

Precautionary Saving and Consumption Fluctuations*

Jonathan A. Parker
Princeton University and NBER

Bruce Preston
Princeton University

April 2002

Abstract

This paper uses data on the expenditures of households to explain movements in the average growth rate of consumption in the U.S. from the beginning of 1982 to the end of 1997. We propose and implement a decomposition of consumption growth into series representing four proximate causes. These are new information, and three causes of predictable consumption growth: intertemporal substitution, changes in the preferences for consumption, and incomplete markets for consumption insurance. Incomplete markets for trading consumption in future states leads to statistically significant and countercyclical movements in expected consumption growth. The economic importance of precautionary saving rivals that of the real interest rate, but the relative importance of each source of movement in the volatility of consumption is not precisely measured.

*For helpful comments, we thank Charles Fleischman, Pierre Olivier Gourinchas, Bo Honore, Andrew Samwick, Chris Sims and seminar participants at the Australian National University, University of Chicago Graduate School of Business, Dartmouth College, Princeton University, the Reserve Bank of Australia, and the 2001 North American Winter Meeting of the Econometric Society. For financial support, Parker thanks the National Science Foundation grant SES-0096076, the Sloan Foundation, and a National Bureau of Economic Research Aging and Health Economics Fellowship through National Institute on Aging grant number T32 AG00186 and Preston thanks the Fellowship of Woodrow Wilson Scholars. Contact information and revised versions available at: www.princeton.edu/~jparker/.

1. Introduction

According to canonical macroeconomic theory, aggregate consumption is the result of the optimal choices of a representative consumer. Consumption growth changes over time due to changes in the risk-free rate of return and changes in the information that the consumer has about current wealth, future income, and future rates of return. Consistent with this approach, consumption growth is largely unpredictable. But contrary to the canonical theory, predictable movements in aggregate consumption growth have little correlation with the risk-free rate of return and significant correlation with predictable movements in income. For the representative agent approach, there are several puzzling features of US consumption data, such as excess sensitivity, excess smoothness, and the equity premium puzzle.¹

This paper measures the relative importance of new information, the real interest rate, the preference for consumption, and precautionary saving for movements in average consumption growth. We develop a method that uses the consumption Euler equation and survey data on the consumption expenditures of households to decompose consumption growth into these four proximate causes, and factors outside the model. Using both estimates and calibrations for a standard utility function, we implement this decomposition and find a significant role for precautionary saving in consumption fluctuations.

We focus particularly on predictable movements in consumption growth caused by changes in the amount of consumption risk faced by households, or precautionary saving. We do this for three reasons. First, because consumers face idiosyncratic risk that they do not completely insure, precautionary saving is potentially a significant determinant of consumption dynamics.² The consumption dynamics of price-taking households in a model economy without complete consumption insurance can be quite

¹Flavin (1981), Shiller (1982), Hall (1988), Campbell and Deaton (1989), and Campbell and Mankiw (1989).

²Complete consumption insurance is rejected by Nelson (1994), Cochrane (1991), and Attanasio and Davis (1996).

different from the behavior of these same households in the equivalent economy with complete markets.³ More specifically, Caballero (1990) and Carroll (1997) argue that many of the existing puzzles of the consumption literature can in theory be explained by precautionary saving. Despite the implications of these models, there is no direct evidence on whether precautionary saving causes movements in expected consumption growth in practice. Such evidence is our main substantive contribution.

Second, the measurement of precautionary saving is not straightforward. The methodological contribution of this paper is the derivation of a robust measure of the predictable fluctuations in aggregate consumption that are due to consumption risk. Consistent estimates of the role of precautionary saving allow us to consistently decompose consumption fluctuations. Finally, measurement of precautionary saving helps gauge its broader significance. To the extent that consumption risk and constraints on intertemporal substitution are important for aggregate consumption dynamics, they are more likely to be important for other areas of research such as asset pricing and the design of optimal policy.⁴

Our method proceeds as follows. We posit a household marginal utility function. We alternately estimate and calibrate the parameters of this function. Given this measure of marginal utility and the consumption Euler equation, we allocate movements in consumption to four proximate causes: new information, changes in the risk-free rate of return, changes in the preference for consumption, and changes in expected risk to consumption or liquidity constraints. In particular, predictable consumption growth due to precautionary saving is the predictable variation in a nonlinear function of the innovation to marginal utility. We deal with three complications. First, we derive a measure that is robust to some misspecification of preferences. Second, measurement error in consumption data implies that mean consumption growth cannot be consistently decomposed, only movements around the mean. Finally, low-wealth households in the

³See Rios-Rull (1994), Krusell and Smith (1998) and Gourinchas (2000).

⁴Such potentially important roles for precautionary saving are considered by Caballero (1991), Constantinides and Duffie (1996) and Aiyagari (1995).

economy may face binding liquidity constraints. We measure consumption growth due to both precautionary saving and liquidity constraints for these households. Aggregating across households, we decompose average consumption growth into its proximate causes. These causes are proximate, rather than structural or exogenous, in the same way the causes of output growth in a Solow decomposition are proximate.⁵

Implementing the decomposition using data from the Consumer Expenditure Survey (CEX), we find that precautionary saving makes a statistically significant contribution to fluctuations in expected consumption growth. The economic importance of each source of predictable variation in consumption growth is not isolated due to covariation among the proximate causes. The share of variance of predictable consumption growth sourced to incomplete markets and that sourced to the real interest rate both vary over similar wide ranges. The pattern of this correlation among sources is itself informative. Consumption growth due to precautionary saving covaries positively with that due to the real interest rate, suggesting that low estimates of the intertemporal elasticity of substitution in linearized models are not due to the fact that these models omit the precautionary saving term. Since consumption growth due to shifts in the preference for consumption covaries negatively with that due to preferences, low observed intertemporal substitution in response to the real interest rate is largely due to omitted preference shifters or misspecification of the consumption optimization problem more generally.⁶

We further find that expected consumption growth due to precautionary savings has declined during the time period of the sample, 1982 to 1997. Consumption growth due to precautionary savings is high following the 1982 recession, consistent with high consumption risk and low levels of consumption during the recession. More generally, con-

⁵In particular, the impact of precautionary saving is not a deep structural parameter, but an endogenous outcome. In order to understand the link between precautionary saving and structural parameters, one must model market completeness, endowment risk, factor prices, etc.

⁶See for example the models of Baxter and Jermann (1999), Basu and Kimball (2000), and Ogaki and Reinhart (1998). Carroll (1997) argues that incomplete markets are an important source of bias, Attanasio and Weber (1995) finds that labor supply is an important shifter of the preference for consumption, and Attanasio and Weber (1995) finds some evidence for both views.

sumption growth due to incomplete markets is countercyclical: expected consumption growth is greater when the unemployment rate is expected to increase. Finally, while there is no statistically significant relationship between money shocks and subsequent consumption growth due to precautionary saving, expected increases in government spending are associated with faster consumption growth due to precautionary saving.

Our research builds most directly on papers that exploit the variation and information in detailed household-level survey data on consumption expenditures to better understand consumer behavior and the dynamics of U.S. consumption (such as Hall and Mishkin (1982), Zeldes (1989a), Attanasio and Weber (1995), and Parker (1999)). Our focus on precautionary saving owes much to Dynan (1993), which estimates the parameters of the utility function from the cross-sectional relationship between consumption growth and the expected volatility of consumption. Our research is also related to previous methodologies for inferring the importance of precautionary saving, such as Skinner (1988), Carroll and Samwick (1997), Banks, Blundell, and Brugiavini (2001), Lusardi (1998), Gourinchas and Parker (2002), Cagetti (1998), and Storesletten, Telmer, and Yaron (2000).

The paper proceeds as follows: section 2 shows how to decompose consumption growth using the consumption Euler equation and discusses the central economic assumptions. Section 3 demonstrates that, in the basic decomposition, precautionary saving is a catch-all residual, and contains all movement in consumption not allocated to the real interest rate or preferences. This section derives a robust measure of expected consumption growth due to precautionary saving that avoids this shortcoming. Section 3 further discusses treatment of some small-sample bias and the presence of households that may be liquidity constrained. Section 4 describes our use of the Consumer Expenditure Survey and the variables that we employ in our analysis. Section 5 presents our results for households that are not liquidity constrained and Section 6 presents our results for all households. A final section concludes and Appendixes contain additional details on the decomposition and estimation of parameters.

2. Theory

This section derives a decomposition of consumption growth from the consumption Euler equation. The economy consists of households that receive uncertain income over time and in each period make both a saving and a portfolio decision. Households are expected utility maximizers and price takers. Given the usual assumptions on preferences, a probability space, and budget constraints, equilibrium consumption for the household obeys the consumption Euler equation for any two periods t and $t + 1$:

$$u'_{i,t}(C_{i,t}) = \delta_i R_{i,t+1}^f E_t [u'_{i,t+1}(C_{i,t+1})] \quad (2.1)$$

where δ_i is household i 's discount factor; $u'_{i,t}$ maps expenditures on consumption, C_t , into the household's marginal utility in period t ; E_t is the expectations operator conditional on the time t information set, common across households; and $R_{i,t+1}^f$ is the gross risk-free rate of return between t and $t + 1$.

We clarify four points implicit in this Euler equation. First, markets in the economy may or may not be complete, in that the available assets may or may not span all possible future states of nature. Second, other than the risk-free rate, $R_{i,t+1}^f$, we do not assume that households have access to the same assets. For example, there might be fixed costs associated with entering the equity market and some households might choose not to participate in this market. Third, equation (2.1) rules out borrowing constraints. Subsequently we will weaken the assumption that every household can borrow or lend freely at rate $R_{i,t+1}^f$. Finally, since $u'_{i,t}(\cdot)$ can depend on any exogenous or endogenous characteristic of the household, equation (2.1) makes no assumptions as to the separability of marginal utility from any of the standard suspects like family structure or leisure. Consequently, equation (2.1) has no scientific content in that it has no testable implications. To proceed, we must place further restrictions on the model.

We restrict the differences in utility across households and time in a way that provides identification, is *ex ante* reasonable, and is consistent with prior research. The discount factor is taken to be an exponential function of a linear combination of L

permanent household characteristics, $\{Q_{i,t}^l\}_{l=1}^L$,

$$\delta_i = \exp(\mathbf{Q}_i^A \boldsymbol{\delta}^A) \quad (2.2)$$

where $\mathbf{Q}_i^A = (1, Q_i^2, Q_i^3, \dots, Q_i^L)$, and $\boldsymbol{\delta}^A = (\delta_1, \delta_2, \dots, \delta_L)'$. Discount rates may differ across the population, in accord with education levels, for example. The marginal utility function is also shifted by $K - L$ other household characteristics, $\{Q_{i,t}^l\}_{l=L+1}^K$. Specifically, marginal utility is the product of a function of these household characteristics and a function of consumption. Flow utility is

$$u_{i,t}(C_{i,t}) \equiv \exp(\mathbf{Q}_{i,t}^B \boldsymbol{\delta}^B) \frac{C_{i,t}^{1-\frac{1}{\sigma}}}{1-\frac{1}{\sigma}} + f(\mathbf{Q}_{i,t}^B)$$

where $\mathbf{Q}_{i,t}^B = (Q_{i,t}^{L+1}, Q_{i,t}^{L+2}, \dots, Q_{i,t}^K)$, $\boldsymbol{\delta}^B = (\delta_{L+1}, \delta_{L+2}, \dots, \delta_K)'$, and σ denotes the (constant) intertemporal elasticity of substitution. The marginal utility of a given level of consumption expenditures is nonseparable from the factors in $\mathbf{Q}_{i,t}^B$; the marginal utility of a given level of consumption may move in response to changes in the size of the family or in labor supply, for example. If these changes are expected, they lead to expected movements in consumption.

Gathering these assumptions, the consumption Euler equation is

$$C_{i,t}^{-\frac{1}{\sigma}} = E_t \left[R_{i,t+1}^f \exp(\mathbf{X}_{i,t+1} \boldsymbol{\delta}) C_{i,t+1}^{-\frac{1}{\sigma}} \right] \quad (2.3)$$

where $\mathbf{X}_{i,t+1} \equiv (\mathbf{Q}_i^A, \mathbf{Q}_{i,t+1}^B - \mathbf{Q}_{i,t}^B)$ and $\boldsymbol{\delta} = (\boldsymbol{\delta}^{A'}, \boldsymbol{\delta}^{B'})'$, a $K \times 1$ vector of preference parameters. Equation (2.3) defines an expectation error for every household

$$\varepsilon_{i,t+1} = R_{i,t+1}^f \exp(\mathbf{X}_{i,t+1} \boldsymbol{\delta}) \left(\frac{C_{i,t+1}}{C_{i,t}} \right)^{-\frac{1}{\sigma}} - 1 \quad (2.4)$$

where $E_t[\varepsilon_{i,t+1}] = 0$.

To move from the dynamics of marginal utility to the dynamics of consumption, take logs and re-organize equation (2.4),

$$\Delta \ln C_{i,t+1} = \sigma \ln R_{i,t+1}^f + \sigma \mathbf{X}_{i,t+1} \boldsymbol{\delta} - \sigma \ln(1 + \varepsilon_{i,t+1}). \quad (2.5)$$

This equation relates actual consumption growth to errors in expectation, changes in preferences, and intertemporal substitution in response to the price of a good delivered with certainty in $t + 1$ in terms of goods at t . To relate consumption growth to precautionary saving, take the conditional expectation of equation (2.5) and add the unexpected part of consumption growth to both sides, to yield our main decomposition:

$$\Delta \ln C_{i,t+1} = \sigma E_t[\ln R_{i,t+1}^f] + \sigma E_t[X_{i,t+1}] \delta + \Delta \ln C_{i,t+1}^{IM} + \eta_{i,t+1}^C \quad (2.6)$$

where

$$\Delta \ln C_{i,t+1}^{IM} \equiv -\sigma E_t[\ln(1 + \varepsilon_{i,t+1})] \quad (2.7)$$

$$\eta_{i,t+1}^C \equiv \Delta \ln C_{i,t+1} - E_t[\Delta \ln C_{i,t+1}]$$

are the contribution of precautionary saving to consumption growth and the contribution of new information to consumption growth respectively. The notation “*IM*” is used for “incomplete markets.” This decomposition is similar to the approximate nonlinear equation used by Dynan (1993) to estimate parameters. We use the exact relationship to decompose consumption growth given parameters.

Equation (2.6) decomposes consumption growth only into proximate causes. It is analogous to a Solow decomposition, which decomposes output growth into growth due to capital, labor and productivity, rather than into growth due to technology or policies. We decompose consumption growth into growth due to the listed factors rather than growth due to technology or policies. As an example, a primitive shock to the economy that alters risk typically changes not only precautionary saving but also wealth accumulation and real interest rates and therefore would appear in both consumption growth due to the real interest rate and that due to precautionary saving. Similarly, at the aggregate level, all four components of consumption growth likely contain some consumption movement driven at the primitive level by preference shifts – changes in consumption allocated directly to preferences measure only the intertemporal substitution due to expected changes in preferences. The interesting question is how important are nonseparabilities and incomplete markets for the propagation of these primitive shocks and so for the dynamics of aggregate consumption.

A related point is that $\Delta \ln C_{i,t+1}^{IM}$ measures only the effect of consumption risk on predictable consumption growth. A primitive shock to the economy that increases future consumption volatility affects both predictable consumption growth and the size of the innovation to consumption growth. Only the former appears in our proximate causes as precautionary saving, although unpredictable and predictable movements are linked through the budget constraint.

How does the presence of liquidity constraints alter this decomposition? Let $\lambda'_{i,t}$ denote the multiplier on the constraint on wealth, then optimal behavior implies

$$\begin{aligned} C_{i,t}^{-\frac{1}{\sigma}} &= E_t \left[R_{i,t+1}^f \exp(\mathbf{X}_{i,t+1} \boldsymbol{\delta}) C_{i,t+1}^{-\frac{1}{\sigma}} \right] + \lambda'_{i,t} \\ \lambda'_{i,t} &\geq 0. \end{aligned} \quad (2.8)$$

Let $\epsilon_{i,t+1}$ be the true innovation to marginal utility implied by equation (2.8). Following the same procedure, consumption growth is given by the decomposition in equation (2.6)-(2.7), but with

$$\Delta \ln C_{i,t+1}^{IM} = -\sigma E_t [\ln(1 + \epsilon_{i,t+1} - \lambda_{i,t})] \quad (2.9)$$

where $\lambda_{i,t} \equiv \lambda'_{i,t} C_{i,t}^{\frac{1}{\sigma}}$. Thus, if some households in the economy face binding liquidity constraints, the effect of these constraints appears as precautionary saving. We refer to equation (2.9) as consumption growth due to incomplete markets, as opposed to equation (2.7) which, for an unconstrained household, measures only the effects of precautionary saving. As before, equation (2.9) measures only the direct or proximate role of market incompleteness in consumption growth. Finally, it is worth emphasizing that little is lost by measuring the impact of liquidity constraints and precautionary saving together. Liquidity constraints are missing markets for transferring consumption over time; if precautionary saving is at all important, it comes from missing markets for trading goods across future states. These effects are substantively and technically closely related.⁷

⁷See Carroll and Kimball (2001), and the similarity in consumption behavior between Carroll (1997) on the one hand and Zeldes (1989b) and Deaton (1991) on the other.

Turning to aggregate consumption growth, cross-sectional averages give

$$\begin{aligned}\Delta \ln C_{t+1}^{Ag} &= \sigma E_t[\ln R_{t+1}^f] + \sigma E_t[\mathbf{X}_{t+1}] \boldsymbol{\delta} + \Delta \ln C_{t+1}^{IM} + \eta_{t+1}^C \\ &\equiv \Delta \ln C_{t+1}^R + \Delta \ln C_{t+1}^\delta + \Delta \ln C_{t+1}^{IM} + \eta_{t+1}^C\end{aligned}\quad (2.10)$$

where $\eta_{t+1}^C = \Delta \ln C_{t+1}^{Ag} - E_t[\Delta \ln C_{t+1}^{Ag}]$, and where for any variable $x_t = \frac{1}{H(t)} \sum_{i \in H_t} x_{i,t}$ with the slight abuse of notation that

$$\Delta \ln C_{t+1}^{Ag} = \frac{1}{H(t)} \sum_{i \in H_t} \Delta \ln C_{i,t+1}, \quad \Delta \ln C_{t+1}^{IM} = \frac{1}{H(t)} \sum_{i \in H_t} \Delta \ln C_{i,t+1}^{IM} \quad (2.11)$$

where H_t denotes the set of $H(t)$ households alive at time t .⁸ According to equation (2.10), consumption growth is composed of four pieces: intertemporal substitution due to variations in rate of return ($\Delta \ln C_{t+1}^R$); intertemporal substitution due to variations in preferences ($\Delta \ln C_{t+1}^\delta$); intertemporal substitution in response to variations in consumption risk or binding liquidity constraints ($\Delta \ln C_{t+1}^{IM}$); and unexpected movements in consumption growth due to innovations to marginal utility (η_{t+1}^C). In this formulation, the innovation to consumption growth is orthogonal to the expected movements in consumption due to preference shifts, the interest rate, and incomplete markets.

3. Estimation methodology

Our goal is to use estimates of preference parameters and survey data on household characteristics and consumption choices to construct empirical counterparts to $\Delta \ln C_{t+1}^{Ag}$ and the conditional moments $\Delta \ln C_{t+1}^R$, $\Delta \ln C_{t+1}^\delta$, $\Delta \ln C_{t+1}^{IM}$ and $E_t[\Delta \ln C_{t+1}^{Ag}]$. This section addresses two potential biases that arise in the measurement of the role incomplete markets. First, in the theoretical decomposition, all predictable movements in consumption not attributable to preferences and the real interest rate are assigned by construction to precautionary saving and liquidity constraints. We choose instead to use a measure of consumption growth due to incomplete markets that is robust to some

⁸We require the expectation in front of the real interest rate series because the household-level real interest rates are not completely risk free.

types of model misspecification. Second, the sample analog to the conditional moment $E_t[\varepsilon_{i,t+1}]$ is zero with probability zero in any finite sample. Since the maintained assumptions imply this term should be zero, and since it would otherwise contribute to precautionary saving, we impose a correlation of zero in the measure of precautionary saving. It is important to note that if there is no misspecification or finite sample correlation, the measure of precautionary saving is not biased by our corrections, having the same asymptotic properties as the unadjusted measure. This section describes the removal of these biases for the economy without liquidity constraints and then formally defines our estimators for the more general economy in which binding liquidity constraints may occur. Appendixes A and B provide further details.

The decomposition requires parameter values, which are chosen alternately by estimation and calibration. When estimating, we apply a generalized method of moments (GMM) procedure based on the nonlinear Euler equation to a synthetic panel. When calibrating, we choose parameters based upon the large previous literature studying consumer behavior. Particularly for the calibration results, some misspecification is likely and this biases estimates of the contribution of incomplete markets to consumption growth. In fact, the measure of consumption growth due to incomplete markets absorbs all unexplained movements in expected consumption growth. To see this, substitute equation (2.4) into the cross-sectional average of equation (2.7):

$$\begin{aligned}
\Delta \ln C_{t+1}^{IM} &= -\sigma \frac{1}{H(t)} \sum_{i \in H_t} E_t [\ln (1 + \varepsilon_{i,t+1})] \\
&= -\sigma E_t \left[\frac{1}{H(t)} \sum_{i \in H_t} \ln \left(1 + R_{i,t+1}^f \exp\{\mathbf{X}_{i,t+1} \boldsymbol{\delta}\} \left(\frac{C_{i,t+1}}{C_{i,t}} \right)^{-\frac{1}{\sigma}} - 1 \right) \right] \\
&= -\sigma E_t \left[\frac{1}{H(t)} \sum_{i \in H_t} \ln R_{i,t+1}^f + \mathbf{X}_{i,t+1} \boldsymbol{\delta} - \frac{1}{\sigma} \Delta \ln C_{i,t+1} \right] \\
&= E_t \left[\Delta \ln C_{t+1}^{Ag} - \Delta \ln C_{t+1}^R - \Delta \ln C_{t+1}^{\delta} \right].
\end{aligned}$$

If the model were correct, then this feature would not be a concern. But in practice, the selection of factors influencing utility and the specification of parameters is not perfect, and assuming it is would label our ignorance or mismeasurement of precautionary

saving.⁹ Therefore, we modify our measure of precautionary saving so that it is robust to some types of model misspecification, and decompose consumption growth into each proximate cause and a time series of expected consumption growth that represents our ignorance.

We consider the following type of misspecification. Suppose that the true model is given by

$$C_{i,t}^{-\frac{1}{\sigma}} = E_t \left[\exp\{\mathbf{X}_{i,t+1}\boldsymbol{\delta}\} \theta_{i,t} R_{i,t+1}^f C_{i,t+1}^{-\frac{1}{\sigma}} \right] \quad (3.1)$$

where $\theta_{i,t}$ is an omitted preference shifter, an incorrect choice of $\boldsymbol{\delta}^B$, or unobserved variation in the actual risk-free real interest rate faced by a given household. Let $\epsilon_{i,t+1}$ denote the true innovation implied by equation (3.1). The innovation used in the decomposition, given by equation (2.4), is

$$\varepsilon_{i,t+1} = \frac{1 + \epsilon_{i,t+1}}{\theta_{i,t}} - 1.$$

It follows that equation (2.7) is an inconsistent estimate of precautionary saving unless $\theta_{i,t} = 1$ since

$$E_t [\varepsilon_{i,t+1}] = \frac{1}{\theta_{i,t}} - 1.$$

To estimate precautionary saving consistently, we replace $\varepsilon_{i,t+1}$ in equation (2.7) with $\tilde{\varepsilon}_{i,t+1}$ defined as

$$\tilde{\varepsilon}_{i,t+1} = \frac{\varepsilon_{i,t+1} - E_t [\varepsilon_{i,t+1}]}{1 + E_t [\varepsilon_{i,t+1}]} \quad (3.2)$$

Since

$$E_t [\tilde{\varepsilon}_{i,t+1}] = 0,$$

⁹An analogy is helpful here. Assume that total factor productivity only changes because of increases in technology and (unmeasured) changes in factor utilization. In this situation, a Solow residual measures changes in productivity due to both sources. Thus, economists develop methods to construct Solow residuals that are purged of the utilization component and measure only actual changes in technology. Similarly, our basic measure of consumption growth due to incomplete markets may contain movements due to unmodelled preference shifts; we develop a method to construct consumption growth due to incomplete markets that is purged of this misspecification and thus isolate actual consumption growth due to changes in consumption risk.

using $\tilde{\varepsilon}_{i,t+1}$ in equation (2.7) yields

$$\begin{aligned}
\Delta \ln C_{i,t+1}^{IM} &= -\sigma E_t [\ln(1 + \tilde{\varepsilon}_{i,t+1})] \\
&= -\sigma E_t \left[\ln \left(1 + \frac{\varepsilon_{i,t+1} - E_t[\varepsilon_{i,t+1}]}{1 + E_t[\varepsilon_{i,t+1}]} \right) \right] \\
&= -\sigma E_t \left[\ln \left(1 + \frac{\frac{1+\varepsilon_{i,t+1}}{\theta_{i,t}} - 1 - \frac{1-\theta_{i,t}}{\theta_{i,t}}}{1 + \frac{1-\theta_{i,t}}{\theta_{i,t}}} \right) \right] \\
&= -\sigma E_t [\ln(1 + \varepsilon_{i,t+1})]
\end{aligned} \tag{3.3}$$

which is the concept of interest. This technique leads to an additional source of movement in average consumption, which we call misspecification and denote $\Delta \ln C_{i,t+1}^\theta$. To be clear, this is not all possible misspecification. In particular, mismeasurement of σ or misspecification of stochastic preference shifters are not addressed by this correction. Finally, if the original Euler equation is not misspecified, that is if $\theta_{i,t} = 1$, then $E_t[\varepsilon_{i,t+1}] = 0$, and this correction does not bias any of our measures.

The second source of bias is the possibility that, in a finite sample, the measured expectation error may be correlated with the innovation to marginal utility. The expectations of households are estimated in the standard manner, as fitted values based on least-squares projections of the variable onto instruments, $\mathbf{Z}_{i,t} = (Z_{i,t}^1, Z_{i,t}^2, \dots, Z_{i,t}^J)'$. Let E_T denote the sample conditional expectation, let I_t denote the set of $I(t)$ households observed at time t , and let $\tilde{\varepsilon}_{i,t+1}$ denote the sample analog to $\tilde{\varepsilon}_{i,t+1}$. The problem we address is that $\tilde{\varepsilon}_{i,t+1}$ does not necessarily have the property

$$E_T \left[\frac{1}{I_t} \sum_{i \in I_t} \tilde{\varepsilon}_{i,t+1} | \mathbf{Z}_t \right] = 0. \tag{3.4}$$

As an example, estimation of parameters by overidentified GMM would assume this condition holds asymptotically, but not impose it in any observed sample.¹⁰ It is also

¹⁰In theory we could estimate σ and δ from equation (3.4), so that this second bias correction, while not causing any biases, would be unnecessary. This procedure would require collapsing all the data into a single time series and using a single set of instruments to estimate all the potential preference heterogeneity in the population. Instead, we estimate heterogeneity in the predictable movements in marginal utility using information available at the household level and panel data. As an example, this

straightforward to see that this condition is not guaranteed in finite samples by our correction in equation (3.2) for expectations calculated as linear projections.

We correct the measure of precautionary saving by subtracting this term, the left hand side of equation (3.4). The robust sample analog to equations (2.11) and (2.7) is

$$\Delta \ln \widehat{C}_{t+1}^{IM} = -\sigma E_T \left[\frac{1}{I(t)} \sum_{i \in I_t} \ln(1 + \tilde{e}_{i,t+1}) - \tilde{e}_{t+1} | \mathbf{Z}_t \right].$$

Consider the expectation of a second-order expansion of the logarithm around $\tilde{e}_{t+1} = 0$:

$$\begin{aligned} -E_t \left[E_T \left[\frac{1}{I(t)} \sum_{i \in I_t} \tilde{e}_{i,t+1} - \frac{\tilde{e}_{i,t+1}^2}{2} - \tilde{e}_{t+1} | \mathbf{Z}_t \right] \right] &= \frac{1}{2} E_t \left[\frac{1}{I(t)} \sum_{i \in I_t} \tilde{e}_{i,t+1}^2 \right] \\ &= \frac{1}{2} \overline{Var}_t [\epsilon_{i,t+1}] \end{aligned}$$

which, up to the second-order approximation, is the cross-sectional average of interest. Note that if $E_T \left[\frac{1}{I_t} \sum_{i \in I_t} \tilde{e}_{i,t+1} | \mathbf{Z}_t \right] = 0$, then corrected and uncorrected measures of precautionary saving are numerically identical. Appendix A shows each adjustment is necessary: removing sample correlation inside the logarithm leads to bias, and removing misspecification outside the logarithm leads to inconsistency.

The possibility of liquidity constraints alters the implementation of these corrections. As described in section 2, for a constrained household, $\epsilon_{i,t+1}$ measures $\epsilon_{i,t+1} - \lambda_{i,t}$. It follows that $E_t [\epsilon_{i,t+1}]$ should not equal zero for constrained households. Therefore, we omit the small-sample correction for possibly constrained households. For the misspecification adjustment, we estimate $E_t [\epsilon_{i,t+1}] = \frac{1}{\theta_{i,t}} - 1$ as a function of characteristics of household i using the sample of households that are not constrained. We denote the sample function as $E_{T,\lambda=0} [\epsilon_{i,t+1} | \mathbf{Z}_{i,t}]$ and use this as the sample analog to $E_t [\epsilon_{i,t+1}]$ in constructing $\tilde{e}_{i,t+1}$. Appendix B shows that as long as $E_t [E_{T,\lambda=0} [\epsilon_{i,t+1} | \mathbf{Z}_{i,t}]] = \frac{1}{\theta_{i,t}} - 1$, the presence of liquidity constraints does not alter the interpretation of our decomposition. If this assumption fails, the interpretation changes: expected changes in preferences that occur for constrained households and that are not predicted by $\mathbf{Z}_{i,t}$ for unconstrained allows us to measure the difference in discount rates among household types in part from cross-sectional variations in consumption growth rather than trying to infer these differences from the time variation in aggregate consumption and the time-varying share of each household type in the population.

households are allocated to the series on liquidity constraints rather than preferences. Both the appendix and the results section contain examples of interpretation.

The estimators of the proximate causes of consumption fluctuations are

$$\Delta \widehat{\ln C}_{t+1}^{Ag} = \frac{1}{I(t)} \sum_{i \in I_t} \Delta \ln C_{i,t+1} \quad (3.5)$$

$$\Delta \widehat{\ln C}_{t+1}^R = \sigma E_T \left[\frac{1}{I(t)} \sum_{i \in I_t} \ln R_{i,t+1}^f | \mathbf{Z}_t \right] \quad (3.6)$$

$$\Delta \widehat{\ln C}_{t+1}^\delta = \sigma E_T \left[\frac{1}{I(t)} \sum_{i \in I_t} \mathbf{X}_{i,t+1} | \mathbf{Z}_t \right] \delta \quad (3.7)$$

$$\Delta \widehat{\ln C}_{t+1}^{IM} = -\sigma E_T \left[\frac{1}{I(t)} \sum_{i \in I_t} \ln(1 + \tilde{e}_{i,t+1}) - \frac{1}{I(t, \lambda=0)} \sum_{i \in I_t, \lambda=0} \tilde{e}_{i,t+1} | \mathbf{Z}_t \right] \quad (3.8)$$

$$\hat{\eta}_{t+1}^C = \Delta \widehat{\ln C}_{t+1}^{Ag} - E_T \left[\Delta \widehat{\ln C}_{t+1}^{Ag} | \mathbf{Z}_t \right] \quad (3.9)$$

where

$$\tilde{e}_{i,t+1} = \frac{\varepsilon_{i,t+1} - E_{T, \lambda=0} [\varepsilon_{i,t+1} | \mathbf{Z}_{i,t}]}{1 + E_{T, \lambda=0} [\varepsilon_{i,t+1} | \mathbf{Z}_{i,t}]} \quad (3.10)$$

The measures of consumption growth due to the risk-free real interest rate, preferences, and new information all follow directly from equation (2.10). The measure of incomplete markets includes growth due to precautionary saving and liquidity constraints, is robust to a class of misspecification, and removes movements due to any in-sample covariation between the innovation to marginal utility and the instruments for unconstrained households. Finally, there is some additional consumption growth, $\sigma E_T [\frac{1}{I_t} \sum_{i \in I_t} \ln(1 + \tilde{e}_{i,t+1}) - \ln(1 + \varepsilon_{i,t+1}) | \mathbf{Z}_t]$. When reporting results based on estimation, for which it is reasonable to believe that this truly is misspecification, we omit this source of variation from the decomposition. When reporting results based on calibrations, we include this source of variation as variation due to preferences.

Three points about the decomposition deserve note. First, both of the corrections described in this section leave the consumption growth due to incomplete markets unbiased if there is no need for them. Second, the methodology delivers a lower bound on the proximate contribution of the interest rate, preferences, and incomplete markets to fluctuations in consumption growth. Unpredictable movements in marginal utility

are estimates of the innovations to marginal utility. To the extent that the information set used to make these predictions is smaller than that available to the household, the role of innovations is overstated and the role of other factors understated. Third, as demonstrated in the next section on the data, because of noise in measurement of household consumption, the sources of average consumption growth cannot be consistently decomposed. Instead, we decompose the fluctuations in these time series, not their means.

4. Data and variables

This section describes our use of the information on consumption expenditures in the CEX, and the implications of measurement error in consumption for inference.

We use data from the Family and Detailed Expenditure files of the Consumer Expenditure Survey (CEX) to make an unbalanced, overlapping panel of households that contains consumption data from December 1981 to February 1998, with some periods omitted due to changes in the survey.¹¹ The Bureau of Labor statistics constructs the CEX data from a series of interviews based on a stratified random sample of the U.S. population. Unless a household attrits, it is interviewed five times, once every three

¹¹The 1997 files include data on household expenditures for all three-month interview periods starting in 1997, so that the data we use cover up to and including February of 1998. We use both the raw data files and SAS files available from Lorna Greening at: <ftp://elsa.berkeley.edu/pub/ices/>. We omit 1980 and 1981 data because it is of significantly worse quality. Due to decennial survey changes, we cannot match consumption growth to the first three months of 1986 or 1996. When decomposing consumption growth we also drop the three observations on consumption growth across 1987 to 1988 and the three across 1995 to 1996 due to large increases in the variance of consumption growth in there periods due to changes in the survey instrument. Finally, when decomposing consumption growth, we drop the observation in which $t + 1$ ends in July 1996 due to several strange factors. It has a variance of consumption growth similar to a survey change, and it has the largest mean movement in several preference categories. This inclusion of this single observation does not substantively alter our inference. This leaves 16 years and three months less two months of observations in levels, for a time dimension of 177 three-month to three-month growth rates and a pre-sample of 6 months.

months. In a household's first interview, the CEX procedures are explained to the members of the household and they are asked to keep track of their expenditures for future interviews. In each subsequent interview, detailed information is collected on the past three months' consumption expenditures. New households replace those that leave the sample on a monthly basis, so that while consumption observations cover three months, the data are effectively monthly. In each family's second and fifth interviews, demographic and income data are collected, including income and earnings information about the previous 12 months. This information is updated if it changes during the course of the survey year. In a household's final interview, a set of questions about assets is asked, including the amount by which wealth changed over the previous 12 months.

We define consumption as the sum of expenditures on food, alcoholic beverages, apparel and apparel services, gasoline and motor oil used in transportation, public transportation, entertainment, personal care, and reading.¹² Data are deflated by the consumer price index for each category of consumption for the Census region in which the household resides.

The real interest rate is the expected return on a three-month U.S. Treasury bill less the constructed, household-specific inflation rates based on the deflation method described above. Thus, estimation uses a household-specific real interest rate even though we consider the same asset for all households. We use monthly data, averaged to generate real interest rates covering three-month to three-month periods.¹³

$E_{T,\lambda=0}[\cdot|Z_{i,t}]$ and $\bar{e}_{i,t+1}$ are estimated from linear regressions on instruments $Z_{i,t}$ that the household might use to forecast future wealth, preferences, and risk, for a sample of households deemed unlikely to be constrained. Households are counted as unconstrained

¹²Other categories that might be included that we do not include are: tobacco, because it is addictive; household operations, because it includes repairs of furniture, appliances, computers, etc. as well as day care expenses including tuition; and utilities, because, apart from phone, these are to a large extent determined by one's housing choice.

¹³As discussed in Hall (1988).

if they have more than 3,000 1982 – 84 dollars in wealth in saving accounts, checking accounts, government bonds, stocks, and mutual funds, less debts, as of the start of their first report of expenditures. We also include any household reporting topcoded wealth in the final interview in this sample. Due to a large amount of missing wealth data, this split leads to only one quarter of the sample being clearly unconstrained – an implausibly small share of the population. Accordingly, we supplement this sample by using prior information about the characteristics of households in the economy that are likely to face binding liquidity constraints. Households are also counted as unconstrained if they have at least some college education and have age greater than 45.¹⁴ These allocation rules impose that half the sample of households are unconstrained.

All instruments are constructed so that they are in the household information set at the start of period t .¹⁵ $Z_{i,t}$ contains: a quadratic function of age, month indicator variables, four education indicator variables, and four aggregate forecast variables – the expected unemployment rate at $t + 1$, its change from t , and expected real and nominal interest rates at $t + 1$, all interacted with five family types.¹⁶ That is, $\tilde{e}_{i,t+1}$ is the fitted value from 5 separate regressions of $\varepsilon_{i,t+1}$ on age, age squared, month and education dummy variables, and the four aggregate variables. The aggregate instruments are forecasts constructed from rolling regressions.¹⁷ In all of our analysis, we eliminate the variation due to month.

$E_T[.|Z_t]$ and $\tilde{e}_{t+1}$ are estimated via regression using the aggregate variables: a constant, the four forecasts of aggregate variables, month indicators, and twice lagged

¹⁴Meghir and Weber (1996) and Japelli, Pischke, and Souleles (1998) suggest that these characteristics are good indicators of lack of liquidity constraints.

¹⁵This allows for the possibility of time aggregation.

¹⁶The education groups are some high school, high school graduate, some college, college graduate. Education is assigned on the basis of the male head if present, otherwise on the basis of the female head. The family types are single, single parent with children, married couple, married couple with children, unrelated individuals.

¹⁷The rolling regressions use the variable's own lags and monthly data for the post-war period to construct $E_T[x_{t+1}]$ for the three month period covered by $t + 1$ given only information available at the start of period $t - 1$.

(uninstrumented) real interest rate series. The seasonal variation is removed in all analyses.

We are forced to drop any observations missing the required consumption data, variables that shift marginal utility, or instruments used to proxy expectations. Rural households are dropped, as are households living in student housing. We drop observations in which the gender of the head of the household remains the same and the head or spouse changes age by more than a year or less than no years in between interviews. Only households for which the age of the head and spouse are less than 85 and greater than 21 are included. We drop observations reporting family size changes greater than 3. The final data set contains 148,117 observations on consumption growth.

One important advantage of the data and one important shortcoming affect the decomposition.

The advantage is that by using household data, we remove composition biases that are present in aggregate consumption data that lead to predictable movements in consumption growth.¹⁸ We analyze average consumption growth across a subset of households,

$$\Delta \ln C_{t+1}^{Ag} = \frac{1}{I(t)} \sum_{i \in I_t} \Delta \ln C_{i,t+1},$$

while aggregate consumption data is,

$$\Delta \ln \bar{C}_{t+1} = \ln \left(\frac{1}{H(t)} \sum_{i \in H_t} C_{i,t+1} \right) - \ln \left(\frac{1}{H(t)} \sum_{i \in H_t} C_{i,t} \right).$$

Thus we omit households who are born or die, and do not incorrectly assign predictable movements in consumption due to births and deaths to precautionary saving or other sources, as could be mistakenly done using aggregate data.¹⁹ Moreover, aggregate consumption data also contain many expenditures which are not truly nondurable consumption. We select items for consumption that better reflect nondurable and services consumption than National Income data.

¹⁸Attanasio and Weber (1993) discuss the importance of this issue.

¹⁹However, there is little evidence that these population dynamics are important for fluctuations in consumption growth (Parker (1999)).

The shortcoming is that the consumption expenditures of each household are measures with a significant amount of error. The standard deviation of the growth of household expenditures on nondurable consumption is 36%. Most researchers believe that the majority of this movement does not represent shifts in preferences or actual risk to marginal utility, but is due to the difficulty in collecting data and the difference between the concept of consumption and the observed expenditures. That is, consumption can vary due to omitted or over-reported expenses. To partly deal with this, we drop the bottom one percent of households in the distribution of real consumption per effective householder.²⁰ In addition, we drop the top and bottom 2.5 percent of households in the distribution of consumption growth, on the grounds that these observations are more likely to represent mismeasurement than actual movements in marginal utility.²¹

Does the mismeasurement of consumption bias our decomposition? Measurement error in consumption does bias upward two statistics: the mean of consumption growth due to incomplete markets and the variance of the movements in consumption due to new information. But measurement error biases only these two statistics. Under plausible assumptions, measurement error does not alter the contribution of precautionary saving, preferences, or the real interest rate to changes in the growth rate of consumption. For this reason, this paper focuses on changes in consumption growth and not mean consumption growth.

To make these claims concrete, let the measurement be classical, additive in logs:

$$\left(\frac{C_{i,t+1}}{C_{i,t}}\right)^{-\frac{1}{\sigma}} = \left(\frac{C_{i,t+1}^*}{C_{i,t}^*}\right)^{-\frac{1}{\sigma}} \left(\frac{\mu_{i,t+1}}{\mu_{i,t}}\right)^{\frac{-1}{\sigma}}$$

where $C_{i,t}^*$ is the true level of consumption, $C_{i,t}$ is the observed measure, and $\mu_{i,t+1}$ is the measurement error. Further, let the measurement error be independent of conditioning

²⁰For this exercise only, effective household size is defined as one, plus one if there is a spouse present, plus 0.4 times each additional family member. The first percentile in the distribution is predicted by a constant and a time trend and observations below this fitted value are dropped. The fitted values range from 190 to 230 real 1982-1984 dollars per effective household member per three months.

²¹These are reasonably standard procedures for dealing with the extreme outliers in the data. See Vissing-Jorgensen (forthcoming) and Brav, Constantinides, and Geczy (forthcoming) for example.

information, and define the nuisance parameter ϕ as

$$\phi \equiv E_t [\Delta \ln \mu_{i,t+1}] .$$

To see that measurement error biases the mean and only the mean of consumption growth due to precautionary saving, plug observed consumption growth into the theoretical measure from equation (2.7)

$$\begin{aligned} \Delta \ln C_{i,t+1}^{IM} &= -\sigma E_t \left[\ln \left(1 + e^{\mathbf{X}_{i,t+1}\delta} \left(\frac{C_{i,t+1}^*}{C_{i,t}^*} \right)^{-\frac{1}{\sigma}} \left(\frac{\mu_{i,t+1}}{\mu_{i,t}} \right)^{\frac{-1}{\sigma}} R_{t,t+1} - 1 \right) \right] \\ &= -\sigma E_t \left[\mathbf{X}_{i,t+1}\delta - \frac{1}{\sigma} \ln \left(\frac{C_{i,t+1}^*}{C_{i,t}^*} \right) - \frac{1}{\sigma} \ln \left(\frac{\mu_{i,t+1}}{\mu_{i,t}} \right) + \ln R_{t,t+1} \right] \\ &= E_t \left[\ln \left(\frac{\mu_{i,t+1}}{\mu_{i,t}} \right) \right] - \sigma E_t \left[\ln \left(1 + e^{\mathbf{X}_{i,t+1}\delta} \left(\frac{C_{i,t+1}^*}{C_{i,t}^*} \right)^{-\frac{1}{\sigma}} R_{t,t+1} - 1 \right) \right] \\ &= \phi - \sigma E_t [\ln (1 + \varepsilon_{i,t+1})] \\ &= \phi + \Delta \ln C_{i,t+1}^{IM*} . \end{aligned}$$

where $\Delta \ln C_{i,t+1}^{IM*}$ denotes the true (correctly measured) consumption growth due to incomplete markets..

To see that the volatility of consumption growth and innovations to consumption are overstated,

$$\begin{aligned} \eta_{t+1}^C &= \frac{1}{H(t)} \sum_{i \in H_t} \Delta \ln C_{i,t+1} - E_t \left[\frac{1}{H(t)} \sum_{i \in H_t} \Delta \ln C_{i,t+1} \right] \\ &= \frac{1}{H(t)} \sum_{i \in H_t} \Delta \ln C_{i,t+1}^* + \frac{1}{H(t)} \sum_{i \in H_t} \Delta \ln \mu_{i,t+1} - E_t \left[\frac{1}{H(t)} \sum_{i \in H_t} \Delta \ln C_{i,t+1}^* \right] - \phi \\ &= \eta_{t+1}^{C*} + \tilde{\mu}_{t+1} \end{aligned}$$

where η_{t+1}^{C*} denotes the correctly measured variance of consumption growth and $\tilde{\mu}_{t+1} \equiv \frac{1}{H(t)} \sum_{i \in H_t} \Delta \ln \mu_{i,t+1} - \phi$ which is an MA(1) variable independent of all lagged instruments. Both arguments also hold in our robust measures. In sum, the measures overstate mean consumption growth and mean consumption growth due to precautionary saving by a factor ϕ , and overstate the volatility of consumption growth and innovations to consumption by $Var_t(\tilde{\mu}_{t+1})$.

5. The proximate causes of consumption growth for unconstrained households

This section studies the consumption growth of households that are not liquidity constrained, as determined by the criteria just described. The contribution of precautionary saving to consumption volatility is statistically significant, but the decomposition leaves a range of uncertainty as to the economic importance of precautionary saving. Our measure of expected consumption growth due to precautionary saving covaries positively with expected changes in government spending. Further, there is some evidence that precautionary saving is countercyclical, consistent with consumption risk being high in recessions.

In order to give a robust picture of the role of precautionary saving in consumption growth, we report results for several different utility functions. First, σ and δ are estimated by GMM using grouped panel-data and the time-series moments implied by the consumption Euler equation. Variables that shift the preference for consumption ($\mathbf{X}_{i,t+1}$) include indicator variables for the month of the year to capture seasonal variations in demand. Since reported consumption declines slightly with each interview that a household participates in, preferences are allowed to vary by interview. We allow discount rates to differ by five-year birth cohorts. The number of family members and the number of children are both included because they shift the marginal utility of a given amount of expenditure. Finally, to control for the possibility that labor supply is non-separable from the marginal utility of consumption, the number of hours that the woman head of household works is included as a preference shifter. The intertemporal elasticity of substitution is estimated relatively precisely as 0.652 with standard error of 0.14, although there is also significant specification uncertainty. Appendix C describes the GMM estimator and estimates in greater detail. Having estimated the utility function, it is possible that misspecification of preferences is minimal, so we present analyses of consumption growth with and without the correction for misspecification.²²

²²Overidentification tests do not reject the model. However, misspecification could potentially remain

The second approach is calibration. We analyze three different intertemporal elasticities of substitution: 0.3, 0.65, and 1 (log utility). In each case we include no preference shifters, so that it is reasonable to believe that the misspecification adjustment is capturing preference variation. Therefore, as noted in section 3, the consumption growth due to misspecification is treated as consumption growth due to preference changes. For the decompositions based on estimation, the variation in consumption growth due to misspecification is omitted. In all cases, seasonal variation is removed.

Table 1 shows, for our sample of unconstrained households, the share of variation in consumption growth due to innovations to consumption growth and the share of predictable variation due to preferences, the real interest rate, and precautionary saving. Before turning to predictable consumption growth, note that 90 percent or more of the variance of average consumption growth in the CEX is due to news or measurement error, evidence that we are not overfitting in predicting consumption growth. This large number stems from the fact that aggregate consumption growth is difficult to forecast with few instruments, and from the fact that measurement error in the CEX leads to unpredictable changes in average consumption growth.

Table 1 demonstrates that precautionary saving causes a statistically significant share of the volatility of expected consumption growth, and that there is a wide range of uncertainty concerning the economic importance of precautionary saving. Innovations to consumption growth are, by construction, orthogonal to the other components of consumption growth; however, the predictable series are not mutually uncorrelated. Therefore the share of variation due to any series depends on the ordering of the series in the variance decomposition. Table 1 presents two alternative orderings for each set of parameters. After the correction for misspecification, the GMM estimates of the percent of variation in predictable consumption growth due to precautionary saving range

a concern as the estimates from this model are used to construct our measure of precautionary savings at an individual household level, rather than at the synthetic cohort level at which estimates of preferences were obtained. That the Euler equation holds for a synthetic cohort is necessary but not sufficient for it to be satisfied by an individual household.

from 2.5 to 95 percent. For the results based on calibration – which treat misspecification as preference variation – estimates of the economic importance of precautionary saving range from 1.4 percent to 37 percent. In contrast, the first row of results, which omit the correction for possible misspecification suggest that precautionary saving is far more economically important, directly causing 70 to 80 percent of variation in expected consumption growth. But as emphasized in section 3, the presence of model misspecification inflates this measure by construction. The results with the correction for misspecification show that the importance of precautionary saving is significantly less, or at least more uncertain, when the predictable component of the error term is removed.²³

Given this range of uncertainty, is there a statistically significant contribution of precautionary saving to the volatility of consumption growth? The last column of Table 1 reports the probability that there is no variation due to precautionary saving; that is the p-value for an F-test that the true coefficients in the regression constructing the expectation in equation (3.8) are all zero. Fluctuations in consumption due to precautionary saving are statistically significant in all cases except the raw residuals from GMM estimation.

The reason for the wide range of uncertainty over the economic importance of precautionary saving is itself informative. Expected consumption growth due to precautionary saving is negatively correlated with that due to preference variation and positively correlated with that due to movements in the real interest rate. These correlations are large, and this is what drives the uncertainty over relative importance. The positive covariance between consumption growth due to the real interest rate and that due to

²³Also, comparing the second and fourth sets of results shows the effect of assuming that expected consumption growth due to misspecification is actually due to preference variation. The GMM estimates assume that estimated preferences capture all true preference variation and omit the misspecification series from the decomposition. The calibration assumes that misspecification captures omitted preference variation. When we include the misspecification as preference variation, the correlation of precautionary saving and preferences decreases and the role of precautionary saving is better pinned down.

precautionary saving implies that omission of the precautionary term from a regression of consumption growth on the expected real interest rate would increase rather than decrease the estimated elasticity of intertemporal substitution. We also find that consumption growth due to shifts in the preference for consumption covaries negatively with that due to the real interest rate. Thus low observed intertemporal substitution in response to the real interest rate, as observed by Hall (1988) for example, is largely due to omitted preference shifters or misspecification of the consumption optimization problem more generally.²⁴ Put differently, since precautionary saving leads to higher consumption growth when the interest rate is higher, precautionary saving cannot be the cause of low observed intertemporal substitution. This supports models in which expected changes in consumption growth are caused by nonseparabilities of nondurable consumption and home production, consumption of durable goods, or leisure.²⁵

Table 1 also shows that the economic importance of precautionary saving is larger the larger the assumed level of household risk aversion and prudence. Lower values of σ , the elasticity of intertemporal substitution, imply a more important contribution of precautionary saving to predictable consumption growth. This occurs because σ also governs relative risk aversion ($\frac{1}{\sigma}$) and relative prudence ($1 + \frac{1}{\sigma}$). A low elasticity of intertemporal substitution implies high risk aversion and high prudence.

Figure 1 displays the three time series of expected consumption growth for both the GMM estimates and the calibration with $\sigma = 0.65$. The data are quarterly averages converted to annual rates, and apart from the series for preferences, are visually similar across panels. The primary difference between the figures is due to the misspecification adjustment, which is included in the preference series in Figure 1b and omitted in Figure 1a. There is a significant amount of predictable variation in consumption growth

²⁴This supports one of the main findings of Attanasio and Weber (1995) and Attanasio and Weber (1993). These papers also argue that correct aggregation is an important part of the difference between linear models in micro and macro data.

²⁵See Baxter and Jermann (1999), Ogaki and Reinhart (1998), and Basu and Kimball (2000) respectively.

“explained” by the nonstructural construction of $E_{T,\lambda=0} [\varepsilon_{i,t+1} | \mathbf{Z}_{i,t}]$, but unexplained by the more structural utility function used in the GMM estimation.

In both panels, we see some evidence that precautionary saving has become less important over the sample, contributing less to consumption growth. This is consistent with the increase in wealth and economic boom of the 1982 – 1997 period and suggests that decreased risk has a role in the consumption boom (see Parker (1999)). We also see that risk contributes substantially to consumption growth after the large 1982 recession, consistent with consumption dropping in recessions due to increased consumption risk in the future. There is no similarly dramatic pattern observed around the 1991 recession, though consumption growth due to precautionary saving rises before the recession, and falls during and after it, with somewhat more accentuated movements in Figure 1b.

These are visual impressions. To be more formal, Table 2 presents the results of investigating precautionary saving by regressing expected consumption growth due to precautionary saving onto: the contemporaneous expected change in the unemployment rate, the contemporaneous expected growth in government spending, and lagged innovations to the federal funds rate.

The first column shows that precautionary saving leads to countercyclical expected consumption growth, although this is statistically insignificant. The point estimates imply that an expected one percent increase in the prime-age male unemployment rate is associated with 0.1 to 0.2 percent faster consumption growth due to consumption risk. The sign of the effect supports our theoretical understanding of precautionary saving – when the unemployment rate is expected to increase there is greater risk and therefore precautionary saving should be high and expected consumption growth should be greater.

The second and third columns of Table 2 present some reduced form evidence on precautionary saving and macroeconomic policy. Theoretically, there are two channels through which economists have considered policy changing consumption risk and precautionary saving. First, expansionary policy can lead to less consumption risk, so that consumption increases when the policy is announced and then is expected to grow more

slowly over time. This effect is not consistent with the observed impulse responses of consumption to monetary policy shocks, for example, which show faster consumption growth for a while after a reduction in the real interest rate (Bernanke and Gertler (1995)). If instead policy is going to cause precautionary saving to increase consumption growth, it must be that expansionary policy leads to increased consumption risk. For example, if policy improves expectations of future income, households increase consumption on announcement reducing their stocks of liquid wealth, and this leads to higher consumption risk and faster consumption growth.²⁶ To evaluate these theories, we ask how consumption growth due to precautionary saving responds to predictable movements in government spending and lagged monetary policy shocks.

An expected one percent higher growth rate of government spending is associated with a one-quarter to three-quarter of a percent faster consumption growth due to precautionary saving. This is consistent with pre-announced increases in government spending leading to faster consumption growth due to precautionary saving. As to monetary policy, we find no statistically significant relationship between consumption growth due to precautionary saving and 12 lags of monthly monetary policy shocks as constructed by Bernanke and Mihov (1998), although the sign of the effect is positive, consistent with the impact of a monetary policy shock on total expected consumption growth.

In sum, we find a statistically significant role for precautionary saving, some evidence that it has declined and is countercyclical, and conclude that its economic impact is both similar to that of the real interest rate and similarly uncertain in magnitude. We turn now from the sample of unconstrained households, in which we are studying the effects of precautionary saving, to the entire sample, in which our measure of consumption growth due to incomplete markets may include the impact of liquidity constraints.

²⁶In models such as Carroll (1997), higher expected income growth leads to higher expected consumption growth through precautionary saving.

6. The proximate causes of consumption growth

This section presents the results from a similar set of exercises as the previous section, but performed on the data for all households. We find that expected consumption growth due to incomplete markets is directly responsible for a slightly larger share of consumption growth in the entire population than for only unconstrained households. Indeed, there is evidence that incomplete markets were a more important determinant of consumption dynamics in the 1991 recession for all households than for unconstrained households alone. Furthermore, consumption growth due to incomplete markets is more countercyclical for all households than for unconstrained households alone, and has no significant comovement with our policy variables. These results must be interpreted with care. There is a large negative correlation between expected consumption growth due to incomplete markets and that due to misspecification. The misspecification adjustment induces significant movement in the incomplete markets series. The economic interpretation of this finding is that, for all households, preference shifters (or misspecification) interact in important ways with incomplete markets. We give a concrete example below.

Table 3 presents the variance decomposition of consumption growth for all households. Similar to the findings in the smaller, unconstrained sample, expected consumption growth due to incomplete markets is highly statistically significant. Slightly more than 90 percent of consumption growth is unpredictable. Unlike for unconstrained households alone, the corrected estimates from the GMM estimation (second set of results) imply that precautionary saving and liquidity constraints together are directly responsible for a reasonably precise 52 to 69 percent of the predictable consumption growth. If however, one views all the predictable variation in the constructed expectation error as due to preferences then the measured effect of incomplete markets is smaller. The calibration results show a lower contribution of incomplete markets, and a greater level of uncertainty as to its role in fluctuations. As is the case for the unconstrained sample, expected consumption growth due to precautionary saving negatively

covaries with that due to preference variation, and this covariance rises when “misspecification” is included as preferences changes. Across specifications, the percent of variation in predictable consumption growth due to incomplete markets ranges from 1.5 to 69 percent.

Figure 2 displays the three time series of expected consumption growth for the GMM estimates (Panel *A*) and the calibration with $\sigma = 0.65$ (Panel *B*). Again, incomplete markets and the real interest rate each lead to similar fluctuations in consumption growth in each panel. Adding the misspecification correction to preferences in Panel *B* leads to volatile consumption growth due to changing preferences that is more negatively correlated with consumption growth due to incomplete markets. Why is this? One answer is simply that $E_{T,\lambda=0}[\varepsilon_{i,t+1}|\mathbf{Z}_{i,t}]$ is estimated only using unconstrained households but used in the construction of $\Delta \ln \widehat{C}_{t+1}^{IM}$ for all households (equation (3.8)). Statistically, the predicted values out of sample create this negative correlation. But a second answer is that there are good economic reasons for these predictions to be poor out of sample: constrained households are different. The presence of a binding liquidity constraint leads a household to adjust margins besides consumption and this can lead to exactly this type of negative covariance. For example, consider a young constrained household that increases labor supply in response to the constraint. As the constraint relaxes, labor supply declines predictably, leading to slower desired consumption growth due to preference shifts, while consumption growth rises predictably due to the constraint relaxing. The large negative correlation suggests this is what is happening: preference shifters and liquidity constraints lead to offsetting movements in expected consumption growth.²⁷ Given the magnitude of the negative correlation, we suspect the negative correlation is not purely due to economic behavior, but is also partly driven by imperfect prediction.

That said, like Figure 1, Figure 2 shows some evidence that precautionary saving and liquidity constraints have become less important over the time period covered by the data, contributing less to consumption growth. Comparing Figures 1 and 2, this

²⁷See also Appendix B.

decline appears slightly larger for all households than for only liquid households. Also, liquidity constraints or precautionary saving of low wealth households seem to have played a role in maintaining consumption growth in the 1991 recession. Again, these are visual impressions. Table 4 reports the results of regressing consumption growth due to incomplete markets onto the same set of variables as Table 2.

First, incomplete markets lead to countercyclical expected consumption growth. The point estimates imply that an expected one percent increase in the prime-age male unemployment rate is associated with just over a one percent higher rate of expected consumption growth rate. This effect is larger than that estimated for unconstrained households alone, implying that precautionary saving is more important for low wealth households. Larger increases in consumption risk for low wealth households could be due to the fact that unemployment falls more heavily on lower income households or that credit constraints are tighter when unemployment is expected to rise, so that borrowing to smooth consumption becomes more difficult. Either way, the estimates imply that incomplete markets amplify business cycle movements in consumption. Consumption risk increases and/or borrowing constraints tighten upon news of a coming recession and therefore incomplete markets for transferring consumption across time and states amplify the decline in consumption that occurs on this news.

The remaining columns of Table 4 show that both our measures of policy have insignificant correlations with expected consumption growth due to incomplete markets. Unlike for unconstrained households, the impact is not consistently of one sign. There is no evidence that the combination of liquidity constraints and precautionary saving amplify or damp the impact of government spending or monetary policy on total consumption growth.

7. Conclusion

This paper proposes and implements a decomposition of consumption growth into its proximate causes. We show how the utility function and consumption data can be used

to source movements in consumption growth and demonstrate how to measure the role of precautionary saving in a manner robust to some misspecification. Using estimates and calibrations for a standard utility function, we implement this decomposition using the CEX survey data on household expenditures. In sum, we find that incomplete markets for trading future consumption cause statistically important movements in expected consumption growth and that the economic importance of precautionary saving rivals that of the interest rate.

For households that are unlikely to be liquidity constrained, expected consumption growth is sourced to movements in the real interest rate, changes in the preference for consumption, and precautionary saving. We find that the impact of precautionary saving on variations in consumption growth is statistically significant. While the magnitude is rather uncertain, the variance of predictable movements in consumption that are due to movements in precautionary saving is if anything slightly less than the variance due to movements in the real interest rate. There is some evidence, although imprecise, that precautionary saving is countercyclical, leading to higher consumption growth when unemployment rates are increasing. And there is some evidence that high expected growth in government spending is associated with greater consumption growth due to precautionary saving.

For average consumption growth, incomplete markets are more important for fluctuations in expected consumption growth. Precautionary saving and liquidity constraints are highly statistically significant, and cause movements in consumption that are positively correlated with movements due to the real interest rate and negatively correlated with movements due to preferences or “misspecification.” Consumption growth due to incomplete markets is countercyclical, and so amplifies recessions, but we find no evidence that it is correlated with past policy.

We suspect that measurement error in consumption growth limits our ability to infer the characteristics of consumption growth. Better measurement (or data in which with longer growth rates can be studied) would allow a tighter decomposition and a more accurate mapping from primitive shocks to subsequent movements in precautionary

saving. Better measurement may also allow analysis of trend growth rates. Of specific interest is the variation across countries in consumption growth rates, which are largely unexplained by variations in real interest rates. Our decomposition suggests a way to use microeconomic data to measure the differences in consumption risk across countries into implied differences in expected growth rates of consumption.

Appendixes

A. The corrections to precautionary saving

This appendix demonstrates that the specification of the corrections described in section 3 are not interchangeable.

Consider first applying the correction for misspecification in the manner of the second correction. Assuming parameter estimates are correct apart from the misspecification term θ , the household-level measure of precautionary saving would be

$$-\sigma \frac{1}{I(t)} \sum_{i \in I_t} E_T [\ln(1 + \varepsilon_{i,t+1}) - \varepsilon_{i,t+1} | \mathbf{Z}_t].$$

Since

$$\varepsilon_{i,t+1} = \frac{1 + \epsilon_{i,t+1}}{\theta_{i,t}} - 1$$

the expectation of this measure is equal to

$$\begin{aligned} & -\sigma E \left[\ln \left(\frac{1 + \epsilon_{i,t+1}}{\theta_{i,t}} \right) - \frac{1 + \epsilon_{i,t+1}}{\theta_{i,t}} + 1 \right] \\ &= -\sigma E [\ln(1 + \epsilon_{i,t+1})] - \sigma \left(-\ln \theta_{i,t} + \frac{1}{\theta_{i,t}} + 1 \right). \end{aligned}$$

Expanding the second logarithm around $\theta_{i,t} = 1$, one can see that this measure is correct to the first order but not higher. Under this alternative, $\theta_{i,t}$ raised to powers would show up in our measure of precautionary saving.

Turning to the correction for small-sample bias, suppose that we were to subtract $\hat{e}_{t+1} \equiv E_T \left[\frac{1}{I(t)} \sum_{i \in I_t} \tilde{e}_{i,t+1} | \mathbf{Z}_t \right]$ from the residual in the manner of the adjustment for misspecification, so that the measure of consumption growth due to precautionary saving is

$$-\sigma E_T \left[\frac{1}{I(t)} \sum_{i \in I_t} \ln(1 + \tilde{e}_{i,t+1} - \hat{e}_{i,t+1}) | \mathbf{Z}_t \right].$$

Taking the expectation of a second-order expansion of the expectation around $\tilde{e}_{i,t+1} - \hat{e}_{i,t+1} =$

0 yields

$$\begin{aligned}
&= -E_t \left[\frac{1}{I(t)} \sum_{i \in I_t} \tilde{e}_{i,t+1} - \hat{e}_{t+1} - \frac{(\tilde{e}_{i,t+1} - \hat{e}_{t+1})^2}{2} \right] \\
&= -E_t \left[\frac{1}{I(t)} \sum_{i \in I_t} \tilde{e}_{i,t+1} \right] + \frac{1}{I(t)} \sum_{i \in I_t} E_t [\hat{e}_{t+1}] + E_t \left[\frac{1}{I(t)} \sum_{i \in I_t} \frac{\tilde{e}_{i,t+1}^2 - 2\hat{e}_{t+1}\tilde{e}_{i,t+1} + \hat{e}_{t+1}^2}{2} \right] \\
&= \frac{1}{2} E_t \left[\frac{1}{I(t)} \sum_{i \in I_t} \tilde{e}_{i,t+1}^2 \right] - \frac{1}{2} \frac{1}{I(t)} \sum_{i \in I_t} E_t [\hat{e}_{t+1}^2] \\
&= \frac{1}{2} \overline{Var}_t [\tilde{e}_{i,t+1}] - \frac{1}{2} \frac{1}{I(t)} \sum_{i \in I_t} E_t [\hat{e}_{t+1}^2]
\end{aligned}$$

which is the object of interest less a term that is always positive and possibly quite large. Taking the true conditional expectation of this measure yields the quantity of interest less a positive term.

B. Liquidity constraints and the decomposition

This appendix shows that the incomplete markets series contains movements in consumption growth due to liquidity constraints and how the presence of binding liquidity constraints alters the interpretation of this series.

Suppose that for a constrained household, omitted preference shifters are predicted by variables in $Z_{i,t}$ with the same coefficients as for an unconstrained household. That is, assume that the true model is given by equation (3.1) and that the function of $Z_{i,t}$, $E_{\lambda=0}[\varepsilon_{i,t+1}|Z_{i,t}] = \frac{1}{\theta_{i,t}} - 1$, independent of λ . For constrained households the (theoretical) robust measure of consumption growth due to incomplete markets is

$$\begin{aligned}
\Delta \ln C_{i,t+1}^{IM} &= -\sigma E_t \left[\ln \left(1 + \frac{\varepsilon_{i,t+1} - E_{\lambda=0}[\varepsilon_{i,t+1}|Z_{i,t}]}{1 + E_{\lambda=0}[\varepsilon_{i,t+1}|Z_{i,t}]} \right) \right] \\
&= -\sigma E_t \left[\ln \left(1 + \frac{\frac{1 + \varepsilon_{i,t+1} - \lambda_{i,t}}{\theta_{i,t}} - \frac{1}{\theta_{i,t}}}{\frac{1}{\theta_{i,t}}} \right) \right] \\
&= -\sigma E_t [\ln (1 + \varepsilon_{i,t+1} - \lambda_{i,t})] .
\end{aligned}$$

which is the quantity of interest.

Is it plausible that $E_{T,\lambda=0}[\varepsilon_{i,t+1}|\cdot]$ is not a function of λ ? Suppose that a binding liquidity constraint causes an increase in labor supply. Suppose further that hours worked is omitted

from $X_{i,t}$ and that hours worked do not otherwise vary with $Z_{i,t}$. Then, this condition fails: the change in the expected growth rate of consumption due to the change in hours is allocated to liquidity constraints. To see this more formally, let $\theta^{-1}(\lambda, Z_{i,t})$ be the predicted value of $\theta_{i,t}^{-1}$ for a household as a function of liquidity constraint and other characteristics.

$$\begin{aligned}\Delta \ln C_{i,t+1}^{IM} &= -\sigma E_t \left[\ln \left(1 + \frac{\varepsilon_{i,t+1} - E_{t,\lambda=0} [\varepsilon_{i,t+1}]}{1 + E_{t,\lambda=0} [\varepsilon_{i,t+1}]} \right) \right] \\ &= -\sigma E_t \left[\ln \left(1 + \frac{\frac{1+\varepsilon_{i,t+1}-\lambda_{i,t}}{\theta(\lambda, Z_{i,t})} - \frac{1}{\theta(0, Z_{i,t})}}{\frac{1}{\theta(0, Z_{i,t})}} \right) \right] \\ &= -\sigma E_t \left[\ln \left(\frac{\theta(0, Z_{i,t})}{\theta(\lambda, Z_{i,t})} (1 + \varepsilon_{i,t+1} - \lambda_{i,t}) \right) \right]\end{aligned}$$

The effect of a binding liquidity constraint on preferences appears in the incomplete markets series since movements in θ caused by λ (unpredictable by $Z_{i,t}$ and unmodelled in $X_{i,t}$) move $\frac{\theta(0, Z_{i,t})}{\theta(\lambda, Z_{i,t})}$ away from unity.

C. GMM estimation of the parameters of the utility function

This appendix describes the estimation of the parameters of the utility function by GMM. The estimator uses grouped CEX data, as in Attanasio and Weber (1995), and the nonlinear theoretical moments implied by equation (2.3), as pioneered using aggregate data by Hansen and Singleton (1982).

C.1. The moments

Given J instruments, the Euler equation implies J moment conditions for each household:

$$E \left[\mathbf{h} \left(\boldsymbol{\delta}, \sigma, \mathbf{X}_{i,t+1}, \frac{C_{i,t+1}}{C_{i,t}}, \mathbf{Z}_{i,t} \right) \right] = 0 \quad (\text{C.1})$$

where $E[\cdot]$ is the mathematical (unconditional) expectations operator and where

$$\mathbf{h} \left(\boldsymbol{\delta}, \sigma, \mathbf{X}_{i,t+1}, \frac{C_{i,t+1}}{C_{i,t}}, \mathbf{Z}_{i,t} \right) \equiv \left[\exp(\mathbf{X}_{i,t+1} \boldsymbol{\delta}) \left(\frac{C_{i,t+1}}{C_{i,t}} \right)^{-\frac{1}{\sigma}} R_{i,t+1}^f - 1 \right] \mathbf{Z}_{i,t}. \quad (\text{C.2})$$

Given ideal data, we would use the estimator

$$(\hat{\boldsymbol{\delta}}, \hat{\sigma}) \equiv \arg \min_{\{\boldsymbol{\delta}, \sigma\}} \mathbf{g}(\boldsymbol{\delta}, \sigma; \mathbf{Y})' \boldsymbol{\Omega} \mathbf{g}(\boldsymbol{\delta}, \sigma; \mathbf{Y}) \quad (\text{C.3})$$

where Y denotes the stacked matrix of all data, Ω is a JI by JI weighting matrix and $g(\delta, \sigma; Y)$ is the JI by 1 column vector of the empirical counterparts to each theoretical moment in equation (C.1). $g(\delta, \sigma; Y)$ consists of I stacked $J \times 1$ column vectors of

$$\mathbf{g}_i(\delta, \sigma, \mathbf{Y}) = \frac{1}{T_i} \sum_{t=1}^{T_i} \mathbf{h} \left(\delta, \sigma, \mathbf{X}_{i,t+1}, \frac{C_{i,t+1}}{C_{i,t}}, \mathbf{Z}_{i,t} \right) \quad (\text{C.4})$$

where T_i is the number of observations available for household i .

We do not directly use equations (C.3) and (C.4) for three reasons.

First, the consistency of this estimator requires that the empirical moments converge to the theoretical moments, which, as is clear from (C.4), happens only as the time dimension on each household becomes large. In the CEX, T_i is fixed at a maximum value of three. Since T is so small, the properties of the estimator proposed in (C.4) are not reliably approximated by the asymptotic distribution. The CEX data does however contain short sequences of overlapping consumption growth of different households covering a long time period. We therefore use a weaker restriction and estimate our model from the restriction that average of the expectation errors across all households average to zero over time,²⁸ formally:

$$E \left[\frac{1}{I(t)} \sum_{i \in I_t} \mathbf{h} \left(\delta, \sigma, \mathbf{X}_{i,t+1}, \frac{C_{i,t+1}}{C_{i,t}}, \mathbf{Z}_{i,t} \right) \right] = 0. \quad (\text{C.5})$$

Condition (C.5) is implied by (C.1) but does not imply (C.1).

Second, household-level characteristics and particularly consumption growth are measured with error in the CEX. When consumption growth is measured with substantial error, and when the model is overidentified so that the moments are not all set to zero, the intertemporal elasticity of substitution, σ , is biased upwards by measurement error in consumption growth. This follows because measurement error increases the variance of the constructed innovation to marginal utility. Since, in finite samples, there is some covariance between the instruments and this residual, the estimation procedure raises σ to lower the variance of the innovation and thus the covariance with the moments. We address this issue and the third issue by partially aggregating the data.

²⁸A better assumption would be that the expectation errors of *similar* households average to zero, but non-linear estimation does not converge with more moments.

Before proceeding however, it is worth noting that the estimate is consistent. Under the assumptions on measurement error used in the main body of the paper, substituting the observed consumption series into our theoretical moment condition yields

$$\begin{aligned}
& E \left[\left(e^{\mathbf{X}_{i,t+1}\delta} \left(\frac{C_{n,t+1}^*}{C_{n,t}^*} \right)^{-\frac{1}{\sigma}} \left(\frac{\mu_{i,t+1}}{\mu_{i,t}} \right)^{-\frac{1}{\sigma}} R_{i,t+1}^f - 1 \right) \mathbf{Z}_{i,t} \right] \\
&= E \left[E \left[\left(\frac{\mu_{i,t+1}}{\mu_{i,t}} \right)^{-\frac{1}{\sigma}} | \mathbf{Z}_{i,t} \right] E \left[e^{\mathbf{X}_{i,t+1}\delta} \left(\frac{C_{i,t+1}^*}{C_{i,t}^*} \right)^{-\frac{1}{\sigma}} R_{i,t+1}^f | \mathbf{Z}_{i,t} \right] \mathbf{Z}_{i,t} \right] - E [\mathbf{Z}_{i,t}] \\
&= E \left[e^\varphi E \left[e^{\mathbf{X}_{i,t+1}\delta} \left(\frac{C_{i,t+1}^*}{C_{i,t}^*} \right)^{-\frac{1}{\sigma}} R_{i,t+1}^f | \mathbf{Z}_{i,t} \right] \mathbf{Z}_{i,t} \right] - E [\mathbf{Z}_{i,t}] \\
&= E \left[\left(e^{\varphi + \mathbf{X}_{i,t+1}\delta} \left(\frac{C_{i,t+1}^*}{C_{i,t}^*} \right)^{-\frac{1}{\sigma}} R_{i,t+1}^f - 1 \right) \mathbf{Z}_{i,t} \right].
\end{aligned}$$

If the true parameters are $\delta^0 = (\delta_1^0, \delta_2^0, \dots, \delta_K^0)'$ and σ_0 , then the probability limits of the parameters that set the moment to zero when the consumption data are mismeasured are $(\delta_1^0 - \varphi, \delta_2^0, \dots, \delta_K^0)'$ and σ_0 . GMM does not estimate consistently δ_1 – the mean level of impatience across households – but it does estimate consistently the remaining structural parameters of interest.

Third, the theoretical moments are too many to feasibly use in estimation. There are J (the number of predetermined variables) times I (the number of households) moment conditions and the predetermined variables may include individual information, aggregate variables, and characteristics of the distribution of individual variables. In an economy in which markets are incomplete, the evolution of the aggregate economy generally depends upon the distribution of wealth and income in the economy. In the extreme case, it may be that the information set that each household uses to form expectations contains its own household-level characteristics and the household-level characteristics of all other households in the economy.

We address these issues by partially aggregating the data. We assume that the economy consists of N groups of households that are able to pool their resources within each group to insure their risks completely from all shocks, except those that affect the group average. This assumption is made because it renders the estimator feasible: groups are observed over long

time periods and averaging reduces the impact of measurement error.²⁹ We construct the expectation error by averaging $X_{i,t+1}$ and $\ln \frac{C_{i,t+1}}{C_{i,t}}$ within each group (denoted by n), thus mitigating the impact of measurement error. As to the subset of instruments, we assume that $Z_{n,t}$ contains only information on group n 's average characteristics and aggregate information.

Our final estimator, based on the theoretical moment given in equation (C.5) but with groups, n , instead of households, is (C.3) where $\hat{\Omega}^{-1}$ is $J \times J$ and $g(\delta, \sigma; Y)$ is the $J \times 1$ column vector

$$g(\delta, \sigma, Y) = \frac{1}{T} \sum_{t=1}^T \left[\frac{1}{N(t)} \sum_{n \in N_t} h(\delta, \sigma, X_{n,t+1}, C_{n,t+1}^{GR}, Z_{n,t}) \right] \quad (C.6)$$

where $Z_{n,t} \equiv \frac{1}{I(n,t)} \sum_{i \in I_{n,t}} Z_{i,t}$, $X_{n,t+1} \equiv \frac{1}{I(n,t)} \sum_{i \in I_{n,t}} X_{i,t+1}$, and $C_{n,t+1}^{GR} \equiv \prod_{i \in I_{n,t}} \left(\frac{C_{i,t+1}}{C_{i,t}} \right)^{\frac{1}{I(n,t)}}$ where $I_{n,t}$ defines the set of $I(n,t)$ households in group n observed in t and $t+1$. We refer to $\frac{1}{T} \sum_{t=1}^T N(t)$ as the number of groups used in the construction of an estimator.

C.2. The weighting matrix and inference

We use the optimal weighting matrix so that the asymptotic distribution of estimator defined by equations (C.3) and (C.6) is normal with variance-covariance matrix given by

$$(G' \Omega G)^{-1}$$

where

$$G = \left(\frac{dg(\delta, \sigma, Y)}{d\delta}, \frac{dg(\delta, \sigma, Y)}{d\sigma} \right)$$

and

$$\Omega = E[g(\beta, \sigma; Y)g(\beta, \sigma; Y)']. \quad (C.7)$$

We construct an empirical counterpart to Ω , denoted $\hat{\Omega}$ that allows for an arbitrary degree of correlation across households within any time period. By allowing for this arbitrary correlation, we make inference consistently when there is an unknown degree of market incompleteness. Since it is possible that the decision period of the household is finer than the

²⁹Even in the presence of true panel data with a large time dimension this assumption may be required. Attanasio and Low (2000) show that group averaging improves the empirical performance of GMM estimation of the linear consumption Euler equation even when the true economy has no insurance of individual income shocks.

period observed in the data, we also allow for arbitrary covariance among all innovations that share any of the months from which consumption growth is calculated. This flexible form also accounts for the adjustment that is necessary to compensate for the temporal correlation for any single group induced by mismeasurement of consumption. Finally, we allow for arbitrary heteroskedasticity across all observations.

Consider a household for whom consumption growth is observed from the three months ending in $t - 3$ to the three months ending in t . This datum is constructed from monthly observations from $t - 5$ through to t . It follows that all household observations for which the last month of an observed growth in consumption falls anywhere between $t - 5$ to $t + 5$ could potentially be correlated with our original household due to a common effect from an aggregate shock or consumption insurance.

For notational simplicity, first define a scalar expectation error for each group as $\varepsilon_{n,t+1} \equiv e^{\mathbf{X}_{n,t+1}\tilde{\delta}} \left(\frac{C_{n,t+1}}{C_{n,t}} \right)^{-\frac{1}{\tilde{\sigma}}} R_{i,t+1}^f - 1$ where $\tilde{\delta}$ and $\tilde{\sigma}$ are estimates from a first-stage estimation that sets $\hat{\Omega} = I$. Second, let N be the number of cohorts; T_j the number of periods for which cohort j is observed; T_{ij} the number of periods over which both cohorts i and j are observed; and $Z_{i,t}$ the instrument vector for individual i at time t .

Our estimate of $\hat{\Omega}$ is:

$$\hat{\Omega} = \mathbf{P}_0 + \alpha_1 \sum_{v=1}^5 \mathbf{P}_v + \alpha_2 \sum_{v=0}^5 (\bar{\mathbf{P}}_v + \bar{\mathbf{P}}_v')$$

where:

$$\mathbf{P}_v = \frac{1}{N} \sum_{j=1}^N \frac{1}{T_j} \sum_{t=v+1}^{T_j} (\mathbf{Z}_{j,t} \mathbf{Z}_{j,t-v}' u_{j,t} u_{j,t-v} + \mathbf{Z}_{j,t-v}' \mathbf{Z}_{j,t} u_{j,t} u_{j,t-v})$$

$$\bar{\mathbf{P}}_v = \frac{1}{N} \sum_{j>i}^N \frac{1}{T_j} \sum_{t=v+1}^{T_j} (\mathbf{Z}_{j,t} \mathbf{Z}_{i,t-v}' u_{j,t} u_{i,t-v} + \mathbf{Z}_{j,t-v}' \mathbf{Z}_{i,t} u_{j,t} u_{i,t})$$

and α_1 and α_2 are weights to ensure the positive definiteness of the estimated variance-covariance matrix. Thus $\mathbf{P}_v$ captures the set of possible correlations within a given cohort (other than variance); and $\bar{\mathbf{P}}_v$ the set of possible correlations between cohorts. Finally we note that if markets are complete and groups fully insure consumption risk across groups, then contemporaneous cross-correlations should be unity apart from measurement error. This

weighting matrix insures that, if this is the case, the optimal estimate and its variance covariance matrix are asymptotically the same as an estimate derived from the model in which all households are assumed part of the same insurance group and $N = 1$.

C.3. Instruments used in estimation

In selecting instruments, $Z_{n,t}$, we choose variables that the household might use to forecast future wealth or preferences and/or that are relevant for predicting income risk. All instruments are constructed so that they are in the household information set at the start of period t for group n , allowing for time aggregation. We assume that households know the group-average change in family size and number of children that will occur between t and $t + 1$, so these variables are included in $Z_{n,t}$. The instrument set includes the indicator variables for month and interview that are included in preferences, as well as the cohort indicators when they are included in preferences. We include a set of average information about the household: two indicator variables for whether a male head is present and whether a female head is present; four indicators for the possible states of married/single and with and without children; three indicator variables for whether there is a working male head, female head, or other; age; age squared; an indicator for older than 45; two indicator variables for whether there is a male or female head who is retired; two indicator variables for whether there is a male or female head who is a government employee; six indicator variables for whether there is a male or female head with less education than a high school degree, a high school degree, and some college but not a degree. We include six aggregate variables. The instruments contain the real interest rate for $t - 1$ and its lag, and 4 variables that are out-of-sample predictions constructed from rolling regressions using monthly data from period $t - 1$: the log of the unemployment rate for $t + 1$ and its first difference; the real and nominal interest rates in the last month of $t + 1$ (deflated using the NIPA aggregate deflator).

C.4. Parameter Estimates

We group households into 5-year birth cohorts, as done in Attanasio and Weber (1995), but estimate a nonlinear Euler equation. At this level of grouping, there are 33 birth cohorts and

1,901 observations. Parameter estimates are summarized in Table A1, and the first column of results is those we use in the paper. The intertemporal elasticity of substitution is estimated to be 0.65 and statistically significant. The coefficients on change in family size, change in number of children and change in female hours worked, the set of cohort dummies, and the set of seasonal dummies, are all economically and statistically significant. The specification tests do not reject the model (the Hansen J-statistic has p-value 0.56). The estimated σ is quite similar if we instead define cohorts by one year birth groups or interact the five-year groups with four education categories. The demographics remain significant, but are less stable. At still finer definitions of cohort, estimated σ rise or the estimator does not converge. The model has similar behavior if cohort effects are not included. If we estimate the model in linearized form (ignoring precautionary saving), estimates of σ typically decline to half their previous values. In the linear model, our estimated coefficients on the demographic variables and labor supply variable are a little smaller than those in Attanasio and Weber (1995), although both sets of estimates vary substantially by specification.

References

- ARIYAGARI, S. R. (1995): "Optimal Capital Income Taxation with Incomplete Markets, Borrowing Constraints, and Constant Discounting," Journal of Political Economy, 103(6), 1158–75.
- ATTANASIO, O., AND S. J. DAVIS (1996): "Relative Wage Movements and the Distribution of Consumption," Journal of Political Economy, 104(6), 1227–62.
- ATTANASIO, O., AND H. LOW (2000): "Estimating Euler Equations," NBER Technical Working Paper 253.
- ATTANASIO, O., AND G. WEBER (1993): "Consumption Growth, the Interest Rate and Aggregation," Review of Economic Studies, 60, 631–49.
- (1995): "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey," Journal of Political Economy, 103(6), 1121–57.
- BANKS, J., R. BLUNDELL, AND A. BRUGIARINI (2001): "Risk Pooling, Precautionary Saving and Consumption Growth," Review of Economic Studies, 68, 757–79.
- BASU, S., AND M. S. KIMBALL (2000): "Long-Run Labor Supply and the Elasticity of Intertemporal Substitution for Consumption," University of Michigan Manuscript.
- BAXTER, M., AND U. J. JERMANN (1999): "Household Production and the Excess Sensitivity of Consumption to Current Income," American Economic Review, 89, 902–20.
- BERNANKE, B. S., AND M. GERTLER (1995): "Inside the Black Box: The Credit Channel of Monetary Policy Transmission," Journal of Economic Perspectives, 9(4), 869–902.
- BERNANKE, B. S., AND I. MIHOV (1998): "Measuring Monetary Policy," Quarterly Journal of Economics, 113(3), 869–902.
- BRAV, A., G. M. CONSTANTINIDES, AND C. C. GECZY (forthcoming): "Asset Pricing with Heterogeneous Consumers and Limited Participation: Empirical Evidence," Journal of Political Economy.

- CABALLERO, R. J. (1990): "Consumption Puzzles and Precautionary Saving," Journal of Monetary Economics, 25(1), 113–36.
- (1991): "Earnings Uncertainty and Aggregate Wealth Accumulation," American Economic Review, 81(4), 859–71.
- CAGETTI, M. (1998): "Estimating Preferences from Wealth Data," Manuscript, University of Chicago.
- CAMPBELL, J. Y., AND A. DEATON (1989): "Why is Consumption so Smooth?," Review of Economic Studies, 56(3), 357–73.
- CAMPBELL, J. Y., AND N. G. MANKIW (1989): "Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence," in N.B.E.R. Macroeconomics Annual, ed. by O. J. Blanchard, and S. Fischer, pp. 185–215. Cambridge: MIT Press.
- CARROLL, C. D. (1997): "Buffer Stock Saving and the Life Cycle/Permanent Income Hypothesis," Quarterly Journal of Economics, 112(1), 1–55.
- CARROLL, C. D., AND M. S. KIMBALL (2001): "Liquidity Constraints and Precautionary Saving," NBER Working Paper No. 8496.
- CARROLL, C. D., AND A. A. SAMWICK (1997): "The Nature of Precautionary Wealth," Journal of Monetary Economics, 40(1).
- COCHRANE, J. H. (1991): "A Simple Test of Consumption Insurance," Journal of Political Economy, 99(5), 957–76.
- CONSTANTINIDES, G. M., AND D. DUFFIE (1996): "Asset Pricing with Heterogeneous Consumers," Journal of Political Economy, 104(2), 219–40.
- DEATON, A. (1991): "Saving and Liquidity Constraints," Econometrica, 59, 1121–1142.
- DYNAN, K. (1993): "How Prudent Are Consumers?," Journal of Political Economy, 101(6), 1104–1113.
- FLAVIN, M. (1981): "The Adjustment of Consumption to Changing Expectations about Future Income," Journal of Political Economy, 89, 974–1009.
- GOURINCHAS, P.-O. (2000): "Precautionary Saving, Life Cycle and

- Macroeconomics," Mimeo, Princeton University.
- GOURINCHAS, P.-O., AND J. A. PARKER (2002): "Consumption Over the Lifecycle," Econometrica, 70(1), 47–89.
- HALL, R. E. (1988): "Intertemporal Substitution in Consumption," Journal of Political Economy, 96.
- HALL, R. E., AND F. S. MISHKIN (1982): "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households," Econometrica, 50, 461–481.
- HANSEN, L. P., AND K. J. SINGLETON (1982): "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," Econometrica, 50, 1269–86.
- JAPELLI, T., J.-S. PISCHKE, AND N. SOULELES (1998): "Testing for Liquidity Constraints in Euler Equations with Complementary Data Sources," Review of Economics and Statistics, 80(2).
- KRUSELL, P., AND A. A. SMITH (1998): "Income and Wealth Heterogeneity in the Macroeconomy," Journal of Political Economy, 106(5), 867–96.
- LUSARDI, A. (1998): "On the Importance of the Precautionary Saving Motive," American Economic Review (Papers and Proceedings), pp. 449–53.
- MEGHIR, C., AND G. WEBER (1996): "Intertemporal Nonseparability or Borrowing Restrictions? A Dissaggregate Analysis Using a U.S. Consumption Panel," Econometrica, 64(5), 1151–81.
- NELSON, J. A. (1994): "On Testing for Full Insurance Using Consumer Expenditure Survey Data," Journal of Political Economy, 102(2), 384–94.
- OGAKI, M., AND C. M. REINHART (1998): "Measuring Intertemporal Substitution: The Role of Durable Goods," Journal of Political Economy, 106(5), 1078–98.
- PARKER, J. A. (1999): "Spendthrift in America? On two decades of decline in the U.S. saving rate," in N.B.E.R. Macroeconomics Annual, ed. by B. Bernanke, and J. Rotemberg, pp. 317–370. Cambridge: MIT Press.
- RIOS-RULL, V. (1994): "On the Quantitative Importance of Market Completeness,"

Journal of Monetary Economics, 34, 462–496.

SHILLER, R. J. (1982): “Consumption, Asset Markets, and Macroeconomic Fluctuations,” Carnegie Mellon Conference Series on Public Policy, 17, 203–238.

SKINNER, J. (1988): “Risky Income, Life-Cycle Consumption, and Precautionary Saving,” Journal of Monetary Economics, 22, 237–55.

STORESLETTEN, K., C. I. TELMER, AND A. YARON (2000): “Consumption and Risk Sharing Over the Life Cycle,” NBER working paper 7995.

VISSING-JORGENSEN, A. (forthcoming): “Limited Asset Market Participation and the Elasticity of Intertemporal Substitution,” Journal of Political Economy.

ZELDES, S. P. (1989a): “Consumption and Liquidity Constraints: An Empirical Investigation,” Journal of Political Economy, 97, 305–346.

——— (1989b): “Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence,” Quarterly Journal of Economics, 104, 275–98.

Table 1: Variance Decomposition of Average Consumption Growth for Unconstrained Households

	Percent of variation that is not predictable	Percent of predictable variation due to			p-value for precautionary saving series
		Real Interest Rate	Changing Preferences	Precautionary Saving	
GMM Estimates (Raw)					
In Order Listed:	96.4	5.3	14.8	79.9	0.167
Order Reversed:		22.9	7.0	70.1	
GMM Estimates					
In Order Listed:	95.4	81.3	16.3	2.5	0.027
Order Reversed:		4.6	0.2	95.2	
Calibration $\sigma=0.3$					
In Order Listed:	89.9	65.9	29.0	5.0	0.000
Order Reversed:		4.5	59.1	36.5	
Calibration $\sigma=0.65$					
In Order Listed:	94.1	36.2	61.4	2.4	0.000
Order Reversed:		8.3	79.1	12.6	
Calibration $\sigma=1$					
In Order Listed:	95.8	6.9	91.6	1.4	0.000
Order Reversed:		13.4	86.6	0.0	

Note: The GMM (raw) series is constructed without applying the either correction for misspecification or for small-sample bias; the remaining rows do apply this correction. The decompositions for the GMM estimates do not include the movement in consumption due to predictable variation in the innovation to marginal utility – the misspecification correction described in the main text; the remaining decompositions include this movement in the preferences series. See text for further details.

Table 2: Properties of Consumption Growth Due to Precautionary Saving

Dependent variable:	Expected change in unemployment rate	Expected change in real government spending	Lagged monetary policy shocks
Standard deviation of variable:	0.0017	0.0041	0.0032
GMM Estimates			
Coefficient:	0.21	0.32	
Standard error:	(0.19)	(0.05)	
Sum of Coefficients:			0.15
P-value:			0.80
Calibration $\sigma=0.3$			
Coefficient:	0.20	0.78	
Standard error:	(0.39)	(0.08)	
Sum of Coefficients:			0.10
P-value:			0.88
Calibration $\sigma=0.65$			
Coefficient:	0.14	0.39	
Standard error:	(0.21)	(0.04)	
Sum of Coefficients:			0.08
P-value:			0.88
Calibration $\sigma=1$			
Coefficient:	0.10	0.26	
Standard error:	(0.14)	(0.02)	
Sum of Coefficients:			0.06
P-value:			0.88

Note: The decompositions for the GMM estimates do not include the movement in consumption due to "misspecification"; the remaining decompositions include this movement in the preferences series. See text for further details. For Monetary policy shocks the table reports the sum of the estimated coefficients. The p-value refers to the f-test that all lags are zero. For the remaining series point estimates and associated p-values from a t-test are reported. Newey-West standard errors are reported, correcting for up to 12 months serial correlation.

Table 3: Variance Decomposition of Average Consumption Growth for All Households

	Percent of variation that is not predictable	Percent of predictable variation due to			p-value for precautionary saving series
		Real Interest Rate	Changing Preferences	Incomplete Markets	
GMM Estimates (Raw)					
In Order Listed:	95.6	0.2	92.1	7.7	0.028
Order Reversed:		1.6	20.2	78.2	
GMM Estimates					
In Order Listed:	92.9	47.1	0.4	52.5	0.013
Order Reversed:		25.2	5.7	69.1	
Calibration $\sigma=0.3$					
In Order Listed:	93.1	47.5	30.2	22.3	0.000
Order Reversed:		41.0	50.8	8.2	
Calibration $\sigma=0.65$					
In Order Listed:	96.1	4.5	53.0	42.4	0.000
Order Reversed:		65.6	32.9	1.5	
Calibration $\sigma=1$					
In Order Listed:	95.5	13.9	46.1	40.0	0.001
Order Reversed:		32.3	62.1	5.6	

Note: The GMM (raw) series is constructed without applying the either correction for misspecification or for small-sample bias; the remaining rows do apply this correction. The decompositions for the GMM estimates do not include the movement in consumption due to predictable variation in the innovation to marginal utility -- the misspecification correction described in the main text; the remaining decompositions include this movement in the preferences series. See text for further details.

Table 4: Properties of Consumption Growth Due to Precautionary Saving and Liquidity Constraints

Dependent variable:	Expected change in unemployment rate	Expected change in real government spending	Lagged monetary policy shocks
Standard deviation of variable:	0.0017	0.0041	0.0032
GMM Estimates			
Coefficient:	1.03	0.00	
Standard error:	(0.16)	(0.13)	
Sum of Coefficients:			-0.27
P-value:			0.78
Calibration $\sigma=0.3$			
Coefficient:	1.19	0.29	
Standard error:	(0.36)	(0.17)	
Sum of Coefficients:			0.03
P-value:			0.32
Calibration $\sigma=0.65$			
Coefficient:	1.09	-0.07	
Standard error:	(0.26)	(0.15)	
Sum of Coefficients:			0.05
P-value:			0.67
Calibration $\sigma=1$			
Coefficient:	1.04	-0.21	
Standard error:	(0.23)	(0.14)	
Sum of Coefficients:			0.05
P-value:			0.71

Note: The decompositions for the GMM estimates do not include the movement in consumption due to "misspecification"; the remaining decompositions include this movement in the preferences series. See text for further details. For Monetary policy shocks the table reports the sum of the estimated coefficients. The p-value refers to the f-test that all lags are zero. For the remaining series point estimates and associated p-values from a t-test are reported. Newey-West standard errors are reported, correcting for up to 12 months serial correlation.

Table A1: Parameter Estimates from Consumption Euler Equations for Cohort Groupings

	Group: 5-year birth cohorts 33 Cohorts	Group: 1-year birth cohorts 165 Cohorts	Group: 5-year birth cohorts and 4 education groups 132 Cohorts
	1	2	3
Number of Observations:	1901	9670	7599
<u>Coefficients:</u>			
Real Interest Rate	0.652 (0.142)	0.639 (0.094)	0.659 (0.168)
Change in Family Size	1.355 (0.250)	-0.004 (0.512)	1.055 (0.587)
Change in Number of Children	-0.405 (0.085)	0.024 (0.181)	-0.275 (0.180)
Change in Female Hours Worked	-0.043 (0.016)	-0.028 (0.008)	-0.064 (0.015)
<u>Specification Test:</u>	24.35 0.56	37.16 0.07	48.49 0.001

Note: Estimates from two-step GMM estimator described in the appendix. All models also include indicator variables for cohort, month and interview.

Figure 1: Predictable Consumption Growth for Unconstrained Households





Figure 2: Predictable Consumption Growth for All Households

