

# Consumption and Saving: Theory and Evidence

Christopher D. Carroll

*Chris Carroll is Professor of Economics at the Johns Hopkins University, and a Research Associate in the National Bureau of Economic Research programs on Monetary Economics and on Economic Fluctuations and Growth.*

Consumption and saving decisions are at the heart of both short-run and long-run macroeconomic analysis (as well as much of microeconomics). In the short run, spending dynamics are of central importance for business cycle analysis and the management of monetary policy. And in the long run, aggregate saving determines the size of the aggregate capital stock, with consequences for wages, interest rates, and the standard of living.

Since the pioneering work of Friedman and of Modigliani and Brumberg in the 1950s, the principal goal of the economic analysis of saving has been to formulate mathematically rigorous theories of behavior. But that goal was difficult until recently because the optimal response of saving to uncertainty was difficult to compute. Research was generally carried out under the assumption that uncertainty might boost saving somewhat, but that behavior in the presence of uncertainty was likely to be broadly similar to optimal behavior in a world in which households had perfect foresight about their future circumstances.

In two papers that grew out of my 1990 dissertation,<sup>1</sup> I showed that the presence of uncertainty could change the nature of optimal behavior in qualitatively and quantitatively important ways. Specifically, I examined the optimal behavior of consumers with standard attitudes toward risk (constant relative risk aversion), who faced income uncertainty of the kind that appears to exist in household-level data sources. The first paper found that target or “buffer-stock” saving can be optimal under some circumstances; the second paper found that, depending on households’ income profiles and their degree of impatience, it can be optimal for average household spending patterns to mirror average household income profiles over much of the life cycle. This was a surprising finding because in models without uncertainty optimizing consumers spend based on their expected lifetime resources without regard to the expected timing of income; that is, spending patterns by age are not intrinsically determined by income patterns by age. (This work, and my subsequent related work, assumes that consumers have successfully solved

any “self-control” problems of the kind that David Laibson and others have so persuasively described).

This paper was related to two other, more abstract, papers. The more fundamental of these,<sup>2</sup> written with Miles Kimball, showed that in the presence of uncertainty, households with low levels of wealth will respond more to a windfall infusion of cash than households with ample resources. The other paper<sup>3</sup> demonstrated that the logic of precautionary saving undermines the standard “Euler equation” method of testing for optimizing consumption behavior.

Mathematical and computational aspects of optimal behavior have remained a theme in my research to the present. A recent paper provides the rigorous foundations for the mathematical methods employed in my earlier work.<sup>4</sup> Another paper with Miles Kimball<sup>5</sup> explores the theoretical implications of borrowing limitations; and a very short new paper describes a conceptual trick that can be used to simplify and accelerate the solution of many kinds of optimal intertemporal choice models.<sup>6</sup> As an aid to other researchers, I have posted on my web page computer software that implements this trick to solve a variety of standard optimization problems; my web page also contains software that reproduces the computational and empirical results in most of my published papers, as well as a set of lecture notes (and associated software) that provide a comprehensive treatment of the methods for solving these models.<sup>7</sup>

In the end, however, mathematical models are useful only insofar as they can be related to empirical evidence about the real world. Toward the end of matching theory and data, Andrew Samwick and I wrote two papers<sup>8,9</sup> whose goal was to get a quantitative sense of the nature and magnitude of household responses to uncertainty. The first of these papers found that a standard source of microeconomic data, the *Panel Study of Income Dynamics*, implied that income uncertainty was very large indeed. According to the benchmark specification, a conservative estimate was that in any given year about a third of households could expect their “permanent” income to rise or fall by as much as ten percent. (“Permanent” changes in income here mean the kind of change associated with a promotion or being laid off and settling for a new lower-paying job). The second paper with Samwick estimated that as much as 40 percent of the wealth held by the typical household represented a response to the fact that some households face greater uncertainty than others.

An important caveat about these results is that many of the wealthiest

households are missing from the PSID dataset on which the estimates are based. Since a large proportion of aggregate wealth is held by the richest few percent of households, these estimates very likely overstate the proportion of aggregate wealth that can be attributed to precautionary motives. Indeed, another paper<sup>10</sup> showed that the theoretical model used in the first paper with Samwick severely underpredicts the wealth holdings of the wealthiest households in the U.S. even if wealthy individuals are assumed to be more patient than others. That paper argued that a bequest motive in which bequests are a “luxury” good is essential to explaining why saving rates of wealthy households are so high. A subsequent paper<sup>11</sup> showed that the “bequests as luxuries” model can also explain a variety of facts about the portfolio choices of wealthy households, particularly their comparatively high tolerance for financial risk.

Another potential problem with my work with Samwick is that we were forced by data limitations to make the assumption that income risk is something over which people have no control. If instead people make employment choices based partly on the riskiness of the different alternatives (for example, if risk-averse people seek civil service jobs while the risk-lovers become entrepreneurs), the estimated effect of uncertainty on saving might be incorrect. The likeliest effect would be to underestimate the importance of precautionary behavior, since the theory tends to suggest that those who dislike risk more will both avoid risky occupations and save more. But in an attempt to get around this problem, Karen Dynan and Spencer Krane and I wrote a paper<sup>12</sup> that used temporary regional variations in unemployment risk (over which individual households have no control) to measure the size of uncertainty. Empirical results in that paper suggested that precautionary motives for saving were more important for people in the upper half of the income distribution, and that precautionary behavior is manifested partly in a reluctance to borrow against home equity when unemployment is high, rather than an explicit accumulation of greater liquid assets.

If uncertainty matters this much for spending decisions on average, it seems plausible that the changes in uncertainty that accompany business cycles might be an important source of fluctuations in consumer spending. Wendy Dunn and I showed<sup>13</sup> that while there does not seem to be any systematic relationship between spending and various measures of households’ financial condition, measures of consumers’ degree of uncertainty (especially their assessment of whether the unemployment rate is likely to rise) have a powerful impact on spending (particularly purchases of big-ticket items like

vehicles and houses).

In fact, the model in that paper suggested that, if anything, the mystery is why uncertainty-driven fluctuations in expenditures on durable goods are not even larger. According to the model, most of the people who were on the verge of buying a car should be willing to postpone their purchase in response to even a very modest increase in uncertainty; while the evidence confirms that durable goods spending is indeed more volatile than spending on nondurables like food, the size of the discrepancy is not as large as the rational optimization model tends to suggest it should be.

This finding seems to fit with the results of an earlier paper with David N. Weil<sup>14</sup> which found that, across countries, the relationship between aggregate saving and aggregate growth is not what would be expected from the standard framework in which spending depends on expectations about future income. The problem is that people living in fast-growing economies should expect their future incomes to be large relative to their current incomes, and should therefore be borrowing to finance their current expenditures, while people in slow-growing economies should anticipate that they may need to save a lot if they wish to maintain their current standard of living in the future. The logic therefore suggests that we should expect to see a negative association between saving and growth.

One objection to this thread of reasoning might be that countries' saving rates differ partly for cultural reasons, and it seems natural to expect that countries whose saving rates are high because of a cultural preference for saving would consequently exhibit high growth. Byung-Kun Rhee and Changyong Rhee and I used data on immigrants to Canada<sup>15</sup> to investigate the possibility that cultural differences explain saving differences. Under the "cultural" theory of saving, one might expect that immigrants from high-saving countries (e.g. Japan) to save more than immigrants from low-saving countries (e.g. Sweden). But we found no evidence of such a pattern, either in Canada or in a subsequent study using Census data from the U.S.<sup>16</sup>

Furthermore, the evidence clearly suggests that the relationship between saving and growth is dynamic, not static: Countries that go through periods of prolonged growth tend to experience *rising* saving rates, while countries that experience sustained economic slowdowns tend to suffer declining saving rates.

Both the sluggish response of spending to uncertainty and the pattern in which increases (or decreases) in growth produce increases (respectively, decreases) in saving might be explained by a model in which spending "habits"

exert a powerful influence on behavior. A paper with Jody Overland and David N. Weil<sup>17</sup> explored how the incorporation of spending habits modifies the predictions of a model of optimal spending behavior. A subsequent paper<sup>18</sup> incorporated both habits and uncertainty, and argued that the broad patterns of saving and growth seen in the East Asian “tiger” economies could be explained in a model where both precautionary motives and habit formation were important. This work meshes with a prominent strand of the macroeconomics and finance literatures over the past decade, which have argued that habit formation can explain a wide range of empirical observations that are difficult to reconcile with standard models without habits.

A new paper with Jirka Slacalek,<sup>19</sup> however, casts doubt on the view that habits are the right explanation for the sluggishness of aggregate spending dynamics. This paper points out that habits imply that spending dynamics should be similar in microeconomic and macroeconomic data. Yet empirical studies using microeconomic data, using exactly the same methods as applied to macroeconomic data, find very different results. While the data hint that there may be some modest habit formation effects in a few categories of spending, models in which habits are a dominant force in microeconomic spending decisions can be decisively rejected.

The new paper relates to another strand of my research, which argues that economists should pay more attention than has been customary to the evidence provided by surveys of households. A 2001 NBER working paper proposed modeling household survey data on inflation expectations using a simple model of disease transmission. The idea is that rather than forming their own independent views of the likely future inflation rate, typical people’s views are formed by exposure to the views of experts as represented in the news media. In this model, households’ forecasts of inflation, while not fully “rational” in the economist’s usual strict sense of the term, do not deviate very long or very far from the experts’ view. The paper presented empirical evidence that information in newspaper reports about inflation seems to filter out to the population gradually rather than instantly. The proposed model can be interpreted as providing a concrete theoretical justification for the model of “sticky expectations” that has become increasingly popular in the macroeconomics literature in recent years. (The NBER working paper was subsequently split into two papers, one containing the empirical evidence and a stripped-down version of the model, and the other examining a detailed examination of the epidemiological modeling framework and its application).<sup>20</sup>

The paper with Slacalek proposes to reconcile the microeconomic and macroeconomic evidence about consumption dynamics by applying the same model of sticky expectations. The essential idea is that people have a very good understanding of the circumstances they face in their own lives (for example, they know whether they have been fired), but they do not pay as much attention to macroeconomic developments (for example, they may not know the latest aggregate unemployment statistic). Since household-specific uncertainty is much greater than aggregate uncertainty (a rough estimate is that household-specific risks are about 100 times larger than macroeconomic risks), it makes sense for busy consumers to pay less than perfect attention to the macroeconomy.

Whether or not this particular explanation for the conflict between microeconomic and macroeconomic data on consumption dynamics is accepted, this conflict seems likely to be a topic of growing attention over the next few years. While great progress has been made in understanding the quantitative implications of alternative models of consumption and saving behavior, much remains to be understood.

## Endnotes

<sup>1</sup>Christopher D. Carroll. The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence. *Brookings Papers on Economic Activity*, 1992(2):61–156, 1992 and Christopher D. Carroll. Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis. *NBER Working Paper No. 5788*, October 1996; the second paper was published in the *Quarterly Journal of Economics*, CXII(1):1–56, 1997. URL <http://econ.jhu.edu/people/ccarroll/BSLCPIH.zip>

<sup>2</sup>

<sup>3</sup>Christopher D. Carroll. Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation). *NBER Working Paper No. 6298*, 1997; published in the

<sup>4</sup>

<sup>5</sup>Christopher D. Carroll and Miles S. Kimball. Liquidity Constraints and Precautionary Saving. *NBER Working Paper No. 8496*, 2001

<sup>6</sup>; forthcoming in *Economics Letters*.

<sup>7</sup>

<sup>8</sup>Christopher D. Carroll and Andrew A. Samwick. The Nature of Precautionary Wealth. *NBER Working Paper No. 5193*, July 1995a; published in the

<sup>9</sup>Christopher D. Carroll and Andrew A. Samwick. How Important Is Precautionary Saving? *NBER Working Paper No. 5194*, July 1995b; published in *Review of Economics and Statistics*, 80(3):410–419, August 1998

<sup>10</sup>Christopher D. Carroll. Why Do the Rich Save So Much? *NBER Working Paper Number 6549*, May 1998; published In Joel B. Slemrod, editor, *Does Atlas Shrug? The Economic Consequences of Taxing the Rich*. Harvard University Press, 2000a

<sup>11</sup>Christopher D. Carroll. Portfolios of the Rich. In *NBER Working Paper No. 7826*. August 2001; published In *Household Portfolios: Theory and Evidence*. MIT Press, Cambridge, MA, 2002.

<sup>12</sup>

<sup>13</sup>Christopher D. Carroll and Wendy E. Dunn. Unemployment Expectations, Jumping (S,s) Triggers, and Household Balance Sheets. *NBER Working Paper No. 6081*, July 1997; published In Benjamin S. Bernanke and Julio J. Rotemberg, editors, *NBER Macroeconomics Annual, 1997*, pages 165–229. MIT Press, Cambridge, MA, 1997

<sup>14</sup>Christopher D. Carroll and David N. Weil. Saving and Growth: A Reinterpretation. *NBER Working Paper No. 4470*, September 1993; published in the *Carnegie-Rochester Conference Series on Public Policy*, 40:133–192, June 1994

<sup>15</sup>

<sup>16</sup>Christopher D. Carroll, Changyong Rhee, and Byungkun Rhee. Does Cultural Origin Affect Saving Behavior? Evidence from Immigrants. *NBER Working Paper No. 6568*, May 1998; published in *Economic Development and Cultural Change*, 48(1):33–50, October 1999

<sup>17</sup>

<sup>18</sup>Christopher D. Carroll. ‘Risky Habits’ and the Marginal Propensity to Consume Out of Permanent Income. *NBER Working Paper No. 7839*, 2000; published in the *International Economic Journal*, 14(4):1–41, 2000b

<sup>19</sup>Christopher D. Carroll and Jiri Slacalek. Sticky Expectations and Consumption Dynamics. *Manuscript, Johns Hopkins University*, 2007

<sup>20</sup>The basic model and evidence are published in ; the full model and its extensions can be found in The Epidemiology of Macroeconomic Expectations. In Larry Blume and Steven Durlauf, editors, *The Economy as an Evolving Complex System, III*. Oxford University Press, 2006.