

This PDF is a selection from an out-of-print volume from the National Bureau of Economic Research

Volume Title: National Saving and Economic Performance

Volume Author/Editor: B. Douglas Bernheim and John B. Shoven, editors

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-04404-1

Volume URL: <http://www.nber.org/books/bern91-2>

Conference Date: January 6-7, 1989

Publication Date: January 1991

Chapter Title: Introduction to "National Saving and Economic Performance"

Chapter Author: B. Douglas Bernheim, John B. Shoven

Chapter URL: <http://www.nber.org/chapters/c5985>

Chapter pages in book: (p. 1 - 14)

Introduction

B. Douglas Bernheim and John B. Shoven

According to official government figures, rates of saving in the United States declined precipitously during the 1980s and are currently much lower than in any other comparable period of our history. During the second half of the last decade, net national savings amounted to a paltry 2.7 percent of net national product, compared with 7.9 percent for the 1970s and 8.6 percent for the 1960s. The United States saved a much smaller fraction of its national income than other industrialized countries throughout the postwar period. While the last decade witnessed a decline in saving throughout the developed world, the United States had the dubious distinction of leading the way.

There is widespread agreement among economists that the consequences of low saving are severe. When an individual fails to save, he jeopardizes his own economic security. Following retirement, serious illness, or involuntary job loss, he may well find that his resources are insufficient to maintain his accustomed standard of living, and at times he may experience significant hardship. Even if his luck holds out during his own lifetime, he will contribute little to the enrichment of his family line.

Inadequate rates of saving have also been blamed for a variety of chronic macroeconomic problems. When a society fails to save, each and every individual may ultimately pay the price for collective profligacy. Traditionally, many economists have been concerned about the link between saving and capital accumulation. If chronically inadequate rates of saving depress investment, then the economy must follow a growth path on which output, income, productivity, and wages are all lower than they would be in a more frugal society. Relatively recent developments have also stimulated interest in the relationships between saving and international economic transactions. Some have argued that low rates of saving have compelled the United States to sell large chunks of its capital stock to foreign investors at "fire sale" prices. Others insist that excessive consumption is primarily responsible for the

staggering U.S. trade deficits of recent years and has indirectly brought about the deteriorating competitiveness of American industry. In short, it can be argued that inadequate saving threatens the foundations of economic prosperity.

The task of restoring acceptable rates of saving in the United States poses a major challenge to those who formulate national economic policy. During the 1980s, widespread concern over various economic woes spawned a series of sporadic attempts to stimulate saving and investment through policies that were designed to make these activities more rewarding. These policies included liberalized individual retirement accounts (IRAs) and Keogh Plans, the special treatment of some reinvested dividends, reductions in capital gains taxes, and increased investment incentives at the corporate level. The results were, to say the least, disappointing. Despite the existence of ample economic incentives, rates of saving continued to decline steadily.

Disillusioned with this approach, Congress eliminated many of these special incentives in the Tax Reform Act of 1986. Unfortunately, rates of saving failed to rebound significantly in the second half of the 1980s. This has generated considerable support for the reinstatement of several special provisions, such as the favorable treatment of capital gains, that were dropped only a few years ago. The Bush administration has even proposed the creation of "family savings accounts," which would considerably broaden the current scope of IRA-like investments.

The 1980s provided a humbling experience for economists and policymakers alike. Certainly, we learned many lessons about the economy, but foremost among them was the realization that we still understand very little about the factors that motivate people to save. There are few areas within the field of economics where the need for continuing research is quite so urgent.

This volume contains papers presented at an NBER conference on saving, held in Wailea, Hawaii, on January 6 and 7, 1989. The conference was part of the NBER's Project on Saving and Investment. The goal of the conference was to further our understanding of the determinants of saving, as well as the relationships between saving and various macroeconomic aggregates. The papers in this volume may be grouped into five areas: (1) the measurement of saving, (2) the effects of corporate saving, (3) the impact of taxation on saving, (4) the relationship between saving and international capital flows, and (5) the relationship between saving and growth. The remainder of this introductory section contains a brief summary of each paper.

The Measurement of Saving

A number of economists are skeptical about the validity of the official data on saving. Some contend that international comparability is a severe problem, despite efforts by the OECD and other organizations to standardize national accounts. Others argue that arbitrary accounting conventions result in the ex-

clusion or mismeasurement of certain forms of saving. Some revisionists have even gone so far as to suggest that, once one adjusts the official figures appropriately, saving rates during the 1980s were not significantly below historical averages.

In "Market Value versus Financial Accounting Measures of National Saving," David F. Bradford argues that notions of saving based on the National Income and Product Accounts (NIPA) are seriously defective. These measures are similar in spirit to the financial accounting concepts that are used to derive the "net worth" of business enterprises. It is well known that accounting concepts provide imperfect measures of economically meaningful variables. In particular, investment (and therefore saving) in NIPA is limited to acquisitions of tangible property, and depreciation is calculated mechanically as a function of historical cost.

Bradford argues that it is more appropriate to measure saving as the change in the market value of net wealth. Certainly, expenditures on intangible assets such as advertising and R&D constitute investment, since they enhance a firm's command over future economic resources. Likewise, since individuals regard capital gains on existing assets as current income, unrealized or reinvested gains should be thought of as saving. Changes in the market valuation of an enterprise reflect both the value of intangible investments and capital revaluations. Bradford therefore concludes that, if one is interested in explaining saving behavior through conventional microeconomic theories, then a market-value notion of saving rather than the NIPA accounting measure, is appropriate.

Bradford considers several potential objections to the use of market-value accounting. Chief among them is the claim that no reliable market-value data exists. In response, he points out that the National Balance Sheets, compiled under the auspices of the Board of Governors of the Federal Reserve, are a largely underexploited resource. In particular, the National Balance Sheets carry land and corporate assets at market value and, in addition, contain information about the stock of consumer durable goods.

Measures of saving based on market-value accounting prove to be much more volatile than conventional accounting measures. Indeed, Bradford's time series for household and aggregate saving bear very little resemblance to the official numbers. He does find evidence of a long-term declining trend in the growth rate of real wealth per capita. However, his calculations also suggest that recent performance has been less disappointing than NIPA figures would lead one to believe. Indeed, the current level of wealth per capita is slightly above its long-term trend.

Bradford acknowledges that there are a variety of problems with the National Balance Sheets data and that his calculations provide imperfect measures of market-value concepts. However, he argues that these calculations represent a distinct improvement over NIPA data and an important step toward more informative measures of national saving.

The Effects of Corporate Saving

Two papers in this volume examine the question of whether the decision by corporations to retain or distribute earnings affects household-level consumption and saving. Do investors pierce the corporate veil and treat retained earnings as if they were their own income? Understanding this issue would certainly be important if one wanted to design policies to increase national saving. If the corporate veil is completely pierced, then any engineered increase in corporate saving would be completely offset at the household level. On the other hand, if it is not, aggregate saving would presumably be positively related to increases in its individual components.

In "Dividends, Capital Gains, and the Corporate Veil: Evidence from Britain, Canada, and the United States," James M. Poterba addresses the question of the corporate veil by estimating aggregate consumption functions for three countries: the United States, the United Kingdom, and Canada. If there were no corporate veil, household consumption should be invariant to firms' policies regarding the distribution of earnings (holding the earnings themselves fixed). Instead of using changes in dividend payments (which might signal improved prospects for the future and not simply a new payout policy) in his consumption regressions, Poterba uses a variable constructed to reflect the relative tax burden faced by dividends and capital gains brought about by retentions. He also examines whether involuntary capital gains resulting from corporate restructurings affect aggregate consumption and saving.

Poterba's results provide some evidence against the no-corporate-veil hypothesis. Under that hypothesis, one expects that increased dividends (holding all else equal) or increased cash received from involuntary capital gains realizations would simply be offset by the additional acquisition of financial assets by the recipients. The new payout behavior would leave consumption unchanged and would not change the composition of the assets of households. Poterba's econometric regressions suggest that increases in dividend payouts increase consumption, particularly when that variable includes the expenditures on household durables. For the United States and the United Kingdom the coefficient reflecting the impact on consumption of reducing the relative tax penalty on dividend payments is consistently positive over a fairly large number of specifications, although only occasionally statistically significantly different from zero. The point estimates are about twice as large when consumption includes the purchase of durables as they are when only nondurables and services are used as the dependent variable. This suggests that much of the additional cash received by the investors is used to purchase household durables, which contradicts the predictions of the pure life-cycle model. The point estimates indicate that nondurable consumption also increases when dividends are raised. The results for Canada show less of an impact of cash receipt on consumer behavior (in fact, the point estimate of the relative divi-

dend tax variable is negative for some specifications), so in the Canadian case the analysis is more consistent with the no-veil hypothesis.

Poterba's analysis of the effects of the involuntary capital gains realizations due to mergers and acquisitions in the United States suggests that consumer spending increases due to these payments, although once again at least half of the increase in spending takes the form of the purchase of consumer durables. For both the United States and the United Kingdom, Poterba finds that between 50 and 60 percent of the proceeds of these cash mergers are spent and not reinvested in financial markets. The implication is that there is a fairly strong "mailbox effect." That is, checks in the mailbox affect spending far more than gains that simply show up on the stock listings in the newspaper.

Poterba's results are admittedly only suggestive. However, what they suggest is that corporate financial policies may have important implications for the aggregate saving rate in the economy. More research on this topic would seem desirable, particularly a more microeconomic evaluation of investor response to cash receipts. In the meantime, policymakers should be aware that the corporate veil may not be completely transparent.

In "Corporate Savings and Shareholder Consumption," Alan J. Auerbach and Kevin Hassett also test for the presence of a corporate veil, again using aggregate macro time-series information. The authors argue that previous papers claiming to have found evidence of a corporate veil have not completely neutralized the analysis from the informational content that may be conveyed by dividend increases. The life-cycle theory does not imply that consumption should not respond to changes in dividends, but only that consumption should be unaffected by wealth-neutral changes in payout policy. They are skeptical as to whether previous investigators have appreciated this distinction and state that Poterba's results in this volume are difficult to interpret because they are unsure whether the switch to a dividend tax preference variable (rather than dividends themselves) solves the problem. Poterba offers no evidence that changes in dividend taxes are independent of wealth changes.

Auerbach and Hassett derive their econometric specification from a representative agent intertemporal utility optimization model. Unlike Poterba, they only examine consumption exclusive of the acquisition of durables. Using quarterly and annual aggregate data from 1947 to 1986, the authors estimate that expected changes in dividends, holding wealth constant, do not cause a statistically significant change in consumption in their Euler equation framework. This is taken as evidence against the presence of a corporate veil and implies that wealth-neutral changes in dividend payout behavior would not affect aggregate national saving. The authors are quick to point out, however, that their tests are not powerful enough to dispose entirely of the possible existence of a corporate veil.

Auerbach and Hassett proceed to investigate the observed excess sensitivity of consumption to short-term changes in income, particularly labor income.

They find that their evidence is consistent with a considerable fraction of the population being subject to liquidity constraints. However, they doubt that the liquidity-constrained consumption case applies to the distribution of corporate earnings because equity ownership is so concentrated among the highest income households in the country.

Finally, Auerbach and Hassett test whether the marginal propensity to consume out of corporate equity wealth is as high as the marginal propensity to consume out of other forms of wealth. If the marginal propensities are different, then a new form of corporate veil is introduced, because a shift in wealth from corporate equity to other forms will change aggregate consumption and saving. In this part of their paper, Auerbach and Hassett come up with a startling result. They find that there is no apparent effect of corporate wealth on consumption. The marginal propensity to consume out of such wealth is not significantly different from zero. This implies that transferring wealth out of corporate equities might indeed increase consumption and decrease saving. The authors concern themselves with the distribution of the holders of equities in trying to explain this result, but it clearly deserves more research and may mean that the corporate veil does exist in a somewhat different form and should be a consideration in the design of prosaving policies.

Taxation and Saving

The next group of papers looks at the effect of taxation on saving. Tax policy is one of the governmental instruments that is most often considered to alter saving behavior. A long literature exists arguing that an income tax discourages saving by taxing it twice: first, earnings that are saved are subject to tax, and second, the return on the savings is also taxed. The elimination of this double taxation is one of the appeals of saving via pension vehicles and individual retirement accounts. If the country adopted a consumption tax rather than an income tax, the double tax would be eliminated.

In "The Saving Effect of Tax-deferred Retirement Accounts: Evidence from SIPP," Steven F. Venti and David A. Wise examine a new source of data regarding households' responsiveness to the availability of tax preferred individual retirement accounts (IRAs). Using the Survey of Income and Program Participants (SIPP) panel data set, the authors document the low levels of financial assets held by most households at all ages and levels of income. The overall median level of household financial assets (including saving and checking accounts, stocks, bonds, etc.) was \$1,600 in 1985. They further show that most IRA accounts are held by households with incomes less than \$50,000 and with only modest amounts of financial assets. They find that families who contributed to IRAs after they became available in 1982 had not, prior to that time, accumulated financial assets at a rate even close to the IRA contribution ceiling. They also find no evidence that IRA contributions have been funded by borrowing. The overall conclusion from a descriptive look at

the SIPP data is that most IRA saving represented new saving and not a re-shuffling of existing accounts.

Venti and Wise propose that IRA saving and other forms of saving are not perfect substitutes, largely because of the liquidity restrictions placed on IRA accounts. They present a model that permits these two forms of saving to be treated as separate goods. The model is estimated with the SIPP data and the perfect-substitutes hypothesis is rejected. They simulate the impact of a \$1,000 increase in the IRA contribution ceiling. This affects only those who are already contributing at the maximum level. For those households, the model of Venti and Wise predicts that the \$1,000 increase in the limit would result in an average increase in IRA contributions of \$856 and an average decrease in other saving of only \$22. Immediate tax proceeds would fall by \$269, but the total effect of the increase in the ceiling would be an increase in the national saving rate.

In "Consumption Taxation in a General Equilibrium Model: How Reliable Are Simulation Results?" B. Douglas Bernheim, John Karl Scholz, and John B. Shoven seek to determine the confidence one can have in the point estimate results of computational general equilibrium models used to evaluate the impact of the U.S. switching from a personal income tax to a consumption or expenditure tax. Most general equilibrium evaluations of this issue have come to the conclusion that such a switch would increase saving and the long-run capital-labor ratio and substantially enhance economic welfare. However, these general equilibrium models require a large number of parameter valuations (particularly behavioral elasticities) that are not known with certainty. The research question that the authors address is how the uncertainty regarding input parameters translates to uncertainty about the model's predictions.

Bernheim, Scholz, and Shoven use the Fullerton-Shoven-Whalley model to examine this question. They present a technique of linearizing the model to get an approximation of the variances in the model's predictions given the variance-covariance matrix associated with the underlying parameters. The consumption tax is modeled as an income tax with complete deductibility of saving (similar to unlimited IRA accounts with no withdrawal penalty). The United States currently does not have a pure income tax, but rather a hybrid tax somewhere between an income tax and a consumption tax. This is due to the fact that at least half of saving is sheltered from taxation either through such vehicles as pension accumulations or through saving that takes the form of investment in owner-occupied housing.

The results of Bernheim, Scholz, and Shoven are mixed. It does appear that the short- and medium-run results regarding the effects of the adoption of a complete consumption tax on saving are tied down fairly precisely. One can at least rule out the possibility that the effect is zero or of the opposite sign with a reasonable degree of certainty. The predicted impact of the tax policy change on the present value of utility or economic welfare is positive, but the point estimate is between one and a half and three times the standard deviation

for this variable. The authors conclude that the results lend some support to the case for a consumption tax, but that they also emphasize the need for more precise econometric estimates of the various key elasticities in the economy (such as the saving and labor-supply elasticities) that are inputs to the general equilibrium models.

In "Taxes and Capital Formation: How Important Is Human Capital?" James Davies and John Whalley study the dynamic effects of taxes on saving and investment. Their analysis departs from most of the preceding literature on this subject by considering both human and nonhuman capital formation.

The accumulation of human capital is of enormous quantitative importance in the U.S. economy. Moreover, human and nonhuman capital may function as substitutes, at least to some extent. For example, when the returns to physical capital are taxed at higher rates, it is conceivable that individuals simply shift their resources toward education and training without actually lowering their overall levels of saving. Thus, the omission of human capital seriously limits the usefulness of many previous models that have been used to study the welfare effects of capital income taxation.

Davies and Whalley simulate the effects of various tax policies, using a fully dynamic, overlapping generations model of the U.S. economy. The structure of the model is conventional, except that the authors have incorporated a process governing the accumulation of human capital. The model is calibrated to a stylized data set that is intended to represent the position of the U.S. economy in the mid-1970s.

For their dynamic general equilibrium model, Davies and Whalley find that the inclusion of human capital increases the short-run impact of taxes on saving. However, they also show that the transition to a steady state is much more rapid than in the absence of human capital, and that there is very little distortion of human capital investment in the steady state. Consequently, they conclude that the incorporation of human capital does not significantly alter the full dynamic welfare effects of most tax reforms. This conclusion stands in sharp contrast to previous partial equilibrium results, which have suggested that the endogenization of human capital substantially increases the welfare effects of various tax policies.

National Saving and International Capital Flows

In a well-known study published in 1980, Martin Feldstein and Charles Horioka documented extremely high correlations between domestic saving rates and domestic investment for industrialized OECD countries. This finding was widely interpreted as evidence of international capital market imperfections. Although the Feldstein-Horioka analysis was subjected to a variety of criticisms, their basic finding appeared to hold up rather well.

During the 1980s, foreign nationals demonstrated a phenomenal appetite for assets in the United States. Although the United States was a net supplier

of capital to the rest of the world in the 1950s, 1960s, and 1970s, net inflows of foreign capital climbed seven-tenths of a percentage point (relative to GNP) in the early 1980s, and then shot up another two-and-a-half points in the late 1980s. At year-end 1988, the Commerce Department estimated that foreigners had accumulated nearly \$1.8 trillion worth of assets in the United States. This number exceeded the value of American-owned foreign assets by more than half a trillion dollars.

These developments imply that the correlation between domestic saving and investment noted by Feldstein and Horioka may have declined significantly during the 1980s. Anecdotal evidence suggests that this may have occurred because international capital markets became increasingly well integrated: governments reduced artificial barriers to capital flows, extensive new markets for hedging exchange-rate risks were developed, and financial institutions became increasingly sophisticated. One might therefore expect national policy regarding saving and capital income taxation to have very different effects in the 1990s than in the 1960s or 1970s. Two of the papers in this volume are concerned with evaluating the Feldstein-Horioka result and its interpretations in light of recent experience.

In "National Saving and International Investment," Martin Feldstein and Philippe Bacchetta examine correlations between domestic saving and investment for the period 1980–86, and compare these results with correlations for earlier periods. They find a substantial decline in the correlation between gross saving and gross investment and a somewhat smaller decline in the correlation between net saving and net investment. In addition, they document significant differences between EEC and non-EEC countries. Specifically, correlations between saving and investment for EEC countries have historically been lower than for non-EEC countries and declined more rapidly between the 1970s and 1980s. This evidence is consistent with the view that capital markets among the EEC countries have become highly integrated over the last decade. Even so, the impact of domestic saving on domestic investment remains substantial. The evidence indicates that, during the 1980s, a one dollar increase in domestic saving added more than 50 cents to domestic investment.

Feldstein and Bacchetta argue that high correlations between saving and investment are likely to persist even with full integration of international capital markets. With perfect integration, each investor would have to receive the same return, contracted in his domestic currency, on all equally risky investments, domestic or foreign. This need not imply equality of real *ex ante* interest rates, where returns on investments in each country are denominated in its own currency, unless expected changes in exchange rates equal the difference between expected inflation rates. Empirical violations of purchasing power parity imply that this condition is often not met. Moreover, even though purchasing power parity might hold in the long run, investors may be highly sensitive to exchange-rate fluctuations and political risks. Consequently, net

international capital flows may respond very little to apparent interest-rate differentials.

Several economists have proposed alternative explanations for the high correlation between saving and investment. It is extremely important to test the validity of these competing explanations, since each has different implications for public policy. One hypothesis, originally advanced by Maurice Obstfeld, is that the Feldstein-Horioka results are spurious and reflect the common influence of economic growth on both saving and investment. Feldstein and Bacchetta's analysis corroborates Obstfeld's claim that this is a theoretical possibility, but their empirical analysis demonstrates that the relationship between investment and saving remains equally strong even when one includes measures of growth. An alternative hypothesis, popularized by Lawrence Summers, is that governments actively seek external balance by manipulating fiscal policy. Feldstein and Bacchetta point out that Summers's evidence on the endogeneity of government deficits is also consistent with the hypothesis that deficits are exogenous, and that they crowd out private investment.

Since the Feldstein-Horioka results concern long-run correlations between saving and investment, it is still possible that international capital flows absorb a substantial fraction of short-run fluctuations in domestic saving rates. Feldstein and Bacchetta investigate this possibility by estimating dynamic adjustment processes for both saving and investment. They find that a gap between domestic saving and investment raises investment in subsequent years but leaves saving unaffected. In particular, a saving-investment gap equal to 1 percent of GDP causes the ratio of investment to GNP to rise by roughly one-quarter of a percentage point in the following year.

In "Quantifying International Capital Mobility in the 1980s," Jeffrey A. Frankel discusses four distinct definitions of perfect capital mobility: the Feldstein-Horioka definition, real interest parity, uncovered interest parity, and closed interest parity. He argues that full integration of financial markets would produce closed interest parity, but would not necessarily yield any of the other three conditions.

Closed interest parity holds when interest rates are equalized across countries for financial contracts that are written in a common currency. This condition would be violated only if there were significant barriers to the flow of financial capital across countries, such as transactions costs, information costs, capital controls (actual or potential), tax laws that discriminate by country of residence, or default risk. Frankel tests this condition by using data on forward exchange rates for 25 countries to construct covered interest rate differentials. His calculations reveal that financial markets became increasingly integrated during the 1970s. Barriers to flows of financial capital remained for a few developed countries until the late 1970s and even mid-1980s. However, by 1988, integration of financial markets had virtually eliminated all covered interest rate differentials for the major industrialized countries.

Uncovered interest parity holds when an investor who has not hedged

against exchange-rate risk expects to receive the same rate of return, denominated in his own currency, on the bonds of all countries. This requires both closed interest parity, as well as risk neutrality with respect to variations in exchange rates. Since investors are probably quite sensitive to exchange-rate risk, it would be rather surprising if this condition was satisfied in practice. More generally, uncovered interest rate differentials measure exchange-rate risk premiums. Frankel's decomposition of interest rate differentials provide a measure of these risk premiums. He finds that they have been both substantial and variable.

Real interest parity holds when real interest rates are equalized across countries. Frankel points out that this requires uncovered interest parity, plus the assumption that there can be no expected real depreciation of a country's currency. This assumption is satisfied only when goods markets are completely integrated. In practice, transportation costs for some goods are high, and many countries impose quotas and tariffs. Thus, the well-documented failure of real interest parity need not have anything to do with the efficiency of financial markets. Frankel provides new evidence on expected changes in real exchange rates and concludes that these expected changes explain a significant fraction of observed real interest rate differentials.

Finally, the Feldstein-Horioka notion of capital mobility requires real interest rate parity, plus the assumption that national saving is uncorrelated with other determinants of national investment. Frankel notes that the validity of this assumption has been disputed but argues that the Feldstein-Horioka results are nevertheless quite robust. He also updates previous estimates of the saving-investment correlation using data from the 1980s. In contrast to Feldstein and Bacchetta, he employs time-series data for the United States rather than a cross section of different countries. His results suggest that the Feldstein-Horioka result has broken down to a much greater extent than is indicated by the work of Feldstein and Bacchetta.

Taxes, Saving, and Growth

The final three papers in this volume deal with national saving and economic growth. In a closed economy, it is saving that funds and permits investment. While there is some debate about the exact magnitude of the contribution, it is universally agreed that one of the key contributors to economic growth is a rapidly growing capital stock. In such a closed economy, the question of the adequacy of saving is the same issue as the adequacy of investment. In an open economy situation, there is no reason why investment and saving in a particular country should be equated. Savers simply would look all over the world for the highest return on their investments, taking account of the various risk factors.

In "A Cross-Country Study of Growth, Saving, and Government," Robert J. Barro develops several theoretical models of the determinants of long-term

growth rates and saving rates for economies. He first presents a model in which public goods and services, jointly with private capital, determine per capita output. The model features an infinitely lived representative agent who maximizes intertemporal utility. Public goods and services are financed by proportional income taxes and the government runs a balanced budget. With some assumptions regarding technology, Barro finds that the growth rate of per capita output initially increases with increases in the level of government investment, but that at higher levels of government investment and taxes, growth is eventually retarded. Government consumption, in contrast to government investment, does not enter into production functions and definitely depresses economic growth and saving rates in his model.

Barro extends his initial model to include endogenous population growth and adds a distinction between physical and human capital. Human capital consists of two components, raw unskilled labor and accumulated human capital. Population growth is in effect a form of saving and investment, as is skill acquisition or extra human capital. Higher rates of population growth require adults to spend more time raising children, which in the model lowers the return to human capital investments.

Barro's empirical work in the paper involves the cross-country estimation of four interdependent endogenous variables. The dependent variables that he is trying to explain are per capita GDP growth, the ratio of physical investment to GDP, human capital acquisition (measured through the rate of secondary school enrollment), and population growth. The independent explanatory variables include five classes of government expenditures (measured as a ratio to GDP), a proxy for the treatment of property rights, dummies for socialist and mixed economies, and one for violent war or revolution.

Barro's results are often consistent with his theoretical models. For example, public consumption spending is systematically inversely related to growth and investment. On the other hand, public investment tends to be positively correlated with growth and investment. The results regarding property rights tend to indicate that they stimulate growth and both physical and human capital investment. Finally, there appears to be a strong negative interaction between population growth and investment in human capital. Barro refers to this as the trade-off between the quantity and the quality of children. He characterizes his paper as a progress report on a large research project to gain a better understanding of the determinants of economic growth.

In "Consumption Growth Parallels Income Growth: Some New Evidence," Christopher D. Carroll and Lawrence H. Summers also deal with cross-country issues of growth and saving, although without Barro's concentration on government expenditures. Carroll and Summers challenge the empirical validity of representative agent life-cycle theories, or the modeling of the economy as infinite horizon optimizers as in the theoretical sections of Barro's paper. They find that the aggregate consumption and saving implications of

these fashionable representations of aggregate behavior are grossly inconsistent with the features of cross-country and cross-section data on consumption and saving. The authors begin with the inference of the infinite horizon representative agent model that consumption growth should depend on the difference between the real interest rate and the rate of time preference. The growth rate of consumption should also be related to the elasticity of substitution of consumption, but the theory says that it should be independent of the rate of growth of income. When Carroll and Summers examine the cross-country data, they find that they do not conform to the predictions of the model. In particular, they observe a very high correlation between the rate of growth of consumption and the rate of growth of income. Further, the predicted positive correlation of consumption growth with real interest rates is not readily apparent.

Carroll and Summers examine several possible ways to reconcile long horizon optimization models with the international empirical evidence. One by one, the attempted reconciliations are rejected. The authors find that consumption tracks income more closely than these theories would predict, both at the aggregate level and at the individual level. The pattern of the consumption of the elderly relative to that of the young across countries also is inconsistent with what the long horizon optimization models would predict. Ultimately, the authors reject all models with lifetime or longer horizons.

Carroll and Summers conclude with two suggestions. First, they think that the evidence favors Milton Friedman's original view that permanent income should not be regarded as lifetime income (or income over even a longer horizon), but simply as the mean or expected income over a much shorter horizon, perhaps several years. They think that a model in which most households are liquidity constrained or only hold a "buffer stock" of saving may be descriptive of the real world. Second, the authors suggest that the bulk of saving in most economies is done by a very small minority of the population and that these households may behave differently than others. That is, they suggest that the consumption and saving behavior of "savers" is possibly quite different than the consumption and saving choices of the vast majority of the population. They offer this insight as a guide for additional research on this topic.

In "Saving Behavior in Ten Developing Countries," Susan Collins documents cross-country differences in rates of saving, as well as within-country trends over time for a sample of developing countries. She also attempts to explain the observed differences of saving behavior. In particular, she explores the roles of economic growth, standard of living, and the age distribution of the population.

Collins's calculations reveal that the experiences of developing countries have been quite diverse. For the period 1960–84, gross saving (expressed as a fraction of GNP) ranged from a low of 11.9 percent in Indonesia, to a high of 24.3 percent in Singapore. Countries with higher rates of saving also

tended to grow more rapidly. Her analysis of the data suggests that there is very little relationship between saving and income inequality, and that standard of living only affects the rate of saving in relatively poor countries.

One striking feature of the data is that four of the 10 countries experienced massive increases in rates of saving during the 1960s, 1970s, and 1980s. Singapore, Taiwan, Korea, and Hong Kong all eventually achieved gross saving rates in excess of one-third of GDP, despite the fact that all saved very little in the early 1960s (e.g., Hong Kong saved only 2% of GNP in 1960). This evidence calls into question the importance of "cultural" determinants of saving, since it indicates that many Asian countries achieved high rates of saving through rather dramatic behavioral changes.

Collins develops a formal model of life-cycle saving and uses it to motivate an econometric analysis of the data. She emphasizes that, according to theory, one needs to control for interactions between the various determinants of saving (such as growth, income, and age distribution). Her estimated equations bear this prediction out. Moreover, they account for a substantial portion of the differences in saving behavior both across countries and over time. On the basis of these estimates, Collins concludes that there appear to be fundamental, structural differences between the determinants of saving in low- and middle-income countries.

Finally, Collins conducts a detailed analysis of rates of saving in Korea. She argues that the trend toward collective frugality has, to a large extent, been driven by the household sector. This conclusion is supported by household survey data that decomposes saving into rural and urban components. Collins interprets the Korean experience in light of her econometric estimates. This allows her to speculate about the special determinants of rising saving in Korea.