

Precautionary Saving and Precautionary Wealth

Christopher D. Carroll¹

ccarroll@jhu.edu

Miles S. Kimball²

mkimball@umich.edu

March 15, 2006

Abstract

This is an entry for *The New Palgrave Dictionary of Economics, 2nd Ed.*

Keywords: Precautionary saving, prudence, consumption function, buffer stock saving

JEL Codes: C61, D11, E21

Archive <http://econ.jhu.edu/people/ccarroll/PalgravePrecautionaryArchive.zip>

Text <http://econ.jhu.edu/people/ccarroll/PalgravePrecautionary.pdf>

¹ Department of Economics, Johns Hopkins University, Baltimore MD
<http://econ.jhu.edu/people/ccarroll>
Corresponding author

² Department of Economics, University of Michigan, Ann Arbor, MI
<http://www.umich.edu/~mkimball>

We would like to thank Luigi Guiso, Arthur Kennickell, Annamaria Lusardi, Jonathan Parker, Luigi Pistaferri, and Patrick Toche for valuable comments on an earlier draft which resulted in substantial improvements.

Contents

1	Introduction	2
2	Strength of the Precautionary Saving Motive	2
3	Buffer Stock Wealth	5
3.1	Concavity of the Consumption Function	8
4	Empirical Evidence	9
4.1	Euler Equation Methods	9
4.2	Structural Estimation Using Micro Data	9
4.3	Regression Evidence	10
4.4	Survey Evidence	11
5	Conclusion	12

1 Introduction

Precautionary saving is additional saving that results from the knowledge that the future is uncertain.

In principle, additional saving can be achieved either by consuming less or by working more; here, we follow most of the literature in neglecting the “working more” channel by treating noncapital income as exogenous.

Before proceeding, a terminological clarification is in order. “Precautionary saving” and “precautionary savings” are often (understandably) confused. “Precautionary saving” is a response of current spending to future risk, conditional on current circumstances. “Precautionary savings” is the additional wealth owned at a point in time as the result of past precautionary behavior. That is, precautionary savings at any date is the stock of extra wealth that results from the past flow of precautionary saving. To avoid confusion, we advocate use of the phrase “precautionary wealth” in place of “precautionary savings.”

2 Strength of the Precautionary Saving Motive

In the standard analysis, precautionary saving is modeled as the outcome of a consumer’s optimizing choice of how to allocate existing resources between the present and the future, originally formulated in a two-period model by Leland (1968), and extended to the multiperiod case by Sibley (1975) and Miller (1976). Additional interest in precautionary saving was stimulated by numerical solution of a benchmark model by Zeldes (1989) and the connection made in Barsky, Mankiw, and Zeldes (1986) between precautionary saving and the effects of government debt. (We assume time-invariant preferences in order to sidestep the the important issues of time consistency recently explored by Laibson (1997) and others. That literature opens up a rich and interesting field of further behavioral possibilities beyond the basic logic outlined here.)

To clarify the theoretical issues, we break down the consumer’s problem into two steps: The transition between periods, and the choice within the period. A consumer who ends period t with assets a_t receives capital income in period $t + 1$ of $a_t r$. The consumer’s immediate resources (‘cash on hand’) in period $t + 1$ consist of such capital income, plus the assets that generated

it, plus labor income y_{t+1} :

$$m_{t+1} = a_t r + a_t + y_{t+1} \quad (1)$$

$$= \underbrace{(1+r)}_{\equiv R} a_t + y_{t+1}. \quad (2)$$

The simplest interpretation of m is as the contents of the consumer's bank account immediately after receipt of the paycheck and interest income ('cash-on-hand'). R is the real interest *factor*, as distinct from the real interest *rate*, lower case r . a_t reflects the consumer's accumulated assets at the end of period t , after the spending decision for period t has been made. The transition from the beginning to the end of period t reflects the fact that spending is paid for by drawing down m :

$$a_t = m_t - c_t. \quad (3)$$

To decide how to behave optimally in period t , the consumer must be able to judge the value of arriving in period $t + 1$ in any possible circumstance. This information is captured by the value function $v_{t+1}(m_{t+1})$. Here, we simply assume the existence of some well-behaved v_{t+1} ; below we show how to construct v_{t+1} .

Standard practice assumes that consumers in period t weight future value by the factor β ; if $\beta = 1$ the consumer today cares equally about current and future pleasure, while if $\beta < 1$ the consumer prefers present to future pleasure. Given β , and assuming that the consumer's period- t beliefs about future distribution of income are captured by the expectations operator \mathbf{E}_t , we can define the value of ending period t with accumulated assets a_t as

$$\omega_t(a_t) = \beta \mathbf{E}_t[v_{t+1}(Ra_t + \tilde{y}_{t+1})], \quad (4)$$

where the \sim over the y indicates that period- $(t + 1)$ income is uncertain from the perspective of period t . Think of $\omega_t(a)$ as the end-of-period value function.

The consumer's goal is to optimally allocate beginning-of-period resources between current consumption and end-of-period assets; the value function for period t is defined as the function which yields the value associated with the optimal choice:

$$v_t(m_t) = \max_{c_t} \{u(c_t) + \omega_t(m_t - c_t)\}. \quad (5)$$

By definition the optimal choice will be a level of c_t such that the consumer does not wish to change his spending. Under standard assumptions this implies that the marginal utility of consumption must be equal to the marginal value of assets:

$$u'(\overbrace{m_t - a_t}^{c_t}) = \omega'_t(a_t), \quad (6)$$

since if this were not true the consumer would be able to improve his well-being (value) by reallocating some resources either from consumption into a or from a into c .

Figure 1 depicts the consumer's problem graphically. For given initial m_t , the consumer's goal is to find the value of a such that (6) holds. The left hand side of (6) is the upward-sloping locus. As for the two downward-sloping loci, the lower one reflects expected marginal value if the consumer is perfectly certain to receive the mean level of income $\mathbf{E}_t[\tilde{y}_{t+1}]$, while the higher downward-sloping function corresponds to the case where income is uncertain.

When the risk is added, the optimal choice for end-of-period assets moves from a^* to a^{**} . Since $c_t = m_t - a_t$, the increase in a in response to risk corresponds to a reduction in consumption. This reduction in consumption is the precautionary saving induced by the risk.

For a given $v_{t+1}(m_{t+1})$, the exercise captured in the diagram can be conducted for every possible value of m_t , implicitly defining a consumption function $c_t(m_t)$.

Kimball (1990) shows that the index of absolute prudence $\frac{-v'''_{t+1}(m_{t+1})}{v''_{t+1}(m_{t+1})}$ and the index of relative prudence $\frac{-v'''_{t+1}(m_{t+1})m_{t+1}}{v''_{t+1}(m_{t+1})}$ are good measures of how much a risk of given size will shift the marginal value of assets curve $\omega'_t(a)$ to the right. For a constant relative risk aversion value function, relative prudence is equal to relative risk aversion plus one. Kimball and Weil (2004) look at the strength of the precautionary saving motive when Kreps-Porteus (1978) preferences are used to break the usual equation $\varsigma = 1/\rho$ where ς is the elasticity of intertemporal substitution and ρ is relative risk aversion. In this more general case, the counterpart to relative prudence \mathcal{P} is given by $\mathcal{P} = (1 + \varsigma\varepsilon)\rho$, where ε is the elasticity with which absolute risk aversion declines and absolute risk tolerance increases.

Note that, given the basic properties $\varsigma > 0$ and $\rho > 0$, a positive wealth elasticity of risk tolerance implies that $\mathcal{P} > \rho$. This is a special case of a much

more general result first hinted at by Drèze and Modigliani (1972). Even for very exotic objective functions, the precautionary saving motive will always be stronger than risk aversion whenever ownership of more a_t due to a small forced reduction in consumption would lead an optimizing investor to bear more risk (a property that Drèze and Modigliani (1972) call “endogenously decreasing absolute risk aversion”). This general result holds because if ownership of extra a_t due to a small forced reduction in consumption would lead an optimizing investor to bear risks she was previously indifferent to, then reduced consumption must be complementary with bearing near-indifferent risks. The symmetry of complementarity then implies that, given a free choice of consumption levels, taking on an additional near-indifferent risk will lead an optimizing consumer to reduce consumption. For example, consider an agent with additive habit formation (as distinct from multiplicative habits, cf. Carroll (2000)), for whom reduced consumption not only increases assets but reduces the size of the consumption habit, and so unambiguously leads to more willingness to bear risks. Such an agent will want to reduce consumption if induced to take on an additional risk by a compensation that makes her indifferent to the risk. The size of the compensation is determined by risk aversion. Yet the compensation for the agent’s risk aversion is not enough to cancel out the precautionary saving effect of the risk.

3 Buffer Stock Wealth

The prior discussion suggested that precautionary behavior can be understood by considering a tradeoff between the present (captured by $u(c_t)$) and the future (captured by $\omega_t(m_t - c_t)$).

That analysis was incomplete in a crucial respect: It took the initial level of resources, m_t , as given exogenously. But arguably the most important question about precautionary behavior is how large an effect it has on the prevailing level of m . This cannot be answered using a framework that treats m as exogenous.

The framework can be extended to address this problem, by defining the problem in such a way that the functions v and ω reflect the discounted value of an infinite number of future periods. This is often accomplished by making assumptions under which optimal behavior in every future period is identical to optimal behavior in the current period; it is then possible to solve for a “consumption function” that provides a complete characterization of

the relationship between resources and spending.

The critical extra assumption is “impatience,” broadly construed as a condition on preferences that prevents wealth (or the wealth to income ratio) from growing to infinity. In the simplest version of the model where income does not grow, the required condition is $R\beta < 1$; for the appropriate condition in models with income growth, see Carroll (2004).

The exact nature of income risk turns out to be less important than the assumption of impatience. Here, we analyze a particularly simple case (which is an adaptation of a model by Toche (2005)). There are two kinds of consumers: workers and retirees. Retirees have no labor income, and must live off their assets. Workers earn a fixed amount of labor income in each period, but face a constant danger of being exogenously forced into retirement. (Exogenous forced retirement is the sole source of risk in the model).

Under these assumptions, if the utility function is of the standard constant relative risk aversion form $u(c) = c^{1-\rho}/(1-\rho)$, optimal behavior for retirees is very simple: They spend a constant fraction of m in each period, where the fraction depends on the degree of impatience and intertemporal substitution ($1/\rho$).

The situation for workers is more interesting; it is depicted in figure 2.

The simplest element of the figure is the line labelled “Perm Inc.” This shows, for any m , the level of spending that would leave expected m unchanged; it is equal to labor income plus the interest on capital income, and is upward sloping because a consumer with more m earns more capital income.

The assumption of impatience is reflected in the fact that the consumption function that would apply if uncertainty did not exist, $\bar{c}(m)$, is everywhere above the level of permanent income (income of the perfect-certainty consumer is adjusted downward so that the reduction in unemployment risk does not cause an increase in mean income). In other words, an impatient consumer facing no uncertainty would choose to spend at a rate that cannot be sustained indefinitely.

The locus with arrows is the consumption function, which indicates the optimal level of spending (in the presence of uncertainty) for any given level of m . Since the difference between $c(m)$ and $\bar{c}(m)$ is purely the consequence of risk, that difference $\bar{c}(m) - c(m)$ constitutes the amount of precautionary saving associated with any specific m .

Standard assumptions about preferences and uncertainty imply that there

will be an intersection between the permanent income locus and the consumption function. (For a proof that there will be only one intersection, see Carroll (2004)). The intersection defines a “target” level for the buffer stock of wealth m : The level such that an employed consumer with this amount of resources today will end up with the same m next period. Dynamics are captured by the arrows, which indicate that, for initial values of m below the target, consumption is below permanent income, so m is increasing and consumption crawls upward along the consumption function toward the target. For initial values of m above the target, consumption is above permanent income, so m is falling. The consumer holds a “buffer stock” of wealth in an attempt to reach the “target” level of wealth as defined above.

The existence of a target level of resources has many interesting implications. Perhaps the most surprising is that in long-run equilibrium the expected growth rate of consumption for employed consumers is unrelated to the interest rate or the degree of impatience.

To understand this point better, and to relate it to the literature, we restate it in a slightly more general form: The equilibrium expected growth rate of consumption for employed consumers is approximately equal to their predictable rate of income growth,

$$\mathbf{E}_t[\Delta \log c_{t+1}^e] \approx g. \quad (7)$$

In many respects the equilibrium equality of consumption growth and permanent income growth seems intuitive. However, it appears to conflict with a standard way of analyzing consumption growth, which relies on the first order condition from the optimization problem (the ‘Euler equation’), which is often approximated by an equation of the form

$$\mathbf{E}_t[\Delta \log c_{t+1}^e] \approx \rho^{-1}(r - \tau) + \phi \quad (8)$$

where ρ is the coefficient of relative risk aversion and τ is the geometric rate at which future utility is discounted (related to the time preference factor β); ϕ is a term that reflects the contribution of precautionary motives to consumption growth.

The resolution of the apparent contradiction is that the precautionary component of consumption growth is endogenous; combining (7) and (8) permits us to solve for the equilibrium value of the precautionary contribution to consumption growth:

$$\phi \approx g - \rho^{-1}(r - \tau). \quad (9)$$

We return to this point below.

3.1 Concavity of the Consumption Function

We can characterize the effect of uncertainty by noting three facts about figure 2: $c(m) < \bar{c}(m)$ (consumption is lower in the presence of uncertainty); $\lim_{m \rightarrow \infty} \bar{c}(m) - c(m) = 0$ (as wealth approaches infinity the effect of uncertainty in labor income vanishes); and $c(m)$ is strictly concave, so that the marginal propensity to consume out of a windfall increase in income, $c'(m)$, is greater for poor people than for rich people.

The concavity of the consumption function bears further comment. Intuitively, it can be understood in a similar light to the effect of liquidity constraints. A consumer who is subject to a currently-binding liquidity constraint is someone for whom a marginal increase in cash will result in an immediate one-for-one increase in spending (a marginal propensity to consume (MPC) of one). However, if the same consumer happened to have a large windfall transfer of cash (say, he wins the lottery), he would no longer be currently constrained, and his MPC would (presumably) be less than one. In the case of precautionary saving, the ownership of an extra unit of wealth relaxes the suppression of consumption due to risk; this relaxation is more powerful for low-wealth consumers living on the edge of (precautionary) fear than for high wealth consumers with plenty of resources. Thus, either liquidity constraints or precautionary motives or both will cause the consumption function to become concave (Carroll and Kimball (2005)). Huggett (2004) shows that consumption concavity in turn implies greater equilibrium wealth.

Empirical evidence indicates that the wealth distribution is highly concentrated. This means that the owners of much of the aggregate capital stock likely inhabit the portion of the consumption function to the far right, where it approaches the linear consumption function that characterizes the perfect foresight solution. Note, however, that this does not necessarily imply that aggregate consumption behavior will resemble that of a perfect foresight consumer, because a large proportion of aggregate consumption is accounted for by households with small amounts of market wealth. Spending of such households is likely determined much more by their permanent income than by their meager wealth, and so it remains possible that a high proportion of consumption is performed by households inhabiting the more nonlinear part of the consumption function.

4 Empirical Evidence

4.1 Euler Equation Methods

The early literature relevant to identifying the strength of precautionary motives tended to rely on Euler equation estimation (see Browning and Lusardi (1996) for a survey), often by estimating regression equations of the form

$$\Delta \log c_{t+1} = \alpha_0 + \alpha_1 \mathbf{E}_t[r_{t+1}] \quad (10)$$

and interpreting the coefficient on the interest rate term as an estimate of the inverse of the coefficient of relative risk aversion (which holds true under time-separable CRRA utility, cf. (8)). However, this analysis did not take into account the dependence of higher order terms like ϕ on the independent variables (see (9)). Some papers like Dynan (1993) attempted to account for precautionary contributions to consumption growth; but see Carroll (2001) for a critique of the whole Euler equation literature (including the second-order approach).

4.2 Structural Estimation Using Micro Data

A new methodology for estimating the importance of precautionary motives was pioneered by Gourinchas and Parker (2002) and Cagetti (2003) (with a related earlier contribution by Palumbo (1999)). Their idea was to calibrate an explicit life cycle optimization problem using empirical data on the magnitude of household-level income shocks, and to search econometrically for the values of parameters such as the coefficient of relative risk aversion that maximized the model's ability to fit some measured feature of the empirical data. Gourinchas and Parker (2002) matched the profile of mean consumption over the lifetime; Cagetti (2003) matched the profile of median wealth. The intensity of the precautionary motive emerges, in each case, as an estimate of the coefficient of relative risk aversion, which Gourinchas and Parker (2002) put at about 1.4 and Cagetti (2003) finds to be somewhat larger. (A value of 1 corresponds to logarithmic utility). One important caution about these quantitative results is that the method's estimates of relative risk aversion depend on the model's assumption about the degree of risk households face. Recent work by Low, Meghir, and Pistaferri (2005) that attempts to correct for measurement problems caused by job mobility suggests that the estimates of the magnitude of permanent shocks in Carroll and Samwick (1997) used

for calibration by Gourinchas and Parker (2002) and Cagetti (2003) may be overstated by as much as 50 percent. Reestimation of the structural parameters using the Low et. al. calibration would generate larger estimates of relative risk aversion.

4.3 Regression Evidence

A separate literature attempts direct empirical measurement of the relationship between uncertainty and wealth. To fix notation, index individual households by i and assume uncertainty for household i in period t can be measured by some variable $\sigma_{t,i}$. Then in its simplest form the idea is to perform a regression of cash-on-hand on its determinants along the lines of

$$\log m_{t,i} = \sigma_{t,i}\gamma + Z_{t,i}\alpha + \epsilon_{t,i} \quad (11)$$

where Z is some set of variables that capture life cycle, time series, and other nonprecautionary effects. In principle, one can then calculate the predicted magnitude of m if everyone's uncertainty were set to zero (or some alternative like the minimum measured value of σ in the population).

In principle this method permits the data to speak in a much less filtered way than the structural estimation approach. A drawback is that even if the magnitude of precautionary wealth could be estimated reliably and precisely, it would not be clear how to translate those estimates into a measure of relative risk aversion or some other set of behavioral parameters that could be used for analyzing policy questions such as the optimal design of unemployment insurance or taxation.

A further disadvantage is that the method does not reliably yield the same answer in different data. Using a measure of subjective earnings uncertainty from a survey of Italian households, Guiso, Jappelli, and Terlizzese (1992) estimate the precautionary component of wealth at only a few percent, while Kazarosian (1997) and Carroll and Samwick (1998) estimate the precautionary component of wealth for typical U.S. households to be in the range of 20-50 percent. Hurst, Lusardi, Kennickell, and Torralba (2005) argue that estimates of α are inordinately sensitive to whether business owners are included in the dataset; and work by Lusardi (1998, 1997) and Engen and Gruber (2001) implies much smaller precautionary wealth. Such large variation in empirical estimates is not plausibly attributable to actual behavioral differences across the various sample populations.

A problem that plagues all these efforts is identifying exogenous variations in uncertainty across households. The standard method has been to use patterns of variation across age, occupation, education, industry, and other characteristics. This runs the danger that people who are more risk tolerant may both choose to work in a risky industry and choose not to save much, biasing downward the estimate of the effect of an exogenous change in risk.

One recent paper attempts to get around this problem by using a natural experiment: Fuchs-Schündeln and Schündeln (2005) show that before the collapse of the Berlin Wall, East German civil servants had similar income uncertainty to that faced by other East Germans. However, after the collapse of Communism, income uncertainty went up dramatically for most East Germans - but not for civil servants, who were given essentially the same risk-free jobs in the new merged government that they had had before the collapse. Fuchs-Schündeln and Schündeln (2005) show that, in accord with a model that includes substantial precautionary effects, saving rates of most East Germans increased sharply after unification, but saving rates of civil servants did not. By contrast, the West Germans—who would have been subject to more selection into jobs based on risk preferences—exhibited little difference in saving rates between civil servants and others with riskier jobs, either before or after reunification.

4.4 Survey Evidence

Given the difficulties of obtaining reliable quantitative measures of precautionary motives using the revealed preference econometric techniques sketched above, some researchers have turned to approaches that involve asking survey participants more direct questions.

Kennickell and Lusardi (2005) find that when respondents for the 1995 and 1998 U.S. *Survey of Consumer Finances* are asked their target level of precautionary wealth, most have little difficulty answering the question; desired precautionary wealth represents about 8 percent of total net worth and 20 percent of total financial wealth. They find that respondents cite a broad array of risks in making their precautionary targets: In addition to labor income risk, they face health risk, business risk, and the risks of unavoidable expenditures (e.g. home repairs). (Consumers are clearly aware of the theoretical point that a given dollar of wealth can provide self-insurance against multiple different kinds of risks, since the risks are not likely to be perfectly correlated with each other).

Carefully designed survey questions can in principle also be used to elicit information on the strength of underlying preferences (like risk aversion) that determine precautionary behavior. The principle that whenever risk-bearing increases with assets, the precautionary saving motive (prudence) must be stronger than risk aversion provides an important theoretical lower bound on the degree of prudence. Survey responses to hypothetical gambles over lifetime income in the *Health and Retirement Study*. Kimball, Sahn, and Shapiro (2005) estimate that relative risk aversion has a median of 6.3 and a mean of 8.2. (Note that because of Jensen's inequality, the mean of relative risk aversion $E\rho$ is larger than the reciprocal of the mean of relative risk tolerance $\frac{1}{E(1/\rho)}$.) These estimates of relative risk aversion imply precautionary saving motives much stronger than those that have been used empirically to match observed wealth holdings. This discrepancy remains unresolved.

5 Conclusion

The qualitative and quantitative aspects of the theory of precautionary behavior are now well established. Less agreement exists about the strength of the precautionary saving motive and the magnitude of precautionary wealth. Structural models that match broad features of consumption and saving behavior tend to produce estimates of the degree of prudence that are less than those obtained from theoretical models in combination with risk aversion estimates from survey evidence. Direct estimates of precautionary wealth seem to be sensitive to the exact empirical procedures used, and are subject to problems of unobserved heterogeneity that have been demonstrated from German data after reunification. Thus, establishing the intensity of the precautionary saving motive and the magnitude of precautionary wealth remain lively areas of debate.

References

- BARSKY, ROBERT B., N. GREGORY MANKIW, AND STEPHEN P. ZELDES (1986): “Ricardian Consumers with Keynesian Propensities,” *American Economic Review*, 76(4), 676–91.
- BROWNING, MARTIN J., AND ANNAMARIA LUSARDI (1996): “Household Saving: Micro Theories and Micro Facts,” *Journal of Economic Literature*, 34(4), 1797–855.
- CAGETTI, MARCO (2003): “Wealth Accumulation Over the Life Cycle and Precautionary Savings,” *Journal of Business and Economic Statistics*, 21(3), 339–353.
- CARROLL, CHRISTOPHER D. (2000): “Solving Consumption Models with Multiplicative Habits,” *Economics Letters*, 68(1), 67–77, <http://econ.jhu.edu/people/ccarroll/HabitsEconLett.pdf>.
- (2001): “Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation),” *Advances in Macroeconomics*, 1(1), Article 6, <http://econ.jhu.edu/people/ccarroll/death.pdf>.
- (2004): “Theoretical Foundations of Buffer Stock Saving,” *NBER Working Paper Number 10867*, <http://econ.jhu.edu/people/ccarroll/BufferStockProofsNew.pdf>.
- CARROLL, CHRISTOPHER D., AND MILES S. KIMBALL (2005): “Liquidity Constraints and Precautionary Saving,” *Manuscript, Johns Hopkins University*, <http://econ.jhu.edu/ccarroll/liquidRevised.pdf>.
- CARROLL, CHRISTOPHER D., AND ANDREW A. SAMWICK (1997): “The Nature of Precautionary Wealth,” *Journal of Monetary Economics*, 40(1), 41–71, <http://econ.jhu.edu/people/ccarroll/nature.pdf>.
- (1998): “How Important Is Precautionary Saving?,” *Review of Economics and Statistics*, 80(3), 410–419, <http://econ.jhu.edu/people/ccarroll/howbig.pdf>.
- DRÈZE, JACQUES H., AND FRANCO MODIGLIANI (1972): “Consumption Decisions Under Uncertainty,” *Journal of Economic Theory*, 5, 308–335.

- DYNAN, KAREN E. (1993): “How Prudent Are Consumers?,” *Journal of Political Economy*, 101(6), 1104–1113.
- ENGEN, ERIC, AND JONATHAN GRUBER (2001): “Unemployment Insurance and Precautionary Saving,” *Journal of Monetary Economics*, 47, 545–579.
- FUCHS-SCHÜNDELN, NICOLA, AND MATTHIAS SCHÜNDELN (2005): “Precautionary Savings and Self-Selection: Evidence from the German Reunification “Experiment”,” *Quarterly Journal of Economics*, 120(3), 1085–1120.
- GOURINCHAS, PIERRE-OLIVIER, AND JONATHAN PARKER (2002): “Consumption Over the Life Cycle,” *Econometrica*, 70(1), 47–89.
- GUIO, LUIGI, TULLIO JAPPELLI, AND DANIELE TERLIZZESE (1992): “Earnings Uncertainty and Precautionary Saving,” *Journal of Monetary Economics*, 30(2), 307–37.
- HUGGETT, MARK (2004): “Precautionary Wealth Accumulation,” *Review of Economic Studies*, 71, 769–781.
- HURST, ERIK, ANNAMARIA LUSARDI, ARTHUR KENNICKELL, AND FRANCISCO TORRALBA (2005): “Precautionary Savings and the Importance of Business Owners,” *Manuscript, Dartmouth College*.
- KAZAROSIAN, MARK (1997): “Precautionary Savings – A Panel Study,” *Review of Economics and Statistics*, 79(2), 241–247.
- KENNICKELL, ARTHUR, AND ANNAMARIA LUSARDI (2005): “Disentangling the Importance of the Precautionary Saving Motive,” *Working Paper, Dartmouth College*.
- KIMBALL, MILES S. (1990): “Precautionary Saving in the Small and in the Large,” *Econometrica*, 58, 53–73.
- KIMBALL, MILES S., CLAUDIA R. SAHM, AND MATTHEW D. SHAPIRO (2005): “Imputing Risk Tolerance from Survey Responses,” *Manuscript, University of Michigan*.
- KIMBALL, MILES S., AND PHILIPPE WEIL (2004): “Precautionary Saving and Consumption Smoothing Across Time and Possibilities,” *Manuscript, University of Michigan*.

- KREPS, DAVID M., AND EVAN L. PORTEUS (1978): “Temporal Resolution of Uncertainty and Dynamic Choice Theory,” *Econometrica*, 46, 185–200.
- LAIBSON, DAVID (1997): “Golden Eggs and Hyperbolic Discounting,” *Quarterly Journal of Economics*, CXII(2), 443–477.
- LELAND, HAYNE E. (1968): “Saving and Uncertainty: The Precautionary Demand for Saving,” *Quarterly Journal of Economics*, 82, 465–473.
- LOW, HAMISH, COSTAS MEGHIR, AND LUIGI PISTAFERRI (2005): “Wage Risk and Employment Over the Life Cycle,” *Manuscript, Stanford University*.
- LUSARDI, ANNAMARIA (1997): “Precautionary Saving and Subjective Earnings Variance,” *Economics Letters*, 57, 319–326.
- (1998): “On the Importance of the Precautionary Saving Motive,” *American Economic Review Papers and Proceedings*, 88(2), 449–453.
- MILLER, BRUCE L. (1976): “The Effect on Optimal Consumption of Increased Uncertainty in Labor Income in the Multiperiod Case,” *Journal of Economic Theory*, 13, 154–167.
- PALUMBO, MICHAEL G (1999): “Uncertain Medical Expenses and Precautionary Saving Near the End of the Life Cycle,” *Review of Economic Studies*, 66(2), 395–421, <http://ideas.repec.org/a/bla/restud/v66y1999i2p395-421.html>.
- SIBLEY, DAVID S. (1975): “Permanent and Transitory Effects of Optimal Consumption with Wage Income Uncertainty,” *Journal of Economic Theory*, pp. 68–82.
- TOCHE, PATRICK (2005): “A Tractable Model of Precautionary Saving in Continuous Time,” *Economics Letters*, 87(2), 267–272, available at <http://ideas.repec.org/a/eee/ecolet/v87y2005i2p267-272.html>.
- ZELDES, STEPHEN P. (1989): “Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence,” *Quarterly Journal of Economics*, 104(2), 275–298.

Figure 1: Marginal Utility of Assets and of Consumption

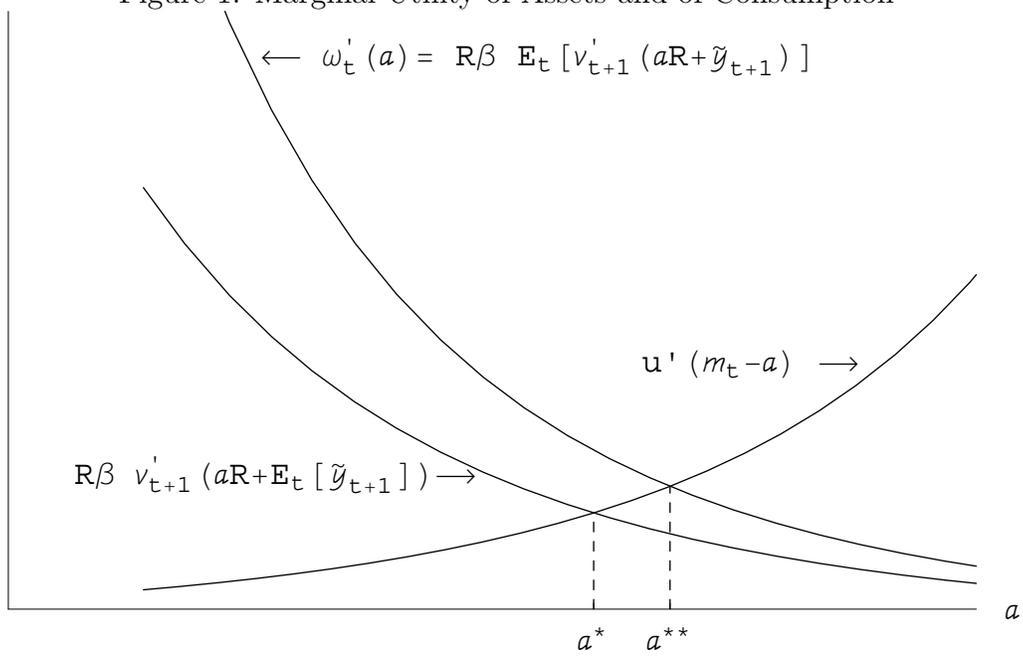


Figure 2: The Consumption Function

