Issues in Political Economy, Vol 20, 2011, 28-44

Determinants of the U.S. Household Saving Rate: An Econometric Analysis

Jeffrey M. Gough, Bellarmine University

Saving, the act of foregoing present consumption, is vital to maintain and expand an economy's capital structure and, in turn, lay the foundation for long run growth. Imagining an economy without saving is rather bleak. Assuming no past capital accumulation, this economy—if it could be called that—would be marked by mere day to day survival and be unable to support anything close to the civilization many know today. For a nation, aggregate saving comes from several sources, but one significant source is the saving contributed by households. Because of its importance to the long run growth of the economy, this paper studies the determinants of the annual household saving rate in the United States and considers some of the policy impacts of its findings. We begin by analyzing the statistical relationships between the household saving rate and other relevant variables. Based on the results of this analysis, we discuss some potential implications for economic policy that may help to strengthen saving patterns.

I. Literature Review

Any theoretical analysis of saving is logically intertwined with consumption theory. In order to research why households save, a researcher would need to understand, at the same time, why households consume. Carroll (2001) provides a helpful overview of the various models of consumption behavior developed by economists since the 1950s. He stresses the explanatory power of Milton Friedman's permanent income hypothesis, a model credited with being one of the first and most important in modern consumption theory. Despite the greater mathematical sophistication and computing power available to later economists, Carroll argues that Friedman's original model, with its emphasis on expectations, stands up well against the conclusions of modern optimization modeling of consumer behavior. The core of Friedman's hypothesis is that consumers will change their consumption patterns only if they expect a *permanent* change in their future income. Otherwise, consumers typically seek to smooth their consumption patterns across time. While risks to income perceived as transitory should have, at most, fleeting effects, a household perceiving a substantial permanent risk to its future income is more likely to consistently decrease consumption and increase its saving rate. Therefore, we expect this latter type of risk to increase household uncertainty, induce the precautionary motive for saving and significantly explain changes in the U.S. household saving rate.

Most research also credits Modigliani and Brumberg's life cycle model as one of the pioneering works in consumption theory. The basic theory behind the model is quite intuitive. Saving patterns differ across households in different stages of their life cycle. In his Nobel prize acceptance speech Modigliani summarizes a basic implication of this model noting that since "the retirement span follows the earning span, consumption smoothing leads to a hump-shaped age path of wealth holding" (Modigliani, 1986: p. 300). In other words, households will tend to save and accumulate during their peak earning years and spend down their accumulated assets in retirement. We will attempt to test this life cycle model by measuring the effect of the retired and youth dependent populations on the U.S. household saving rate.

The closest measure of the household saving rate in the United States is a measure of the *personal* saving rate.[1] From roughly the mid 1980's through the mid 2000's the U.S. personal saving rate followed a pronounced downward trend. In their article, "The Decline in the U.S. Personal Saving Rate: Is It Real and Is It a Puzzle?," Guidolin and La Jeunesse (2007) set out to address whether or not this trend was a legitimate movement or simply an illusion created by the data.[2] To do this, they compare and contrast the two main data sets tracking the personal saving rate. One data set is provided by the Bureau of Economic Analysis (BEA) while the other is estimated from the Flow of Funds Accounts released by the Federal Reserve (Fed). Methodological imperfections aside, the authors accept that the decline in personal saving was legitimate because both the BEA and Fed measures clearly revealed the same trend. They note that any differences between the BEA and Fed data sets are effectively negligible since they are so strongly correlated (p. 503 n 18). (Because of this high correlation, we have decided to use the more popular measure reported by the BEA for our analysis.)[3] The end of the article outlines various economic theories that may help explain why the personal saving rate declined so dramatically. Some of the explanations include wealth effects, easier access to credit, Ricardian equivalence, and macroeconomic stability, all of which we will incorporate into our model.

Ewing and Payne (1998) investigate the relationship between consumer sentiment and the personal saving rate. One highlight of Ewing and Payne's article is its plotting of monthly consumer sentiment data in the same chart with the personal saving rate from 1959-1997. This visual suggests, as expected, a negative relationship between the two variables. After several econometric manipulations the authors confirm this suspicion and conclude that, although short-run results are weaker, there exists a statistically significant negative long run relationship between consumer sentiment and the personal saving rate in the U.S. This long run conclusion is derived using a method known as cointegration. In further analysis, when the authors included other variables in a regression, they found a negative relationship between changes in disposable income and the personal saving rate and a positive relationship between the interest rate and the personal saving rate. All three of these findings will be helpful to evaluate and compare against our own results.

Saving rates in other countries have been analyzed as well. Doshi (1994) incorporates cross-sectional data from 129 different countries to study the determinants of saving rates internationally. His regression model includes five independent variables: percentage of population 14 years old or younger, percentage of population 65 years old or older, average life expectancy, GNP per capita, and average growth rate of real GNP. For high-income countries, including the U.S., he found that the average growth rate of real GNP was the only variable not significant at the 10% level. All of the other four variables were significant at the 1% level. Per capita GNP was calculated to have a positive relationship with a country's saving rate while life expectancy and the population variables, meant to be proxies for young dependency and retired individuals, were found to have negative relationships with the saving rate. The major differences between this paper and Doshi's analysis are that Doshi does not report U.S.-specific information, he uses cross-sectional as opposed to time-series data, and his dependent variable is a measure of *gross* national saving, which includes business and government saving.

Of the literature consulted for this paper, Kim's article (2010) "The Determinants of Personal Saving In The U.S." is certainly the most similar in setup. Kim uses annual U.S. data from 1950 to 2007 for his econometric study. He regresses fourteen different independent variables on personal saving in a self-described “general-to-specific modeling procedure” (p. 38.) His analysis finds that personal income has a positive relationship and personal income tax has a negative relationship with personal saving, significant at the 1% level. At the 5% level, Kim finds the civilian employment-population ratio to have a positive relationship with personal saving while at the 10% level consumer credit outstanding has a negative relationship. Indeterminate relationships include the real interest rate and Social Security, as well as the old and young dependency ratios. Although the article’s main structure and focused topic is very similar to this paper a few points remain ambiguous. For instance, it is unclear whether Kim measures personal income tax as total revenue or top marginal rate or something else. Second, his varied use of the terms personal saving and personal saving *rate* leave us wondering if personal saving is measured in a total aggregate amount or as a percentage of personal disposable income. Despite the lack of clarity on these items, however, this article will serve as another useful comparison with our results.

II. Hypothesis and Initial Model Specification

We will utilize OLS linear regression techniques to investigate the statistical significance of quantitative relationships between the U.S. household saving rate (as measured by the BEA’s ‘personal’ saving rate) and several theoretically important variables. After surveying research in this area, we hypothesize the most significant variables will be household uncertainty, the annual percentage change in real personal disposable income, the real interest rate, and the retired population. We begin with a set of annual time-series data from 1964 to 2006.[4] Our initial model is specified in Model Specification *(1)* below.

Model Specification *(1)*

$$SVGRT_t = \beta_0 + \beta_1 UNCER_t + \beta_2 RINC_t + \beta_3 HHNW_t + \beta_4 RINT_t + \beta_5 CRACCESS_t + \beta_6 RETPOP_t + \beta_7 YNGD_t + \beta_8 CLGPOP_t + \beta_9 SOCSEC_t + \beta_{10} CGTAX_t + \beta_{11} INCTAX_t + \beta_{12} RGOVDEF_t + \varepsilon_t$$

Where:
SVGRT = Household Saving Rate (Annual Aggregate % of After-Tax Disposable Income Saved)
UNCER = Proxy for Household Uncertainty (Consumer Sentiment Index)
RINC = Real Personal Disposable Income (Annual Percentage Change)
HHNW = Household Net Worth (Annual Percentage Change)
RINT = Real Interest Rate (Inflation-Adjusted Interest Rate)
CRACCESS = Access to Consumer Credit (Annual Percentage Change in Consumer Credit)
RETPOP = Retired Population (% of U.S. Population Age 65+)
YNGD = Dependent Youth Population (% of U.S. Population under age 18)
CLGPOP = College Graduate Population (% of U.S. Population with 4 or more years of college)
SOCSEC = Social Security benefits proxy (Annual Percentage Change)
CGTAX = Capital Gains Tax (Annual Average Effective Rate)
INCTAX = Personal Income Tax (Annual Average Effective Rate for Median Income Household)
RGOVDEF = Real Federal Government Deficit (Millions of Inflation-Adjusted Dollars)

III. Hypothesized Effects and Coefficient Signs

The initial specification includes twelve independent variables that we broadly categorize as economic, demographic, and government policy measures.[5]

A. *Economic Measures*

We will be using consumer sentiment as a proxy for economic uncertainty. Since high economic uncertainty should be reflected via low consumer sentiment, we predict a negative relationship between our proxy for uncertainty and the personal saving rate. When the future looks grim or becomes less certain, a household's precautionary motive for saving should increase in response. We assume household net worth and credit access will have a negative relationship with the personal saving rate. Higher household net worth should induce the wealth effect, encouraging households to consume more and save less. Credit should serve as a substitute for savings to some extent. Thus, when households have easier access to credit we assume they would be less likely to save.

For the measures of real personal disposable income and the real interest rate we expect a positive relationship with the personal saving rate. With higher real interest rates, households should be induced to save more because the opportunity cost of not saving is higher. With higher real disposable income—although this is much less clear than other variables—we will assume households obtain the means to save more.

B. *Demographic Measures*

We assume all of the measures of demographics, the retired and youth population as well as the population of college graduates, will have a negative relationship with the personal saving rate. Following the life cycle model, the retired population is expected to be dissaving while young dependents should cause households to have higher consumption levels and thus lower saving rates. The population of college graduates has a much less clear effect. We assume that the expense of college serves as a substitute for saving and that student loan debt typically incurred for college attendance diminishes graduates' ability to save.

C. *Government Policy Measures*

To test the Ricardian equivalence theorem we assume households will save in anticipation of future tax increases when the federal government runs a deficit. Since higher deficits will lead to more negative values, we expect the government deficit variable to have a negative relationship with the personal saving rate. For the other government policy measures, Social Security benefits and personal income and capital gains tax, we expect a negative relationship with the personal saving rate as well. If Social Security retirement benefits increase, we expect this will indicate to the population nearing retirement that they will not have to save as much to supplement retirement income. Since personal income tax captures interest payments and the capital gains tax captures realized gains, an increase in these rates amounts to a lower net return for savers, providing a disincentive to save.

We summarize the hypothesized coefficient signs for each of the independent variables in Table 1 below.

Table 1. Hypothesized Coefficient Signs for Independent Variables

Variable	Hypothesized Sign
UNCER	(-) Negative
RINC	(+) Positive
HHNW	(-) Negative
RINT	(+) Positive
CRACCESS	(-) Negative
RETPOP	(-) Negative
YNGD	(-) Negative
CLGPOP	(-) Negative
SOCSEC	(-) Negative
CGTAX	(-) Negative
INCTAX	(-) Negative
RGOVDEF	(-) Negative

IV. Initial Results

The regression results from our initial specification can be seen in Sample Regression Line 1 below. This regression yields an adjusted R^2 of 0.943 while eight of our twelve independent variables are estimated to have coefficient signs consistent with our hypothesis. One variable is significant at 10%, UNCER, while two are significant at 5%, RINC and RETPOP, and two at a 1% level, CLGPOP and CGTAX. Perplexingly, the results for RETPOP, the retired population variable (measured by a proxy as the percentage of the U.S. population age 65 or older), indicate that a 1% increase in the retired population leads to a 2.05% *increase* in the household saving rate holding constant the other independent variables. Of course, this is the exact opposite of what the life cycle model predicts! We scrutinize the robustness of this highly unexpected result as we test for multicollinearity, superfluous variables, heteroskedasticity, and serial correlation.

Sample Regression Line 1

Predicted $SVGRT_t$ =	**12.15**	**– 0.038$UNCER_t$**	**+ 0.244$RINC_t$**	**– 0.032$HHNW_t$**
t-statistic		-1.732	2.664	-1.005
p-value		0.094	0.012	0.323
		– 0.001$RINT_t$	**– 0.037$CRACCESS_t$**	**+ 2.048$RETPOP_t$**
t-statistic		-0.016	-1.22	2.296
p-value		0.987	0.232	0.029
		– 0.338$YNGD_t$	**– 1.179$CLGPOP_t$**	**+ 0.033$SOCSEC_t$**
t-statistic		-1.667	-5.973	0.935
p-value		0.106	1.502×10^{-6}	0.357
		– 0.219$CGTAX_t$	**+ 0.106$INCTAX_t$**	**– 0.000$RGOVDEF_t$**
t-statistic		-2.847	0.489	-0.472
p-value		0.008	0.629	0.64

N = 43; Adjusted R^2 = 0.943

V. Superfluous Variable Tests

Including superfluous variables in a model tends to increase the variances of coefficient estimates. The increased variances lead to higher standard errors for the coefficients, decreasing the absolute value of t-statistics, increasing p-values, making it more likely we will incorrectly deem variables insignificant in explaining household saving rates. Thus, if a variable is superfluous we would expect its exclusion to decrease the standard errors of the remaining variables as well as increase or negligibly change our model's adjusted R^2. With this as our criteria, we conduct superfluous variable tests by running new regressions excluding a particular variable each time and analyzing its effects on the model. Each variable whose theoretical basis is border-line—should it be in or out of our model?—and whose p-value from our initial regression is high, viz. a p-value greater than or equal to 0.300, is subjected to this test. The condensed results shown in Table 2 below indicate why we deem RGOVDEF, INCTAX, SOCSEC, and HHNW to be superfluous variables.

Table 2. Results of Superfluous Variable Tests on Initial Specification

Excluded Variable	Effect on Adjusted R^2	Effect on Coefficient Standard Errors
RGOVDEF	Increase 0.002	All Decrease
INCTAX	Increase 0.001	All Decrease
SOCSEC	Approximately 0	All Decrease
HHNW	Increase 0.001	All Decrease

VI. Multicollinearity Tests

We now test our model for multicollinearity. After reviewing the correlation matrix of independent variables, we are encouraged to find none of the correlations between variables indicate perfect multicollinearity (absolute value equal to one), which would violate a classical assumption.[6] However, simple correlations of -0.97, 0.94, and -0.92 exist between our population variables RETPOP, YNGD, and CLGPOP. This indicates a high likelihood that our population variables are introducing severe multicollinearity into our model. To test for this more formally, we calculate the variance inflation factors (VIFs) for each independent variable. The resulting VIF calculations can be seen in Table 3 below.

Table 3. VIF Calculations from Initial Regression

Variable	Initial VIF	Variable	Initial VIF
UNCER	6.224	YNGD	61.295
RINC	2.463	CLGPOP	71.037
HHNW	1.718	SOCSEC	3.453
RINT	3.237	CGTAX	7.8
CRACCESS	2.011	INCTAX	12.735
RETPOP	111.351	RGOVDEF	1.924

Many researchers recommend using a VIF calculation greater than ten as a threshold for determining whether or not a model suffers from severe multicollinearity.[7] From our initial regression four of the twelve variables have VIF calculations greater than ten. The most concerning results are not surprising; RETPOP, YNGD, and CLGPOP have VIFs of 111, 61, and 71 respectively. Without correction, our model's estimates of statistical significance could be woefully unreliable.

One potential method for correcting multicollinearity is to drop superfluous variables. Thus, we revise our specification dropping the four superfluous variables identified above and calculate the VIFs for the remaining variables. Table 4 presents the results below. This method succeeds in decreasing the severity of multicollinearity, but with VIFs for our population variables still several multiples greater than ten this method is clearly insufficient.

Table 4. VIF Calculations without Superfluous Variables (w/o SVs)

Variable	VIF(Initial)	VIF(w/o SVs)	Variable	VIF(Initial)	VIF(w/o SVs)
UNCER	6.224	3.45	RETPOP	111.351	101.134
RINC	2.463	2.37	YNGD	61.295	45.633
RINT	3.237	2.345	CLGPOP	71.037	28.676
CRACCESS	2.011	1.579	CGTAX	7.8	6.529

Another way to potentially correct multicollinearity issues is to modify the model specification. After considering possible modifications, we decide to combine the variables

RETPOP and YNGD, to create a new variable RETPOP+YNGD. This combined variable represents the percentage of the U.S. population that is under age 18 or age 65 or older. In following the life cycle model we determine this to be a theoretically sound modification since the new variable will serve as a proxy for the percentage of the U.S. population that is nonworking and, thus, is unlikely to save. Consequently, we expect a negative coefficient for RETPOP+YNGD.

VII. Specification Revisions

Compelled by multicollinearity issues we revise our model's specification to Model Specification *(2)* below. The results from this new specification are summarized in Sample Regression Line 2.

Model Specification *(2)*

$$SVGRT_t = \beta_0 + \beta_1 UNCER_t + \beta_2 RINC_t + \beta_4 RINT_t + \beta_5 CRACCESS_t + \beta_6 (RETPOPt + YNGD_t) + \beta_8 CLGPOP_t + \beta_{10} CGTAX_t + \varepsilon_t$$

Sample Regression Line 2

Predicted $SVGRT_t$ =	**51.077**	**– 0.061$UNCER_t$**	**+ 0.235$RINC_t$**
t-statistic		-3.451	2.405
p-value		0.001	0.022
		+ 0.131$RINT_t$	**– 0.028$CRACCESS_t$**
t-statistic		2.151	-0.943
p-value		0.038	0.352
		– 0.731($RETPOP_t$ +$YNGD_t$)	**– 0.906$CLGPOP_t$**
t-statistic		-5.17	-9.88
p-value		9.589×10^{-6}	1.162×10^{-11}
		+ 0.005$CGTAX_t$	
t-statistic		0.151	
p-value		0.881	
N = 43; Adjusted R^2 = 0.932			

The most notable result from this new specification is that our variable RETPOP+YNGD, significant at 1%, now has a negative coefficient sign, consistent with the life cycle model. All other coefficients, with the exception of CGTAX, also have the hypothesized sign. CGTAX and CRACCESS are no longer significant at even a 20% level; however, the other five variables are significant at either 5% or 1%.

A. *Further Superfluous Variable Tests*

Using the same procedures from section V above we perform tests to identify any potentially superfluous variables in our revised specification. The results summarized in Table 6 below indicate why we also deem CGTAX and CRACCESS to be superfluous variables.

Table 5. Results of Superfluous Variable Tests on Specification *(2)*

Excluded Variable	Effect on Adjusted R^2	Effect on Coefficient Standard Errors
CGTAX	Increase 0.002	All Decrease
CRACCESS	Increase 0.006	All Decrease

Thus, we revise the model once more to arrive at our final specification expressed by Model Specification *(3)*. By dropping CGTAX we are capable of expanding our sample data set to the years 1964-2008, an increase of two observations.

Model Specification *(3)*

$$SVGRT_t = \beta_0 + \beta_1 UNCER_t + \beta_2 RINC_t + \beta_4 RINT_t + \beta_6 (RETPOPt + YNGD_t) + \beta_8 CLGPOP_t + \varepsilon_t$$

B. *Further Multicollinearity Tests*

To test our final specification for multicollinearity issues we recalculate VIFs for each of the remaining five variables. Table 6 below illustrates that we have successfully diminished the severity of multicollinearity; each variable's VIF is now less than ten. Our modifications have corrected one of the most concerning issues with our initial model.

Table 6. VIF Calculations for Final Specification

Variable	VIF	Variable	VIF
UNCER	2.105	RETPOP+YNGD	7.907
RINC	2.067	CLGPOP	7.5
RINT	1.473		

VIII. Heteroskedasticity Tests

Given the method of our variable measurements we do not anticipate (nor do any of our residual plots suggest) heteroskedasticity issues for our final specification. Nonetheless, to further evaluate the reliability of our model we formally test for heteroskedasticity by conducting the Park test. Since some of the data points for the variables RINC and RINT are negative we are only capable of conducting this test for three of our variables. As Table 7 illustrates, our variables are not significant in explaining the variability of the error term. Thus, we conclude our model has successfully avoided heteroskedasticity issues.

Table 7. Final Specification Park Test Results

Variable	t-statistic	p-value
ln(UNCER)	0.360	0.721
ln(RETPOP+YNGD)	-0.230	0.820
ln(CLGPOP)	0.754	0.455

IX. Serial Correlation Test

Since we are utilizing time-series data for our research, it is imperative to test the model for serial correlation. Beginning with a cursory observation of the graph of residuals, we are encouraged to find no obvious patterns amongst the error terms.[8] More formally, we calculate the Durbin-Watson d-statistic, testing for pure, first-order serial correlation. The results of our calculation can be seen in Table 8 below.

Table 8. Durbin-Watson d-statistic

$\sum (e[t]-e[t-1])^2$	24.91
$\sum (e[t])^2$	17.78
D-W d-statistic	**1.40**

Since our calculated d-statistic is greater than 1.29 but less than 1.78 it lies within the inconclusive range for determining whether or not the model suffers from pure, first order serial correlation (Studenmund, p. 591.)[9] Therefore, we are unable to either confirm or deny its existence by analyzing the d-statistic. In following Studenmund, who does "not recommend the application of a remedy for serial correlation if the Durbin-Watson test is inconclusive," we decide to forgo any modifications to our specification because of this result (p. 317).

X. Omitted Variable Bias Test

We will choose UNCER, the variable measuring consumer sentiment as an inverse proxy of household uncertainty, to conduct an omitted variable bias test. After omitting UNCER from our model we will determine whether or not its omission causes bias in the coefficient of RINC. We will also analyze its effect on the regression's adjusted R^2. If UNCER is indeed an important explanatory variable for our model, as we presume, its omission should decrease the adjusted R^2 as well as decrease the magnitude of RINC's coefficient.[10] As the results in Table 9 illustrate, these two results occur as anticipated. Thus, we verify UNCER as an important explanatory variable; it will remain in our model.

Table 9. Omitted Variable Bias Test - UNCER

	With UNCER	**Without UNCER**
RINC Coefficient	0.212	-0.027
Adjusted R^2	0.940	0.897

XI. Final Specification Results

Although we are unable to conclusively rule out the existence of serial correlation, the cumulative results of the tests conducted on our final specification are encouraging. We are confident that our specification revisions have helped us create a reliable model whose coefficient estimates are best, linear, unbiased estimators (BLUe). Our model excludes superfluous variables, diminishes the severity of multicollinearity, and avoids heteroskedasticity. Therefore, constructed from the annual U.S time-series data set 1964-2008, we summarize our final OLS linear regression results in Sample Regression Line 3 below.

Sample Regression Line 3

Predicted $SVGRT_t$ =	**48.033**	**– 0.070$UNCER_t$**	**+ 0.212$RINC_t$**
t-statistic		-5.46	2.427
p-value		2.906×10^{-6}	0.02
		+ 0.159$RINT_t$	**– 0.653($RETPOP_t$ + $YNGD_t$)**
t-statistic		2.985	-5.935
p-value		0.005	3.308×10^{-7}
		– 0.852$CLGPOP_t$	
t-statistic		-12.447	
p-value		3.395×10^{-16}	

N = 45; Adjusted R^2 = 0.940

XII. Analysis of Results

A. Results and Comparison with Earlier Research

All five of the variables in our final specification are found to have coefficient signs consistent with our hypothesis. Four variables, UNCER, RINT, RETPOP+YNGD, and CLGPOP, are found to be significant at a 1% level while RINC is significant at a 5% level. By way of the household net worth variable (HHNW), we are unable to observe the wealth effect from our data. In addition, compared to our original hypothesis, one of the more surprising results is the significance and magnitude of CLGPOP. If the percentage of the U.S. population that is college educated increases 1% our model predicts a 0.85% decrease in the household saving rate (holding the other variables constant.) This effect occurs in the direction we hypothesized but its magnitude is stronger than anticipated.

In terms of the variable for household uncertainty (UNCER) our coefficient estimate (-0.07) may suggest that, despite its statistical significance, it has a very minor effect on the annual saving rate. This would be too quick a conclusion, however, because the values of this variable, measured by consumer sentiment, have ranged anywhere from roughly 63.8 to 107.6 in the period analyzed. Thus, for example, if the consumer sentiment index were to increase 10 units, which is by no means an unprecedented annual fluctuation, our model predicts a 0.70% decrease in the annual saving rate (holding the other variables constant), an effect we would not classify as minor. Therefore, our results also suggest a significant effect derived from the precautionary motive for saving.

In general, our results, measured by the nonworking population variable (RETPOP+YNGD), are consistent with the life cycle model and, more specifically, are consistent with Doshi's (1994) findings on the effects of young and old population ratios. The sign and significance of household uncertainty in our research is consistent with Ewing and Payne's (1998) investigation into the relationship between the consumer sentiment index and the saving rate. Additionally, the positive relationship we found for real personal disposable income agrees with Kim (2010) although runs counter to Ewing and Payne (1998). Kim (2010) also reports significant relationships for personal income tax and credit access, two results we are unable to recreate in this study.

B. Potential Policy Implications

If policymakers consider strengthened saving patterns to be a priority for economic growth, the results of our household saving rate analysis may lead to potential policy implications. One implication stems from the strong negative correlation we find between the percentage of the population that is college educated and the saving rate. Since education can be viewed as an investment in human capital, one purported benefit of increased public expenditures on college education is its long run effect on productivity. A more educated workforce is a more productive workforce. Of course, all other things equal, this is true. What our study suggests though is that college education may be a strong substitute for saving, another important investment in (physical or financial) capital. Thus, while college education certainly has economic benefits, any thorough analysis of public education expenditures should strongly consider the offsetting opportunity costs with respect to saving.

Another potential policy implication stems from the life cycle model. Our analysis finds statistical evidence suggesting the percentage of the population that is either retired or younger than 18 has a strong negative effect on the saving rate. For policymakers looking to strengthen saving patterns, they may seek to adopt policies mitigating these demographic effects. One potential policy is to encourage younger workers to enter the job market earlier in life. Relaxation of minimum wage restrictions and weakening restrictions on various internship opportunities could diminish some of the artificial barriers precluding younger workers from entering the labor force, at least as part time labor. The immediate effects on the saving rate may be minimal, but with more young workers gaining work experience at earlier ages this could not only increase the labor force, but help a segment of the labor force become more productive and thus more capable of saving. Another policy could be to incentivize older workers to delay retirement. Lengthening the full retirement age of the Social Security program could not only strengthen the fiscal position of Social Security and bring it more in line with contemporary life expectancies but also incentivize workers considering early retirement to remain in the labor force. Additionally, another possible way to mitigate the life cycle effects on the saving rate may be to loosen immigration restrictions. Although effects of family size and foreign transfer payments must be considered, from a general perspective allowing more workers to enter the labor force from abroad could help to offset increases in the percentage of the population that is nonworking.

XIII. Conclusion

Saving is an essential element of economic productivity and growth. Since a key part of a nation's aggregate saving is contributed by households, it is valuable to investigate what determines household saving patterns and consider possible economic policy implications stemming from this investigation. With this motivation in mind, we have analyzed the quantitative relationship between several variables and a measure of the annual household saving rate in the United States. Ultimately, using a sample of time-series data from 1964-2008, we conclude that household uncertainty, the percentage of the U.S. population that is college educated, the percentage of the U.S. population that is nonworking, the annual percentage change in real personal disposable income, and the real interest rate are all statistically significant in explaining changes in the saving rate. While our results are consistent with the life cycle model and find a significant effect from precautionary motives for saving, they fail to find a significant wealth effect explaining changes in saving patterns. For policymakers these results may suggest a need for stronger consideration of the substitution effect between education and saving. Additionally, they may suggest adopting policies to mitigate the demographic effects on the saving rate by way of loosening minimum wage, internship, and immigration restrictions as well as increasing the full retirement age for Social Security.

XIV. Endnotes

[1] However, some authors, like Reinsdorf (2007), argue that stripping out the information pertaining to nonprofits serving households from the BEA rate would create a better measure of household saving.

[2] Notably, this decline in personal saving rates was not limited to the United States, e.g. Japan and Canada also witnessed drastic reductions over this same time period. See Garner (2006: pp. 8-10.)

[3] For example, Ewing and Payne (1998) and Kim (2010) both use saving data from the BEA. Garner (2006) notes the BEA's personal saving rate is the "most commonly cited measure of personal saving" (p. 8).

[4] Although measurements of several of our variables date back further, we are restricted to using 1964 as our first observation. This is the first year the U.S. Census Bureau reports consistent annual data on educational attainment levels, used to measure our variable CLGPOP.

[5] With the exception of capital gains tax, CGTAX, whose theoretical connection to the saving rate is outlined in section III-*C*, our choice of independent variables closely follows the consulted research. From Guidolin and La Jeunesse (2007) we find mention of the wealth effect and Ricardian equivalence theorem, which motivated us to create the household net worth, HHNW, and real federal government deficit, RGOVDEF, variables to measure each respectively. Ewing and Payne (1998) provide part of the impetus for including the proxy for household uncertainty, UNCER, real personal disposable income, RINC, and the real interest rate, RINT. In addition, for the variables personal income tax, INCTAX, Social Security benefits, SOCSEC, and access to consumer credit, CRACCESS, we find justification from Kim (2010). The life cycle model along with Doshi's (1994) article lent credence to analyzing the retired population, RETPOP, and the dependent youth population, YNGD. Lastly, Doshi's (1994) discussion of education as human capital investment compelled us to include the college graduate population variable, CLGPOP.

[6] The consulted correlation matrix can be seen in the Appendix, Figure A1.

[7] See Introduction to SAS. Chapter 2: Regression Diagnostics from UCLA ATS for reference to a VIF threshold greater than ten for indication of severe multicollinearity.

[8] A graph of the final specification residuals can be seen in the Appendix, Figure A2.

[9] Studenmund: p. 591. N=45, k=5: Table B-4: Critical Values of the Durbin-Watson Test Statistics dL and dU: 5-Percent One-Sided Level of Significance (10-Percent Two-Sided Level of Significance)

[10] First, the hypothesized relationship between UNCER and our dependent variable, SVGRT, is negative for the reasons explained in section III.A. Second, as real personal disposable income (RINC) increases we expect consumer sentiment to increase as well indicating a positive relationship between these two variables. Thus, if UNCER is an important explanatory variable the expected bias of its omission on RINC should be negative, (–) x (+) = (–), i.e. we expect the coefficient of RINC to decrease.

XV. Appendix

	UNCER	*RINC*	*HHNW*	*RINT*	*CRACCESS*	*RETPOP*	*YNGD*	*CLGPOP*	*SOCSEC*	*CGTAX*	*INCTAX*	*RGOVDEF*
UNCER	1.00											
RINC	0.53	1.00										
HHNW	0.03	0.05	1.00									
RINT	0.33	0.24	0.00	1.00								
CRACCESS	0.24	0.42	0.08	0.03	1.00							
RETPOP	0.19	-0.37	-0.02	0.25	-0.19	1.00						
YNGD	-0.07	0.41	-0.04	-0.24	0.17	-0.97	1.00					
CLGPOP	0.27	-0.31	-0.05	0.12	-0.22	0.94	-0.92	1.00				
SOCSEC	-0.54	0.06	0.10	-0.28	-0.07	-0.71	0.64	-0.70	1.00			
CGTAX	0.20	-0.24	-0.15	0.11	-0.11	0.67	-0.53	0.48	-0.50	1.00		
INCTAX	-0.45	-0.11	0.24	0.27	0.23	-0.15	0.09	-0.43	0.29	0.12	1.00	
RGOVDEF	0.16	0.32	-0.07	-0.03	0.27	-0.43	0.47	-0.42	0.24	-0.12	0.16	1.00

Figure A1. Correlation Matrix of Independent Variables From Initial Specification

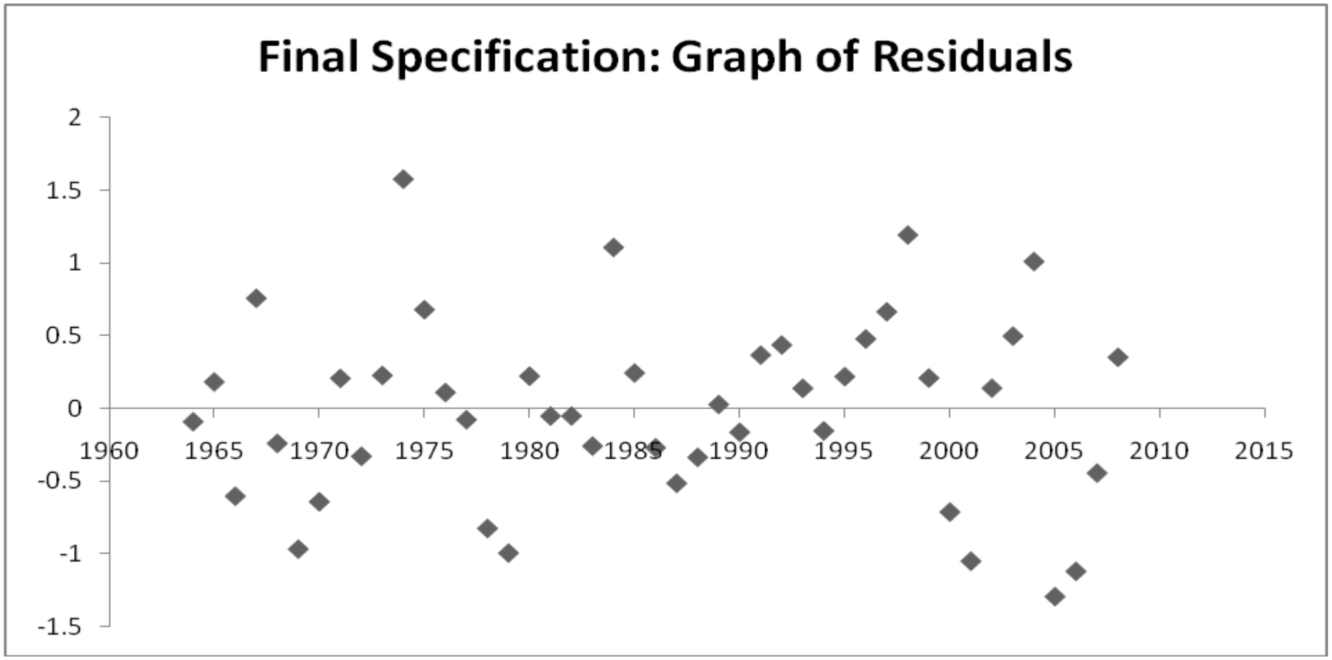

Figure A2. Graph of Residuals for Final Specification

XVI. References

Carroll, Christopher D. 2001. "A Theory of the Consumption Function, With and Without Liquidity Constraints" *Journal of Economic Perspectives*, 15(3): 23–45.

Doshi, Kokila. 1994. "Determinants of the Saving Rate: An International Comparison." *Contemporary Economic Policy*, 12(1): 37-45.

Ewing, Bradley T. and Payne, James E. 1998. "The Long-Run Relation Between The Personal Savings Rate And Consumer Sentiment." *Financial Counseling and Planning*, 9(1): 89-96.

Garner, C. Alan. 2006. "Should the Decline in the Personal Saving Rate Be a Cause for Concern?" Federal Reserve Bank of Kansas City *Economic Review*, Second Quarter Issue: 5-28.

Guidolin, Massimo and La Jeunesse, Elizabeth A. 2007. "The Decline in the U.S. Personal Saving Rate: Is It Real and Is It a Puzzle?" Federal Reserve Bank of St. Louis *Review*, 89(6): 491-514.

Introduction to SAS. Chapter 2: Regression Diagnostics. 2010. UCLA: Academic Technology Services, Statistical Consulting Group. http://www.ats.ucla.edu/stat/sas/webbooks/reg/chapter2/sasreg2.htm

Kim, Myeong Hwan. 2010. "The Determinants Of Personal Saving In The U.S." *Journal of Applied Business Research*, 26(5): 35-43.

Modigliani, Franco. 1986. "Life Cycle, Individual Thrift, and the Wealth of Nations." *American Economic Review*, 76(3): 297-313.

Reinsdorf, Marshall B. 2007. "Alternative Measures of Personal Saving." Bureau of Economic Analysis *Survey of Current Business*, 87(2): 7-13.

Studenmund, A.H. 2011. *Using Econometrics: A Practical Guide.* 6th Ed. Boston, MA: Pearson Education, Inc.