

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

Volume Title: A Theory of the Consumption Function

Volume Author/Editor: Milton Friedman

Volume Publisher: Princeton University Press

Volume ISBN: 0-691-04182-2

Volume URL: <http://www.nber.org/books/frie57-1>

Publication Date: 1957

Chapter Title: Summary and Conclusion

Chapter Author: Milton Friedman

Chapter URL: <http://www.nber.org/chapters/c4411>

Chapter pages in book: (p. 220 - 240)

CHAPTER IX

Summary and Conclusion

THE central theme of this monograph can be illustrated by a simple hypothetical example. Consider a large number of men all earning \$100 a week and spending \$100 a week on current consumption. Let them receive their pay once a week, the pay days being staggered, so that one-seventh are paid on Sunday, one-seventh on Monday, and so on. Suppose we collected budget data for a sample of these men for one day chosen at random, defined income as cash receipts on that day, and defined consumption as cash expenditures. One-seventh of the men would be recorded as having an income of \$100, six-sevenths as having an income of zero. It may well be that the men would spend more on pay day than on other days but they would also make expenditures on other days, so we would record the one-seventh with an income of \$100 as having positive savings, the other six-sevenths as having negative savings. Consumption might appear to rise with income, but, if so, not as much as income, so that the fraction of income saved would rise with income. These results tell us nothing meaningful about consumption behavior; they simply reflect the use of inappropriate concepts of income and consumption. Men do not adapt their cash expenditures on consumption to their cash receipts, and their cash expenditures on consumption may not be a good index of the value of services consumed—in our simple example, consumption expenditures might well be zero on Sunday.

Lengthening the period of observation from a day to a week would eliminate entirely the error introduced into our simple example by the use of inappropriate concepts of income and consumption. It is the central theme of this monograph that the use of a period as long as a year does not render the error in actual data negligible, let alone eliminate it entirely. The results obtained from such annual data conform in broad outline to those of our simple example: recorded consumption is on the average positive when recorded income is zero, and the fraction of income saved rises with income. If the thesis of this monograph is correct, these results are to be explained in the same way. They too reflect the use of inappropriate concepts of income and consumption.

SUMMARY AND CONCLUSION

Our analysis accordingly distinguishes sharply between income as recorded—which we term measured income—and the income to which consumers adapt their behavior—which we term permanent income—and, similarly, between measured consumption and permanent consumption. The concept of permanent income is easy to state in these general terms, hard to define precisely. Permanent income cannot be observed directly, it must be inferred from the behavior of consumer units. And this is equally true of permanent consumption and its relation to permanent income.

The wide range of empirical material examined in this monograph turns out to be consistent with a rather simple relation between permanent consumption and permanent income suggested by purely theoretical considerations, namely, a ratio between permanent consumption and permanent income that is the same for all levels of permanent income but depends on other variables, such as the interest rate, the ratio of wealth to income, and so on. The widespread belief that the ratio of consumption to income declines as income rises can be explained entirely by the considerations stressed in our example.

Our conclusion about the meaning of permanent income cannot be stated so simply. We can think of the factors affecting the consumer's receipts as having a range of time dimensions: some factors affect his receipts only for a day, others for a week, a year, two years, and so on. We have approximated this continuum by a dichotomy. Effects lasting less than a certain time period are considered transitory, those lasting for a longer time, permanent. The length of this time period we call the consumer unit's horizon. A number of different pieces of evidence support the highly tentative conclusion that the horizon so defined is about three years.

On our interpretation of the evidence, the transitory components of a consumer unit's income have no effect on his consumption except as they are translated into effects lasting beyond his horizon. His consumption is determined by longer-range income considerations plus transitory factors affecting consumption directly. The transitory components of income show up primarily in changes in the consumer units' assets and liabilities, that is, in his measured savings.

This approach to the interpretation of consumption data and the particular hypothesis to which it has led have far-reaching implications. The rest of this chapter states the hypothesis more formally, summarizes the evidence adduced in support of it, lists generalizations about consumer behavior derived from it, and outlines some of its implications for research, economic understanding, and economic policy.

SUMMARY AND CONCLUSION

1. *Summary Statement of Hypothesis*

The permanent income hypothesis can be summarized in a system of three simple equations for the individual consumer unit:

$$\begin{aligned} (\alpha) \quad & c_p = k(i, w, u)y_p, \\ (\beta) \quad & y = y_p + y_t, \\ (\gamma) \quad & c = c_p + c_t. \end{aligned}$$

Equation (α) asserts that planned or permanent consumption (c_p) is a fraction (k) of planned or permanent income (y_p) that does not depend on the size of permanent income but does depend on other variables, in particular, the interest rate (i), the ratio of nonhuman wealth to income (w), and other factors affecting the consumer unit's tastes for current consumption versus accumulation of assets (u), such as the degree of uncertainty attached to the receipt of income, the consumer unit's age and its composition, and objective indexes of cultural factors like race or national origin. This is the simplest equation that seems consistent with the pure theory of consumer behavior as presented in Chapter II.

Equations (β) and (γ) assert that measured income (y) and measured consumption (c) can each be regarded as the sum of two components: (1) the permanent component that enters into (α), and (2) a transitory component reflecting the influence of factors regarded as chance or random by the consumer unit, as well as errors of measurement. As they stand, these equations have no substantive content; they are purely definitional.

The permanent components of income and consumption can never be observed directly for an individual consumer unit; we can only observe *ex post* what it spends and what it receives. We can, however, make inferences about the permanent components for groups of families from observed data if we accept certain assumptions about the relation between permanent and transitory components. The particular assumptions I have made are that the transitory components of consumption and income can be taken to be uncorrelated with the corresponding permanent components and with each other; these are an essential part of the hypothesis presented in this monograph. In addition, I have on occasion assumed the mean transitory components of consumption and income to be zero. This is not essential to the hypothesis and has been done only for the convenience or simplicity of the particular application.

These assumptions breathe substantive content into equations (β) and (γ). Equations (α), (β), and (γ) then imply an observed regression of measured consumption on measured income for which the ratio

SUMMARY AND CONCLUSION

of consumption to income declines as measured income increases—as in our simple introductory example and for the same reasons. They imply also a computed elasticity of measured consumption with respect to measured income that is proportional to the fraction of the total variance of income for the group concerned that is contributed by the permanent component (P_v), and a height of the regression that depends on the mean level of the permanent and transitory components of income and consumption and on the variables affecting k . The regression is shifted upward by a rise in mean permanent income and by an increase in k . Thus changes in neither the elasticity nor the height of the observed regressions need imply any changes in consumer tastes and preferences for current consumption versus accumulation of wealth, or in opportunities for exchanging the one for the other. They may instead reflect simply changes in certain characteristics of the income distribution. The appearance of changing consumer behavior may simply be a disguised reflection of the fact of changing income structure.

For simplicity of exposition, the hypothesis has been described in its arithmetic form. A variant is to retain (α) but to replace (β) and (γ) by similar expressions in the logarithms of the various terms, and to assume zero correlation between the logarithmic transitory components of income and consumption and each of these and the corresponding logarithmic permanent component. This logarithmic variant seems to fit the empirical evidence better than the arithmetic variant and is the one that has been used in most of the empirical work of the preceding chapters. Its implications are essentially the same as those of the arithmetic variant, since the one can be regarded as a first order approximation to the other, and most verbal statements of the implications apply equally to both; its advantage is that the implications hold over a wider range.

The relation between aggregate consumption and aggregate income depends not only on the consumption function for individual consumer units but also on the distribution of consumer units by the variables affecting their behavior. Under simplifying assumptions, however, the aggregate function has the same form as the individual function and can likewise be described by (α), (β), and (γ), with the exception that the variables determining the ratio of permanent consumption to permanent income (designated k^* for aggregate data) are different. They are now the distribution of consumer units by i , w , and u , or such summary measures of these distributions as their means and variances. Given the same assumptions of zero correlation between transitory and permanent components and between transitory components of consumption and income, the

SUMMARY AND CONCLUSION

hypothesis then has the same implications for the regression of consumption on income computed from aggregate data as for the regression computed from data for individual consumer units. In neither case is stability of the observed regression a necessary consequence of stability in consumer behavior with respect to current consumption and current saving.

The hypothesis has many empirical implications in addition to those already stated about the regression of measured consumption on measured income. For example, it can be used to decompose the dispersion of measured income, and also of measured consumption, into the parts attributable to transitory and permanent components. It implies that if consumer units are classified by the change in income from one year to another, the regressions of consumption on income for such groups will, under plausible conditions, be parallel and differ in height by amounts that can be specified in advance; and that the common slope will be steeper than the slope of the regression for all units combined by an amount that can be calculated from a characteristic of the income distribution for the group as a whole. It can be used to predict the correlation between the ratio of measured saving to measured income of the same units in different years. For aggregate data for a country like the United States that has been experiencing secular growth, it implies that the elasticity of consumption with respect to measured income computed from time series will be higher, the longer the period spanned by the data, and the longer the elementary time unit of observation; that it will also be higher when computed from data on aggregate consumption and income than from per capita data and when computed from data in current prices than from data in constant prices.

2. Evidence on the Acceptability of the Permanent Income Hypothesis

The implications of the permanent income hypothesis explain the major apparent anomalies that arise if the observed regression between measured consumption and measured income is interpreted, as it generally has been, as a stable relation between permanent components—though, of course, this is not the name that has been attached to the measured magnitudes. On such an interpretation the observed regression of consumption on income for a single group of consumer units implies (a) that inequality of income will increase over time—since consumption exceeds income for low income units and is less than income for high income units so apparently the poor are getting poorer and the rich richer; (b) that savings must have become an increasing fraction of income over time in the United

SUMMARY AND CONCLUSION

States and similar countries—since real income has been increasing more or less steadily; and (c) that regressions computed from budget studies made at widely spaced dates will not differ systematically. Yet there is ample evidence that (a) inequality of income has, if anything, decreased over time in the United States, (b) savings have been a roughly constant fraction of income over time in the United States, (c) computed regressions have steadily been higher, the later the date of the budget study. All three observations are entirely consistent with the permanent income hypothesis presented in this monograph.

The consistency of the hypothesis with these broad facts is only a small part of the evidence in its favor. In addition, the hypothesis is consistent with numerous detailed findings about consumption behavior that have accumulated from analyses of both budget data for individual consumer units and time series data on aggregate consumption and income; in particular, each of the implications listed in the preceding section has been compared with observation and no serious discrepancy has been found in either qualitative or quantitative elements of the implications. Perhaps the two most striking pieces of evidence for the hypothesis are, first, its success in predicting in quantitative detail the effect of classifying consumer units by the change in their measured income from one year to another; and, second, its consistency with a body of data that have not heretofore been used in analyzing consumption behavior or, indeed, even regarded as relevant to consumption behavior, namely, data on the measured income of individual consumer units in successive years. An estimate of the fraction of the variance of measured income contributed by permanent components (i.e. of P_y) can be made from such data by techniques that I developed much earlier for another purpose, namely, the analysis of the stability of relative income status. On the permanent income hypothesis the measured income elasticity of consumption is also an estimate of this same fraction.¹ These two estimates are derived from two largely independent bodies of data. Comparison of them for a variety of groups of consumer units show that they are highly correlated and approximately of the same order of magnitude.

There is some leeway in the hypothesis in the precise meaning to be assigned to the permanent component of income. The broadest definition would regard this component as attributable to any factors whose influence extends over more than one elementary time unit (a year, in most studies). Successively narrower definitions would include

¹ If computed from an arithmetically linear regression, at the mean income and for zero mean transitory components of income and consumption.

SUMMARY AND CONCLUSION

only factors affecting income in three or more years, four or more years, and so on until the narrowest definition would identify the permanent component with expected lifetime income. The comparisons mentioned in the preceding paragraph as well as a number of other pieces of evidence suggest that the empirically appropriate definition is to regard the permanent component as reflecting the influence of factors affecting income for a period of three or more years. But this must still be regarded as a highly tentative conclusion.

A number of hypotheses have been suggested in recent years to explain the contradictions mentioned above between the available evidence and the hypothesis that consumption expenditures are a stable function of absolute income. The chief such hypotheses can be regarded as special cases of the permanent income hypothesis under special conditions. This is true of the hypothesis that the ratio of consumption to income for a consumer unit depends on the relative income position of the consumer unit as measured by either the ratio of its income to the mean income of the group of which it is regarded as a member or its percentile position in the income distribution. It is equally true of the hypothesis that aggregate consumption depends not only on current aggregate income but also on the highest previous income, which has been considered a special case of the relative income hypothesis. Regarded as an alternative theory, the relative income hypothesis has fewer empirical implications than the permanent income hypothesis, so is less fruitful; in addition, such empirical evidence as I have examined, for circumstances when the implications of the two hypotheses differ, favors the permanent income hypothesis rather than the relative income hypothesis.

3. Generalizations about Consumer Behavior Based on the Hypothesis

Empirical evidence has been considered in this study primarily from the standpoint of its consistency with the permanent income hypothesis rather than of its contribution to the understanding of consumer behavior. In the process of using the evidence to test the hypothesis, however, we have necessarily been led to use the hypothesis to extract generalizations from the evidence; these are two sides of the same coin. It may illuminate these tests and this evidence if we summarize here this by-product, taking for granted that the agreement of the hypothesis with the available evidence is sufficient to justify its tentative acceptance. It should be emphasized that the generalizations that follow are all for personal consumption and personal savings; they do not cover corporate savings or governmental savings.

SUMMARY AND CONCLUSION

a. We have found no evidence of any structural change in the behavior of consumer units in the United States with respect to spending and saving over at least the past sixty years. The data for this period all conform to the pattern defined by our hypothesis, including not only the general equations (α), (β), and (γ) but also the more specific assumptions about the lack of correlation between the transitory components of income and of consumption and between each of these and the corresponding permanent component. And they conform to this pattern in the sense not only that the general functional relations apply but also that the parameters of the relations seem to have been unchanged over the period in question. A horizon of about three years seems to have characterized the outlook of consumer units, though it should be noted that the results are not very sensitive to the length of the horizon.

b. Over this period, k , the ratio of permanent consumption to permanent income has been decidedly higher for wage earners than for entrepreneurial groups; from .90 to .95 for wage earners, from .80 to .90 for entrepreneurs, probably close to the lower end of this range for nonfarm entrepreneurs and to the middle or upper end for farmers.

The difference between entrepreneurial and nonentrepreneurial groups in the size of k seems larger and better established than any other we have examined. The value of k is perhaps a trifle higher for Negroes than for whites; this difference, which is small and not well established, is opposite in direction from that which has been inferred from conventional analysis of the data. The value of k is probably higher for large than for small families, but again this cannot be regarded as well established.

c. At least part of the reason why k is lower for entrepreneurial than for nonentrepreneurial groups is the greater uncertainty of income prospects for the former, which makes the need for a reserve against emergencies greater.

For nonentrepreneurial consumer units living in urban communities of at least moderate size, the dispersion of transitory components of income is about 20 to 25 per cent of their average income; that is, about two out of three will in any year be within plus or minus 20 or 25 per cent of what they regard as their permanent position.

For nonentrepreneurial groups in small cities and villages, it seems likely that the transitory component is even less widely dispersed, though the evidence for this statement is very limited.

For entrepreneurial consumer units, whether farm or nonfarm, the relative dispersion of transitory components of income seems to be upwards of 40 per cent, perhaps as high as 50 per cent, of average

SUMMARY AND CONCLUSION

income, or something like twice as great as for nonentrepreneurial groups.

For the various nonfarm groups combined, including entrepreneurs and others, the average relative dispersion of transitory components is about 30 per cent.

d. Part of the reason why k is lower for entrepreneurial than for nonentrepreneurial groups may be the ability of the entrepreneurial group to earn a higher rate of return on accumulated capital; a similar difference between nonfarm and farm entrepreneurs may also explain why k is lower for the nonfarm entrepreneurs. However, this conclusion is highly conjectural and cannot be regarded as well established.

e. In terms of permanent income status, farmers are less dispersed than nonfarm groups as a whole, though perhaps about as dispersed as nonfarm wage and clerical workers. Nonfarm entrepreneurial groups are more widely dispersed than either farmers or other nonfarm groups. Our estimates of the size of the dispersion vary from about 60 to 70 per cent for farmers to 80 to 90 per cent for nonfarm entrepreneurial groups.

f. The variability in permanent income status has accounted for something like 80 to 85 per cent of the variability of measured annual income for broad nonfarm groups in the United States; for a much smaller fraction, for farm groups. It has apparently accounted for a larger fraction in Great Britain and Sweden than in the United States. This means that the distribution of annual income exaggerates the inequality of long-run income status by more for farmers than for nonfarmers, and by more for the United States than for Great Britain or Sweden.

g. There is limited evidence that transitory components are much less important for consumption than for income, having a relative dispersion of the order of perhaps 10 per cent instead of the 30 per cent recorded for income.

h. There is no evidence of a lag in the adjustment of consumer expenditures to changes in circumstances beyond that which is implicit in the idea that consumers adapt their expenditures to longer-run income status as measured by permanent income rather than to their momentary receipts. The effects of changes in measured income on consumer expenditures can all be accounted for in this way.

i. The ratio of aggregate consumption to aggregate income for the United States (k^*) has remained roughly constant for more than half a century at about .88 for a definition of consumption that excludes expenditures on major consumer durable goods and includes

SUMMARY AND CONCLUSION

their estimated use value. Accumulation of durables has accounted for an increasing fraction of savings, so the ratio of consumption to income would be slightly higher and would show a moderately rising secular trend for measures of consumption that treated expenditures on durables as consumption. These conclusions are supported by both budget data and time series data.

j. The constancy of the numerical value of k^* , though consistent with the permanent income hypothesis, is not required by it, even in the absence of structural change and even when the conditions are satisfied for the aggregate function to be described by equations like (α) , (β) , and (γ) . The constancy of k^* means that the variables determining k and the distribution of consumer units by these variables have been either constant or offsetting in their effects. Probably the two major offsetting forces have been (1) the declining relative importance of farming, which would tend to raise k^* , and (2) the declining size of family, which would tend to lower k^* . A third major factor, the changing role of the state in the provision of security, has itself had offsetting effects on k^* as it is measured from the available statistics.

k. Permanent income for the community as a whole can be regarded as a weighted average of current and past measured incomes, adjusted upwards by a steady secular trend and with weights declining as one goes farther back in time. The average time span between the measured incomes averaged and current permanent income is about $2\frac{1}{2}$ years.

The corresponding aggregate consumption function is

$$c^*(T) = k^* \beta \int_{-\infty}^T e^{(\beta - \alpha)(t - T)} y^*(t) dt$$

where c^* is aggregate or per capita consumption and y^* aggregate or per capita income, T designates the time unit in question, t designates time in general and is simply a variable of integration that does not appear in the final function, and k^* , α , and β are the parameters of the function. k^* is to be interpreted as the ratio of permanent consumption to permanent income, α as the secular rate of growth of income, and β as the damping coefficient which describes the process of forming estimates of expected or permanent income from current and past measured income; the higher β , the more rapidly the weights decline as one goes back in time, and the shorter the average lag between permanent income and the incomes averaged. For the period 1905 to 1951, and Raymond Goldsmith's data on deflated per capita savings and deflated per capita personal income,

SUMMARY AND CONCLUSION

the estimated values of the parameters are

$$k^* = .88$$

$$\alpha = .02$$

$$\beta = .40.$$

The values of α and β should be fairly insensitive to the precise data used, whereas k^* will be quite sensitive. The value cited is for a concept of consumption that includes only the use value of major consumer durable goods, treating accumulation in the stock of consumer durables as savings, and that includes additions to social security reserves as personal savings and income.

4. *Implications of the Hypothesis for Research*

The broader implications of acceptance of the permanent income hypothesis affect two very different areas of human effort: (1) research into consumption behavior and income structure and (2) economic understanding and policy.

A major part of the effort in consumption research, both with respect to total consumption and consumption expenditures on particular categories, has been directed toward determining the regression of consumption on income. This emphasis reflects the belief that current income is the major determinant of current consumption expenditures and that comparison of regressions is a way of eliminating the influence of income and so isolating the effects of other factors affecting consumer behavior. This partial correlation approach underlies most of the data collection and presentation; it explains alike why the Study of Consumer Purchases, perhaps the largest and most carefully planned budget study ever undertaken, collected expenditure data from a controlled rather than representative sample; and why measured income is the major, and often the only, variable used to classify consumer units in tabulations of budget data. This partial correlation approach characterizes also the bulk of the analytical research into consumer behavior, from Engel's original enunciation of his famous laws which led to his name being attached to regressions of consumption on income, to current self-consciously complex econometric research; from the examination of data for a small group of consumer units, to the calculation of demand functions from a combination of time series and budget data for a nation as a whole and for many separate commodities. Sophistication has taken the form of adding more and more variables, and of using more refined statistical techniques to estimate their effects and to allow for sampling and measurement errors; it has not changed the basic orientation or direction of the research.

SUMMARY AND CONCLUSION

Acceptance of the permanent income hypothesis implies that much or most of this research has been misdirected. What has been held constant is not income in the sense that is relevant to consumption behavior but a more or less arbitrary mixture of income in this sense and accidental elements. Statistical measures that have been taken to reflect the influence of differences in consumer behavior in fact reflect features of the income distribution. In consequence, the discovery of regularities in consumer behavior has been bedeviled by the confounding of such behavior with unrelated though not irrelevant features of the distribution of income. The result has been the introduction of increasing complexity into the analysis in an effort to rationalize the data within the same general framework.

This complexity is frequently pointed to with pride by workers in the field as evidence of the subtlety of their analysis. It is tempting to make a virtue of necessity by asserting that the consumer *is* a complex creature who is influenced by everything under the sun and hence that only an analysis in terms of a large number of variables can hope to extract a consistent pattern from his behavior. In fact, the necessity of introducing many variables is a sign of defeat and not of success; it means that the analyst has not found a truly fruitful way of interpreting or understanding his subject matter; for the essence of such a fruitful theory is that it is simple. The consumption analyst, as it were, has been priding himself on his success in adding yet more epicycles. The possibility of dispensing with these does not, of course, mean that his empirical findings are in error, that the variables he finds related to consumer behavior are not related to it, any more than acceptance of the Copernican view rendered non-existent the astronomical movements that it was necessary to introduce additional epicycles to explain. What it does mean is that these empirical relations can all be inferred from a much simpler structure, that they can all be regarded as manifestations in different guise of a single and simpler set of forces rather than as the result of largely irreducible ultimate variables.

Acceptance of the permanent income hypothesis means that much less emphasis should be attached to the regressions of consumption on income, especially in the analysis of total consumption and savings. The principal task in this area at the present stage of knowledge is to find the major determinants of k and to measure their influence. The data needed for this purpose, at least for the first attack on the problem, are average consumption and average income for groups of consumer units for which transitory components can largely be expected to average out; for example, communities, or moderately homogeneous occupational groups within cities. I know, myself, of

SUMMARY AND CONCLUSION

only one study, by Dorothy Brady, which has used the community as the unit of observation and has dealt with relations among averages.² We need to determine whether and how the ratio of such averages, which we have been calling the average propensity to consume, is connected with the variables our hypothesis leads us to put into the forefront: the rate of interest, the relative dispersion of transitory components of income and of consumption, the ratio of wealth to income, the age and composition of consumer units. Much can be done along these lines with existing data, though thanks to their different orientation, community averages are either not available for many such data or can be computed only with difficulty. With respect to future collection of data, acceptance of this approach in many ways simplifies the problem, by enabling emphasis to be put almost entirely on samples giving good estimates of means. This can be done with a smaller sample and perhaps a simpler sampling design than is required to get good estimates of multivariate relations including current income of the consumer unit as a major variable.

From the point of view of the permanent income hypothesis, the regression of consumption on income not only should receive much less exclusive emphasis in consumption research than it has heretofore, it also serves a different function and use. Its function is primarily to provide a means to decompose the total variation in income into the parts contributed by permanent and transitory components. The result is useful for consumption research in providing an estimate of one variable that may be expected to influence k , namely, the relative dispersion of transitory components. Its main use, however, is not in consumption research at all but in analyzing the distribution of income. It enables the mass of consumer budget data to be used to interpret data on the distribution of income and to convert them into estimates of the distribution of permanent income status.

Curiously enough, while the hypothesis converts the regression of consumption on income into a tool for analyzing income distribution, it gives the regression of income on consumption, which has heretofore been almost entirely neglected, significance for consumption research. For this regression enables us to decompose the total variation in consumption into the parts contributed by permanent and transitory components and so to estimate the relative dispersion of transitory components of consumption. The regression of income on consumption needs to be computed from any existing data for which it is possible to do so, and in future studies, the two regressions should be treated symmetrically.

A final implication of the hypothesis for research that deserves

² "Family Savings in Relation to Changes in the Level and Distribution of Income."

SUMMARY AND CONCLUSION

mention is the importance it confers on data on the consumption or income of the same consumer units in different years, especially on such data giving both the consumption and income of the same units. It is likely that data of this kind now exist which have not been exploited, and these are one of the kinds of consumption data that should receive highest priority in future collection of data.

5. *Substantive Implications of the Hypothesis*

Acceptance of the permanent income hypothesis necessarily has implications for any problem of economic understanding or policy in which the determinants of savings play a significant role. At least in recent years, there have been two main classes or problems of this kind: those connected with the process of economic development, particularly of so-called underdeveloped economies, and those connected with economic fluctuations.

a. ECONOMIC DEVELOPMENT

Students of economic development tend to give a major role to the availability of resources for capital formation. One source is, of course, domestic savings. Their availability has, in turn, been taken to depend largely on the level of real income, on the one hand, and the inequality of income on the other.

The level of real income has been regarded as playing a dual role. First, the level of income defines the total amount available for consumption and savings; if, by some criterion, the total is low, so is the potential amount available for either purpose. This is, of course, a purely arithmetical truism and is unaffected by the hypothesis accepted about the factors determining the division of the total between consumption and savings. Second, acceptance of the absolute income hypothesis led to the belief that a low real income was unfavorable to savings in the further sense that it made for a relatively low ratio of savings to income.

The relative income hypothesis, which has received increasing acceptance in recent years, removed the direct connection between low real income and a low savings ratio but substituted an indirect connection. True, it argued, in an isolated community, the level of real income would have no effect on the savings ratio. But in a community connected with the rest of the world it would. According to the most widely accepted theoretical justification for the relative income hypothesis, that of Duesenberry, relative income is important within a community because of emulation and the demonstration of the availability and usefulness of superior goods. But these same effects work as well between communities. The "demonstration

SUMMARY AND CONCLUSION

effect" of the level of consumption in high income countries, or of citizens of high income countries resident in underdeveloped countries, tends, it has been argued, to lead the citizens of the underdeveloped countries to devote an unduly high percentage of their low level of income to current consumption and, especially, to use in that way any increases in income.

Acceptance of the permanent income hypothesis removes both the direct and this particular indirect connection between low real income and a low savings ratio. According to it, the savings ratio is independent of the level of income. Relative income, as measured, is empirically related to the savings ratio within a country not because of emulation or the demonstration effect but because relative measured income is a biased index of relative permanent income status. If the emulation and demonstration effects are not present within a community, there is no reason to expect them to operate between communities. It may be that a country or group with a relatively low real income will also have a low aggregate ratio of savings to income; but it may also have a high ratio, and in either case the explanation is to be sought not in the level of income but in other factors.

Although, on the permanent income hypothesis, a low level of real income does not make for a low savings ratio, a rapid rate of rise in income, whatever the level, may do so. The reason is that a rise expected to continue tends to raise permanent income relative to measured income and so to raise consumption relative to measured income. I conjecture that whether this effect shows itself is likely to depend critically on the source of the rise in real income. If it reflects development financed at least in part from domestic capital in an environment which makes for a high rate of return on domestic capital, the high rate of return to savings is an offset to the high ratio of permanent to current income and may well be more important. On the other hand, if the rise in income reflects primarily an external stimulus that gives little or no role to domestic capital, there may be no offset and one might expect the savings ratio to fall. An example of this second possibility might be the rising income among Okinawan natives as a result of its development as a United States military base, though I know too little about the details to be confident that it is. At any rate, if some examples of the two kinds of developments could be found, they might offer a rather nice test of the present analysis and of the applicability of the permanent income hypothesis across countries.

Thanks to the widespread acceptance not only of the absolute income hypothesis but of a very special form of it, the inequality of

SUMMARY AND CONCLUSION

the distribution of income has been regarded as a major factor explaining the aggregate savings ratio.³ Wide inequality of income is thought to tend toward a high savings ratio, and an approach toward equality, toward a low savings ratio. This consideration has frequently been a major argument offered in defense of inequality by people who are in other respects egalitarians: inequality in an underdeveloped country, they say, is a necessary evil since there is no other way to generate the savings needed for economic development.

According to the permanent income hypothesis, the effect of inequality depends critically on the source of the inequality. Insofar as the inequality is attributable to differences in permanent income status, it has no effect on the savings ratio. Insofar as it is attributable to differences in transitory components, it does, because inequality then means uncertainty about income prospects and hence increases the need for a reserve against emergencies. What is favorable to a high savings ratio is not inequality per se but uncertainty, provided, of course, it is uncertainty of a kind that does not reduce the average rate of return on capital—a qualification that is entered to allow for the clearly unfavorable effect on savings of increased uncertainty about the security of property such as might arise from fears of confiscation or close regulation by government.

This distinction between the sources of inequality seems to me of great importance. If I may speculate on the basis of utterly inadequate knowledge in the hope of provoking further study by better qualified students, it seems to me that the kind of inequality characteristic of many so-called underdeveloped countries is precisely the kind that is irrelevant to the savings ratio. Such countries frequently have rigid social systems, sharp separations between classes, great stability in the membership of classes within generations and from generation to generation; in short, wide inequality in permanent income status. The process of development, of industrialization, breaks down these rigid class distinctions; historically, it tends not only to produce a smaller degree of inequality in measured income but, what is more important for our purposes, also to substitute inequality arising from transitory factors for inequalities of permanent income status. The reduction of the inequality of permanent income status, whatever its importance in other connections, is neutral with respect to the savings ratio. This kind of inequality, which interestingly is generally the kind that is most distasteful to egalitarians, cannot be defended

³ If consumption is a linear function of absolute income, the aggregate savings ratio depends only on the mean income and not on its distribution, although savings are a larger fraction of income the higher the income, so long as the intercept of the consumption function is positive. For the relation described in the text to hold, the consumption function must be concave downward on the average.

SUMMARY AND CONCLUSION

as required to generate savings. On the other hand, the fluidity introduced into relative income status, the emergence of fresh possibilities of moving from one class to another, of possibilities of large gains and large losses over short periods of time—changes which the egalitarian may welcome as increasing equality of opportunity even if resulting in inequality of outcome—these changes are favorable to the savings ratio.

To continue these speculations outside my own field of competence, I wonder whether undue attention has not been given to the magnitude of the savings ratio at the expense of the form that savings take. Savings may well have been at least as large a fraction of income in the Middle Ages as in modern times; they then in considerable measure, perhaps in major part, took the form of cathedrals, which, however productive of ultimate satisfaction and of social security in more than one sense of that term, were not productive of worldly goods. I understand that budget studies for India, which at first sight seem to give very different results from corresponding studies for the United States, are found largely to duplicate the latter if the category “ornaments” is interpreted as savings or, in the jargon of budget studies, as “net changes in assets and liabilities.” The East was for long regarded as a “sink” for the precious metals, surely evidence both of substantial savings and of the particular form that it took. Perhaps the crucial role that has been assigned to the savings ratio in economic development should be assigned instead to the factors determining the form in which wealth is accumulated; to the investment rather than saving process, as it were.

b. ECONOMIC FLUCTUATIONS

There has been widespread acceptance in recent years of explanations of economic fluctuations that interpret them as primarily a resultant of the interaction of unstable investment and a relatively stable relation between consumption and current income. While I do not myself accept this income-expenditure theory as a valid and tested interpretation of experience, the acceptance of the permanent income hypothesis clearly has important implications for it that are worth recording.

The combination of this interpretation with a belief in a shortage of investment opportunities and in a rising ratio of savings to income as real income rises led, particularly in the United States in the late 1930's, to a fear of “secular stagnation”: “mature” economies, it was argued, tend to have limited investment opportunities and high savings ratios at full employment. Acceptance of the permanent income hypothesis removes completely one of the pillars of the

SUMMARY AND CONCLUSION

"secular stagnation" thesis; there is no reason to expect the savings ratio to rise with a secular rise in real income. In addition, it destroys the case for one proposed remedy. To counter the danger of secular stagnation, it was argued, requires raising the average propensity to consume at a full employment income. Acceptance of a particular version of the absolute income hypothesis led to the belief that one way to do so was to reduce the inequality of income, so the fear of secular stagnation was used as an argument in favor of income redistributive measures. Acceptance of the permanent income hypothesis means that, whatever may be the merits or demerits of raising the consumption ratio, changes in the inequality of income, at least of permanent income, cannot be expected to have this result. Apparently the permanent income hypothesis is evenhanded—if it removes the justification for inequality as a necessary evil to produce required savings, it also removes the justification for reducing inequality as a means of reducing attempted savings.

Postwar expansion and apparently widespread investment opportunities have on occasion led some proponents of the income-expenditure theory to fear "secular exhilaration." Consistency would have required them to favor measures designed to increase inequality as a means of reducing the ratio of consumption to income at full employment and so reducing the danger of inflation. However, to the best of my knowledge, none has done so. Clearly, on the permanent income hypothesis no such conclusion would be justified.

As was noted in Chapter I, the permanent income hypothesis has relevance not only to these arguments about inequality based on particular empirical judgments, but also to the underlying Keynesian theoretical structure, in both its long-run or structural, and its short-run or cyclical, aspects.

In its long-run aspect, the central analytical proposition of the structure is the denial that the long-run equilibrium position of a free enterprise economy is necessarily at full employment; there may be, it is asserted, no monetary equilibrium at all, unless some *deus ex machina* such as rigid nominal wage rates is introduced to produce one; and the "real" equilibrium may be at a less than full employment position. Acceptance of the permanent income hypothesis in its most general form does not render these propositions invalid, for they do not depend on the proposition that savings is an increasing ratio of income as income rises and could be valid even if savings were a constant ratio. What does render these analytical propositions invalid is acceptance of the special feature of the hypothesis that w , the ratio of wealth to income, is a significant variable affecting k , the

SUMMARY AND CONCLUSION

ratio of permanent consumption to permanent income, and that an increase in w tends to raise k .

In its short-run aspect, as an interpretation of cyclical fluctuations, the central role in the income-expenditure theory is played by the relation between consumption and current income. The permanent income hypothesis has, so far as I can see, no implications for the empirical validity or acceptability of this interpretation of cyclical fluctuations; that must be decided by comparing its predictions with the predictions of alternative theories. But it does have important implications for the form of the consumption function and, in consequence, for the cyclical characteristics of an economy for which the income-expenditure explanation of fluctuations holds. The permanent income hypothesis leads to an aggregate consumption function like that presented above in point k of section 3, in which current consumption is largely determined by past incomes. One need not accept this particular form; the general result follows simply from the idea that current consumption is adapted to some measure of longer-run income status rather than to current receipts. The effect is almost certain to be a much smaller estimate of the marginal propensity to consume out of current income than would be obtained from a function that makes consumption dependent on current income alone. To put it in other terms: it means that a much larger part of current consumption is interpreted as autonomous and a much smaller part as dependent on current income and hence, through the multiplier process, on investment. The result is a smaller investment multiplier, and an inherently cyclically more stable system. For the particular consumption function we have estimated from the data, the multiplier of personal disposable income with respect to autonomous expenditures is only about 1.4, and this takes no account of the stabilizing effects of the progressive personal tax structure, corporate taxation and savings, and the like.⁴ To avoid misunderstanding, I hasten to repeat that these are not intended to be assertions about the actual empirical characteristics of our economy; they are conditional assertions and dependent for their validity on the prior acceptance of the income-expenditure theory as an explanation of economic fluctuations.

An enumeration of the implications of acceptance of a new hypothesis can never hope to be exhaustive. Indeed, one of the main implications is that it will stimulate people to think in new directions

⁴ A more extended treatment of this point is contained in Milton Friedman and Gary Becker, "A Statistical Illusion in Judging Keynesian Models," *Journal of Political Economy*, February, 1957.

SUMMARY AND CONCLUSION

and new ways that cannot possibly be specified in advance. Even though I have ventured in this section well beyond my empirical evidence and the areas of my own competence, I have no doubt omitted more of the ultimate implications of the acceptance of the permanent income hypothesis than I have included. This is at once the appeal and the justification of what we flatter ourselves by calling "pure" research.

