does not necessarily indicate advance information. A positive regression coefficient could result from the properties of the income process alone. Intuitively, the larger the correlation induced by the income process, the smaller the contribution of advance information to the regression coefficient. While we specify i.i.d. transitory shocks, Blundell et al. (2008) consider MA(1) transitory shocks. A concern might be that MA(1) transitory shocks could account for a larger fraction of the regression coefficient than i.i.d. shocks. However, our findings indicate the opposite. An income specification with i.i.d. transitory shocks explains slightly more of the regression coefficient than one with MA(1) transitory shocks. Thus, our estimates in Table 5 can be interpreted as conservative estimates of advance information. Further details on this robustness exercise can be found in Appendix F.

In Appendix G, we conduct another robustness experiment. We suppress partial insurance by setting $\theta = 0$ and extend the signal structure to allow for separate signals for the persistent and transitory components of income, with different levels of informativeness, κ_P and κ_T, respectively. We find that households have significantly more information about the persistent part of their income, with $\kappa_P = 0.470$ and $\kappa_T = 0.060$. This version of the model is also able to closely approximate the coefficients of the consumption growth regression, with $\beta_{\Delta y_t} = 0.188$ and $\beta_{\Delta y_{t+1}} = 0.045$. The implications for consumption insurance are consistent with the benchmark, with $\beta_{INS} = 0.268$ (with $\beta_{INS,P} = 0.413$ and $\beta_{INS,T} = 0.109$), while the reduction in overall forecast errors is slightly lower, $AI = 0.125$. We also investigate consumption insurance with respect to persistent and transitory shocks separately and find that, in line with the permanent income hypothesis, the pass-through of persistent shocks is significantly higher than that of transitory shocks. Still,

[20] In our baseline estimation, we do not include the residual standard deviation in Eq. (9), 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 $\sigma_c = 0.278$, similar to the value of 0.36 found in Guvenen and Smith (2014).

[21] The pass-through coefficients for persistent and transitory shocks are derived from the following regression equation: Δc_{it} = constant + $\beta_{INS,P}\left\{\log(y_{P,it}) - \log\left(\mathbb{E}_{t-1}[y_{P,it}]\right)\right\} + \beta_{INS,T}\left\{\log(y_{T,it}) - \log\left(\mathbb{E}_{t-1}[y_{T,it}]\right)\right\} + \epsilon_{it}$. For more details, see Appendix G.

Table 6
Counterfactual scenarios.

	Parameter estimates			Implications					
	κ	θ	β	AI	β_{INS}	$\beta_{INS,P}$	$\beta_{INS,T}$	$\beta_{\Delta y_t}$	$\beta_{\Delta y_{t+1}}$
Benchmark	0.391	0.158	0.944	0.153	0.269	0.395	0.093	0.184	0.045
No AI — Scenario 1	0.000	0.741	0.948	0.000	0.201	0.323	0.030	0.151	0.025
No AI — Scenario 2	0.000	0.620	0.947	0.000	0.217	0.340	0.045	0.163	0.024
No PI — Scenario 3	0.405	0.000	0.943	0.164	0.287	0.413	0.110	0.193	0.046

this difference is remarkably similar to the corresponding one in our benchmark calibration. Overall, the implications of the model for the degree of advance information and consumption insurance are not sensitive to this variation.

We also investigate, in Appendix H, the possibility that the degree of insurance and advance information differs across education levels. We divide our baseline sample into two groups, college and no-college, based on the educational attainment of the household head. We then leverage the partial-equilibrium nature of our model and replicate our baseline experiment for each subsample separately. We find stark differences: college households exhibit substantially more advance information, with a forecast error reduction (AI) of 0.199, roughly double the AI of 0.099 observed for the no-college group.[22] Additionally, college households are better able to insure against income shocks, with a pass-through coefficient (β_{INS}) of 0.176 compared to 0.276 for the no-college sample.

7. Counterfactual scenarios and policy implications

In the following, we first study counterfactual scenarios in which we suppress advance information. We then investigate the implications of advance information for consumption insurance and the effect of policy across the wealth distribution.

7.1. Counterfactual scenarios

We compute two scenarios that show that neglecting advance information leads to downward-biased estimates of the income-consumption pass-through, indicating too much consumption insurance. In a third scenario, we suppress partial insurance but retain advance information, finding that consumption insurance remains similar to the benchmark. Table 6 summarizes our findings on economies with and without advance information.

Scenario 1: No advance information and targeting a downward-biased regression coefficient. Suppose advance information is assumed to be absent, with $\kappa = 0$, and that the downward-biased regression coefficient $\widetilde{\beta}_{\Delta y_t} = 0.151$ is targeted to estimate θ. In this case, we estimate $\theta = 0.741$, which is substantially higher than in our baseline estimation with advance information. This occurs for two reasons. First, targeting the lower coefficient $\widetilde{\beta}_{\Delta y_t} = 0.151 < 0.184 = \beta_{\Delta y_t}$ requires a larger amount of informal partial insurance. Second, without advance information, the entire response of consumption to income changes is attributed to insurance, both informal and self-insurance, which increases the estimate of θ.[23] 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%. For persistent shocks, the pass-through coefficient drops to 0.323 (from 0.395 in the benchmark). Further, only 3% of all transitory shocks pass through to consumption, compared to 9% in the benchmark with advance information.

Scenario 2: No advance information but targeting the coefficients $\beta_{\Delta y_t}$, and $\beta_{\Delta y_{t+1}}$. Our benchmark estimate for the informal partial insurance parameter, $\theta = 0.158$, is lower than the corresponding value found in the related work of Guvenen and Smith (2014), who estimate $\theta = 0.451$. Their model, like in Scenario 1, does not allow for advance information on future income shocks.[24] However, unlike Scenario 1, their auxiliary model does allow consumption to respond to future income changes. In an experiment that is closer to Guvenen and Smith (2014), we continue to assume households do not possess advance information and estimate θ by targeting the two regression coefficients $\beta_{\Delta y_t} = 0.184$ and $\beta_{\Delta y_{t+1}} = 0.045$. We find $\theta = 0.620$, which is closer to Guvenen and Smith (2014)'s estimate.[25] As expected, the model without advance information struggles to capture the consumption response to future income changes but accurately reproduces the response to current income changes. Nevertheless, the model still predicts excessive insurance, with a pass-through of 0.217, which is 20% lower than our baseline estimate of 0.269 obtained in the model with advance information. The same pattern holds when examining pass-throughs for both income components separately. For persistent shocks, we now estimate $\beta_{INS,P} = 0.339$ and for transitory shocks $\beta_{INS,T} = 0.044$, compared to 0.395 and 0.093 in the benchmark.

[22] These results are consistent with Ozkan et al. (2023), who find that low-wage workers experience higher job-ladder risk.

[23] The overestimation of insurance due to ignoring advance information can be understood within the simple model of Section 2. Both κ and θ reduce the first coefficient in Eq. (2). Thus, restricting $\kappa = 0$ leads to a higher estimate for θ, which subsequently reduces the pass-through in Eq. (3).

[24] 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.

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

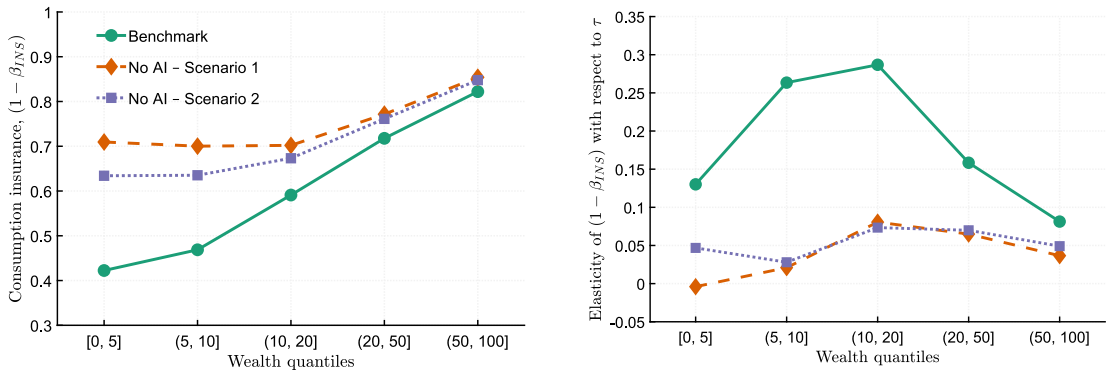

(a) Consumption insurance

(b) Effect of public insurance

Fig. 2. Insurance by quantiles of wealth. *Notes:* The left panel plots consumption insurance, measured by $(1-\beta_{INS})$, conditional on the wealth quantiles. The right panel plots the elasticity of $(1-\beta_{INS})$ with respect to τ.

Scenario 3: No informal partial insurance but targeting the coefficients $\beta_{\Delta y_t}, \beta_{\Delta y_{t+1}}$. Given that informal partial insurance is less precisely estimated, it is natural to investigate whether its omission affects the estimation of consumption insurance and advance information. For this exercise, we set $\theta = 0$ and estimate κ by targeting the two regression coefficients $\beta_{\Delta y_t} = 0.184$ and $\beta_{\Delta y_{t+1}} = 0.045$. The corresponding results are displayed in the fourth row of Table 6. We estimate $\kappa = 0.405$, which is only slightly higher than the baseline estimate of $\kappa = 0.391$. The resulting reduction in forecast errors is $AI = 0.164$ compared to 0.151 in the baseline. The income-consumption pass-through is somewhat higher, with $\beta_{INS} = 0.287$ compared to 0.269. The pass-through coefficients for persistent and transitory shocks are also moderately higher than in the benchmark economy. The two regression coefficients are well-approximated, given that only one parameter, κ, is used to match the two targets. The model slightly overstates the consumption response to current income growth, which explains the higher income-consumption pass-through.

7.2. Insurance across the wealth distribution and implications for policy

Fig. 2 shows how consumption insurance varies across the wealth distribution under the different scenarios considered above. As expected, households with higher levels of wealth are better able to insure their consumption, resulting in a lower pass-through and a higher measure of consumption insurance, $(1-\beta_{INS})$. However, the degree to which wealthier households are better insured is more pronounced in the benchmark scenario. As a result, the overestimation of consumption insurance from ignoring advance information is significantly magnified at the bottom of the wealth distribution. For the 5-percent poorest households, 42% of all shocks do not pass through to consumption in the benchmark, compared to 71% in No AI — Scenario 1. This discrepancy is not due to more households being borrowing constrained in the economies with advance information (in fact, the opposite is true). The key mechanism is that with advance information, households save less when they receive a positive signal about their future income and are then particularly unprepared if the signals are incorrect. In other words, the worst surprise in an economy without advance information is receiving a negative income shock, whereas in an economy with advance information, it is receiving a positive income signal followed by a negative income shock.

To investigate the policy relevance of these results, we also estimate the improvement to consumption insurance, measured by $(1-\beta_{INS})$, resulting from an increase in public insurance. To model public insurance parsimoniously, we consider a simple tax and transfer scheme,

$$y_{t,\text{new}} = (1-\tau)y_t + \tau\bar{y},$$

with an income tax τ that is rebated via lump-sum transfers, $\tau\bar{y}$, where $\bar{y}$ denotes the average income, ensuring the policy is budget-neutral. Since the income y_t already denotes after-tax income, a positive τ should be interpreted as providing additional progressivity relative to the current US tax system.

In Fig. 2(b), we report the elasticity of $(1-\beta_{INS})$ with respect to τ for the same calibration scenarios and wealth quantiles. Increases in public insurance are significantly more effective in improving consumption insurance when advance information is present. For households with wealth between the 5th and 10th percentiles, the improvement in consumption insurance is higher by an order of magnitude. Public insurance is generally the least effective in No AI — Scenario 1, where advance information is neither included in the model nor taken into account in the consumption growth regression used to discipline the model. This policy experiment shows that ignoring advance information leads to a substantial underestimation of the effects of public insurance schemes.

8. Conclusion

In this paper, we have investigated whether US households possess advance information about their future income. In a simple income-fluctuation model, we first show that when households possess advance information, they adjust consumption in response to both 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 using 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 of abstracting from advance information results in an overestimation of consumption insurance. On average, 27% of all unexpected income changes are passed through to consumption with advance information, whereas 25% fewer income shocks pass through without advance information. This overestimation of consumption insurance is more pronounced for households at the bottom of the wealth distribution. Suppressing 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. This overestimation of consumption insurance without advance information leads to a significant underestimation of the effect of additional public insurance. The elasticity of consumption insurance with respect to public insurance is underestimated by an order of magnitude for some of the poorest wealth quantiles.

We believe the methods developed in this paper can be extended in several interesting directions. One could explore additional data moments that could enable the proper identification of additional parameters such as the precision of information households have about the persistent and transitory components of income. It would also be interesting to introduce other relevant trade-offs that interact with the household consumption-savings decision. Obvious examples include labor supply and household composition decisions.

Online Appendix

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jmoneco.2025.103748.

Data availability

Data will be made available on request.

References

Arellano, Manuel, Blundell, Richard, Bonhomme, Stéphane, 2017. Earnings and consumption dynamics: a nonlinear panel data framework. Econometrica 85 (3), 693–734.
Arellano, Manuel, Blundell, Richard, Bonhomme, Stéphane, Light, Jack, 2024. Heterogeneity of consumption responses to income shocks in the presence of nonlinear persistence. Journal of Econometrics 240 (2), 105449.
Attanasio, Orazio, Pistaferri, Luigi, 2014. Consumption inequality over the last half century: Some evidence using the new PSID consumption measure. Amer. Econ. Rev. 104 (5), 122–126.
Barsky, Robert B., Sims, Eric R., 2011. News shocks and business cycles. J. Monetary Econ. 58 (3), 273–289.
Blundell, Richard, Pistaferri, Luigi, Preston, Ian, 2008. Consumption inequality and partial insurance. Amer. Econ. Rev. 98 (5), 1887–1921.
Campbell, David, Carruth, Alan, Dickerson, Andrew, Green, Francis, 2007. Job insecurity and wages. Econ. J. 117 (518), 544–566.
Commault, Jeanne, 2022. Does consumption respond to transitory shocks? Reconciling natural experiments and semistructural methods. Am. Econ. J.: Macroecon. 14 (2), 96–122.
Dominitz, Jeff, 1998. Earnings expectations, revisions, and realizations. Rev. Econ. Stat. 80 (3), 374–388.
Dominitz, Jeff, Manski, Charles F., 1997. Using expectations data to study subjective income expectations. J. Amer. Statist. Assoc. 92 (439), 855–867.
Ejrnæs, Mette, Browning, Martin, 2014. The persistent–transitory representation for earnings processes. Quant. Econ. 5 (3), 555–581.
Guvenen, Fatih, Karahan, Fatih, Ozkan, Serdar, Song, Jae, 2021. What do data on millions of U.S. workers reveal about lifecycle earnings dynamics? Econometrica 89 (5), 2303–2339.
Guvenen, Fatih, Smith, Anthony A., 2014. Inferring labor income risk and partial insurance from economic choices. Econometrica 82 (6), 2085–2129.
Hendren, Nathaniel, 2017. Knowledge of future job loss and implications for unemployment insurance. Amer. Econ. Rev. 107 (7), 1778–1823.
Hryshko, Dmytro, Manovskii, Iourii, 2018. On the heterogeneity in family earnings and income dynamics in the PSID. AEA Pap. Proc. 108, 292–296.
Jappelli, Tullio, Pistaferri, Luigi, 2010. The consumption response to income changes. Annu. Rev. Econ. 2 (1), 479–506.
Jordà, Òscar, Knoll, Katharina, Kuvshinov, Dmitry, Schularick, Moritz, Taylor, Alan M, 2019. The rate of return on everything, 1870–2015. Q. J. Econ. 134 (3), 1225–1298.
Kaplan, Greg, Violante, Giovanni L., 2010. How much consumption insurance beyond self-insurance? Am. Econ. J.: Macroecon. 2 (4), 53–87.
Kaufmann, Katja, Pistaferri, Luigi, 2009. Disentangling insurance and information in intertemporal consumption choices. Amer. Econ. Rev. 99 (2), 387–392.
Mueller, Andreas I., Spinnewijn, Johannes, Topa, Giorgio, 2021. Job seekers' perceptions and employment prospects: Heterogeneity, duration dependence, and bias. Amer. Econ. Rev. 111 (1), 324–363.
Ozkan, Serdar, Song, Jae, Karahan, Fatih, 2023. Anatomy of lifetime earnings inequality: Heterogeneity in job-ladder risk versus human capital. J. Political Econ. Macroecon. 1 (3), 506–550.
Stoltenberg, Christian A., Uhlendorff, Arne, 2025. Consumption choices and earnings expectations: Empirical evidence and structural estimation. Mimeo, University of Amsterdam.
Tauchen, George, Hussey, Robert, 1991. Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models. Econometrica 59 (2), 371–396.