Precautionary Saving and Aggregate Shocks

Jeanne Commault*

(Previously circulated as 'An Exact Analysis of Precautionary Consumption Growth')
February 2026

Abstract

Using surveys about the distribution of future consumption growth, I find that on average 8% of expected consumption growth from a year to the next is explained by precautionary behavior. This precautionary component is sensitive to employment shocks, while the non-precautionary component of consumption growth is not. It is not very sensitive to monetary shocks. This is consistent with precautionary behavior playing an important role in business cycles through fluctuations in employment, and with monetary policy affecting consumers mostly indirectly, through its general equilibrium effect on GDP. Standard approximated expressions of consumption growth do not capture this large contribution, mostly because they are based on the approximations of mathematical identities—a relation of the form $x = f^{-1}(f(x))$. Empirically, using such an approximation underestimates precautionary growth by more than half.

Key Words: Precautionary Saving; Heterogeneity; Wealth Inequality; Income Risk

*Department of Economics, Sciences Po, 28 Rue des Saints-Pères, 75007 Paris, France; jeanne.commault@sciencespo.fr.

1 Introduction

Heterogeneous-agent models, pioneered by Bewley (1986), Imrohoroglu (1989), Huggett (1993), Aiyagari (1994), have become a leading macroeconomic framework for studying business cycles and stabilization policies. This is because, since the Great Recession, these models can account for stylized facts that representative agents models have difficulties matching, in particular the sharp fall in consumption at the onset of the recession (see the review of Kaplan and Violante 2018).

The key innovation of these models is that, contrary to representative agent models, consumption is not linear in resources. Consumers face uninsurable income risk, so they engage in precautionary saving. The amount of precautionary saving they make, and the sensitivity of this saving to shocks, depends on their wealth and income. As a result, different people respond differently to the same shock, and the impact of an aggregate shock depends on the wealth and income distribution of the population affected by the shock. In addition to making consumption responses heterogeneous, precautionary saving also amplifies all these consumption responses, beyond the response that would take place under perfect foresight, which is typically small.

Now, the older literature estimating Euler equations found that incorporating second-order variance terms has little effect on the estimation results (see the review of Browning and Lusardi 1996). Standard numerical simulations with precautionary saving only also had trouble replicating the large consumption response to shocks. As a result, heterogeneous-agent models shifted to incorporate other sources of heterogeneity—borrowing constraints, taste for wealth, or behavioral biases.

Yet, recent surveys show that consumers would be willing to forgo substantial amounts of consumption to avoid business cycle fluctuations (Georgarakos et al. 2025), suggesting that precautionary behavior may play a significant role. In addition, there is new evidence that people face a substantially wider distribution of shocks (Guvenen and Smith 2014, Guvenen, Karahan, Ozkan, and Song 2021, Braxton, Herkenhoff, Rothbaum, and Schmidt 2025) and substantially more frequent shocks than previously thought (Ganong, Noel, Patterson, Vavra, and Weinberg 2025).

In this paper, I revisit the empirical measurement of the importance of precautionary saving, and the response of precautionary saving to shocks. I argue that precautionary behavior is a significant determinant of consumers' response to shocks. To motivate my new take on the question, I discuss the approximation method behind the log-linearized

Euler equation at the heart of the Euler equation estimation literature. I note that the method does not actually approximate the consumption growth implied by a consumer model but simply approximates a mathematical identity. Thus, the problem with the log-linearized Euler equation is not that it is an approximation but that it approximates the wrong relation. I then propose a method to empirically estimate the contribution of precautionary behavior to consumption growth that is based on survey data eliciting the respondents' expected distributions of future consumption. I find that the contribution of precautionary behavior to consumption growth is relatively large and, importantly, sensitive to employment shocks, contrary to the other component of consumption growth. It is also relatively unresponsive to monetary shocks.

My first contribution is to identify and explain a problem with the way consumption and log-consumption is approximated in the micro Euler equation literature. My starting point is that there is a discrepancy between the assumptions that the micro-level literature requires for consumption (or log-consumption) to evolve as a random walk, and the assumptions that the macroeconomic literature requires. Microeconomic studies approximate consumption as a random walk with a drift, thus eventually as a linear function of resources, around small realized shocks without constraining the uncertainty that people expect. Some papers interpret the terms in the resulting expression of consumption as driven by precautionary behavior, for instance Blundell, Pistaferri, and Preston (2008) discuss the amount of self-insurance through precautionary saving in the random walk expression they obtain (paragraph 'Self-Insurance' p1897).

In contrast, the macroeconomic literature recognizes that consumption or log-consumption can only be approximated as a random walk (with a drift) around a risk-free steady-state, where people expect no shocks. Because of this strong assumption, the literature is moving away from it. Frontier methods are now to approximate around a risky steady-state (Coeurdacier, Rey, and Winant 2011), where people expect shocks although the realized shocks are close to zero, or around a point with no expected aggregate shocks but some expected (and realized) idiosyncratic shocks (see Auclert, Bardóczy, Rognlie, and Straub 2021).

I consider the standard consumer's problem and I show that the discrepancy arises because the seminal consumer's problem study of Hall (1978), on which micro papers have built, approximates a mathematical identity—a relation of the form $x = f^{-1}(f(x))$—rather than the solution of the consumer's problem. This can arbitrarily yield random walk expressions. This generates two problems: (i) a misinterpretation of assumption required to obtain a random walk expression of consumption; (ii) different second order terms than the

ones that would arise around an approximation of the first order condition of the model.

While this may explain why this literature typically finds little role for precautionary behavior or for second-order terms, what is true one? My second contribution is to examine the size of the true contribution of precautionary behavior to consumption growth. Measuring precautionary saving is typically difficult: it is not just a share of saving, since people in debt can be making precautionary saving if they are taking less debt than they would absent uncertainty, and it is defined relative to a counterfactual world with no uncertainty that is not observed. To measure it, I rely on the New York Fed Survey of Consumer Expectations (SCE). I use the waves between April 2015 and December 2019. Crucially, the survey includes information about the respondents' distribution of expected consumption growth. Furthermore, other questions in other modules of the SCE make it possible to recover the level of consumption. Combining both, I eventually obtain the distribution of expected consumption. Under some assumptions about the utility function, this makes it possible to measure both the expected marginal utility of consumption and the marginal utility of expected consumption, thus to recover the contribution of precautionary behavior to consumption growth.

I find that precautionary behavior explains 8% of consumption growth under an isoelastic utility with a relative risk aversion of 2. This means that, absent uncertainty, expected consumption growth would be 8% lower.[1] Precautionary behavior still explains 4% of expected consumption growth with a relative risk-aversion of 0.5, and 17% with a relative risk-aversion of 5. The distribution of the precautionary component of consumption growth is strongly right-skewed, despite the fact that I winsorize the top and bottom 0.5% of the data. The precautionary component of consumption growth also correlates positively with the level of consumption and the level of consumption growth. As a result, its contribution to aggregate consumption growth is larger than its average contribution to expected consumption growth across respondents.

I regress this precautionary component against two shocks. The first one is the monetary policy surprise constructed by Bauer and Swanson (2023). The second one is an employment shock that I build as the experience of finding a job over the next three months while reporting low expectations of doing so. I find that the monetary policy surprise has little impact on precautionary growth. In contrast, the employment shock reduces precautionary

[1] Uncertainty can be removed in different ways, for instance by keeping expected consumption the same vs by keep expected income the same. However, the way in which it is remove does not change the impact of uncertainty on consumption growth—it only changes the impact of uncertainty on the level of consumption.

growth significantly. The magnitude corresponds to a 40% reduction at the average precautionary growth. This is consistent with the simulations of heterogeneous-agent models, in which monetary policy affects consumers mostly indirectly, through its general equilibrium effect on GDP. The non-precautionary component of consumption growth does not respond. This other component captures the consumption that would take place absent any uncertainty nor friction, but also the effect of liquidity constraints or of other constraints (at the existing level of uncertainty not how they would be absent uncertainty). The fact that they do not respond suggest that precautionary behavior is an important driver of the consumption response during business cycles.

2 The consumers' problem

Consumers' problem. I consider the standard consumers' problem with constant relative risk aversion. This is the one commonly found in macroeconomic frameworks and the one considered in Hall (1978). At each period t, a consumer i chooses their current consumption and the distribution of their future consumption as the solution of the following intertemporal optimization problem:

$$\max_{c_{i,t},\dots c_{i,T}} \quad \sum_{s=0}^{T-t} \left(\frac{1}{1+\delta}\right)^{t+s} E_t\left[u(c_{i,t+s})\right] \tag{2.1}$$

$$s.t. \quad a_{i,t+k+1} = (1+r)a_{i,t+k} - c_{i,t+k} + y_{i,t+k} \qquad \forall\, 0 \leq k \leq T-t, \tag{2.2}$$

$$a_{i,T} \geq 0. \tag{2.3}$$

Consumers are finite-lived with T the length of their life. They have time-separable utility, and at each period t they derive utility from their contemporaneous consumption expenditures $c_{i,t}$. The period utility function $u(c)$ is isoelastic, which is equivalent to saying that it displays constant relative risk aversion (CRRA). Its functional form is $u(c) = (c^{1-\rho}-1)/(1-\rho)$, with ρ the relative risk aversion and $1/\rho$ the intertemporal elasticity of substitution. Future utility is discounted by the factor β, which is constant and known in advance by the household. At each period t, the consumer i earns the stochastic amount $y_{i,t}$, which is uncertain prior to period t. The period budget constraints (2.2) state that, to store their wealth from one period to another, consumers only have access to a risk-free asset that delivers the certain interest rate r, where $a_{i,t}$ denotes the level of this asset at the beginning of period t (or at the end of period $t-1$). The terminal condition on wealth (2.3)

states that the consumers cannot die with a strictly positive level of debt.

First order condition The first order condition of the maximization problem of the household, known as the Euler equation, is as follows:

$$u'(c_{i,t}) = E_t[u'(c_{i,t+1})] \underbrace{\beta(1+r)}_{R}. \tag{2.4}$$

The term $R \equiv \beta(1+r)$ is the factor accounting for the deterministic intertemporal substitution motives. This first order condition states that an optimizing household chooses its consumption path so that current and future consumption delivers the same expected marginal utility, after weighting these marginal utilities to account for deterministic intertemporal substitution motives.

Precautionary consumption growth. Defining φ_t the equivalent premium such that $E_t[u'(c_{t+1})] = u'(c_{t+1} - \varphi_t)$, the first order condition implies:

$$u'(c_{i,t}R^{1/\rho}) = E_t[u'(c_{i,t+1})], \tag{2.5}$$

$$u'(c_{i,t}R^{1/\rho}) = u'(E_t[c_{i,t+1}] - \varphi_t), \tag{2.6}$$

$$c_{i,t}R^{1/\rho} = E_t[c_{i,t+1}] - \varphi_t. \tag{2.7}$$

From Jensen's inequality, the premium φ_t is strictly positive because $u'(.)$ is decreasing and convex. Intuitively, when marginal utility is convex ($u'''(.) > 0$), the effects of negative and positive shocks are asymmetric: a shock that reduces future consumption raises the value of one additional unit of future consumption more than a shock that raises future consumption reduces the value of one additional unit of future consumption. Because of this asymmetry, on average, the effect of the negative shocks dominates, and the presence of mean-zero shocks to future consumption increases the expected marginal utility of future consumption above the marginal utility of expected future consumption: $E_t[u'(c_{i,t+1})] > u'(E_t[c_{i,t+1}])$. Since marginal utility is decreasing ($-u''(.) > 0$), this induces the consumers to set their current consumption below their expected future consumption for the marginal utility of their current consumption to be as high as the expected marginal utility of their future consumption. The amount by which the consumers must decrease their current consumption coincides with the risk premium φ_t. Kimball (1990) notes that φ_t, which he calls the precautionary premium, corresponds to precautionary saving in a two-period model. In this

more general multiperiod model, it simply corresponds to the contribution of precautionary behavior to consumption growth.

Precautionary saving I derive an expression of precautionary saving in the multiperiod framework (not closed-form since φ are endogenous in a multiperiod framework). Iterating forward on equation (3.1), consumption growth between t and any future period $t+s$ is a weighted sum of the precautionary premiums between these two dates: $E_t[c_{t+s}] = c_t R^{s/\rho} + \sum_{k=1}^{s} E_t[\varphi_{t+k-1}] R^{(s-k)/\rho}$. I plug these expressions into the intertemporal budget constraint that arises from the combination of (2.2)-(2.3). The resulting intertemporal budget constraint states that the net present value of current and future expected consumption equals current assets plus the net present value of current and future expected income. I rearrange and obtain

$$c_t = \underbrace{\frac{1}{l_{t,0}} \left((1+r)a_t + \sum_{s=0}^{T-t} \frac{E_t[y_{t+s}]}{(1+r)^s} \right)}_{\substack{\text{consumption under perfect foresight} \\ \frac{1}{l_{t,0}} W_{i,t}}} - \underbrace{\frac{1}{l_{t,0}} \left(\sum_{s=1}^{T-t} \frac{l_{t,s} E_t[\varphi_{t+s-1}]}{(1+r)^s} \right)}_{\substack{\text{precautionary saving} \\ PS_{i,t}}}, \tag{2.8}$$

with $\frac{1}{l_{t,k}} = (\sum_{s=k}^{T-t} \frac{R^{s/\rho}}{(1+r)^s})^{-1}$ the weight put on consumption at period t relative to consumption between t and $T-k$.[2] This expression encompasses the perfect foresight case, in which the consumers simply consume a fraction $\frac{1}{l_{t,0}}$ of their total expected lifetime resources—the sum of their net assets, current income and expected future income—denoted W_t (for wealth). In the presence of uncertainty, to be able to implement the additional precautionary consumption growth they desire, the consumers take out their total expected precautionary growth from their total expected resources and consume a constant share of what remains. Precautionary saving is defined as the difference between actual consumption and the consumption that would be chosen under perfect foresight everything else being equal. Thus, it writes as a weighted sum of the current and future expected value of the precautionary premiums φ.

In the remainder of the paper, I drop the household index i.

[2] When consumers are neither patient nor impatient ($\beta = \frac{1}{1+r}$ so $R = 1$), $l_{t,0}$ tends toward $\frac{r}{1+r}$ as T approaches infinity.

3 Comparisons of consumption growth approximations

3.1 The discrepancy between macro and micro approximations

Exact expression The Euler equation leads to expression (3.1)

$$c_{i,t}R^{1/\rho} = E_t[c_{i,t+1}] - \varphi_t. \tag{3.1}$$

Consumption evolves as a random walk with drift $R^{1/\rho}$ if and only if the precautionary premium φ_t is zero, that is, when the consumer faces no uncertainty about their future consumption. With shocks at $t+1$, I define the innovation ε^c_{t+1} of realized consumption c_{t+1} defined as $\varepsilon^c_{t+1} = c_{t+1} - E_t[c_{t+1}]$. Plugging this expression in (3.1), realized consumption growth is

$$c_{t+1} = \varphi_t + c_t R^{1/\rho} + \xi_{t+1}, \text{ with } E_t[\xi_{t+1}] = E_t[(c_{t+1} - E_t[c_{t+1}])] = 0. \tag{3.2}$$

Combining this expression with the budget constraints, consumption is not linear in shocks to wealth and income because φ_t varies with income and wealth—it does not necessarily do so for all utility functions but under the isoelastic utility we assume here it does.

Macroeconomics approach In macroeconomics, it is typical to approximate the Euler equation around the expected distribution of future consumption c_{t+1}. For simplicity of exposition, I assume there is a finite number S of states of the world at $t+1$. The probability of each state s is p^s. The Euler equation rewrites

$$\begin{aligned} u'(c_t R^{1/\rho}) &= \big(\sum_{s\in S} p^s u'(c^s_{t+1})\big), \\ c_t R^{1/\rho} &= (u')^{-1}\big(\sum_{s\in S} p^s u'(c^s_{t+1})\big) \end{aligned}$$

Macroeconomic analyses that are not concerned with uncertainty take a first order approximation around a deterministic steady-state. To do the same here, I consider the deterministic state where consumers expect that their future consumption will take the same value in each possible state s: $c^s_{t+1} = \bar{c}$ for each s. I take a first-order Taylor expansion of c^s_{t+1}

around $\bar{c}$ for each s

$$c_t R^{1/\rho} \approx (u')^{-1}\big(u'(\bar{c})\big) + \sum_{s\in S}(c_{t+1}^s - \bar{c})p^s u''(\bar{c})((u')^{-1})'(u'(\bar{c})),$$
$$c_t R^{1/\rho} \approx \bar{c} + (E_t[c_{t+1}] - \bar{c})u''(\bar{c})((u')^{-1})'(u'(\bar{c})).$$

For $\bar{c} = E_t[c_{t+1}]$, which means the deterministic state considered is the one in which the consumer expects their future consumption to always be equal to its weighted average value across states

$$c_t R^{1/\rho} \approx E_t[c_{t+1}] + \underbrace{(E_t[c_{t+1}] - E_t[c_{t+1}])}_{=0} u''(E_t[c_{t+1}])((u')^{-1})'(u'(E_t[c_{t+1}]))$$
$$c_t R^{1/\rho} E_t[c_{t+1}].$$

Consumption evolves as a random walk with drift $R^{1/\rho}$. This means that, approximating in first order around a deterministic state without uncertainty yields the exact same expression as in the total absence of uncertainty. With shocks at $t+1$, realized consumption is

$$c_{t+1} \approx c_t R^{1/\rho} + \xi_{t+1}, \text{ with } E_t[\xi_{t+1}] = E_t[(c_{t+1} - E_t[c_{t+1}])] = 0.$$

Combining this expression with the budget constraints, consumption is linear in income and wealth. All consumers with the same age, discount factor and interest rate respond in the same way to a change in income or wealth, even if their wealth and income differ.

Approximating the expected distribution of future consumption in first order around any risky state, where consumption is not the same in all states, consumption no longer evolve as a random walk approximation. This is because the precautionary premium φ_t only disappears when consumers expect no shocks. Any uncertainty about future consumption generates a non-zero φ_t.

Hall (1978) approach: approximation of a mathematical identity around a risky point While in the macroeconomic approach researchers approximate the expected distribution of future consumption, in the case of isoelastic utility, Hall (1978) approximates the realized marginal utility of consumption at $t+1$, $u'(c_{t+1})$, around its expected value at t, $E_t[u'(c_{t+1})]$. It does not impose restrictions on the consumer's expectations and still obtains a random walk. However, realized consumption (or the realized marginal utility of

consumption) is not in the Euler equation (2.4): this equation defines c_t as a function of variables observed at t, specifically the expected distribution of future consumption at t. Thus, what Hall (1978) approximates is not the Euler equation but the mathematical identity $c_{t+1} = (u')^{-1}(u'(c_{t+1}))$ that it plugs in the Euler equation. Formally, what the paper of Hall (1978) does can be decomposed as follows

$$
\begin{aligned}
u'(c_t R^{1/\rho}) &= E_t[u'(c_{t+1})] \\
u'(c_t R^{1/\rho}) &= u'(c_{t+1}) + (E_t[u'(c_{t+1})] - u'(c_{t+1})) && (3.3) \\
c_{t+1} &= (u')^{-1}(u'(c_{t+1}) + \cancel{E_t[u'(c_{t+1})]} - \cancel{u'(c_t R^{1/\rho})}) && (3.4) \\
c_{t+1} &\approx (u')^{-1}(E_t[u'(c_{t+1})]) + \big(u'(c_{t+1}) - E_t[u'(c_{t+1})]\big)((u')^{-1})'(E_t[u'(c_{t+1})]) && (3.5) \\
c_{t+1} &\approx c_t R^{1/\rho} + \underbrace{\big(u'(c_{t+1}) - E_t[u'(c_{t+1})]\big)}_{\text{innovation, zero expected value at } t} \big(u''(c_t R^{1/\rho})\big)^{-1}. && (3.6)
\end{aligned}
$$

Consumption also writes as its past value multiplied by a drift, plus a term that is zero in expectation at t. The Euler equation, however, is inessential in the derivation of the expression itself: equation (3.4) shows that one could get the same final expression by starting from $c_{t+1} = (u')^{-1}(u'(c_{t+1}))$ and taking the same first order Taylor expansion around $u'(c_{t+1}) = E_t[u'(c_{t+1})]$. The Euler equation only ensures, after the approximation, that the first term in the final expression coincides with consumption at t. Note that one could equivalently take a first order Taylor expansion if $c_{t+1} = (u')^{-1}(u'(c_{t+1}))$ around $u'(c_{t+1}) = u'(c_t R^{1/\rho})$. In that case, one obtains $c_{t+1} \approx c_t R^{1/\rho} + \big(u'(c_{t+1}) - u'(c_t R^{1/\rho})\big)$ and the Euler equation simply ensures that the second term is an innovation with zero expected value at t.

As a result of this method, contrary to the macroeconomic approach or the exact Euler equation without uncertainty, one does not get that expected consumption equals current consumption times a drift. Here, one directly obtains that realized consumption equals current consumption times a drift plus an innovation.

Obtaining any random walk expression from mathematical identities The problem with approximating mathematical identities is that one can often approximate an identity of a variable x as the sum of any desired value y plus an innovation with mean zero by selecting the right mathematical identity. To see this, consider first the very trivial identity $x = x$. Because the variable x is on both sides, it can be directly approximated as the desired

value y plus a first-order term

$$x \approx y + (x - y)$$

But the first order term, that is, the second term on the left-hand side, is not an innovation that is zero in expectation. Yet, if one then selects *the right function* $f(.)$ such that $E[f(x)] = f(y)$, then approximating the identity $x = f^{-1}(f(x))$ around $f(x) = f(y)$ yields the desired expression

$$x \approx f^{-1}(f(y)) + (f(x) - f(y))(f^{-1})'(f(y)),$$

with $E[f(x) - f(y)] = 0$. This can often be done, in the sense that it can be done under well-behaved, uniform expectations about x. In that case, because the premium φ^x associated with $f(.)$ and the uncertainty about x increases with the curvature of the function $f(.)$, it should be possible for any x and y to find a function f such that $E[f(x)] = f(E[x] - \varphi^x) = f(y)$.

3.2 Second order approximations

Second order version of Hall's approximation. Doing a second-order Taylor expansion of $c_{t+1} = (u')^{-1}(u'(c_{t+1}))$ around $u'(c_{t+1}) = E_t[u'(c_{t+1})]$ yields

$$\begin{aligned} c_{t+1} = c_t R^{1/\rho} &+ \big(u'(c_{t+1}) - E_t[u'(c_{t+1})]\big) \frac{1}{u''(c_t R^{1/\rho})} \\ &+ \big(u'(c_{t+1}) - E_t[u'(c_{t+1})]\big)^2 \frac{u'''(c_t R^{1/\rho})}{(-u''(c_t R^{1/\rho}))^3} + o(u'(c_{t+1}) - E_t[u'(c_{t+1})]). \end{aligned} \tag{3.7}$$

However, the consumption literature has mostly considered second-order expansion of log-consumption growth.

Log-linearized expressions of consumption. Dating back to Campbell and Mankiw (1989), it is common in the consumption literature to consider the first or second order log-linearized

Euler equations.

$$\Delta ln(c_{t+1}) = \underbrace{\frac{1}{\rho} ln(R)}_{\text{Deterministic}} - \underbrace{\frac{1}{\rho}\zeta_{t+1}}_{\substack{\text{Innovation,}\\ \text{orthogonal}\\ \text{to var. at } t}} + o(\zeta_{t+1}), \tag{3.8}$$

$$\Delta ln(c_{t+1}) = \underbrace{\frac{1}{\rho} ln(R)}_{\text{Deterministic}} + \underbrace{\frac{E_t[\zeta_{t+1}^2]}{2\rho}}_{\text{Precaution}} \underbrace{-\frac{1}{\rho}\zeta_{t+1} + \frac{\zeta_{t+1}^2 - E_t[\zeta_{t+1}^2]}{2\rho}}_{\substack{\text{Innovation, orthogonal}\\ \text{to var. at } t}} + o(\zeta_{t+1}^2). \tag{3.9}$$

with $\zeta_{t+1} = (u'(c_{i,t+1}) - E_t[u'(c_{i,t+1})])/E_t[u'(c_{i,t+1})]$ the percentage innovation to the marginal utility of consumption between t and $t+1$ over the marginal utility of current consumption.

The method commonly used to derive them, described for instance in the literature review on consumption and saving of Browning and Lusardi (1996) and after that in the consumption textbooks , follows the same principle as in Hall 1978: it is based on the approximation of an identity relating c_{t+1} to itself around small realized shocks. One starts from the Euler equation, takes the log of each side, adds $ln(c_{t+1})$ to each side, rearranges to obtain an expression of $ln(c_{t+1})$ as a function of $\zeta_{t+1} = (u'(c_{t+1}) - E_t[u'(c_{t+1})])/E_t[u'(c_{t+1})]$, and approximates it around $\zeta_{t+1} = 0$.[3] What is approximated is the identity $ln(c_{t+1}) = ln\big((u')^{-1}\big(u'(c_{t+1})\big)\big)$.

Implications for measuring the importance of precautionary behavior from these second-order terms. My finding that this expression is the approximation of a mathe-

[3]The exact procedure is here

$$ln(c_{t+1}) = ln\big((u')^{-1}\big(u'(c_{t+1})\big)\big) \tag{3.10}$$

$$= ln\left((u')^{-1}\left(E_t[u'(c_{t+1})]\frac{u'(c_{t+1}) - E_t[u'(c_{t+1})]}{E_t[u'(c_{t+1})]} + E_t[u'(c_{t+1})]\right)\right) \tag{3.11}$$

$$= ln\left(E_t[u'(c_{t+1})]^{-1/\rho}\left(\zeta_{t+1}+1\right)^{-1/\rho}\right) \tag{3.12}$$

$$= \underbrace{-\frac{1}{\rho} ln(E_t[u'(c_{t+1})])}_{=ln(c_t)+\frac{1}{\rho}ln(R)} - \frac{1}{\rho} ln(1+\zeta_{t+1}). \tag{3.13}$$

Approximating this expression around $\zeta_{t+1} = 0$ yields $ln(c_{t+1}) = ln(c_t) + \frac{1}{\rho}ln(R) - \frac{1}{\rho}\zeta_{t+1} + o(\zeta_{t+1})$ in first order and $ln(c_{t+1}) = ln(c_t) + \frac{1}{\rho}ln(R) - \frac{1}{\rho}\zeta_{t+1} + \frac{1}{2\rho}\zeta_{t+1}^2 + o(\zeta_{t+1})$. Again the Euler equation intervenes only to substitute $E_t[u'(c_{t+1})] = u'(c_t R^{1/\rho})$ in $ln(E_t[u'(c_{t+1})])$. It does not constitute the relation that is approximated since ζ_{t+1} is not initially in the Euler equation.

matical identity means that the second-order terms are unlikely to be meaningfully capturing actual precautionary behavior. The fact that the second order term is a variance, that would be zero absent any uncertainty, does not automatically mean it coincides with the contribution of precautionary behavior. My finding that these expressions result from approximations of identities also offers an explanation for Carroll 2001's finding that even the second-order log-linearized Euler equation performs badly when tested against numerical simulations of the exact model it approximates. The paper attributes to the importance of higher-order moments.What I show is that the problem is not so much that the second order expression is an approximation, but that it is an approximation of a mathematical identity rather than an approximation of the relations implied by the model.

Blundell, Low and Preston (2013) and Blundell Pistaferri and Preston (2008)'s derivation of a random walk in second order approximation. The influential paper of Blundell, Pistaferri, and Preston 2008 claims to have derived a random walk expression of consumption even in second-order approximation around small realized shocks, and not just in first-order approximation. The derivation of the paper actually follows the one in Blundell, Low, and Preston 2013, presented as the detailed version of the derivation. I note that, besides approximating an identity, they additionally equate the innovation to log-consumption and log-consumption growth net of deterministic motives. Equating the two means assuming that log-consumption growth is the sum of a deterministic term plus an innovation, which is what they find. It leads the paper to make the second order term in the approximation a little-o of the innovation term.

In addition, because Blundell, Pistaferri, and Preston 2008 believes the expression they obtain captures self-insurance from precautionary behavior, they interpret some of the terms that would be present under perfect foresight as the contribution of precautionary behavior (paragraph 'Self-Insurance' p1897).

3.3 Hansen and Singleton (1983)'s assumption of future consumption log-normality as making exogenous variables endogenous

Method. Another way to express log-consumption as a random walk is to assume that future log-consumption is normally distributed (or, if the interest rate is stochastic, that future consumption and the future interest rate are joint log-normal). This method is initially developed in Hansen and Singleton (1983). It is presented as the main way to derive an

expression of log-consumption growth in the early consumption textbooks, e.g. Deaton 1992 and Attanasio 1999, but the most recent textbook of Jappelli and Pistaferri 2017 cites it second, after presenting the Taylor approximation method. The derivation relies on the assumption that $ln(c_{t+1})|I_t$ is Gaussian with $\mathcal{N}(E_t[ln(c_{t+1})], var_t(ln(c_{t+1})))$ its distribution. The Euler equation, which recalling that $x = exp(ln(x))$ can be rearranged as $E_t[exp\big(-\rho(\ln(c_{t+1}) - ln(c_t)) + ln(R_{t,t+1})\big)] = 1$, becomes:

$$exp\big(-\rho E_t[ln(c_{t+1})] + \frac{\rho^2}{2} var_t(ln(c_{t+1})) + \rho ln(c_t) + ln(R_{t,t+1})\big) = 1. \tag{3.14}$$

This is because of the property of the normal distribution that if $x \sim \mathcal{N}(\mu, \sigma^2)$, then $E[exp(x)] = exp(\mu + \sigma^2/2)$. Taking the logarithm of each side:

$$-\rho E_t[ln(c_{t+1})] + \frac{\rho^2}{2} var_t(ln(c_{t+1})) + \rho ln(c_t) + ln(R_{t,t+1}) = 0. \tag{3.15}$$

Eventually:

$$\Delta ln(c_{t+1}) = \underbrace{\frac{1}{\rho} ln(R_{t,t+1})}_{\text{Deterministic}} + \underbrace{\frac{\rho}{2} var_t(ln(c_{t+1}))}_{\text{Precaution}} + \underbrace{ln(c_{t+1}) - E_t[ln(c_{t+1})]}_{\text{Orthogonal to var. at } t} \tag{3.16}$$

The expression obtained is similar to the second-order log-linearized Euler equation.

Contradictory implications. The problem is that consumption is endogenous and the assumption that $ln(c_{t+1})|I_t$ is Gaussian has implications that are in contradiction with the model. Indeed, if given each I_t, there exists a distribution of exogenous variables—income in this simple case—, such that $ln(c_{t+1})|I_t$ is Gaussian, the distribution of exogenous variables that makes $ln(c_{t+1})|I_t$ Gaussian is different for different I_t. Now, it happens that the variables on which one must condition to make future log-consumption Gaussian includes decisions variables at t, in particular consumption at t. It is not sufficient to condition on the realized exogenous variables at t. For instance, in a simple two-period version of the model (2.1)-(2.3), one needs $ln(c_2) = ln(y_2 + a1 + y_1 - c_1)$ to be normal, thus $y_2 + a1 + y_1 - c_1$ to be log-normal. This implies that y_2 must be three-parameter log-normal with any mean μ, any variance σ, but a threshold parameter $-(a1 + y_1 - c_1)$. Thus, the required distribution of y_2 varies with c_1. This means that, for the assumption to hold, the income process at $t+1$ that a consumer expects at t depends on its consumption decision at t. But this effect

is not in the problem that consumers solve: it's in contradiction with the consumer problem stated at the beginning.

Now, even assuming that consumers do not internalize the effect of their decision on the distribution of their future income (or any other source of shocks e.g. health in a model with health variations), this implies income is not exogenous. The microfoundation for why an increase in consumption at t would shift down the distribution of log-income at $t+1$ by this particular amount is unclear. It is therefore difficult to use this reasoning to justify that log-consumption growth evolves according the second order log-linearized expression.

4 Measuring precautionary behavior

The previous section establishes that approximating consumption as a random walk requires assuming people expect no shocks, not just that consumers expect shocks but the realized shocks happen to be small, as commonly presented in the micro literature. This is because approximations around small realized shocks are approximations of mathematical identities. It follows that the second order terms in these approximations, including the second-order terms in the common log-linearized expressions of consumption, do not capture precautionary behavior. While it is unlikely that consumers have no uncertainty about their future consumption, it's unclear how large is the wedge between actual expected consumption growth and the consumption growth predicted by a random walk. That's what I measure in this section.

4.1 The Survey of Consumer Expectations

Survey. To measure an empirical counter part to φ, I use data from the Survey of Consumer Expectations (SCE) of the Federal Reserve Bank of New York (2015-2019). It is a monthly online survey with a rotating panel of about 1,300 household heads based in the United States. A household head is defined as a person in the household who owns, is buying, or rents the home. A household may have multiple co-household heads. Respondents stay on the panel for up to twelve months before rotating out of the panel. The survey started in June 2013. While the Core Survey takes place monthly, its topical modules only take place either every four months or every year. The Spending Survey module, which asks about the probability of future consumption growth and about the composition of consumption expenditures, takes place every four months. The Housing module, which

contains information about the level of housing expenditures, and the Household Finance survey, which contain information about wealth, both take place once a year.[4] My period of observation is between April 2015 and December 2019. I stop in December 2019 so my sample is not affected by covid shutdowns and their effects on consumption expectations.

Distribution of future consumption growth. To build the distribution of future consumption, I use a set of questions in the Spending Survey module of the SCE about the probabilities that the respondent assign to experiencing an annual consumption growth of different values. These values are: less than -12%, between -10% and -8%, between -8% and - 4%, between -4% and -2%, between -2% and 0%, between 0% and 2%, between 2% and 4%, between 4% and 8%, between 8% and 12%, or more than 12%. I consider these categories as ten different states of the world. In each state of the world, consumption grows at the average value of the range (i.e. 6% for the range between 4% and 8%). For the lowest and highest categories, I consider consumption grows at -15% and 15% in each state. In a robustness check, I remain agnostic about the values below and above 12%, and estimate them assuming a normal distribution for each individual's expectations.

Consumption. I build consumption from a combination of questions in the Spending module and in the Housing module. Indeed, there is no direct question about the household's *level* of consumption expenditures in the SCE. However, the Spending module reports information about the share of their total monthly spending allocated to different consumption categories in a typical month. It includes the share allocated to housing, defined as mortgage, rent, maintenance and home owner/renter insurance. The Housing module further reports information about the level of typical monthly spending on housing, where spending on housing is defined in the same way.I recover the level of household's total typical monthly spending with a proportionality rule, using the level of housing spending and the share of monthly spending devoted to housing. I also recover the level of spending on different consumption categories (utilities, food, clothing, transportation, medical care, entertainment and education), using the share of monthly spending devoted to each of the other consumption categories. My measure of yearly consumption is the typical monthly consumption spending multiplied by 12. Because this measure is based on multiple answers from different modules, I can only build it for a limited set of observations. I deflate

[4] See Armantier, Topa, Klaauw, and Zafar (2017) for technical background information on the SCE, and `www.newyorkfed.org/microeconomics/sce.html` for additional information.

this measure of consumption with a CPI index and express it in 2015$.

Distribution of future consumption. Since I know the distribution of future annual consumption growth, and since I know what the current annual consumption of each respondent is, I can compute the distribution of their future annual consumption.

Employment, probability of non-employment at the next quarter, and income variability. I observe whether respondents are currently employed (including self-employed) or unemployed, but have been employed at least once in the past. I also observe the probabilities they assign to the events of being unemployed or non-employed in the next four months. I measure their overall household income variability from a question in the Spending module of the SCE. This question asks the respondent whether their household income before taxes is 'more or less constant, slightly variable, somewhat variable or highly variable from month to month?'.

Selection and winsorizing. I select out respondents who allocate more or less than 100 percentage points to the possible future consumption growth values. I also winsorize the top and bottom 0.5% of the distribution of consumption: this means replacing the values above the top 0.5% with the value at the 0.5th percentile, and the values below the bottom 0.5% with the value at the 99.5th percentile.

4.2 Contribution of the precautionary component

Identification. The definition of the wedge between exact expected consumption growth and random-walk consumption growth is

$$\varphi_t = E_t[c_{t+1}] - (u')^{-1}(E_t[u'(c_{t+1})]).$$

For a given parametrization of the utility function, data on the expected distribution of future consumption make it possible to build (i) the average expected value of future consumption; (ii) the average expected value of the marginal utility of future consumption. I can then compute the empirical counterfactual to φ_t.

Precautionary consumption growth φ. Because I can measure $E_t[c_{t+1}]$ and $E_t[u'(c_{t+1})]$ under some assumptions about risk-aversion ρ, I can measure $\varphi_t = E_t[c_{t+1}] - (u')^{-1}(E_t[u'(c_{t+1})])$. Table 1 shows that only 21% of the sample reports 100% certainty that their annual con-

ρ	0.5	1	2	5
Share with $\varphi_{t,i} = 0$	21.3%	21.3%	21.3%	21.3%
Precautionary growth (\$)	110	146	220	437
Expected consumption growth (\$)	2,603	2,603	2,603	2,603
Contribution to aggregate growth $\frac{\sum_{i\in I}\varphi_{t,i}}{\sum_{i\in I}(E_t[c_{i,t+1}]-c_{i,t})}$	4.2%	5.6%	8.4%	16.8%
Average contribution to growth $\frac{1}{I}\sum_{i\in I}\frac{\varphi_{t,i}}{(E_t[c_{i,t+1}]-c_{i,t})}$	0.7%	0.9%	1.3%	2.8%
$\text{Cov}(\varphi_{i,t}, E_t[c_{i,t+1}] - c_{i,t})$	1.0e+06	1.4e+06	2.3e+06	5.4e+06
$\text{Cov}(\varphi_{i,t}, E_t[c_{i,t+1}] - c_{i,t} - \varphi_{i,t})$	-2.9e+05	-9.1e+05	-3.0e+06	-1.5e+07
Observations	7,853	7,853	7,853	7,853

Table 1: Average precautionary consumption growth for different values of risk-aversion

sumption will grow by one of the range suggested. I treat it as complete certainty here because the categories make it impossible to measure uncertainty more finely, but consumers may still be uncertain. For instance, a respondent being sure that their consumption will grow between 2% and 4% maybe unsure whether it will be 2%, 3% or 4%. As a result, what I obtain is a lower bound on the contribution of precautionary behavior φ.

The second line presents the average values of φ under different assumptions about ρ. I denote I the number of the respondents observed. The third line puts these values into perspective. It shows the resulting contribution of aggregate precautionary growth $\sum_{i\in I}\varphi_{t,i}$ to aggregate expected consumption growth $\sum_{i\in I}(E_t[c_{i,t+1}] - c_{i,t})$. This contribution ranges from 6% to 24%. At the common value of $\rho = 2$, precautionary behavior represents 12% of expected consumption growth.

The fourth line shows that the average contribution per capita is smaller. At the common value of $\rho = 2$, the average contribution of precautionary behavior to expected consumption growth is only 1.3%. This would be the case if people with the largest precautionary behavior also have the largest consumption growth.[5] The next line shows that, for values of

[5] I denote $\bar{E}[x] = \frac{1}{I}\sum_{i\in I} x$ the operator that takes the average over the population of respondents. We have $\bar{E}\left[\frac{\varphi_{t,i}}{(E_t[c_{i,t+1}]-c_{i,t})}\right] = \bar{E}[\varphi_{t,i}]\bar{E}\left[\frac{1}{(E_t[c_{i,t+1}]-c_{i,t})}\right] + cov(\varphi_{t,i}, \frac{1}{(E_t[c_{i,t+1}]-c_{i,t})}) > \frac{\bar{E}[\varphi_{t,i}]}{\bar{E}[(E_t[c_{i,t+1}]-c_{i,t})]} + cov(\varphi_{t,i}, \frac{1}{(E_t[c_{i,t+1}]-c_{i,t})})$. The last inequality is because $f(x) = 1/x$ is convex. Because empirically we observe that $\bar{E}\left[\frac{\varphi_{t,i}}{(E_t[c_{i,t+1}]-c_{i,t})}\right]$ (fourth line in Table 1) is much smaller than $\frac{\bar{E}[\varphi_{t,i}]}{\bar{E}[(E_t[c_{i,t+1}]-c_{i,t})]}$ (third line in Table 1), in the inequality above, the covariance $cov(\varphi_{t,i}, \frac{1}{(E_t[c_{i,t+1}]-c_{i,t})})$ must be negative empirically. Thus, the covariance $cov(\varphi_{t,i}, E_t[c_{i,t+1}] - c_{i,t})$ must be positive.

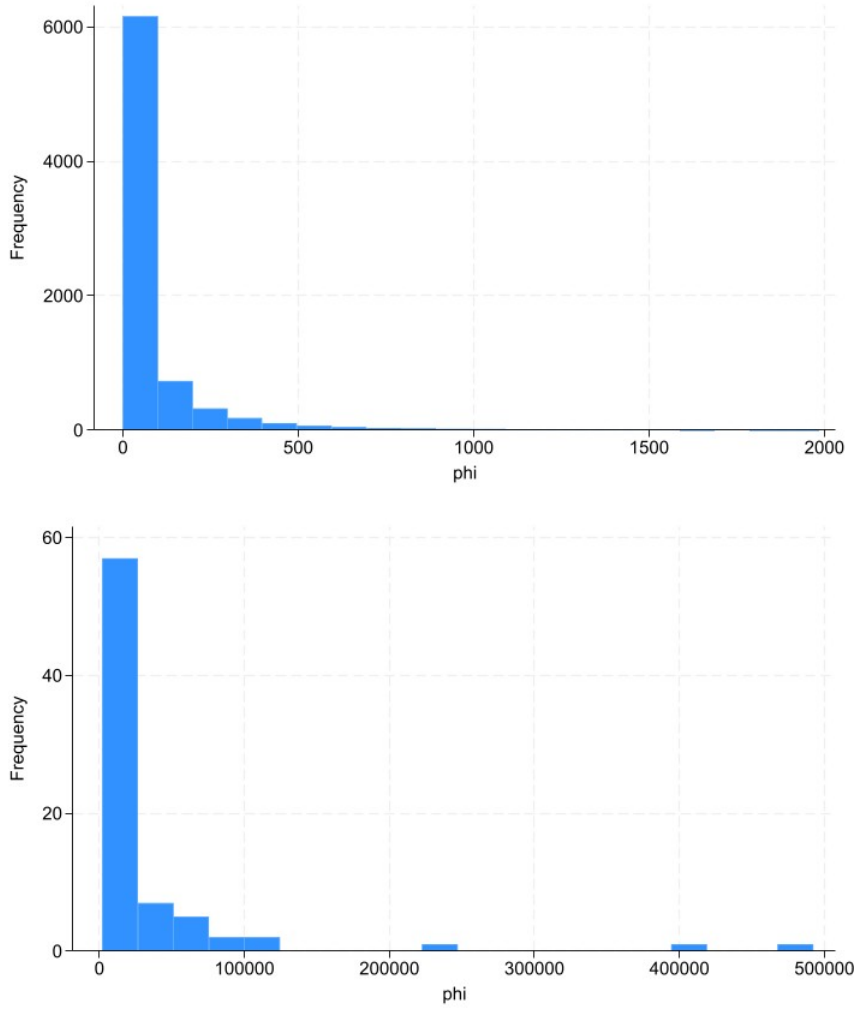


Figure 1: Histogram of φ below \$2,000 (top panel) and above \$2,000 (bottom panel)

$\rho \leq 2$, even the expected consumption growth that is not driven by precautionary behavior covaries positively with φ.

This indicates that precautionary behavior is relatively important for aggregate outcomes, even when its contribution is small at the individual level on average. To further emphasize that respondents at the top display large precautionary growth that matters for the aggregate, Figure 1 shows that the average value of φ_t is very right-skewed. The two panels show that the distribution has a similar fat right tail below and above \$2,000.

Comparisons with precautionary terms in approximated log-consumption expressions.

Approximation	Exact	Hall	Log-linearized	Hansen & Singleton
Growth	$\Delta E_t[c_{t+1}]$	$\Delta E_t[c_{t+1}]$	$\Delta E_t[ln(c_{t+1})]$	$\Delta E_t[ln(c_{t+1})]$
ρ	2	2	2	2
Precautionary growth ($ or ln($))	220	1.38e-09	0.00112	0.00109
Expected consumption growth ($)	2,603	2,603	.	.
Expected log-consumption growth (ln($))	.	.	0.03072	0.03072
Contribution of precautionary growth	8.5%	5.30e-13%	3.6%	3.6%
Observations	7,853	7,853	7,853	7,853

Table 2: Precautionary consumption growth for different values of risk-aversion

Table 2 compares the true value of φ with its approximations in the literature. These approximations substantially underestimate the true contribution of precautionary saving. Expanding Hall's approximation to the second order, the second order term raises expected consumption growth by a minuscule amount. Considering the more commonly used log-linearized consumption expression, the contribution of the precautionary term is about one-third of its actual contribution, at 3.6% instead of 11.8%. This confirms that, even when it includes second-order terms, the log-linearized Euler equation remains a poor fit for actual log-consumption growth.

This underestimation holds even with data on individual-level variance of future consumption, which captures idiosyncratic risk. This risk is recognized to be larger than aggregate consumption risk. I also find it to larger than the variance of consumption within a population with similar characteristics. It is therefore unsurprising that the log-linearized Euler equation in which the individual variances are proxied by the variance of aggregate consumption over time or by the cross-sectional variance of consumption find a very limited role for precautionary behavior.

Precautionary saving. Equation (2.8) shows that precautionary saving writes as a weighted sum of all current and future expected precautionary premiums φ. I do not observe the premiums φ at future periods, nor people's expectations about them. However, to compute the level of precautionary saving, I need to make assumptions, not just about the relative risk aversion ρ, but also about the consumers' discount factor, the interest rate they face, and

their life-expectancy. I also need to get a proxy for $E_t[\varphi_{t+s}]$. To get an order of magnitude of precautionary saving, I make the simplifying assumptions that (i) $\beta = 0.97$, $(1+r) = 0.01$, and $\rho = 2$ for everyone, (ii) $T = 90$ for everyone (iii) $E_t[\varphi_{t+s}] = \varphi_t$ for everyone. I assume the latter because I show in Table 3 that φ does not strongly respond to changes in age: over the age range of the respondent it would decrease then increase back but the coefficients are not significant. This could be because new risks to consumption, such as health shocks, arise in old age.

Under these simplifying assumptions, I find that average precautionary saving corresponds to 4.5% of average consumption. This is consistent with the finding of Georgarakos et al. (2025) that consumers would be willing to renounce to 5% of their life-time consumption to avoid business cycles fluctuations.

4.3 Response of the precautionary component to shocks

Regardless of the average size of precautionary saving, to understand its importance for business cycles, what matters is its response to shocks. In heterogeneous agent models, consumers do not respond much directly to interest rate shocks but respond strongly to employment shocks. I measure the response of the precautionary premium to employment and interest rate shocks. I compare it to the response of expected consumption growth net of precautionary growth.

Shocks. I consider two shocks. The first one is the monetary policy surprise built by Bauer and Swanson (2023). They measure monetary policy surprises using interest rate changes over 30-minute windows around Federal Open Market Committee (FOMC) announcements. They aggregate the measure into monthly interest rate surprises. I sum all the shocks The second one is an employment shock. It characterizes those who, four months prior, were unemployed and put a probability of 70% or more to be either unemployed or out of the labor force now, while in fact they are now employed.

Specification. I estimate the following specifications

$$\varphi_t = a_0 + a_1 MPS_t + a_2 emp_t + a_3 age_t + a_4 age_t^2 + a_5 income\ variability_t \quad (4.1)$$

$$+ a_6 children_t + a_7 period_t \varepsilon_t, \quad (4.2)$$

$$\Delta E_t[c_{t+1}^{nophi}] = a_0 + a_1 MPS_t + a_2 emp_t + a_3 age_t + a_4 age_t^2 + a_5 income\ variability_t \quad (4.3)$$

$$+ a_6 children_t + a_7 period_t + \varepsilon_t. \quad (4.4)$$

The term MPS_t measures the interest rate surprise on month t. The term emp_t is a dummy that takes value one if the respondent was unemployed four months before t and putting a more than 70% probability on being non-employed at t, but is actually employed at t. The term age_t is the age of the respondent at t. The term $income\ variability_t$ measures whether the respondent reports that their household income (not just earnings) is 'more or less constant, slightly variable, somewhat variable or highly variable from month to month?'. I use this as a control for overall household income risk. The variable $period_t$ is a dummy for the month-year t. The term $\Delta E_t[c_{t+1}^{nophi}]$ measures expected consumption growth minus φ_t.

Implementation. I estimate (4.2) and (4.4) with an OLS. Because there is a small panel dimensions, and two observations used in the estimation can come from the same respondent, I cluster the standard errors at the respondent level.

Results. Table (3) presents the results of the estimation of (4.2) and (4.4). The first two columns presents regressions that only include the shocks, without controls. They show that interest rate shocks have no effect on neither the precautionary nor the non-precautionary component of expected consumption growth. However, an employment shock significantly reduces the precautionary component of expected consumption growth. The point estimate is -83, significant at the 1% level. This means those moving from unemployment to employment while they did not expect it reduce their precautionary growth by \$85, which is about 40% of the average precautionary growth.

The third and fourth columns presents the results after including demographic controls. The effect of an employment shock on precautionary remains significant. The point estimate is similar, at -79. The demographics have no precisely measured effect on the precautionary component. Looking at the determinants of the non-precautionary component of consumption growth, the shocks have no effect. Having two children in the household raises the expected consumption growth compared to having none.

Component	Precautionary	Non-precautionary	Precautionary	Non-precautionary
Policy surprise	2195.6	-10812.9	-3197.5	-9046.3
	(2417.2)	(7599.6)	(2806.3)	(57867.7)
Employment shocks	-82.9***	-145.9	-79.3**	-113.6
	(31.9)	(530.0)	(39.5)	(570.1)
Income variability			60.1	-277.3
			(37.7)	(289.0)
Age			-13.0	35.5
			(12.7)	(58.6)
Age^2			0.2	-0.1
			(0.1)	(0.6)
Children=1			28.4	545.4
			(41.9)	(468.7)
Children=2			32.0	1125.0*
			(32.4)	(579.4)
Children=3			216.4	-529.5
			(151.6)	(1700.9)
Children=4			-6.3	1178.5
			(36.9)	(1431.9)
Constant	143.1***	2452.6***	196.5	436.1
	(23.1)	(211.9)	(265.5)	(1479.1)
Period FE	No	No	Yes	Yes
Observations	4,960	4,960	4,820	4,820

Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3: The response of the precautionary and non-precautionary components of consumption growth to shocks

This is consistent with the predictions of HANK models: taking into account uncertainty, households respond significantly to changes in employment but are not directly sensitive to changes in interest rate. The response to employment shocks is exclusively driven by the precautionary component of expected consumption growth. The non-precautionary component, which includes the effect of constraints, does not respond to these shocks.

5 Conclusion

In this paper, I show that the Euler equation literature misinterprets the conditions required for precautionary behavior not to matter in first order approximation, and misspecifies the second-order terms that would capture it.

I design a way to measure the contribution of precautionary behavior from the survey questions of the SCE. It explains 8% of expected consumption growth at the standard risk-aversion value of 2. Importantly, it significantly responds to employment shocks, while the non-precautionary component of consumption growth does not. No component of consumption growth responds to interest rate shocks.

References

Armantier, Olivier, Topa, Giorgio, Klaauw, Wilbert Van der, and Zafar, Basit (2017). "An overview of the Survey of Consumer Expectations." *Economic Policy Review* 23-2, pp. 51–72.

Attanasio, O. (1999). "Consumption." *Handbook of Macroeconomics*. Ed. by Taylor, J. B. and Woodford, M. Vol. 1. Handbook of Macroeconomics. Elsevier. Chap. 11, pp. 741–812.

Auclert, Adrien, Bardóczy, Bence, Rognlie, Matthew, and Straub, Ludwig (2021). "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models." *Econometrica* 89.5, pp. 2375–2408. DOI: 10.3982/ECTA17434.

Bauer, Michael D. and Swanson, Eric T. (2023). "A Reassessment of Monetary Policy Surprises and High-Frequency Identification." *NBER Macroeconomics Annual 2022, vol. 37*. Ed. by Eichenbaum, Martin, Hurst, Erik, and Ramey, Valerie. University of Chicago Press, pp. 87–155. DOI: 10.1086/723574.

Blundell, R., Low, H., and Preston, I. (2013). "Decomposing changes in income risk using consumption data." *Quantitative Economics* 4.1, pp. 1–37.

Blundell, R., Pistaferri, L., and Preston, I. (2008). "Consumption Inequality and Partial Insurance." *American Economic Review* 98.5, pp. 1887–1921.

Braxton, J. Carter, Herkenhoff, Kyle F, Rothbaum, Jonathan L, and Schmidt, Lawrence (2025). "Changing Income Risk across the US Skill Distribution: Evidence from a Generalized Kalman Filter." *American Economic Review*.

Browning, M. and Lusardi, A. (1996). "Household Saving: Micro Theories and Micro Facts." *Journal of Economic Literature* 34.4, pp. 1797–1855.

Campbell, J. and Mankiw, G. (1989). "Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence." NBER Chapters, pp. 185–246.

Carroll, C. (2001). "Death to the Log-Linearized Consumption Euler Equation (And Very Poor Health to the Second-Order Approximation)." *The B.E. Journal of Macroeconomics* 1.1, pp. 1–38.

Coeurdacier, Nicolas, Rey, Hélène, and Winant, Pablo (2011). "The Risky Steady State." *American Economic Review, Papers and Proceedings* 101.3, pp. 398–401. DOI: 10.1257/aer.101.3.398.

Deaton, A. (1992). *Understanding Consumption*. OUP Catalogue 9780198288244. Oxford University Press. ISBN: ARRAY(0x4e725868).

Federal Reserve Bank of New York (2015-2019). *Survey of Consumer Expectations.*

Ganong, Peter, Noel, Pascal J, Patterson, Christina, Vavra, Joseph S, and Weinberg, Alexander (2025). *Earnings Instability*. Working Paper 34227. National Bureau of Economic Research. DOI: 10.3386/w34227. URL: http://www.nber.org/papers/w34227.

Georgarakos, Dimitris et al. (2025). *How Costly Are Business Cycle Volatility and Inflation? A Vox Populi Approach*. Working Paper 33476. National Bureau of Economic Research (NBER). DOI: 10.3386/w33476.

Guvenen, F. and Smith, A. (2014). "Inferring Labor Income Risk and Partial Insurance From Economic Choices." *Econometrica* 82.6, pp. 2085–2129. ISSN: 1468-0262.

Guvenen, Fatih, Karahan, Fatih, Ozkan, Serdar, and Song, Jae (2021). "What Do Data on Millions of U.S. Workers Reveal About Lifecycle Earnings Dynamics?" *Econometrica* 89.5, pp. 2303–2339.

Hall, R. (1978). "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence." *Journal of Political Economy* 86.6, pp. 971–987.

Hansen, Lars Peter and Singleton, Kenneth J. (1983). "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns." *Journal of Political Economy* 91.2, pp. 249–265.

Jappelli, T. and Pistaferri, L. (2017). *The Economics of Consumption: Theory and Evidence*. 9780199383146. Oxford University Press. ISBN: 9780199383146.

Kaplan, G. and Violante, G. (2018). "Microeconomic Heterogeneity and Macroeconomic Shocks." *Journal of Economic Perspectives* 32.3, pp. 167–94.

Kimball, M. (1990). "Precautionary Saving and the Marginal Propensity to Consume." mimeo.