

Aggregate Consumption in the USA

While there is a huge literature on aggregate consumption much of which is based on US data, this document attempts a fresh and also naïve exploration of the raw data. The motivation is the sense that the literature based on testing the stochastic implications of the life cycle permanent income hypothesis (PIH) which followed from Hall's seminal work has fulfilled its potential but also left key questions unanswered. My perception is that it is generally (if not universally) agreed that the PIH is inconsistent with the data, and in particular rejected in the way one would expect if some consumers are liquidity constrained. In other words, I think it is agreed that excess sensitivity of consumption to income can't be explained by relaxing auxiliary assumptions about utility functions and, in particular, by relaxing the assumption of additive separability. Nonetheless, there is a huge amount of ongoing work with models with non liquidity constrained optimizing consumers with utility functions which are separable in consumption leisure and money holdings. The reaction to rejection of the null seems to be that it is a mistake to treat a model as a hypothesis – that models are false by definition and that the interesting question is whether the model usefully captures aspects of reality and not whether the hypothesis that it is reality can be rejected. My perception is that the criticism of testing models as hypotheses has not lead to an effort to decide if the, admittedly false, PIH model usefully captures aspects of reality. It seems to me that the critique of the application of the Neyman Person framework has not lead to enough work on alternative approaches to model evaluation and has rather been generally used to protect well liked models from data.

Let me be the first to write that I am, at least in part, demonstrating my own ignorance. Let me ignore the preceding sentence.

That typed, I propose the following question – does the permanent income model explain features of the data which would otherwise be very puzzling. Roughly, the question is whether the PIH glass is at least half full. It is very clear that most economists would agree that the answer is yes, that the PIH explains enough that it would be extremely unwise to ignore it and explore the data without the PIH in mind. There are three key striking features of the data which lead many economists to this conclusion.

First there is the fact that, as predicted by the PIH, consumption growth is smoother than GDP growth. This can be explained by arguing that aggregate permanent income is smoother than aggregate current income. Much more striking is the apparent contrast between the relationship between household consumption and household income, which appears to show an income elasticity of consumption well below one, and the aggregate time series in which the ratio of consumption to income does not trend down. This pattern can be explained if one assumes that the regression coefficient of permanent income on current income is less than one. There are many other cross sectional stylized facts which can be explained by the PIH – for example that for given household income African American households have lower consumption than White American households. These stylized facts seem to have reassured economists that inter-temporal optimizing models of consumption savings choices are useful even if the hypothesis that they are the truth is rejected by the data. This is a bit odd, since Stephen Marglin explained decades ago that the same patterns can be reconciled with complete myopia if there is habit formation or a need to learn how to consume. The argument is elementary. Habit formation causes lagged income to affect current consumption. The consumption of a myopic agent with habits depends on a moving average of current and lagged income. The consumption of a rational forward looking agent without habit formation depends on a moving average of current

and future income. The cross sectional relationship between a current income and a forward looking moving average is the same as the relationship between current income and a backward looking moving average. Similarly, habit formation can explain why the variance in the rate of growth of aggregate consumption is lower than the variance of the rate of growth of GDP. It is now agreed that it is necessary to assume habit formation in order to explain the otherwise excess smoothness of consumption. The case for adding forward looking expectations to a model which already includes habit formation receives no support at all from the stylized facts which have convinced many economists that the PIH captures useful insights.

It is easier to criticize an approach than to defend one. To decide if the PIH influenced modeling is inferior to an alternative approach, an alternative approach must be described and considered. The PIH was explicitly a response to Keynesian economics, so a natural first place to look for a competing view is The General Theory of Employment Interest and Money chapters 8 and 9. In particular, there is a clear statement of the view that there is nothing to the PIH. I quote Keynes on “The principal objective” factors which affect the propensity to consume” (that is on some things other than aggregate current income which affect aggregate consumption)

The principal objective factors which influence the propensity to consume appear to be the following: [skip]

(2) A change in the difference between income and net income. We have shown above that the amount of consumption depends on net income rather than on income [skip]

(3) Windfall changes in capital-values not allowed for in calculating net income. [skip] The consumption of the wealth-owning class may be extremely susceptible to unforeseen changes in the money-value of its wealth. [skip]

(4) Changes in the [skip] ratio of exchange between present goods and future goods. [skip]

The usual type of short-period fluctuation in the rate of interest is not likely, however, to have much direct influence on spending either way. There are not many people who will alter their way of living because the rate of interest has fallen from 5 to 4 per cent, if their aggregate income is the same as before. [skip]

(5) Changes in fiscal policy. [skip]

a change-over from a policy of Government borrowing to the opposite policy of providing sinking funds (or vice versa) is capable of causing a severe contraction (or marked expansion) of effective demand.

(6) Changes in expectations of the relation between the present and the future level of income. — We must catalogue this factor for the sake of formal completeness. But, whilst it may affect considerably a particular individual’s propensity to consume, it is likely to average out for the community as a whole. Moreover, it is a matter about which there is, as a rule, too much uncertainty for it to exert much influence.

To translate into contemporary American, Keynes does not believe in Ricardian equivalence, he expects real interest rate fluctuations to have a very small effect on consumption, he notes that consumption depends on wealth as well as on income and, especially, he thinks expectations about future aggregate income have a negligible role in determining aggregate consumption.

In contrast, models of consumption based on inter-temporal optimization can imply Ricardian equivalence and, at least, imply lower spending given “net income” (that is disposable income) if there are budget deficits. If reconciled with historical data on real interest rates and growth rates, the simplest inter-temporal models implied a large effect of fluctuations in real interest rates on the rate of growth of consumption. Current income has no special role as a source of current wealth, so even much more emphasis is placed on current financial wealth. Most of all, the relation between present and future income is critically important in such models.

On topic (4), there is no longer disagreement between Keynes and later economists. It is very clear that short term fluctuations in real interest rates have almost no relation with the rate of growth of consumption. Thus the first simple inter-temporal models have been replaced by models with habit formation. This is a point for Keynes.

Keynes discussion on points (2) and (5) can be formalized as a testable hypothesis that, given disposable personal income, public budget deficits and surpluses are not correlated with consumption. The contemporary non-Keynesian alternative is that given disposable income, higher budget deficits imply lower consumption. It is not necessary for there to be full Ricardian equivalence for Keynes's hypothesis to be demonstrably false. However, as shown below, Keynes's totally non-Ricardian null is not rejected by US data (in particular the US data available at FRED)

Keynes specifically argued only that expectations about future aggregate income had a negligible effect on aggregate consumption. The cross sectional evidence for the PIH has no bearing on Keynes's hypothesis. Rather (even ignoring the argument about habit formation above) it is an argument with a strawman Keynes who didn't write “it may affect considerably a particular individual's propensity to consume“. This means that Keynes's hypothesis can be tested only with aggregate data. Now the alternative is that expected future income affects consumption (and usually this means the objective conditional expected value). However, actual future income is correlated with the expected value of future income. The PIH can be reconciled with Keynes if one assumes that the best conditional forecast of future income has very low variance. This would be consistent with his “too much uncertainty” (although I suspect he had a different argument in mind). Insignificant correlation between achieved and expectable future income would not imply that people aren't rational intertemporal utility maximizers, but it would imply that Keynes was right.

The test is so obvious that I'm sure it has been reported in the literature (I don't know where). It is a regression of the ratio of consumption to disposable income on some ratio of future income to current income. The claim that the PIH captures important aspects of reality is the claim that the two variables are positively correlated. Below I find essentially no correlation in US data. This is the main result of this paper.

Finally note that lumped capital gains due to reinvested profits in along with other capital gains. This isn't clearly stated in “The General Theory ...” But Keynes seems to have sharply distinguished dividends (part of capital income) from reinvested profits (which might or might

not cause capital gains). Much contemporary empirical work uses GDP rather than personal income treating reinvested and distributed profits identically. This fits with the assumptions of rational consumers and value maximizing firms, since the reinvested profits belong to shareholders. I think it is fairly clear that Keynes would have preferred the older approach of relating consumption to aggregate personal income not to national income.

Keynes definitely argues that financial wealth affects consumption and suggests that fluctuations in the ratio of consumption to disposable personal income should be correlated with asset prices.

The brief reading of one passage of *The General Theory* suggests obvious empirical questions. Most importantly, does the relation between present and future aggregate income help explain the ratio of consumption to disposable personal income? There are less crucial questions. Which is more stable, the ratio of consumption to aggregate personal income or to national income? Do public sector deficits help explain the ratio of consumption to disposable personal income (contradicting Keynes)? Do asset prices help explain that ratio (as predicted by Keynes)?

Statistical Methods and Data

It may sometimes seem that the idea of this paper is to imagine someone who learned of the then new PIH from Friedman and decided to explore whether it seemed useful and to imagine that that person had access to data up to 2012 but only knew 1950s era econometrics. A justification for this approach is that the discussion began with the observation that models are not hypotheses and are false by definition, so formal tests of implications of models have limited impact. The question of whether a model is useful in understanding striking features of the data suggests looking at graphs. I couldn't resist estimating OLS regressions. The reader is asked to pay no attention to standard errors and t-like statistics. In particular, it is clear that the time series in the regression are not stationary so the t-like statistics do not have a t-distribution under any null hypothesis. The signs of the coefficients are, at least, slightly interesting. Also the R-squared and plots of fitted values give some sense as to how well crude simple models fit the data. The true explanation (excuse) is that it is very hard to look at the scatter plots without running a regression.

All data were downloaded from FRED <http://research.stlouisfed.org/fred2/>. Annual and quarterly time series were analysed. The annual data are (with FRED's code and definition)

Yr = Year AD.

rcons = PCECCA, Real Personal Consumption Expenditures (PCECCA), Billions of Chained 2009 Dollars, Annual.

Cons = PCECA; Personal Consumption Expenditures (PCECA), Billions of Dollars, Annual.

Rinc = A067RX1A020NBEA, Real disposable personal income (A067RX1A020NBEA), Billions of Chained 2009 Dollars, Annual.

inc = A067RC1A027NBEA, Disposable personal income (A067RC1A027NBEA), Billions of Dollars, Annual.

fsurplus = FYFSGDA188S, Federal Surplus or Deficit [-] as Percent of Gross Domestic Product (FYFSGDA188S), Percent of GDP, Annual.

Sp500 = SP500, S&P 500 Stock Price Index (SP500), Index, Annual

housep = USSTHPI, All-Transactions House Price Index for the United States (USSTHPI), Index 1980 Q1=100, Annual.

cpi = CPIAUCNS, Consumer Price Index for All Urban Consumers: All Items (CPIAUCNS), Index 1982-84=100, Annual.

fdebt = FYGFD, .Gross Federal Debt (FYGFD), Billions of Dollars, Annual.
GS10;10-Year Treasury Constant Maturity Rate (GS10), Percent, Annual
GS1;1-Year Treasury Constant Maturity Rate (GS1), Percent, Annual
GS5;5-Year Treasury Constant Maturity Rate (GS5), Percent, Annual
GS3;3-Year Treasury Constant Maturity Rate (GS3), Percent, Annual

From which the following variables are calculated

consinc = the ratio of real consumption to real disposable personal income.
Sp500inc = the ratio of the S&P 500 index to personal disposable income.
fdebtinc = The ration of gross federal debt to personal disposable income.
housscpi = the ratio of the all transactions house price index to the consumer price index

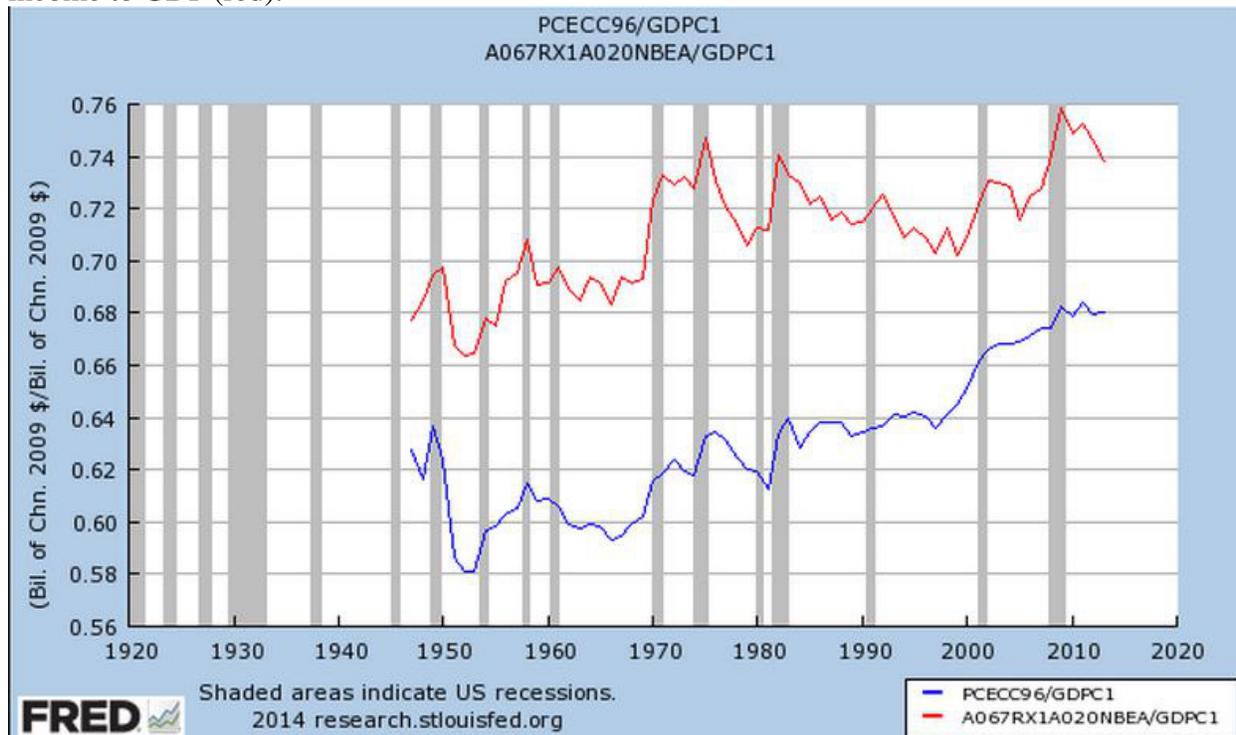
An indicator "rationing" is one for 1942, 1942, 1944, 1945 and 1946 years in at least part of which consumption was rationed.

Quarterly data (with FRED'scode and definition) are

qtr = year AD plus (quarter of the year -1)/4
cons = PCECC96;Real Personal Consumption Expenditures (PCECC96), Billions of Chained 2009 Dollars, Quarterly, Seasonally Adjusted Annual Rate.
cons = PCEC;Personal Consumption Expenditures (PCEC), Billions of Dollars, Quarterly, Seasonally Adjusted Annual Rate
rinc = DPIC96;Real Disposable Personal Income (DPIC96), Billions of Chained 2009 Dollars, Quarterly, Seasonally Adjusted Annual Rate.
inc = DPI; Disposable Personal Income (DPI), Billions of Dollars, Quarterly, Seasonally Adjusted Annual Rate.
Rgdp = GDPC1;Real Gross Domestic Product (GDPC1), Billions of Chained 2009 Dollars, Quarterly, Seasonally Adjusted Annual Rate.

Some Graphs

The fact that consumption is smoother than GDP is one of the main pieces of evidence that permanent income is a useful concept. Also just reading Keynes leads to the question of whether the ratio of consumption to aggregate personal income is more stable than the ratio of consumption to GDP. Both issues are illustrated by Figure 1 which plots annual time series of the ratio of personal consumption expenditures to GDP (blue) and of the ratio of personal disposable income to GDP (red).



The blue curve shows the familiar evidence for the usefulness of the PIH, the ratio of consumption to GDP rises during recessions. But the red curve shows a very similar pattern for the ratio of disposable personal income to GDP. The ratio rises in recessions when profits, including undistributed profits, are a very small share of GDP and public budget deficits are a large fraction of GDP. The correlation of the two graphs is striking.

In fact, for what it's worth, the correlation in quarterly data between the ratio of consumption to GDP and the ratio of disposable personal income to GDP is 0.778. In striking contrast, the correlation between the ratio of consumption to disposable personal income and the ratio of GDP to disposable personal income is actually negative -0.077. There is no sign that retained earnings or the budget surplus (or government purchases of goods and services and depreciation of capital) affect consumption.

This simple graph makes it hard to understand why time series econometricians attempt to model consumption and GDP without introducing personal disposable income. It also supports the old fashioned view that consumption is better explained with personal disposable income than with national income. The pattern of a rising ratio of consumption to GDP during recessions may be interpreted as a sign that sophisticated consumers know that recessions are temporary (at least

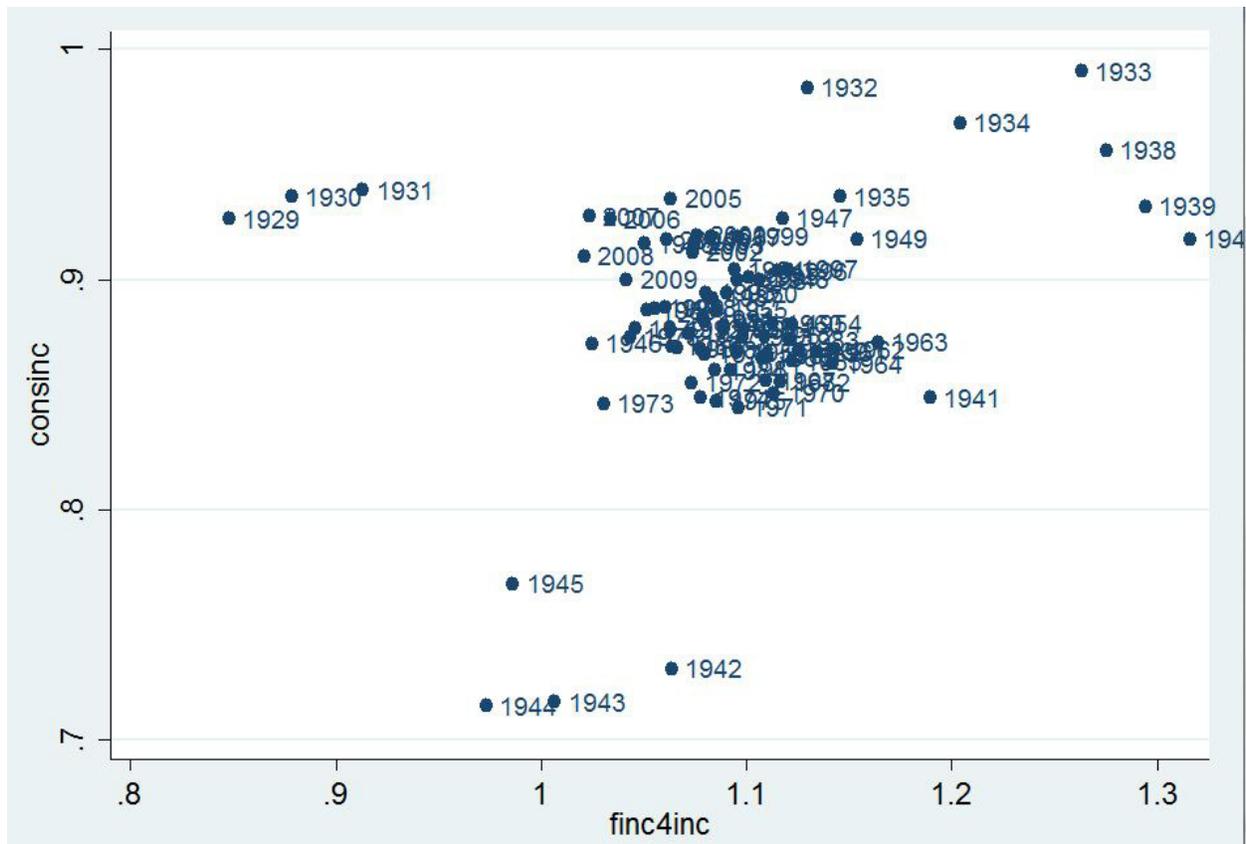
from 1933 through 2008). It could also be interpreted as a sign that extremely unsophisticated consumers don't understand that they own the firms and they or their heirs will be held responsible for the public debt. Finally, it can be interpreted as a sign that extremely unsophisticated economists insist on assuming that consumption can be understood with reference to a representative consumer when, in fact, most shares belong to people so rich that their consumption is limited by a time constraint not a budget constraint. In any case, the time series would come as no surprise to, and create no intellectual problems for, Keynes or paleo Keynesians. In contrast, astonishing ignorance of the striking correlation embarrassed at least one contemporary macro-economist (at your service).

In the simplest inter-temporal model of consumption, in which utility is logarithmic in consumption and additively separable in consumption and everything else, consumption is proportional to permanent income. Even in this case, income available to consumers is not equal to GDP, because GNP is not equal to GDP (although they are very similar for the USA), because income is net of depreciation and finally, because of government consumption. Ricardian equivalence, however, is a very simple implication of inter-temporal models without liquidity constraints. Even if some consumers are liquidity constrained, a government budget deficit should imply lower consumption than one would predict given disposable income. This implies a prediction that the ratio of the government surplus to GDP should be positively correlated with the ratio of consumption to disposable income. When the WWII years with huge deficits and rationed consumption are considered this correlation is positive. Otherwise it is tiny and negative -0.024 for annual data from 1946 through 2013 and -0.084 for 1929 through 2013 but excluding 1942, 1943, 1944 and 1945. Models with Ricardian equivalence are regularly used in spite of the complete absence of evidence for any Ricardian effect in US time series data. Again, there is no evidence against Keynes's null in the annual US data.

The main point on which Keynes disagreed with new Keynesians was his assertion that “Changes in expectations of the relation between the present and the future level of income.” don't have much influence on consumption. The PIH is essentially the assertion that such changes are critically important. It is easy to confront Keynes's guess with the data if one is willing to assume rational expectations or even that expectations are not grossly irrational. Under rational expectations future income is subjective expected future income plus an innovation which must be uncorrelated with subjective expected future income. This means that subjective expected future income must be positively correlated with achieved future income. More generally, this correlation must be positive if expectations are even close to rational expectations. Keynes must predict that the actual ratio of future to present income is not very useful in fitting consumption. It is unfair to Keynes to pretend that he asserted that future income will not add useful information to that contained in current personal disposable income and current financial wealth. After chapter 8 comes chapter 9 in which he discusses habit formation and precautionary saving among other things. In contrast, it is more than fair to Friedman to assert that he only claimed that the ratio of aggregate consumption to current aggregate personal disposable income is positively correlated with the ratio of future aggregate personal disposable income to current aggregate personal disposable income.

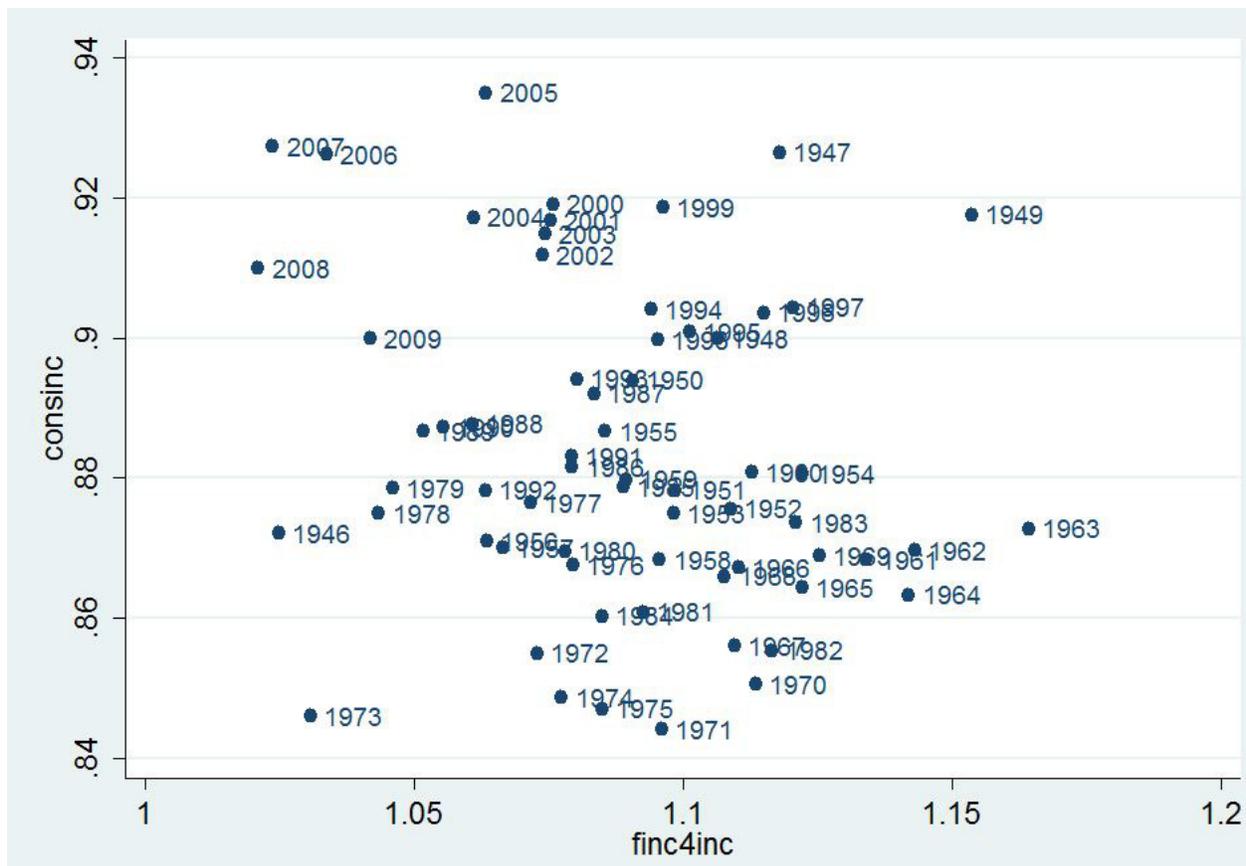
The only challenge in testing that claim is to decide how far into the future to look for “future income”. In figure 2 consinc – the ratio of real US consumption expenditures (PCECCA) to US real disposable personal income (A067RX1A020NBEA) both from Fred is plotted on finc4inc -- the ratio of the average of US real disposable personal income over the next four years (t+1,

t+2,t+3 and t+4) to current year real disposable personal income (A067RX1A020NBEA).



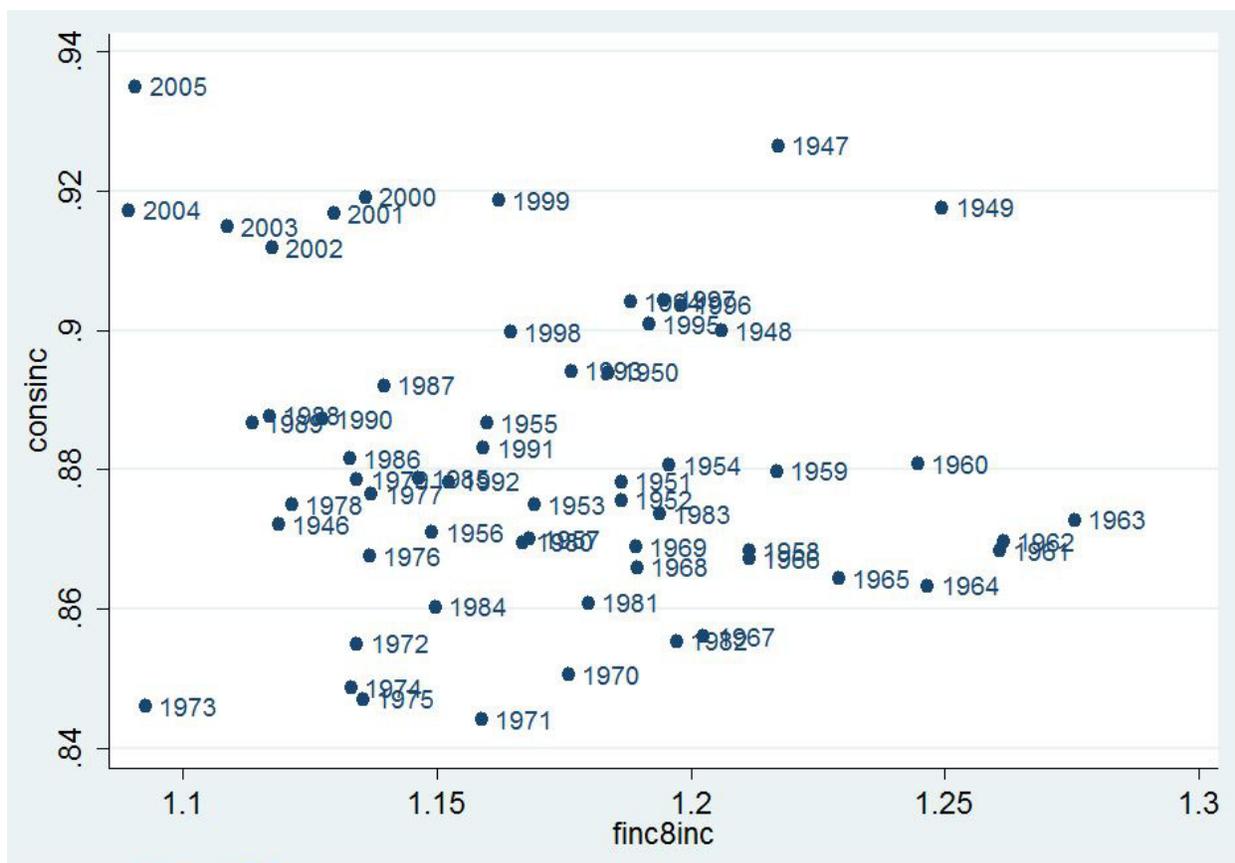
Three sets of outliers are clearly visible. Consumption was very low when it was rationed during world war II, and the ratio of future to current income was very low when Hoover was president and it was very high when Roosevelt was president. Aside from that, there is no clear pattern. Certainly the two variables are not strongly correlated.

Figure 3 shows the same variables for the post world war II sample 1946-2008. There is no sign of correlation. The figure provides no hint that any useful insights into aggregate consumption can be obtained from the PIH.



The correlation of consinc and finc4inc is actually negative being -0.207. There is no sign at all that consumption is influenced by the relationship between future and current income.

Since this is the key issue, it makes sense to consider a variety of measures of “future income.” finc8inc is the ratio of the average of real disposable personal income over the next 8 years to current real disposable personal income. The correlation of consinc and finc8inc over 1946-2004 is -0.139 – again actually negative. Figure 4, the scatter for 1946-2004 is noticeably different from the scatter based on income over the next four years, but it too shows no hint of any association of the two variables.



Summarizing the pattern more briefly, table 1 shows the correlation between the ratio of consumption to real personal disposable income (consinc) and the ratio of the average of real disposable personal income over the next N years to current real personal disposable income fincNinc for different N

Table 1

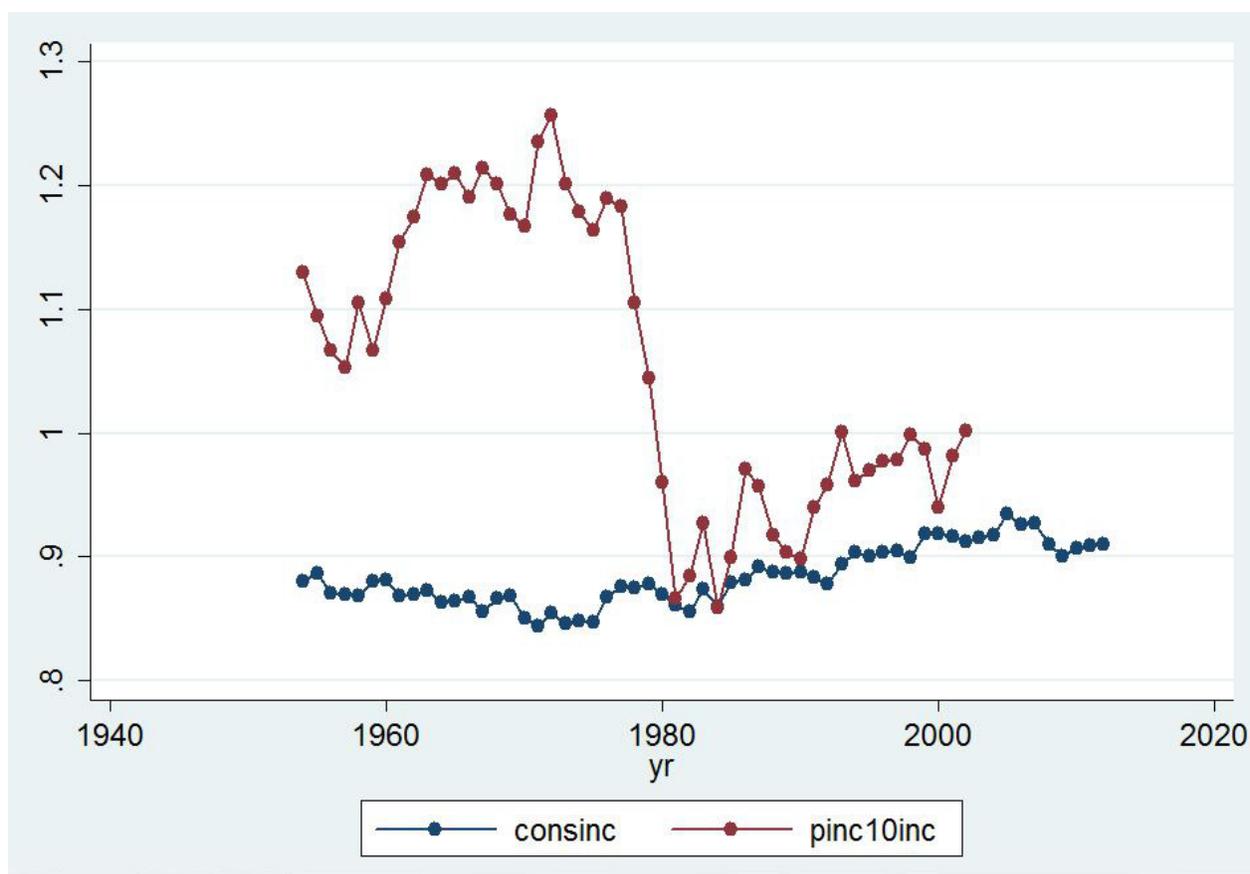
N	Correlation of consinc and fincNinc
1	-0.051
2	-0.146
3	-0.373
4	-0.207
5	-0.352
6	-0.336
7	-0.311
8	-0.139
12	-0.161
16	0.010

The absence of evidence of a positive correlation between the consumption income ratio and future income growth is very impressive.

The PIH and, in general, modern efforts to understand aggregate consumption with models of inter-temporal optimization do not imply that consumption depends on the simple average of expected future income. The importance of discounting and infinite stream of income is stressed and the discount factors play a fundamental role in the theory (even if there is notoriously little evidence in the data that they matter). There are not long time series of real interest rates, because Treasury Inflation Protected Securities were introduced relatively recently. Fortunately, if one is willing to use the achieved present discounted stream of income on the grounds that it should be correlated with the expected present discounted stream of income, then one can use future nominal income discounted with nominal interest rates. Even with this trick, FRED only has constant maturity interest rate series from 1954 on, so some data points are lost. FRED does not provide the entire yield curve. The 2,4,6,7,8 and 9 year interest rates are estimated by linear interpolation.

Of course, the forward looking calculation of the achieved present discounted value of income implies that recent data on consumption can't be used and, in any case, the infinite horizon sum of discounted income can't be calculated. I consider discounted income over the current year plus the next four, eight or ten years. Just to fit the series on the same graph, the four, eight and ten years discounted streams are divided respectively by 5, 10 and 12. So the four year achieved discounted income (pinc14) is one fifth of the sum of current nominal income plus next years income discounted plus income in 2 years discounted plus income in 3 years dicounted plus income in four years discounted.

The resulting time series have no connection with consumption or even macroeconomists's perceptions of the consumption possibilites that we learn ex post would have been available to a representative consumer with perfect foresight. Fluctuations in the time series of discounted income are dominated by fluctuations in interest rates – the sharp increase in ex post realized real interest rates in the early 80s imply a huge drop in the present discounted value of the stream of income. Figure 5 shows the time series for the of the discounted stream of nominal ncome for the current and next ten years to divided by 12 times current nominal income (pinc10inc). For comparison the ratio of real consumption to real disposable income (consinc) is plotted in figure 5.



In the 60s the combination of rapid real disposable personal income growth and low nominal interest rates imply a high present value of future income even with modest inflation. In the 70s high inflation implied low to negative ex post achieved real interest rates and a high value of discounted income in spite of low real income growth. The huge ex post achieved real interest rates of the 80s imply extremely low discounted income. Not only does the time series have no detectable correlation with consumption – it has no correlation with consumption possibilities even perceived with the benefit of hindsight.

Imagine a dialog in which Modern Macroeconomist tells Joe Consumer that, while no doubt he has rational expectations, he must have had very limited information, since with the benefit of hindsight they can see that Joe Consumer would have had much higher utility if during the 60s and 70s he had borrowed on the order of 10% to 20% of his disposable income at the treasury rate. I imagine Joe Consumer say “professor Modern Macroeconomist, you believe I could have built up a huge debt borrowing at the treasury rate. Do you also believe in the tooth fairy ?”

Figure 5 is grossly inconsistent with the model in which consumption is proportional to permanent income. That model is a very special case in with a utility function which is time separable, additively separable in consumption and everything else and logarithmic in consumption. Even with time separability, with an intertemporal elasticity of substitution less than one a higher real interest rate implies a higher ratio of consumption to permanent income. Figure 5 illustrates the small effect of interest rates on consumption which normally shows up in a regression of consumption growth on the expected real interest rate (instrumented with information available to the consumer). Figure 5 illustrates the same fact in levels.

[in principle it is possible to get a better index of permanent income using consumption in 10 years to stand for income in year 10 and later. But it's a hassle. I don't know if I want to do it.]

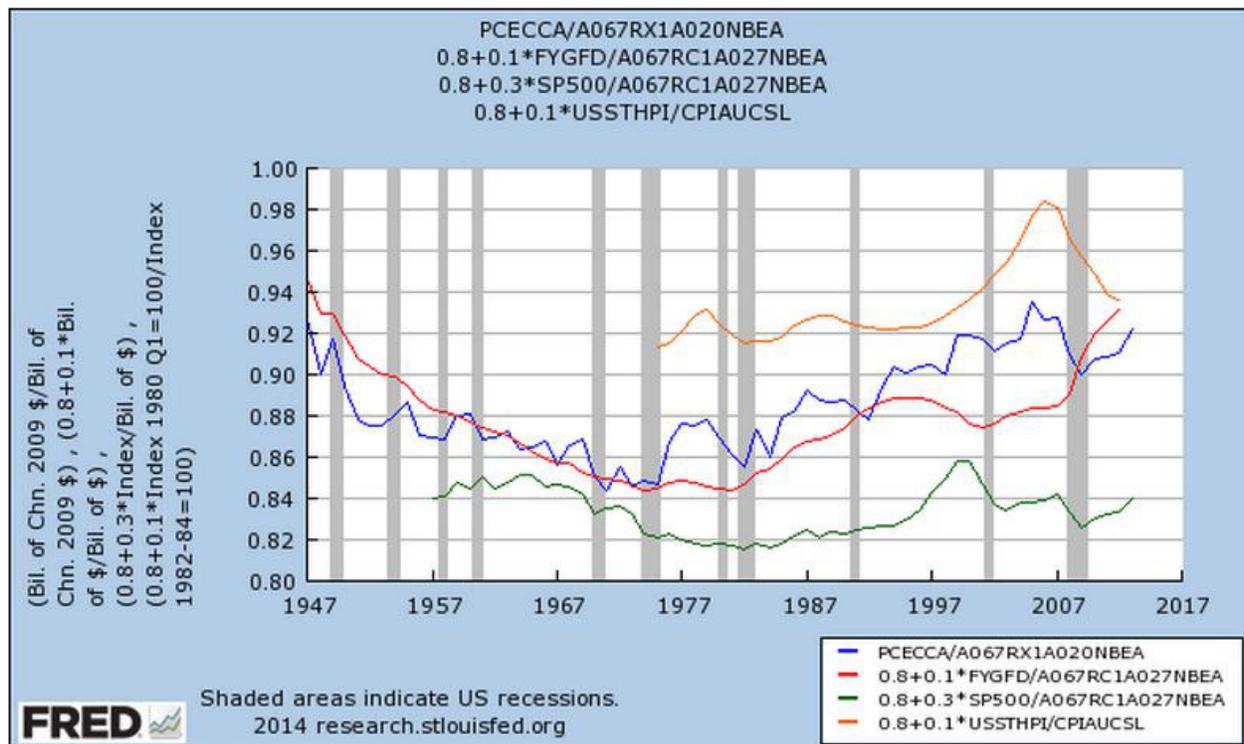
So What Does Explain the Ratio of Consumption to Disposable Personal Income ?

Wealth is a remaining candidate explanation for the ratio of consumption to disposable income. In principle financial wealth appears already in permanent income as the present value of capital income is equal to capital. However stock indexes fluctuate dramatically. The natural interpretation of this is that the subjective expected value of the discounted stream of dividends fluctuates. Someone might even once have imagined that stock indexes are a better indicator than realized income of the objective expected value of future income. However, as noted by Shiller the actual achieved present value of dividends grows almost smoothly over time and doesn't justify stock market fluctuations *ex post*. It is still possible that stock prices reflect evidently irrational fluctuations in the subjective expected growth rate of dividends. It is also possible that even people who own no stock have similar expectations for the growth of wages, since, in fact, real dividend growth and real wage growth are correlated. In any case the ratio of a stock market index to personal disposable income is a natural explanatory variable.

It is well known that a high relative price of housing is correlated with high consumption. On the face of it, this is a bit odd higher house prices imply expanded consumption possibilities for people who plan to sell their house and buy a smaller house, but also reduced consumption possibilities for those who plan to buy a first house or move to a larger house. The effect on measured consumption could be explained as due to substitution of other goods and services for services from housing, but periods of high house prices are, in fact, periods of unusually high housing investment. An explanation consistent with rationality is that high house prices relax liquidity constraints as homeowners can finance consumption with home equity loans. Finally high relative prices of housing may reflect high expected GDP growth as the relative price of land should increase as the supply of goods increases, since the supply of land is, of course, fixed. Here the variable of interest might be the ratio of the value of the housing stock to aggregate disposable income. However, the available data are house price indices (and at FRED only from 1975 on) so it is necessary to make do with the ratio of the housing price index to the consumer price index.

Finally public debt may be mistaken for net wealth. This can be due to the rational belief that the debt will not be a burden on the consumer and his or her heirs (for example because it will be discharged with taxes paid by people who have not yet immigrated). It is much more likely that people simply don't know how large public debt is or what share will be paid by them and their heirs. The fact that many people think the US budget deficit increased over a period when it roughly halved makes it hard to believe that consumers include an accurate estimate of their share of the public debt in their subjective inter-temporal budget constraint (the fact that I wrote roughly halved and not decreased from so many dollars and so many cents to some roughly half as large number demonstrates that I sure don't).

Figure 6 shows the ratio of personal consumption expenditures to personal disposable income (consinc = blue), the ratio of the gross Federal debt to personal disposable income (fdebtinc = red), the ratio of the S&P 500 index to personal disposable income (sp500inc = green), and the ratio of the all transactions house price index to the cpi (housepcip = orange) all ruthlessly scaled to fit on the same graph.



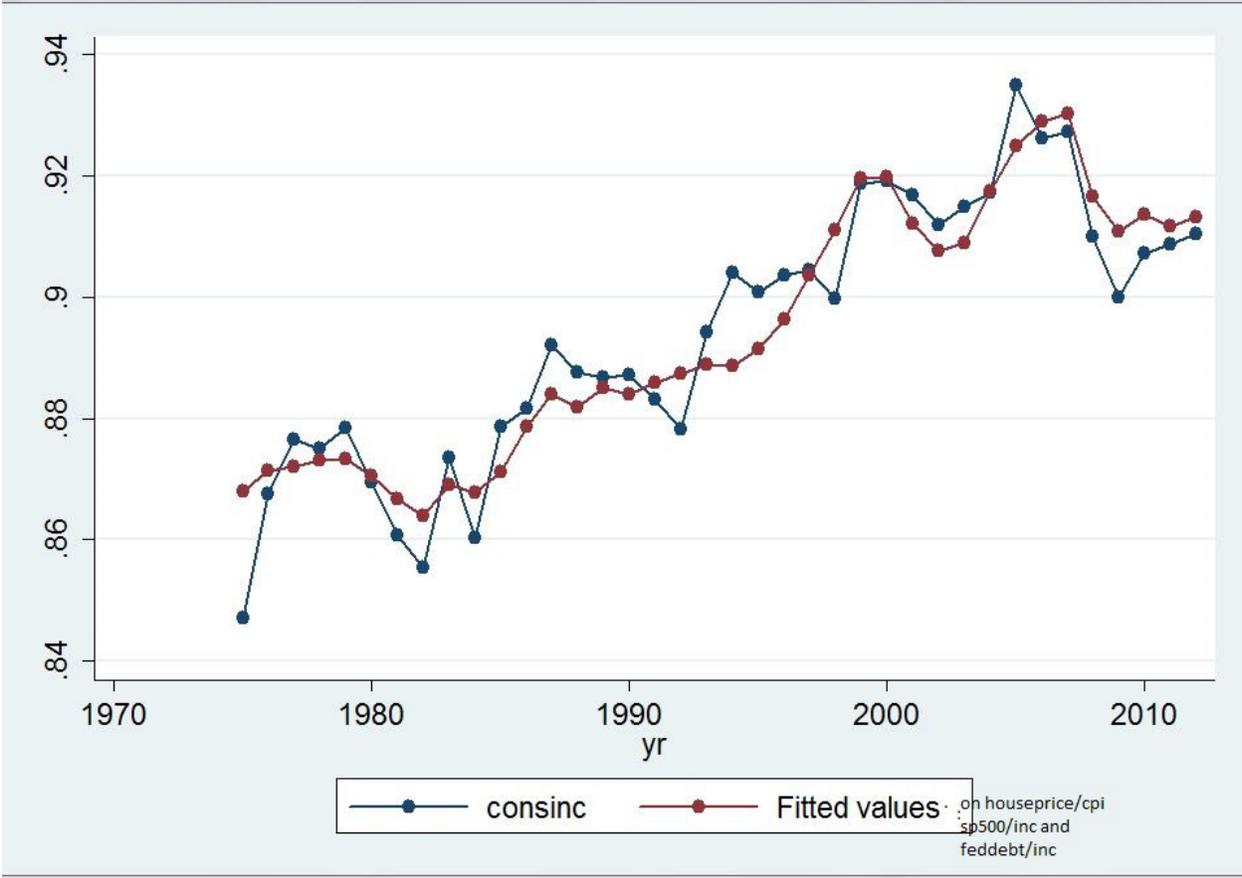
As noted at the times, in the late 1990s there were extremely high stock prices and a high ratio of consumption to disposable income, while in the early 2000s there were high house prices and a high ratio of consumption to disposable income. The similarly u shaped patterns of the ratios of gross Federal debt to GDP and consumption to disposable income have not been so widely noted. Clearly it might be a coincidence, but it is additional evidence against Ricardian equivalence.

After resisting regressing one non stationary time series on another (except for simple correlations) so long, I give in. Table 2 shows the regression of consinc on the wealth variables. The numbers under the little t are not proper t-statistics. Since the time series are clearly non stationary, they do not have a t distribution under the null that the time series are independent. The useful part of the ugly STATA output is that R-squared of over 89%. This confirms what was obvious in figure 6 – fluctuations in the ratio of consumption to personal disposable income are almost exactly fit using the three explanatory variables.

Source	SS	df	MS			
Model	.015772556	3	.005257519	Number of obs =	38	
Residual	.001912385	34	.000056247	F(3, 34) =	93.47	
				Prob > F =	0.0000	
				R-squared =	0.8919	
				Adj R-squared =	0.8823	
Total	.017684941	37	.000477971	Root MSE =	.0075	

consinc	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]	
sp500inc	.2651509	.0401713	6.60	0.000	.183513	.3467888
fdebtinc	.0266599	.0066072	4.03	0.000	.0132325	.0400873
housepci	.0486931	.0076434	6.37	0.000	.0331599	.0642264
_cons	.7821719	.0087799	89.09	0.000	.7643289	.8000149

Figure 7 plots time series of consinc and the fitted values from the regression shown in Table 2.



There is clearly almost no residual variance to be explained by interest rates or the relationship between future and current income.

Since the house price index is available only from 1975 on it is worth looking at the regression using data from 1957 through 2012 excluding housecpi. It is striking that the coefficient on sp500inc is much smaller. This might reflect increased exposure to stock market risk as pensions shift from defined benefit to 401(k) or it might reflect the invalidated of attempted inference based on regressions of one non stationary time series on another.

```
. reg consinc sp500inc fdebtinc
```

Source	SS	df	MS			
Model	.015999034	2	.007999517	Number of obs =	56	
Residual	.015174136	53	.000286304	F(2, 53) =	27.94	
Total	.03117317	55	.000566785	Prob > F =	0.0000	
				R-squared =	0.5132	
				Adj R-squared =	0.4949	
				Root MSE =	.01692	

consinc	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]	
sp500inc	.02161	.0603455	0.36	0.722	-.0994278	.1426477
fdebtinc	.0814063	.0115163	7.07	0.000	.0583076	.104505
_cons	.8234952	.0093506	88.07	0.000	.8047402	.8422503

Finally, since FRED only reports the S&P 500 from 1957 on, consider a regression of consinc on the ratio of gross federal debt to personal disposable income (fdebtinc) for the entire postwar period from 1946 through 2012

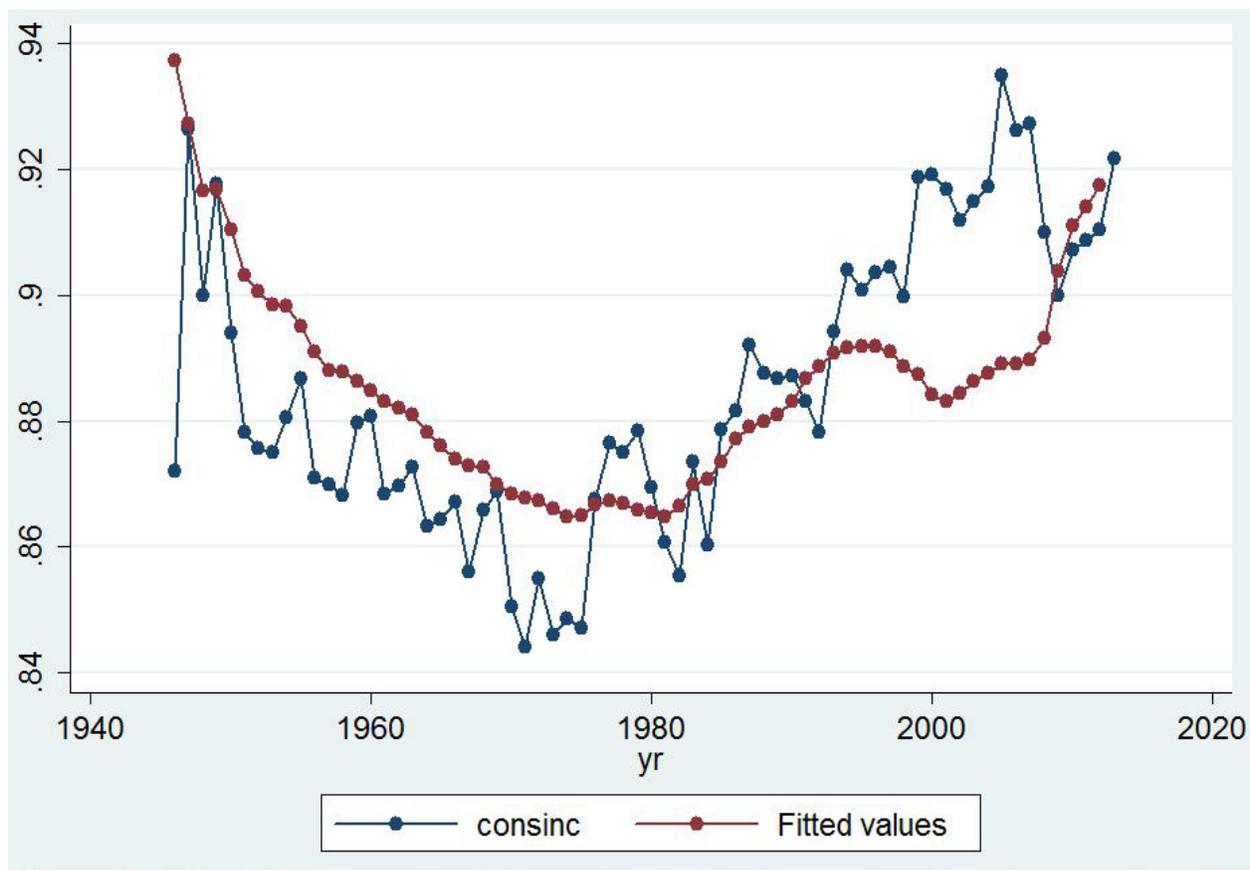
```
. reg consinc fdebtinc if yr>1945
```

Source	SS	df	MS			
Model	.01147868	1	.01147868	Number of obs =	67	
Residual	.023450201	65	.000360772	F(1, 65) =	31.82	
Total	.034928882	66	.000529225	Prob > F =	0.0000	
				R-squared =	0.3286	
				Adj R-squared =	0.3183	
				Root MSE =	.01899	

consinc	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]	
fdebtinc	.0490205	.0086906	5.64	0.000	.0316643	.0663768
_cons	.8462111	.0072345	116.97	0.000	.8317627	.8606594

The not low R-squared increases to 0.427 if 1946 is excluded. War time rationing was extended part way through 1946, which might explain the very low ratio of consumption to disposable income in spite of the huge ratio of illusory wealth to income due to the bonds sold to finance the war effort.

Figure 8 plots the fitted values (just a constant plus a constant times $fdebtinc$) and $consinc$. The fit is impressive except for 1946 as mentioned, the late 1990s when stock prices account for the high consumption and the early 2000s when high house prices account for the high consumption.



This graph should be highly embarrassing for people who believe in Ricardian equivalence (one of the main implications of the PIH). Both figure 7 and figure 8 raise the question of what feature of the data is not explained by income and wealth, that is, what remains to provoke speculation about the possible relevance of future income or interest rates to aggregate consumption.

Conclusion

I think that an economist armed only with 1950s statistical techniques but with data through 2012 (in some cases even 2013) might reasonably have told Milton Friedman that he had a nice idea and made an eloquent case, but there was (or would come to be) no sign in the data that his ideas have any practical application to macroeconomics. In contrast, Keynes's guesses from 1936 read like a good concise summary of empirical work to date, that is up until 2014 not up until 1936.