Labor Income Uncertainty and Macroeconomics

May 21, 2014

Christopher D. Carroll[1]

Abstract

This is a response to email queries from Larry Christiano about my views on the role of uncertainty, and in particular unemployment risk, in business cycle fluctuations.

Keywords Uncertainty, Unemployment, Business Cycles, Consumer Sentiment

JEL codes D81

PDF: http://econ.jhu.edu/people/ccarroll/papers/papers/ChristianoBufferStock.pdf
Web: http://econ.jhu.edu/people/ccarroll/papers/papers/ChristianoBufferStock/
Archive: http://econ.jhu.edu/people/ccarroll/papers/papers/ChristianoBufferStock.zip

[1]Contact: ccarroll@jhu.edu, Department of Economics, 440 Mergenthaler Hall, Johns Hopkins University, Baltimore, MD 21218, http://econ.jhu.edu/people/ccarroll, and National Bureau of Economic Research.

1 Introduction

Larry,

You wrote:

> It seems like one way to quantity the impact of unemployment uncertainty on saving would be to work it our in an Aiyagari model. One way to do it would be two compare two steady states: 6 percent of the population is randomly selected to be unemployed in one and 10 percent is randomly selected to be unemployed in the other. Unemployment spells for individuals would have an expected duration that corresponds to the data. No aggregate uncertainty, so easy to work out. The impact of unemployment on the saving rate could be computed as the difference between the two steady states.
>
> Alternatively, the fraction of unemployment could fluctuate between 6 and 10 percent according to a persistent two-state Markov chain. That would be harder because of the aggregate uncertainty.
>
> Have people done things like this?
>
> My gut feeling is that changes in the unemployment rate are too small to produce a sizable effect on the saving rate. That is, to get a one percent rise in the unemployment rate to produce an appreciable rise in the saving rate, you would need a high degree of risk aversion. But, that same high risk aversion would lead people to hold, on average, a counterfactually high level of assets.
>
> I'm sure I'm wrong. Actually, I want to be wrong on this.
>
> Best, Larry

First, I do not think an Aiyagari model is the right way to think about this, at least insofar as Aiyagari (1994) is distinct from Bewley (1977) and my own paper that I forwarded to you, Carroll (1992). The only conceptual (as opposed to quantitative/calibration) distinction between Aiyagari (1994) and Bewley/Carroll is that Aiyagari's (constant) real interest rate is calibrated to be equal to the steady-state general equilibrium marginal product of capital. Bewley (1977) just fixes an interest rate and requires impatience relative to it. Carroll (1992) targets the median wealth/income ratio rather than the aggregate wealth/income ratio (with corresponding interest rate).

I'm with Bob Hall, who has been heard to remark that the only context in which that kind of general equilibrium really needs to be insisted upon is in the treatment of a medium-run model of the entire world. (The short run is not well captured by such models because all sorts of frictions, discontinuities, market failures, expectational errors, capital market frictions, and other inconveniences intervene. The long run is best analyzed with growth models.) Given the freedom with which capital moves around the world these days, even the United States is arguably best modeled as a small open economy (as in Bernanke's "Global Savings Glut" argument, Greenspan's "conundrum" and similar arguments).

The reason that this matters is that when you constrain the model to match a capital/output ratio of 3 or 4, and you assert that all agents have the same time preference rate, you end up with a model in which most of the individual microeconomic agents also have wealth-to-income ratios in the range of 3-4.

But people that rich are not the ones whose consumption is likely to react violently to a change in unemployment risk. They have enough wealth to ride out even quite a long unemployment spell.

Instead, you need a model that is calibrated to match the behavior of the typical *consumer*, not the *mean* holder of *wealth*.

The reason that this is not a distinction without a difference is that, when microeconomic risk is properly accounted for, the consumption function is not linear. (Carroll and Kimball (1996) prove the concavity of the consumption function in quite generic circumstances: HARA utility (which encompasses all the usual specific functional forms for utility usually considered).)

This point is illustrated in figure 1, taken from Carroll (2001), which depicts the optimal level of consumption (c) as a function of 'cash-on-hand' x for a consumer with CRRA utility and a microeconomically realistic calibration of uncertainty.

The lower "No Borrowing" consumption function reflects optimal behavior for a consumer who is prohibited from borrowing at all. The higher, "Can Borrow Up to 0.3" function is optimal for a consumer who can borrow an amount up to 30 percent of a year's permanent income.

In either case, though, the point is that the response of spending to a transitory income shock (which, in this framework, translates directly into an incremental amount of x) depends critically on the consumer's cash-on-hand. At low levels of cash-on-hand, the marginal propensity to consume is (optimally) 1: Each extra dollar in resources translates one-for-one into an extra dollar of spending. Carroll (2004) proves that as x approaches infinity, the MPC approaches the value that it takes in a perfect foresight model.

So, average consumption behavior will depend on how consumers are distributed across different levels of x.

Returning to Aiyagari (1994) versus Carroll (1992), the key difference is that I calibrated my model so that it matched the wealth of the *median* consumer, while Aiyagari (1994) calibrated his model so that it matched mean wealth rather than the median. Thus, his model effectively puts most consumers way way way out to the right in the figure, while mine puts them much closer to the steep part of the consumption function.

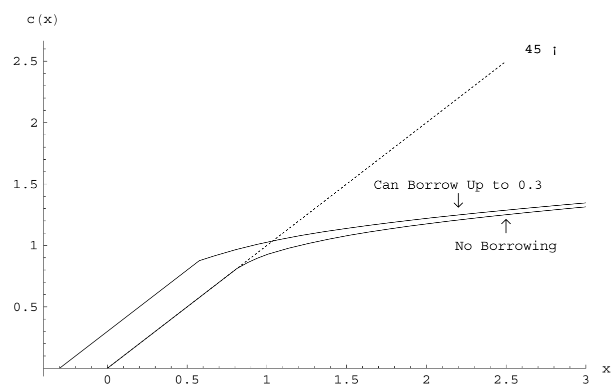

Figure 1 Strict and Looser Liquidity Constraints

Fast-forwarding to today, people working on the microfoundations of aggregate consumption behavior like Kaplan (2012) or Kaplan and Violante (2011) are increasingly taking the view that I got it right (from the standpoint of macroeconomic analysis), because if you want to study consumption then you want to study the typical *consumer* not the typical *wealth holder*. (The two will be identical only if consumption is strictly proportional to wealth, $c = \gamma x$ for some γ – which is what Kimball and I proved is NOT TRUE except in a set-of-measure-zero set of knife-edge cases).

Returning, now, to the essence of your question, which I will reframe as follows: "Suppose we calibrate a [Bewley-Carroll] model to exhibit fluctuations in unemployment risk of the magnitude that we typically see in recessions. Does that model imply that those fluctuations in risk should result in fluctuations in spending that are large enough to make an important contribution to business cycle dynamics?"

That's a question that I think has not been fully addressed by the existing literature. But there are good reasons to think that the answer is "easily."

The first important point to make here is that your quantitative intuitions are likely not to be very good on such a question, because they are likely mostly formed from linear approximations around steady states. The key insight is that the effect of uncertainty on behavior is nonlinear, and even low-probability events can have a profound effect.

This point is illustrated in figure 2, taken from Carroll and Toche (2009), which shows a comparison of the consumption function that is optimal in a perfect foresight model $\bar{c}(m^e)$ with the consumption function that is optimal when that model is modified to introduce the possibility that the consumer might become permanently unemployed.[1]

As you can see, the effect on the level of consumption is both large and nonlinear, being larger for people at lower levels of wealth (because their consumption is greater risk since they have lower resources to buffer their consumption against unempoyment).

Another piece of the answer is contained in Carroll, Slacalek, and Tokuoka (2013). That paper asks and answers two questions:

1. What modifications to a standard [Bewley/Carroll] model are needed to permit that model to do a reasonably good job of matching the distribution of wealth in the U.S. *Answer:* A modest amount of heterogeneity in the time preference rate is sufficient. Specifically, if annual time preference rates are uniformly distributed in the interval $[\beta - 0.02, \beta + 0.02]$, the model-implied distribution of wealth is remarkably close to the empirical distribution

[1] The income process for the consumer facing the risk is adjusted so that the introduction of the risk constitutes a mean-preserving spread; so, *all* of the change in the level of consumption is the effect of risk – none is attributable to the average level of income being lower for consumers subject to the unemployment risk.

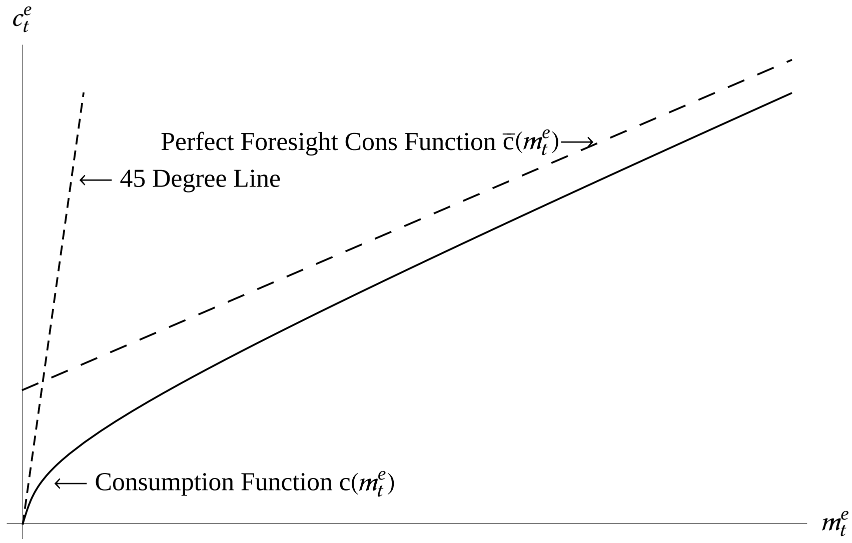

Figure 2 Consumption Function With and Without Unemployment Risk

2. Does the model that matches the distribution of wealth have different implications than one that gets the mean right *a la* Aiyagari (1994) but fails to match the distribution. *Answer:* Yes. For a version of the model that matches the distribution of total net worth, the aggregate Marginal Propensity to Consume is about 0.23 (that is, a manna-from-heaven stimulus check of $1 to every American would result in $0.23 of extra spending in the next year). For a version of the model that matches the distribution of liquid assets, the model-implied MPC is about 0.45. (The micro empirical literature on the MPC typically measures it as being somewhere in the range of 0.2 to 0.7).

A final, closely related, question is how to think about the distinction between "expenditures" (in the NIPA sense) and "consumption" (in the model). The main difference between these is in the treatment of durable goods. NIPA personal consumption expenditures include, e.g., purchases of new motor vehicles, but those are not the kinds of nondurable spending that the theoretical model is designed to model.

The beginnings of an answer are in Carroll and Dunn (1997). That paper explicitly incorporates unemployment risk into an sS model of the purchase of durable goods. The main quantitative problem with the model is that when we change the unemployment probability (in exactly the kind of Markov-switching way you describe above), the impact of a change from the low-unemployment to the high-unemployment state is much too LARGE. (That is, expenditures in the model move even more than NIPA durable goods expenditures move over the business cycle). Roughly speaking, this is because, if the economy moves from expansion to recession, almost all of the people who were just about to buy a car (they were on the margin) decide that they want a bigger precautionary buffer and so they put off their car purchase for 6-9 months. But if EVERYBODY who was about to buy a car puts it off, spending goes to zero. Of course, it would not be too hard to damp this effect by, e.g., saying that some kinds of people are in jobs where unemployment risk does not change much when a recession comes along, and some people have to buy a new car because they totalled their old one, and so on. But the takeaway is that, far from the effects of fluctuating uncertainty being too small to make much difference, the problem is the opposite.

References

Aiyagari, S. Rao (1994): "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics*, 109, 659–684.

Bewley, Truman (1977): "The Permanent Income Hypothesis: A Theoretical Formulation," *Journal of Economic Theory*, 16, 252–292.

Carroll, Christopher D. (1992): "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity*, 1992(2), 61–156, http://econ.jhu.edu/people/ccarroll/BufferStockBPEA.pdf.

——— (2001): "A Theory of the Consumption Function, With and Without Liquidity Constraints," *Journal of Economic Perspectives*, 15(3), 23–46, http://econ.jhu.edu/people/ccarroll/ATheoryv3JEP.pdf.

——— (2004): "Theoretical Foundations of Buffer Stock Saving," *NBER Working Paper No. 10867 (Status: Revise and Resubmit,* Quantitative Economics*, 2014)*, Latest version available at http://econ.jhu.edu/people/ccarroll/BufferStockTheory.pdf.

Carroll, Christopher D., and Wendy E. Dunn (1997): "Unemployment Expectations, Jumping (S,s) Triggers, and Household Balance Sheets," in *NBER Macroeconomics Annual, 1997*, ed. by Benjamin S. Bernanke, and Julio Rotemberg, pp. 165–229. MIT Press, Cambridge, MA, http://econ.jhu.edu/people/ccarroll/macroann.pdf; Methodological Appendix: http://econ.jhu.edu/people/ccarroll/methods3.pdf; Empirical Results and Simulation Programs: http://econ.jhu.edu/people/ccarroll/cdfiles.html;.

Carroll, Christopher D., and Miles S. Kimball (1996): "On the Concavity of the Consumption Function," *Econometrica*, 64(4), 981–992, http://econ.jhu.edu/people/ccarroll/concavity.pdf.

Carroll, Christopher D., Jiri Slacalek, and Kiichi Tokuoka (2013): "The Distribution of Wealth and the Marginal Propensity to Consume," mimeo, Johns Hopkins University, At http://econ.jhu.edu/people/ccarroll/papers/BSinKS.

Carroll, Christopher D., and Patrick Toche (2009): "A Tractable Model of Buffer Stock Saving," *NBER Working Paper Number 15265*, http://econ.jhu.edu/people/ccarroll/papers/ctDiscrete.

Kaplan, Greg (2012): "Inequality and the Life Cycle," *Quantitative Economics*, 3(3), 471–525.

Kaplan, Greg, and Giovanni L. Violante (2011): "A Model of the Consumption Response to Fiscal Stimulus Payments," *NBER Working Paper Number W17338*.