DISTRIBUTION EFFECTS
IN THE
AGGREGATE CONSUMPTION FUNCTION

TESTING FOR THE EXISTENCE OF DISTRIBUTION EFFECTS IN THE AGGREGATE CONSUMPTION FUNCTION

By

SAYYID TAHIR, B.A., M.Sc., M.A.

A Thesis
Submitted to the School of Graduate Studies
in Partial Fulfilment of the Requirements
for the Degree
Doctor of Philosophy

McMaster University

January, 1981

DOCTOR OF PHILOSOPHY (1981) (Economics)

McMASTER UNIVERSITY Hamilton, Ontario

TITLE: Testing for the Existence of Distribution Effects in the Aggregate Consumption Function

AUTHOR: Sayyid Tahir, B.A. (University of the Punjab, Pakistan)

M.Sc. (University of Islamabad, now named Quaid-e-Azam University, Pakistan)

M.A. (McMaster University)

THESIS-SUPERVISOR: Dr. A.L. Robb

MEMBERS SUPERVISORY COMMITTEE: Dr. A.J. Harrison, and Dr. W.M. Scarth

NUMBER OF PAGES: xiii, 168

TO:

My MOTHER---The most beautiful known gift of ALLAH

My Late FATHER---The spring of love, kindness and generosity, all of which translate into this work

Rabbe Irhum-homā Kamā Rabbayānee Sagheera (Amen)

ABSTRACT

This thesis addresses a long-standing puzzle in empirical econometrics: Does the size distribution of income matter in the aggregate consumption function? Current opinion on whether distribution matters is divided. There is also a lack of consensus (among those who believe distribution effects exist) on the nature of such effects; that is, whether a decrease or an increase in income inequality is needed to stimulate aggregate demand. In this thesis, the previous or existing tests are challenged on the grounds that they do not properly take into account the causal link between the variability of the marginal, not the average, propensity to consume (with respect to the income level) and the existence of distribution effects. This particular link is taken care of, however, if one tests for the linearity (in income) of the micro relation underlying one's aggregate consumption function. The rejection of the linearity hypothesis will establish the existence of distribution effects. *Ex post*, if the nonlinear relation is such that the marginal propensity to consume declines with income, it also follows that an equalization in the income distribution produces greater aggregate consumption. The theoretical contribution of this thesis lies in the clarification of these issues.

On the empirical side, this thesis cautions against the casual use of the term "distribution effects". In the current income-current expenditure framework of the Keynesians, it refers to "the effect of a redistribution of real disposable income" on aggregate real consumers' expenditure. In the Permanent Income Hypothesis framework, however, it could mean <u>either</u> "the effect of a redistribution of real disposable income" <u>or</u> "the effect of a redistribution of real permanent income" on aggregate real consumption. In this thesis, the distributions of real disposable income and real permanent income are alternatively assumed to follow the lognormal density, and two conclusions are empirically determined:

I. The distribution of real disposable income matters in the current income-current expenditure framework---this result is statistically significant at a 10% level after the correction for serial correlation and simultaneity bias. In particular, the estimates indicate that the marginal propensity to consume declines with the level of real disposable income and, hence, a decrease in inequality would stimulate aggregate demand.

II. The elasticity of consumption out of real permanent income is unity; therefore, the distribution of real permanent income does not matter in the Permanent Income Hypothesis framework---this result is statistically significant at all conventional levels of significance both before and after the correction for serial correlation.

Both findings are based on aggregative time-series data for Canada. The consumer unit in this thesis is an individual income-recipient, and the data period is 1947-1976. Maximum-likelihood procedures have been used in the estimation, with proper allowance for across-parameter constraints. In the event of correction for serial correlation, the autocorrelation coefficient is constrained to the open-interval (+1,-1). The results are also double-checked by examining many avenues that might affect the nature of the outcomes.

Another contribution of this study is the compilation of data on the distribution of pre-tax personal income (in current dollars) in Canada under the lognormality hypothesis. The parameters of this distribution are determined using the minimum chi-square method. Estimates of the variance (of logarithms of income) parameter show a slight increase in income inequality over the period 1946 to 1976. The data on this parameter are used to approximate the variance of logarithms for the distribution of real disposable income (while establishing result I) and also the same for the distribution of real permanent income (while establishing the result II).

ACKNOWLEDGEMENTS

Allahumman laka al-hamd wa laka al-shukr

The process underlying this thesis spans the better part of seven years. My parents, sisters and brothers are the people who made this work at all possible by their constant encouragements and prayers, not to mention the sacrifices they have made. It is not enough to utter a note of thanks and, thereby, equal their support in my otherwise lone journey. All I can say to them, however, is: May Allah increase you in prosperity and happiness both in this world and in the Hereafter (Amen).

My single greatest intellectual debt goes to my thesis-supervisor Dr. A.L. Robb. He has been associated with my work since Winter 1975 when I took a course in Labour Economics with him. This thesis is an out-growth of the work started during that period. Since then, every single point discussed in this study owes something to him. His keen observation and perseverence have been my assets all along. Thanks baba.

I also sincerely appreciate the help of the other two members of my thesis-committee: Dr. W.S. Scarth and

Dr. A.J. Harrison. They too contributed a great deal to my work. Professor F.T. Denton and Dr. D.W. Butterfield also generously gave their time to clarify numerous econometric matters. To all of them, and many others not mentioned in person, I offer my thanks.

In the end, I acknowledge that all the errors that might still be present in this work go to my credit.

CONTENTS

****

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

PROLOGUE

Aggregate consumption is considered by most economists to be important---by those interested in the short-run because of its importance for aggregate demand, and, hence, production and employment levels; and by neoclassical economists because its complement is savings which determines the capital stock and productivity of the economy. Aggregate income has long been known to be its primary determinant. An important secondary consideration is whether aggregate consumption is affected, in addition, by the size distribution of this aggregate income. This is the issue addressed in this thesis.

The importance of such a study can hardly be underscored. Personal expenditure on consumer goods and services is the largest component of final demand in any country; this fraction runs around sixty percent in Canada. A successful effort to establish the role of the size distribution of income in the consumption sector has obvious pay-offs. For example, it would justify making the size distribution of income a part of our economy-wide models. Such models, in turn, should enable the policy-makers to quantify the impact of main economic variables, such as unemployment, inflation and fiscal actions of the government, on the distribution of income. Moreover, they would also contribute to the understanding and the identification of feedbacks from the distribution to the level of economic activity. An in-depth analysis is also likely to provide fruitful insights into econometric modelling of distribution effects. But these considerations are by no means the only reasons for our undertaking this project. The following empirical puzzle has also been instrumental in making our decision.

Cross-section studies on the consumption function are unanimous in concluding that people with lower incomes tend to have a greater average propensity to consume than those with higher incomes. This observation when supplemented by the popular belief that the poor have relatively more suppressed needs than the rich, yields an interesting proposition: the marginal propensity to consume is expected to become progressively smaller as the level of income rises. Supposing that this is in fact the case, there are far-reaching implications for the redistribution of income. If a dollar of income is transfered from the rich to the poor, the decline in the consumption of the former will be more than offset by the increase in that of the latter. Thus an income-equalizing redistribution will produce a net increase in aggregate consumption. By the same argument, aggregate consumption will decline if an income transfer from the poor to the rich takes place. This argument which we shall refer to as <u>the conventional wisdom,</u> suggests that income inequality ought to be inversely related to aggregate consumption at any point in time. Since a time-series is nothing but a string of cross-sections, the conventional wisdom implies that the inverse relation between aggregate consumption and income inequality should also follow from time-series data. By the same token, this result should also emerge from inter-regional or cross-country data. But the empirical literature on this subject contains mixed results. Staehle (1937) and Van Doorn (1975) report results that conform to the conventional wisdom. However, the findings of Polak (1939), Metcalf (1972) and Blinder (1975) contradict it. The work of Della Valle and Oguchi (1976) contains both types of conclusions, one contradicting and the other favouring the conventional wisdom. These conflicting results have quite opposite policy implications. For example, Staehle's result suggests that a more equal distribution of purchasing power is necessary to stimulate aggregate demand in times of economic recession; whereas Blinder's finding supports the opposite course of action. What is the reason for these conflicting results? Do they follow from the data employed by various authors? Or, are they the product of different methodologies? This situation, in our opinion, calls for a new study to review the past empirical approaches, and to settle the question on the nature of these distribution effects.

The existing works on the subject employ aggregative data; in some cases the analysis is based on inter-regional or cross-country data, while in the others time-series data are used. We also use aggregative, time-series data (for Canada) in this study. Thus, on this aspect of methodology, our work should be directly comparable with the existing studies. We also hope to show that the conflicting results in the literature about the nature of the distribution effects are not due to the differences in data.

The general question addressed in this thesis is:

> Does the distribution of income matter in the aggregate consumption function?

There are three principal theories of the consumption function: the current income-current expenditure approach attributed to Keynes, the Permanent Income Hypothesis (PIH) of Friedman, and the Life-cycle Hypothesis of Modigliani, Brumberg and Ando. But we find that in the literature on the distribution-consumption subject, only two variants of the above-mentioned question have been asked:

A. Does the distribution of real personal disposable income matter in the current income-current expenditure framework?

B. Does the distribution of real permanent income matter in the Permanent Income Hypothesis (PIH) framework?

We intend to study these two questions separately---A in Chapter 4, and B in Chapter 5. Both questions emphasize different things. In the former, consumption is synonymous with personal expenditure on consumer goods and services, including consumer durables; whereas in the latter, consumption refers to the value of services actually consumed. The difference in the income concept is quite obvious; unlike current income, permanent income of a consumer unit is its expected income based on its present wealth holdings (both human and nonhuman) and projected future income-receipts. The nature of the distributions and, therefore, the distribution effects involved in both cases is also understandable. Their emphases aside, however, both questions are identical from the methodological point of view. This is perhaps the reason that we could not find any fundamental differences attributeable to methodology in previous studies that have addressed these questions.

In Chapter 2, section 2.1, we survey the present state of the art of testing for the existence of distribution effects. We define the issue by demonstrating that there is a causal link between the variability of the marginal (not the average) propensity to consume and the existence of distribution effects. This is also the message one gets while rationalizing the conventional wisdom, as we did on page 2. Our general impression, however, is that the majority of existing studies do not address this issue in their tests. This raises a question as to the validity of their results. Only Blinder (1975), in a part of his work, and Van Doorn (1975) explicitly consider this point. However, in each of these two studies the hypotheses tested are imprecise, and improvement in their respective methodologies is desireable.

We conclude our discussion with the following observation in section 2.2: any test should be for the "nonexistence", versus the "existence", of distribution effects---not necessarily against the alternative hypothesis of the distribution effects conforming to the conventional wisdom. This can be done by testing for the linearity of the micro model underlying one's aggregate consumption function, as against the nonlinearity of the same. Ex post, however, if the marginal propensity to consume turns out to decline with the income level, one can infer support for distribution effects conforming to the conventional wisdom. The implementation of this test from aggregative data requires complete knowledge of the functional form of the distribution of income---the distribution of real personal disposable income for Question A, and the distribution of real permanent income for Question B.

The subject of the functional representation of the distribution of income, let alone of real personal disposable income or real permanent income, has received little attention over the past 30 years. In the Canadian case, there is still room for pioneering work. Data constraints shape the course of our empirical work. The micro unit in our statistical illustrations is an "income-recipient" instead of a family or a household as is often used. Our assumption is that a spender of income is actually the one who earns it in the first place. With this in mind, in Chapter 3 we compile evidence on the distribution of pre-tax personal income (in current dollars) in Canada using the two-parameter lognormal distribution. The necessary data are taken from Taxation Statistics (the Green Book of the Department of National Revenue, Taxation Division---now Revenue Canada, Taxation). The appropriateness of these data for the said objective is dicussed in Appendix A. In section 3.1 we provide a statistical-cum-technical background to our parameter estimates. This involves the explanation of (1) our implementation of the minimum chi-square method for estimation, and (2) the adjustments in the number of income-groups in our data, which we make. In section 3.2 we report on the lognormal parameter estimates for 31 years, from 1946 to 1976. The advantage of this exercise is that it gives us a consistent series on the variance of the logarithm (abbreviated as 'variance of logs' throughout this thesis) of pre-tax personal income in Canada. This variance series indicates the pattern of inequality in the said distribution; we discuss it in section 3.3. This statistical information is used twice in the remaining chapters, once in answering Question A in Chapter 4, and a second time in answering Question B in Chapter 5.

In Chapter 4 our working hypothesis is that in every year, from 1947 to 1976, the distribution of real personal disposable income (among the income-recipients) follows the pattern of a two-parameter lognormal distribution. This proposition does not mathematically follow from our assumption of lognormality about the distribution of pre-tax personal income; it is invoked solely on grounds of statistical convenience. Necessary data for the arithmetic mean of the distribution of real disposable income are drawn from the Canadian National Accounts; but for the variance parameter, we use the measure constructed for the pre-tax personal income (in current dollars) in Chapter 3. Whereas the use of the variance of logs measure constructed from nominal data can be justified, by invoking the assumption of a one-good framework, nothing categorical can be said about the lognormality hypothesis as such. So the results are given a conditional interpretation; that is, they are considered valid subject to the accuracy of the lognormality hypothesis about the distribution of real personal disposable income among the income-recipients.

The argument in Chapter 4 runs as follows. There is no <u>definitive</u> micro theory in the current income-current expenditure framework of the consumption function. This is true equally of the basic real disposable income-real consumer expenditure relation and the stochastic factors affecting the consumer units (both at a point in time and across time). So we start at a very elementary level. We define our problem at the micro level in section 4.1. The issue of the stochastic factors is addressed in section 4.2 where we consider a number of alternative assumptions about these, and look at the aggregative regression models that emerge. In our empirical work we use two of these regression models---one in which the stochastic factors enter in an additive fashion, and the other in which they appear in a multiplicative manner. We also supplement these two models by a stationary first-order autocorrelation scheme about the stochastic error term. Using maximum likelihod procedures, these models are estimated and tests of the hypothesis of "no distribution effects" are performed. The matters relating to data are explained in section 4.3. The parameter estimates and results of our tests for the afore-mentioned null hypothesis are discussed in section 4.4. In section 4.5, we study three aspects of our experiments: (1) the use of a proxy in place of the variance of logs of real disposable income; (2) the specification of our micro model; and (3) the issue of the simultaneity bias due to the use of a single-equation model. The final result of Chapter 4 is both interesting and confusing. The interesting part is that the numerical magnitude of the distribution effects conforms to the conventional wisdom. That is, a redistribution of real disposable income from the rich to the poor is expected to lead to an increase in the aggregate real consumer expenditure, and vice versa . The confusion arises when we address the statistical significance of this result. If the correction for both autocorrelation and simultaneity bias are ignored, the evidence against the "no distribution effects" hypothesis is overwhelming, both at the 5% and 10% significance levels. However, if an allowance is made for the impact of either serial correlation or simultaneity bias (or both), the null hypothesis is rejected at the 10% level of significance but not at the 5% level. Thus on the whole, the odds are 90% in favour of the conclusions that the distribution effects exist in the current income-current expenditure framework, and that they conform to the conventional wisdom.

The analysis of Question B, regarding the role of the distribution of real permanent income in the PIH framework, is the theme of Chapter 5. This analysis turns out to be interesting for several reasons. Friedman's micro specification of the consumption-income relationship does not have a constant term; it asserts that permanent consumption is strictly proportional to permanent income, or that the elasticity of permanent consumption out of permanent income is unity. During the past two decades, the challenge to this proposition has revolved around demonstrating the existence of a constant term in the permanent income consumption function. However, we are able to use the approach developed in this thesis to test for the unitary elasticity of consumption without concerning ourselves with the constant term. Our test also happens to be a test for the nonexistence of permanent income distribution effects. This is because the unitary elasticity of consumption, in our experiments, is synonymous with the linearity of Friedman's micro model.

The structure of our discussion in Chapter 5 does not parallel that in Chapter 4. The reason is that the PIH contains a well-defined micro theory; we restate it in section 5.1. This restatement is intended to bring out the central argument of the PIH. In this manner we are able to define the scope of our argument, and to put the issue in proper perspective. The issue, as mentioned above, is the test of the unitary elasticity of consumption out of real permanent income, or the nonexistence of the distribution of real permanent income effects. This is formally stated in section 5.2. A meaningful test of these hypotheses requires _a priori_ knowledge of the functional form of the distribution of real permanent income and the properties of the stochastic errors. The PIH is silent on these issues. We, therefore, make certain assumptions in addition to those made by Friedman. For the distribution of real permanent income, we invoke the lognormality hypothesis. Again, as in Chapter 4 (in that case regarding the distribution of real disposable income), this is done without any theoretical justification. As a consequence, we again recommend caution in the interpretation of the results. This assumption, together with the others made in order to put estimation in the proper perspective, is discussed in section 5.3. The reader should be fully aware that data are a major problem in any test of the PIH. We use _two_ proxies for aggregate real consumption, _one_ for the total number of income-recipients (the consumer unit in our analysis) and _two_ for aggregate real measured income. The income data serves as an input to construct aggregate real permanent income which is determined as part of the estimation. With these data, we are able to define per income-recipient magnitudes for real consumption and real permanent income. For the variance of logs of real permanent income, we use the series constructed in Chapter 3 for the variance of logs of pre-tax personal income (in current dollars). In section

5.4, theoretical issues relating to data are discussed. All technical matters, such as the generation of aggregate real permanent income, are explained as part of the estimation and tests. These are reported in section 5.5. We find that the elasticity of consumption (out of real permanent income) is almost unity, and indiscernible from it on the statistical tests. This finding supports Friedman's specification of the consumption function. In section 5.6, the results are subjected to further tests. These ensure that our conclusions are insensitive to the problem of serial correlation, the potential existence of a constant term in the data, and the presence of the consumer asset component in our consumption data. We conclude our discussion with the inference that the distribution of real permanent income plays no role in determining the level of aggregate real consumption.

Chapter 6 is the last of six chapters in this study. It contains no new findings. Rather, here we look at our work in retrospect, and also indicate what directions future research might follow.

CHAPTER 2

TESTING FOR THE EXISTENCE OF DISTRIBUTION EFFECTS: A CRITICAL SURVEY

The issue of the role of the size distribution of income in the consumption function has almost as long a history as the subject of the aggregate consumption function itself. Staehle (1937), within one year of the publication of Keynes's General Theory, is a good example. But unfortunately very little work has been done on this topic since then. The most likely reason for this neglect, in our opinion, is the lack of data on the size distribution of income. The dating of the studies reviewed in this chapter corroborate this assessment; these are, besides Staehle (1937), Polak (1939), Metcalf (1972), Blinder (1975), Van Doorn (1975), and Della Valle and Oguchi (1976).[1] It is only since the mid-sixties that serious efforts to compile distribution data on a regular basis got underway in most countries; these coincide with the revival of academic interest in welfare and equity matters. All of the six studies discussed here address statistically the question: does the size distribution of income matter in the aggregate consumption function? A partial list of works not included in this survey is given in footnote 2. On the presumption that works analysed here, especially the more recent ones, reflect the approaches followed in the past, the coverage of our survey should be comprehensive for all practical purposes.

2.1 State of the Art---A Synthesis

The term size distribution of income refers to the relative-frequency distribution of income among families, households or income-recipients, whichever (of these) happens to be the primary unit in a given distribution situation. Income inequality is just an economically meaningful aspect of this distribution. This perception of the size distribution of income implies that its primary function is to facilitate consistent aggregation, with respect to income, of micro relations. Put differently, the size distribution provides the theoretical link between one's aggregate consumption function and its underlying micro model. Income-inequality considerations are a second, though by no means a less important, matter. Average (and, given the number of consumer units, total) income is long known to be a determinant of aggregate consumption. But this average is the first moment, about the origin, of the relative-frequency distribution of income. Thus when the issue of distribution effects is raised, the question being asked is: when do the dispersion parameters of the relative-frequency distribution of income enter the aggregate consumption function?[3]

To answer this question, we assume that the consumption function for a typical consumer unit is $c=g(y)$, and the relative-frequency distribution of income is given by $f(y)$, for all $y>0$. For analytic simplicity we will also assume that $g(0)>0$, $g(.)$ is a well-behaved continuous function over the entire income range, $f(y)>0$ for all y's, and $\int f(y)dy=1$, the integration is over the domain of y. These assumptions about $f(y)$ imply that it may be treated as a probability density function of the continuous type. We will denote the successive moments of $f(y)$, about $y=0$, by m_1, m_2, m_3, and so on. According to the Taylor's series expansion of $g(y)$ about $y=0$, we may write

$$c = A_0 + A_1 y + A_2 y^2 + A_3 y^3 + ---, \tag{2.1}$$

where A_0, A_1, A_2, A_3,... are the appropriate Taylor's series coefficients, each defined with respect to $y=0$. The aggregate consumption function for $g(y)$ will be defined in per capita terms as

$$\begin{aligned}\bar{c} = E(c) &= \int cf(y)dy,\\ &= A_0\int f(y)dy + A_1\int yf(y)dy + A_2\int y^2 f(y)dy\\ &\quad + A_3\int y^3 f(y)dy + ---,\end{aligned}$$

all integrations are over the domain of y. Using the notation introduced earlier, we get

$$\bar{c} = A_0 + A_1 m_1 + A_2 m_2 + A_3 m_3 + ---.$$

This can also be rewritten in the following manner:

$$\bar{c} = A_0 + (A_1+A_2m_1)m_1 + A_2(m_2-m_1^2) + A_3m_3 + ---.$$

But by definition $m_1=\bar{y}$, the mean income, and $(m_2-m_1^2)=var(y)$, the variance of incomes. Therefore,

$$\bar{c} = A_0 + (A_1+A_2\bar{y})\bar{y} + A_2var(y) + A_3m_3 + ---. \tag{2.2}$$

It should be noted here that, given the mean income ($\bar{y}$), income inequality notions relate to var(y), m_3 and other higher-order moments of f(y). This illustration allows us to discuss a number of interesting points relevant for the existence of distribution effects. If g(y) were linear in the first place, the coefficients A_2, A_3,... will be zero in equation (2.1). In that event

$$c = A_0 + A_1y,$$

and, from equation (2.2),

$$\bar{c} = A_0 + A_1\bar{y}.$$

That is, if consumption is linearly related to income at the micro level, then aggregate consumption will be completely specified by mean (or, given the number of consumer units, total) income alone. The dispersion and skewness in the distribution of income will not affect aggregate consumption. But as soon as the linearity proposition is dropped, notions relating to income inequality begin to appear in the aggregative functional form, that is, equation

(2.2). Their exact nature will depend, of course, on the shape of g(y), the micro consumption relation, and f(y), the relative-frequency distribution of income. A related point may also be noted here. For a linear micro consumption relation, the average propensity to consume (apc) will vary with the level of income, unless $A_0=0$, though the marginal propensity to consume (mpc) will not. On the other hand, in the case of a nonlinear micro consumption function the mpc does vary, in addition to the apc, with the level of income. Given this, it is understandable that in the empirical literature, justification for the role of distribution effects has been given in terms of the mpc.[4]

As mentioned earlier, we review a total of <u>six</u> works: Staehle (1937), Polak (1939), Metcalf (1972), Blinder (1975), Van Doorn (1975), and Della Valle and Oguchi (1976). The terms of reference in Staehle (1937), Polak (1939) and Van Doorn (1975) relate to the current income-current consumer expenditure framwork, so these studies qualify as attempts to answer Question A of Chapter 1. Blinder sets out to answer Question B. The dependent variable in his work is consumption (relevant for the Permanent Income Hypothesis (PIH)-related studies). However, he uses the distribution of current income, rather than real permanent income, in his analysis; a justification is, of course, provided (Blinder, 1975, footnote 14). Metcalf's case is somewhat complicated.

The dependent variable in his analysis is aggregate real consumers' expenditure. But he includes both current income and lagged consumption terms in his consumption function suggesting that he is thinking on the lines of the PIH. His distribution variables, however, relate to the distribution of current income (without any justification, such as given by Blinder). His consumption function is just one part of a macroeconometric model with the distribution of current income incorporated. Thus, it may be that he was guided by considerations of empirical fit (rather than those involved with testing a consumption theory) in the selection of the functional form for his consumption function. Della Valle and Oguchi (1976) provide little in the way of a theoretical argument for their experiments. Their dependent variables also are based on aggregate consumers' expenditure. Apart from these differences in their respective theoretical contexts and, therefore, the implications of their results, all the authors test for the effects of the size distribution of current income. Blinder tests for these effects with respect to "consumption", in the sense of the value of services actually consumed. All the others test for the same while focussing on "personal expenditure on consumer goods and services (including that on consumer durables)", as in the National Accounts. Further, the differences in the dependent variables aside, the methodological approaches are fairly common in all these studies, so we review them together.[5]

Of the six works that we review, Staehle (1937), Polak (1939), Metcalf (1972) and Blinder (1975) are time-series studies; Van Doorn (1975) is based on a regional cross-section; and Della Valle and Oguchi (1976), as mentioned above, is a cross-country analysis. In the time-series case, aggregate consumption (or consumers' expenditure) is compared across time. In the two other cases, similar comparisons are performed across different regions in a country or across different countries. From the methodological point of view the following discussion applies equally to all the approaches covered in these six works. But differences in the nature of the dependent variables, as pointed out by Blinder (1975,p.455), do recommend some caution in comparing their results. By definition, we see that

$$C = CE - CD + UD,$$

where C = consumption; CE = personal expenditure on consumer goods and services; CD = personal expenditure on consumer durables; and UD = use-value of the stock of consumer durables. From this, it follows that

$$\frac{\partial^2 CE}{\partial Y^2} = \frac{\partial^2 C}{\partial Y^2} + \frac{\partial^2 CD}{\partial Y^2} - \frac{\partial^2 UD}{\partial Y^2},$$

where Y is disposable income. Quite obviously, it is not necessary that $(\partial^2 CE/ \partial Y^2)$ and $(\partial^2 C/ \partial Y^2)$ should have the same sign. Suppose that the former is negative, but the latter is positive; that is, the marginal propensity to spend declines with an increase in disposable income, but at the same time the marginal propensity of consumption increases. In such a situation a redistribution of disposable income from the rich to the poor will increase aggregate CE, but decrease aggregate C. The response of aggregate CE will correspond to the conventional wisdom, but that of aggregate C, though perfectly legitimate, will not. This exercise demonstrates the possibility of two contradictory results with the same income and distribution data, but a different consumption variable. Among the six works at hand, only Blinder (1975) is based on the consumption concept; the rest of the studies, as mentioned earlier, focus on consumer expenditure. We hope to show, however, that the differences in the results in these studies, excluding Van Doorn (1975), are not due to the differences in the consumption concept. Hence the distinction between consumption and consumer expenditure will not be maintained in the following discussion, and both terms will be used inter-changeably.

All the studies, except Van Doorn (1975) which we examine below, have two mutually exclusive parts. The first part contains some justification for having a distribution proxy (or proxies) in the aggregate consumption function. Sometimes this justification is purely mechanical, as in the case of Staehle (1937), and sometimes purely economic, as with Polak (1939), Metcalf (1972), and Della Valle and Oguchi (1976). At times it a mixture and quite sophisticated too; Blinder (1975) is a good example of this. But when it comes to actual experimentation, the <u>second</u> of the two parts of each study, <u>ad hoc</u> approaches and convenience seem to take over. As a result, in each case (excluding a part of Blinder's work, as we will see below) a distribution variable (or a set of such variables) is simply added to the list of regressors. We are not claiming that the correct reason for the possibility of distribution effects is overlooked. On the contrary, it can be shown that all the authors are aware, either explicitly or implicitly, of the necessity of the variability of the mpc (with respect to the income level) for distribution effects. But, excepting Van Doorn (1975), this point does not properly manifest itself in the statistical analysis. This issue may be illustrated with the help of the following example.

Suppose that at the micro level

$$c_i = Ay_i^B, \tag{2.3}$$

where c_i and y_i refer to consumption and income of the ith consumer unit, and A and B are parameters. Let us further assume that the size distribution of income, $f(y_{it})$, is given by the lognormal density:

$$(2.4)\qquad f(y_i) = \frac{1}{y_i \sigma \sqrt{2\pi}} e^{(-0.5/\sigma^2)(\ln y_i - \mu)^2}, \quad y_i > 0;$$
$$= 0, \quad \text{otherwise}.$$

μ, the logarithm of geometric mean income, and σ^2, the variance of logarithms of incomes, are the two parameters of $f(y_i)$. σ^2 also serves as an income-inequality index; an increase (a decrease) in its value is synonymous with an increase (a decrease) in income inequality.[6] Aggregating the micro relation (2.3) with respect to (2.4), we get[7]

$$\bar{c} = Ae^{B\mu + 0.5B^2\sigma^2},$$

where $\bar{c}$ is average consumption. Making further use of the relation $\bar{y}$ (average income)$=\exp(\mu+0.5\sigma^2)$, or $\mu=\ln \bar{y}-0.5\sigma^2$, which is true for the lognormal distribution, this becomes

$$(2.5)\qquad \bar{c} = Ae^{B\ln \bar{y} + 0.5B(B-1)\sigma^2}.$$

Application of the logarithmic transformation to (2.5) yields the following aggregative relation, in logarithmic form, corresponding to (2.3):

$$(2.6)\qquad \ln \bar{c} = \ln A + B\ln \bar{y} + 0.5B(B-1)\sigma^2.$$

σ^{-2} is the distribution variable in this aggregative model.

In this illustration two points are noteworthy. First, for our micro model (2.3)

$$mpc_i = (dc_i/dy_i) = ABy_i^{B-1},$$

$$\text{and} \quad (dmpc_i/dy_i) = AB(B-1)y_i^{B-2}.$$

That is, the mpc for a typical consumer unit is a function of income (of course, under the assumption that $B \neq 1$). Second, given the micro model (2.3), the distribution variable σ^{-2} is not independent of the income term $\ln \bar{y}$ in (2.6). In fact, this illustration suggests that if the individual mpc's are a function of income (which is necessary for the existence of distribution effects), then there is a link between the coefficients that apply to the income and distribution variables. This link depends on the functional form of the micro model at hand, and it becomes apparent when one follows the route of consistent aggregation. Obviously this link is obscured if one directly jumps to the macro model.

In Staehle (1937), Polak (1939), Metcalf (1972), and Della Valle and Oguchi (1976), the regression models are formulated directly at the aggregate level. Blinder initially maintains the distinction between the micro and macro models (Blinder, 1975, sections IIIC, IVB and IVC);

but in the last set of experiments reported in section IVD of Blinder (1975), his approach is similar to that in the afore-mentioned studies. In all of these cases (including Blinder, 1975, section IVD), the authors directly estimate regressions such as

(2.7) $$C = a_0 + a_1 D + a_2 Z;$$

or

(2.8) $$(C/Y) = b_0 + b_1 D + b_2 Z;$$

where C = aggregate real consumers' expenditure (consumption in Blinder, 1975)

(C/Y) = aggregate propensity to consume (definition subject to the nature of C and the aggregate income variable Y);

D = some distribution proxy (or proxies, in which case a_1 and b_1 will be vectors of parameters); and

Z = income plus any other variables (in which case a_2 and b_2 will be vectors of parameters) which the researchers think to be important explanatory variables.

For the reasons cited in the last paragraph, these studies are, in fact, simple correlation analyses. They cannot be judged directly as tests for the existence of distribution effects.

As mentioned above, the empirical work in sections IIIC, IVB and IVC of Blinder (1975) is based on micro foundations. In these sections Blinder reports <u>two</u> sets of experiments. In the first, reported in sections IIIC and

IVB, he tests for different mpc's for separate quintiles of families; in the second, reported in section IVC, he constrains the behaviour of the mpc across the quintiles, and tests for distribution effects, using an aggregative income-inequality index. In the case of his first experiment, Blinder's micro model is of the following type:[8]

$$c_j = \text{const} + k_j y_j, \tag{2.9}$$

where j is the quintile index. Equation (2.9) implies that in the jth quintile, current consumption (c_j) depends on current income (y_j) and factors embedded in k_j, such as the rate of interest---these factors do not include y_j, however. The mpc is thus constant within each quintile but variable between quintiles. This assumption is at best an approximation, and, in fact, a worse approximation the more nonlinear is the true underlying function. This may partly account for his somewhat confusing results. According to his best results, the short-run mpc's are 0.26 for the second and the fourth quintiles and 0.78 for the first, third and fifth quintiles; the corresponding long-run figures are 0.36 and 1.09. Blinder himself acknowledges that this zig-zag pattern of the mpc's is not very illuminating as to the nature of the distribution effects, that is, whether or not they conform to the conventional wisdom (see Blinder, 1975,p.461).

Blinder's second experiment is based on the following micro model:[9]

$$\text{(2.10)} \qquad c_j = \Upsilon_j + (k_0^j + k_1^j r) y_j + \lambda c_{j,-1},$$

where $k_0^j = m_0 + m_1 j$; $k_1^j = n_0 + n_1 j$; r is the rate of interest; and c_{-1} is the last period's consumption---again this model is specified at the quintile level. Here

$$mpc_j = m_0 + n_0 r + (m_1 + n_1 r) j.$$

Given m_0, m_1, n_0 and n_1, this specification constrains the way the mpc varies for movements along the income-scale. To see this, one should recall that in the quintiles' construction, the micro units are pre-arranged in ascending order of their incomes. Thus going from one quintile to the next means moving up along the income ladder. Suppose that all of the coefficients m_0, m_1, n_0 and n_1 are positive. In this case, given a positive r, the mpc should be increasing as one goes from the lower to the higher quintiles. On the other hand negative m_1 and n_1, together with positive m_0, n_0 and r, would imply a declining mpc with the level of income. The aggregative model used by Blinder, corresponding to (2.10), is:

$$\text{(2.11)} \quad (C/Y) = (\Upsilon/Y) + m_0 + n_0 r + m_1 D_B + n_1 r D_B + \lambda(C_{-1}/Y),$$

(Blinder, 1975, equation 18).[10] Here the distribution variable is $D_B = \sum j(y_j/Y)$, the summation is from j=1 to j=5. An increase in the value of D_B implies an increase in income inequality, and vice versa. This approach is desireable in the sense that the idea of the variability of the mpc, with movements along the income-scale, is incorporated in the analysis. But the problem is that it still ignores the constraint (on the distribution variable's coefficient) that is implied if consumption is specified as a nonlinear function of income to begin with. Consequently we cannot place too much reliance on this experiment of Blinder (1975) and the associated results.

So far we have been discussing the methodological aspects of five works: Staehle (1937), Polak (1939), Metcalf (1972), Blinder (1975), and Della Valle and Oguchi (1976). As far as testing for the existence of distribution effects is concerned, our stance is that the methodologies in these studies, with the exception of the first experiment in Blinder (1975), require serious reconsideration. If our explanations of their methodological errors are correct, one still needs to explain the asymmetry in their results. For example, Staehle reports a negative coefficient for his distribution variable (in line with the conventional wisdom). On the other hand, Polak ends up with a negative coefficient for Pareto's α while a positive one is expected. As we mentioned in the previous chapter, these results, if both correct, have contradictory implications for a redistribution policy. Staehle's finding recommends a decrease in income inequality in order to increase aggregate consumers' expenditure, whereas Polak's result supports the opposite course of action. Among the other researchers, Metcalf and Blinder come up with results similar to that of Polak, while Della Valle and Oguchi report mixed results. Do we attribute these conflicting results to the use of different consumption and distribution data? Or, should we look for economic reasons to explain the diversity in results? Either of these, perhaps, would be the appropriate course to follow, had it been that the correct methods of analyses were used. However, we think that the reason for this asymmetry is statistical.

We will use two of the models in Della Valle and Oguchi (1976, equations (1) and (2)) to explain this point. These models are

$$APC = \beta_0 + \gamma_1 D + \text{error term},$$

and $$APC = \beta_0 + \beta_1 Y + \gamma_2 D + \text{error term},$$

where APC is average propensity to consume, Y is aggregate income, and D refers to the distribution of income. According to the simple least-squares formulae

$$\hat{\Upsilon}_1 = (s_P/s_D)r_{PD},$$

$$\text{and} \quad \hat{\Upsilon}_2 = \frac{1}{1 - (r_{PD})^2}(r_{PD} - r_{PY}r_{YD})(s_P/s_D),$$

where s and r, respectively, stand for the standard deviation and the correlation coefficient of variables listed as their subscripts ('P' is an abbreviation of APC used in the subscripts only). Each r lies between -1 and +1. For simplicity, we assume that all the correlations are positive. Now suppose that in a particular data set, there is positive correlation between APC and D ($r_{PD}>0$). A regression of APC on D, as the first one listed above, with these data will necessarily produce a positive coefficient for D ($\hat{\Upsilon}_1>0$). What would happen if one simply adds another regressor to the model, as Y in the second regression listed above? In this case, the sign of distributional variable's coeffient, that is, $\hat{\Upsilon}_2$ would be positive or negative depending on whether the product '$r_{PY}r_{YD}$' is less or greater than r_{PD}. This idea can be extended to any number of regressors, with either the average propensity to consume or aggregate consumption in place of APC. Only the balance of different correlations will decide the sign of the distribution variable's coefficient, if an OLS regression procedure (as in these works) is followed. This conclusion and the other afore-mentioned objections, however, do not apply to Van Doorn (1975), because in this study the distribution variable does not enter independently (of the income variable) in the regression.

Van Doorn's work is commendable, and, in our view, a step in the right direction. But two errors, one methodological and the other a technical oversight, marr his work. His micro model is

$$(2.12) \qquad \ln c_i = \alpha + \beta \ln y_i + \gamma \ln z_i,$$

where c_i, y_i and z_i, in this order, refer to consumer expenditure, disposable income and size of the ith household. He perceives the aggregative relation as

$$(2.13) \quad (1/N)\{\sum \ln c_i\} = \alpha + \beta(1/N)\{\sum \ln y_i\} + \gamma(1/N)\{\sum \ln z_i\};$$

N is the total number of consumer units, and each summation runs over i=1,2,...N. Next he assumes that each of the variables c_i, y_i and z_i has a lognormal density. If a variable x_i has a lognormal density given as $\Lambda(x_i)$, then the continuous equivalent of $(1/N)\{\sum \ln x_i\}$ would be the definite integral $\int \ln x_i \Lambda(x_i) dx_i$ or μ_x, the logarithm of the geometric mean of x_i's. He also notes that for a lognormal distribution $\bar{x}$, the mean of the x_i's, equals $\exp(\mu_x + 0.5\sigma_x^2)$, where σ_x^2 is the variance of logs of x_i's. Using these points, he writes down the following aggregative model for (2.12):

$$\text{(2.14)} \qquad \ln \bar{c} - 0.5\sigma_c^2 = \alpha + \beta(\ln \bar{y} - 0.5\sigma_y^2) + \gamma(\ln \bar{z} - 0.5\sigma_z^2),$$

where each σ_x^2 applies to the variable listed as its subscript.[11] He then assumes $\sigma_c^2=\gamma\sigma_z^2$. This yields the following aggregative model:

$$\text{(2.15)} \qquad \ln \bar{c} = \alpha + \beta(\ln \bar{y} - 0.5\sigma_y^2) + \gamma \ln \bar{z}.$$

He tests for distribution effects by comparing the results for equation (2.15) with those for the following:

$$\text{(2.16)} \qquad \ln \bar{c} = \alpha + \beta \ln \bar{y} + \gamma \ln \bar{z}.$$

Here, he is simply assuming $\sigma_y^2=0$ for all the regions.[12]

But this approach of testing for the existence of distribution effects is wrong. We notice that for his micro model $mpc_i=\beta(c_i/y_i)$, and $(\partial mpc_i/\partial y_i)=\beta(\beta-1)(c_i/y_i^2)$. If $\beta=1$, the mpc_i will be constant, and, in principle, there should be no distribution effects. Thus a test for the existence of distribution effects ought to be a test for a non-unitary β in the context of (2.12). This is neither Van Doorn's objective, nor possible with his method of comparing results for (2.15) and (2.16). This is what we call the methodological error in his work.

As for the technical oversight, we notice that in writing equation (2.14), the maintained hypothesis is that

y_i and z_i are independently distributed as lognormal variates---though Van Doorn (1975) does not contain an explicit statement to this effect. c_i is related to y_i, and z_i via equation (2.12). If y_i and z_i are two random lognormal-variates, it can be shown that c_i will also be lognormal; moreover, the relation between σ_c^2, σ_y^2 and σ_z^2 will be as follows:[13]

$$\sigma_c^2 = \beta^2\sigma_y^2 + \gamma^2\sigma_z^2.$$

This means that the basic aggregative model should not have been (2.15), but instead

$$\ln \bar{c} = \alpha + \beta \ln \bar{y} + 0.5\beta(\beta-1)\sigma_y^2 + \gamma \ln \bar{z} + 0.5\gamma(\gamma-1)\sigma_z^2. \tag{2.17}$$

At this stage, of course, one can invoke some assumption about σ_z^2 due to data limitations. Van Doorn's assumption $\sigma_c^2 = \gamma\sigma_z^2$ led him to forget the "$0.5\beta^2\sigma_y^2$" term which is also relevant for income distribution effects. In Van Doorn's study, the constraint on the coefficient of σ_y^2 is -0.5β, as in equation (2.15), instead of $+0.5\beta(\beta-1)$. This error may account for the result that occasionally he obtained β greater than unity (see Van Doorn, 1975, Table 1). $\beta>1$ implies that the mpc for a typical consumer unit increases with the level of income. This, in turn, suggests that an increase in income inequality should increase $\bar{c}$. But the fact that σ_y^2 is preceded by a negative sign in equation

(2.15) suggests otherwise. This point is also overlooked by Van Doorn.

2.2 Concluding Observations

Variability of the mpc with respect to the income level, or equivalently, nonlinearity (in income) of the micro model at hand, is the key to the existence of distribution effects. None of the studies to date effectively handle this point.

Let income be the only determining factor in consumption decisions of the micro units. In a situation such as this, the model (2.3) can be a useful tool in testing for distribution effects. In the context of (2.3), the hypothesis of "no distribution effects" is synonymous with $B=1$. The distribution effects alternative may be conveyed more generally by considering $B \neq 1$. However, $B<1$ may be the appropriate alternate hypothesis, if one has strong *a priori* reasons that the conventional wisdom is the only rational possibility. In our opinion, it should be left to the data to decide whether B is less or greater than unity.

A priori knowledge of the functional form of the income distribution, rather than summary measures of income inequality, is necessary to test the above-mentioned propositions from aggregative data.

FOOTNOTES TO CHAPTER 2

1. Musgrove (1980) was published after this thesis was substantially complete. We briefly comment on it in footnote 14.

2. The works excluded are principally those which do not statistically test for distribution effects. Brady and Friedman (1947), Lubell (1947), Johnson (1952), and Alamgir (1976) are examples of such works. Also excluded are studies which approach the problem indirectly, for example, by incorporating the factor shares as regressors in the aggregate consumption relations---see Blinder (1975,pp.453-55) for a useful discussion of these.

3. Stating the problem thus, we are in effect assuming that income is the sole factor influencing the consumption decisions of the micro units. In cases when other determining variables are also present, and the distribution of consumer units according to income is independent of their distribution according to these other variables, this would still be the relevant question. To simplify our exposition, we assume income to be the only determining factor.

4. It may also be mentioned that the possibility of a nonlinear consumption function has been examined by Husby (1971,1974). He does recognize the link between the nonlinearity and the existence of distribution effects. However, he does not formally incorporate distribution variables into his model.

5. There are also consumer unit-related differences in these studies. For example, the focus in Staehle (1937), Blinder (1975) and Van Doorn (1975) is on wage-earners, families and households, respectively. But these differences are trivial as far as testing for distribution effects is concerned. What matters, as we hope to show in somewhat greater detail, is the nonlinearity of the micro relation underlying one's aggregative model.

6. This follows from the relation between σ^2 and the Gini coefficient of concentration, as demonstrated by Aitchison and Brown (1957,pp.111-13).

7. See Aitchison and Brown (1957,p.8).

8. In explaining Blinder's work, we suppress all the details not necessary for our point. For example, Blinder defines (2.9) as a relation between permanent consumption and permanent income---see Blinder (1975, equation 12). We modify it as in our text. Also on this and the later occassions, we will not explain the parameters not relevant for the mpc.

9. See Blinder (1975, equation 17). We change the subscript "i" to "j", and drop both the time subscript and the stochastic term.

10. See footnote 9. Υ is the aggregate equivalent of Υ_j's. We further modify Blinder's D to D_B.

11. For the afore-mentioned reasoning, see Van Doorn (1975,p.420).

12. See Van Doorn (1975,equations 4 and 5).

13. See Aitchison and Brown (1957,p.11).

14. The reader should consider this footnote in conjunction with footnote 1. All the studies (reviewed in section 2.1) have been shown to include the distribution variable in the aggregative model without proper allowance for the link between it and the income variable(s). Musgrove (1980) does exactly the opposite of this. In his study, the necessary precautions seem to have been taken in building the aggregative model. But when it comes to testing for distribution effects, the distribution variables are omitted in the tests without regard for the inter-relationship between these and the income variables. The aggregative model is

$$\text{(i)} \quad APC = b_0 + b_1(1/\bar{y}) + b_2(1/\bar{y})^2 + b_3(G(y)/\bar{y}) + b_4W + b_5(W/\bar{y}),$$

where APC is average propensity to consume, $\bar{y}$ is per capita income, G(y) is the Gini coefficient (of concentration in the distribution of y), and W is a measure of asymmetry (see Musgrove,1980,pp.510-12). The coefficients b's are related in the following manner:

$$
\text{(ii)} \begin{cases} b_1 = -2(1-b_0)(b_5/b_4), \\ b_2 = -(1-b_0)(b_5/b_4)^2, \\ \text{and} \quad b_3 = +(1-b_0)(b_5/b_4)^2. \end{cases}
$$

Musgrove's model under the null hypothesis of "no distribution effects" is

$$\text{(iii)} \qquad APC = b_0 + b_1(1/\bar{y}) + b_2(1/\bar{y})^2;$$

see Musgrove (1980,p.516). The principal results are: "the distributional variables seldom have coefficients distinguishable from zero, and the group of three variables never adds sufficiently to R^2 to pass an F-test at the 95 percent confidence level. The hypothesis that the distribution of income has nothing to do with the consumption propensity cannot be rejected ..." (Musgrove,1980,p.516). On the same page, we are also told: "The APC is clearly related to the level of income, either $1/\bar{y}$ or $(1/\bar{y})^2$ or both being significant in every equation."

The reader may note that in going from (i) to (iii), Musgrove treats b_3=b_4=b_5=0. But in that event b_1 and b_2 are either zero or undefined, in the light of (ii). Therefore, (iii) does not seem to be the appropriate model against which to test the said null hypothesis. Moreover, in the light of our discussion in section 2.1, the conclusion of no-distribution effects cannot be reconciled with the evidence of a nonlinear relation between the APC and the income variable. In our opinion, the appropriate constrained model in Musgrove's framework would be

$$\text{(iv)} \qquad APC = b_0,$$

and not (iii). Equation (iv) asserts that consumption is proportional to, or a linear function of, income. This, of course, corresponds to the point of view expressed throughout this chapter.

CHAPTER 3

THE CANADIAN DISTRIBUTION OF PRE-TAX PERSONAL INCOME UNDER THE LOGNORMALITY HYPOTHESIS

The lognormality proposition about the functional form of the distribution is assumed to hold for each of the 31 years from 1946 to 1976. The primary unit in our study is an "individual income-recipient", and the income concept is synonymous with that of personal income (in current dollars) in the National Accounts. Necessary statistical information is taken from Revenue Canada's Taxation Statistics. The adequacy of these data for drawing inferences about the distribution of pre-tax personal income (in current dollars) in Canada, and other related matters, are discussed in Appendix A.

As mentioned before, the lognormal density has two parameters: μ, the logarithm of geometric mean income, and σ^2, the variance of logs of incomes. σ^2 can be used to discuss issues related to income inequality. The parameters are estimated using the minimum chi-square (MCS) method. In the process, however, we have to make some adjustments in the number of groups in our data. All these matters are discussed in section 3.1. The results of our estimation, MCS estimates of μ and σ^2, are reported in section 3.2. This is followed by an overview of the pattern of income inequality (implied by our estimates of σ^2) in section 3.3.

3.1 A Statistical-cum-Technical Background to the Parameter Estimates

Several methods have been proposed for the estimation of μ and σ^2. To name a few, we may cite the method of moments, the quintile method, the graphical method, the maximum likelihood (ML) procedure, the minimum chi-square (MCS) method, and the ordinary least-squares (OLS) approach.[1] Each method offers certain conveniences in the estimation process. In their performance, however, they are not all equal. For example, McDonald and Ransom's results show differences of 25 to 40 percent in estimates of σ^2 derived from the MCS and OLS methods.[2] Of the six methods mentioned above, the ML and MCS methods are asymptotically equivalent and superior to the rest in that both satisfy the Fisherian criteria of consistency, efficiency and sufficiency.[3] McDonald and Ransom (1979) provide empirical evidence to reaffirm their asymptotic equivalence and their superiority to the OLS alternative. We use the MCS method in our estimation.

Let the domain $(0,\infty)$ of the random variable Y (income) be classified into k continuous, mutually exclusive and exhaustive groups:

$$y_0(=0)-y_1, \quad y_1-y_2, \quad y_2-y_3, \quad \ldots, \quad y_{k-1}-y_k(=\infty).$$

Suppose that we have a sample of N observations (on incomes) of which n_j belong to the jth income group, where $j=1,2,..k$; moreover, all n_j's sum to N. Under the hypothesis that Y has a lognormal density, the theoretical relative-frequency distribution of Y, $p(y;\theta)$, is given as

$$(3.1)\quad p(y;\theta) = \frac{1}{y\sigma\sqrt{2\pi}} e^{(-0.5/\sigma^2)(\ln y - \mu)^2}, \quad 0<y<\infty;$$
$$= 0, \text{ otherwise;}$$

where $\theta=(\mu,\sigma^2)$ is an unknown population parameter vector. If the proposition that $\theta=\tilde{\theta}$ (or, $\mu=\tilde{\mu}$ and $\sigma^2=\tilde{\sigma}^2$) is true, the theoretical relative-frequency of an income-recipient belonging to the jth income group will be given by

$$(3.2)\qquad \phi_j(\tilde{\theta}) = \int_{y_{j-1}}^{y_j} p(y;\tilde{\theta}).dy$$

The associated expected frequency of the jth income group, under the lognormality hypothesis, will be $N\phi_j(\tilde{\theta})$. The MCS method of estimating θ is based on the minimization of

$$(3.3)\qquad X^2 = \sum_{j=1}^{k} \frac{(n_j - N\phi_j(\tilde{\theta}))^2}{N\phi_j(\tilde{\theta})},$$

with respect to $\tilde{\theta}$. One can obtain MCS estimates by solving the two normal equations associated with (3.3), one for $\tilde{\mu}$ and the other for $\tilde{\sigma}^2$. These normal equations will be nonlinear in $\tilde{\mu}$ and $\tilde{\sigma}^2$. An alternative approach is to conduct a direct search for the minimum X^2 over a pre-specified grid of $\tilde{\mu}$ and $\tilde{\sigma}^2$, and thereby determine the MCS values for μ and σ^2. We follow this grid-search procedure. It is admittedly time-consuming, but its advantage is that a global minimum for (3.3) is virtually assured.[4]

It should be noted that the statistic (3.3) has been proposed to test the goodness of fit of a given functional form to the sample at hand. But it is highly sensitive to the sample size. This fact implies that the chi-square test will usually reject the null hypothesis if the sample size is very large.[5] To show this, we rewrite (3.3) as

$$(3.3') \quad X^2 = NQ; \quad \text{where} \quad Q = \sum_{j=1}^{k} \frac{(f_j - \Phi_j)^2}{\Phi_j}, \quad \text{and} \quad f_j = (n_j/N).$$

$f_j=(n_j/N)$ is the sample-based relative-frequency of the jth group. For 20 degrees of freedom, the theoretical value of X^2 at the 5% level of significance is 30.4104 (Christ, 1966, Table B-2). Suppose that on the basis of a given relative-frequency distribution of income records, and an assumption about the functional form of income distribution, Q equals 0.2500. Were it that N=100, the calculated value of X^2 would be 25.00, and this would imply a vote of confidence for the assumed functional form. But for N=140 the same relative-frequency distribution of records (that is, the same Q) would yield X^2 equal to 35.00 (>30.4104), and a rejection of the null hypothesis would be warranted. In our case, the sample sizes are in hund-reds of thousands. So, we want to emphasize that, the phenomenally high X^2 values reported in this chapter should not alarm the reader. Testing the goodness of fit of the lognormal distribution to our data is not our main objective. In fact, we do not rule out the possibility of specification error. Our interest is primarily in consistent and efficient estimates of μ and σ^2, and the MCS method is employed toward this end.

The second point that needs elaboration in this section, is associated with the groupings used in the tabulation of income distribution statistics. Both the number and the size of income groups, according to which the distribution data are available, are not the same for all the years (1946 to 1976). To be more specific, the changes have taken the form of collapsing details into fewer income groups at the lower levels of income, and offering more disaggregated information in the middle and upper ranges of income. Whereas smaller intervals are beneficial from the standpoint of numerical integration, they imply a nearly flat distribution of records in the large group (of which they are sub-groups). This phenomenon is likely to inflate the value of X^2 in the years for which we have more detailed groupings. Ultimately, the juggling with the number of income groups may also affect the magnitude of the parameter estimates, although each set of estimates will be consistent in its own right.

To define identical groupings for all the years requires adjustments which would result in only 31 income groups. This would involve a tremendous loss of information. Except for the first six years (1946-1951) the least number of groups is 51, and the number rises steadily to over 60 in the later years. Moreover, such a wholesale modification is not guaranteed to bring us closer to the truth, with respect to the time pattern of the σ^2 values. But the matter cannot be left untouched either, because there is evidence that the lognormal function does not fit the tails of the income distribution well (Lydall,1968). We follow a middle course. This involves a few alterations both at the lower and upper ends of the distribution in some years. We shall explain these modifications separately, starting with those at the left-hand tail of the distribution.

For the years 1946, 1947 and 1948, initially the first ten groups in our data happen to be "Under $100, $100-$200, $200-$300,... $800-$900 and $900-$1000". Collapsing these ten into one large group "Under $1000" is found to lead to a substantial improvement in the value of X^2---for example, for 1946 X^2 is reduced to 44332.95 from 81044.05.

A similar point is noted for 1958, 1959 and 1960, when the first <u>two</u> groups "Under $500" and "$500-$1000" are replaced by <u>one</u>.[7] Table 3.1 contains statistical evidence to support these points. As a result of this observation, the first income group for all the years is set as "Under $1000" (or "$1-$1000", in accordance with our empirical implementation of the MCS method). In addition to the afore-mentioned six years, we also introduce this modification to the data for 1949, 1950, 1951, 1952, 1953, 1954, 1955, 1956, and 1957.

At the other end of the distributon, our cut-off point is <u>$50,000</u>. However, adjustments are restricted to a few selective years. Three things together constitute the criteria for correction for the number (and, therefore, the structure) of income groups:

(i) The time pattern of σ^2's;

(ii) AX^2, the chi-square value <u>adjusted</u> for the differences in the sample size and the number of income groups, which indicates the degree of fit;[8] and

(iii) the number of income groups themselves.

The inter-temporal pattern of σ^2's reflects the course of income inequality over time. A σ^2 noticeably off the trend suggests that there may be something questionable about the data in a particular year, since the distribution of income is unlikely to change very rapidly. A second look is also warranted if, after the adjustment for sample size and income-group differences, the associated chi-square is unusually different from those for the preceding and (or) the succeeding year. Of course, the number and the structure of income groups at the right-hand end of the distribution, are treated as possible sources of such abnormalities. However, in cases where the number of groups in the $50,000+ income range are identical, potential sources of errors have to be looked for elsewhere, such as noise in the data. 1946, 1947 and 1952 are the only years for which adjustment is made on these criteria.

As for 1946 and 1947, with the first group redefined as "Under $1000", initial results show σ^2's for these years to be markedly off the trend. These values are .76083 for 1946 and .74771 for 1947 as opposed to .68553, .67609 and .68390 for 1948, 1949 and 1950, respectively. At this stage the only noticable difference among data for these two sets of years, except the sample size, pertains to the last group. For 1946 and 1947 there are two groups "$50,000-$100,000" and "$100,000 and over", whereas for the other three years it is a single "$50,000 and over" group. Suppressing the detail brings both the number and the structure of income groups for 1946 and 1947 in line with those for the other three years: 35 in all. Subsequent re-estimation affirms our initial reservations about the income-group structure. New X^2's reveal a much better fit for both 1946 and 1947---compare 33940.77 with the old 44332.95 for 1946, and 46871.15 with 61348.72 for 1947.

Again for 1952 σ^2, when estimated together with μ on the basis of 52 income groups, exhibits a healthy jump; both before and after 1952, σ^2's are lower than this estimate (=.71884) for 1952. Moreover, the implied value of AX^2 (=.001747) for 1952 is much higher than those for the next three years (given in column 7 of Table 3.2). At this point the 1952 distribution has just one more group, at the upper end of the income-scale, than those for 1953, 1954 and 1955. We collapse the last two income groups ("$50,000-$100,000" and "$100,000 and over") for 1952 into one ("$50,000 and over"), and make the grouping similar to that for the other three years.[9] This alteration produces a σ^2 for 1952 which is slightly lower than those for 1951 and 1953, quite contrary to the initial observation. But both the new X^2 and the corresponding AX^2 for 1952 are more in line with their counterparts for the other years (given in Table 3.2).

Of the remaining years, in 1956-1960 the $50,000+ income bracket is split into two groups "$50,000-$100,000" and "$100,000 and over". For 1961-1976 a further sub-division of the last group into "$100,000-$200,000" and "$200,000 and over" is also available. There is nothing unusual in the pattern of AX^2's (reported in column 7 of Table 3.2) which can be attributed to the introduction of these details in 1956 and 1961. So the matter is not pursued further.

The real increase in the number of income groups can be traced to increasingly disaggregated information in the \$5000-\$50,000 income range (\$3000-\$50,000 for the period 1946-1951). 1952 can be called the turning point in our data period. It is the only year in which we have a radical change in the structure of income groups over the previous year. Since 1952 there has been a steady increase in the number of income groups. But it is hard to show that this occurrence has had any special influence on the parameter estimates. As an example of our point, let us consider two years 1952 and 1969. In 1952 AX^2 equals .001448 with 51 income groups, and in 1969 the same statistic is .001436 with 60 income groups. The AX^2 values imply that the degree of fit of the lognormal function is more or less the same in these two years, despite the differences in the number and the structure of income groups. On the other hand, for 1968 the group structure is identical to that for 1969; but AX^2 is .001593 which is much higher than .001436 (for 1969)---a clearly superior fit in 1969. In the light of these two examples, nothing definite can be established as to the effect of the varying group structures in the \$5000-\$50,000 income range since 1952. It seems more appropriate to attribute any remaining errors in the fit of the lognormal function to noise in the data for various years.

3.2 The Lognormal Parameter Estimates

Table 3.2 contains our final estimates of μ and σ^2 for all the 31 years and the related information on the number of income groups, the X^2 values, the sample size (N) and the AX^2's. A parallel look at columns 5 and 6 of this table supports our earlier claim that there is a correlation between the sample size and the value of X^2. Accordingly, we are not in a position to comment on the issue of the specification error in any year. However, it is possible to check the degree of fit of the lognormal distribution across different years using AX^2. In our case, the AX^2 values (column 7, Table 3.2) fall into two categories. The first consists of AX^2's in the neighbourhood of .0020, and the second those clustured around .0014. The first category includes the results for the years 1946 to 1951, while the second those for the remaining years (1952 to 1976). The pattern of AX^2's within and between these two time periods brings to light two aspects of the performance of the lognormal distribution:

I. The degree of empirical fit offered by the lognormal function is fairly stable within each of the two periods 1946-1951 and 1952-1976. This implies that it may be preferable to attribute any differences in the fit <u>among</u> the years <u>(within</u> each of the two aforementioned periods) to sampling errors in the data for those years, rather than to the functional form of the distribution.

II. Empirically, the lognormal function performs better during the years 1952-1976 than it does in the first six years after the World War II.

The apparent discrepency in the performance of the lognormal distribution in the two periods can be attributed to the number (and, therefore, the structure) of income groups. This number is constant at <u>35</u> during the first time-span, while it is <u>51 or more</u> in the second. In any event, this discrepency should not prevent a comparison of the parameter values <u>across</u> the two time periods. Statistically speaking, the parameter estimates in each year (in both the time periods) are consistent and efficient. The across-period comparison is also justified because of an earlier observation that the distribution of income does not change very rapidly; moreover, our parameter estimates show no break between 1951 and 1952.

μ is the location parameter of the distribution of Y. As mentioned earlier, this is the logarithm of the geometric mean income. According to our estimates given in column 2 of Table 3.2, μ rises steadily through time, with the exception of 1954. This general pattern is compatible with the observation that in our growing world, incomes (in current dollars) have been growing secularly. As for 1954, μ_{1954}=7.73638 which is fractionally less than μ_{1953}=7.73824. But this decline is not to be blamed on the inappropriateness of the lognormality hypothesis, on the structure of income groups, or on our estimation method (in its inability to reach the global minimum of the X^2 function). In fact, 1954 is an unusual year in Canadian economic history. We also find a drop in the mean income of all return-filers (our data base) as well as in the mean personal income (in current dollars) of all Canadians aged 15 years and over.[10]

σ^2, the variance of logs of income, is the parameter of real interest from the inequality point of view. A rise in its value implies an increase in the degree of overall income inequality, and a fall indicates the converse. According to the estimates in column 3 of Table 3.2, σ^2 is fairly high immediately following the World War II. But apparently as the economy readjusts to peacetime, income inequality follows a declining trend until 1952. After this point in time one cannot mistake a clear upward trend in the inequality in the income distribution. The year 1949 is an exception to the declining trend, and the years 1972 and 1973 are exceptions to the latter upward trend. As for 1949, we do not have an explanation. However, we observe that the parameter values tend to over-estimate the true mean income (from Taxation Statistics) more in 1949 than in any other year.[11] This situation may, perhaps, be blamed on unusual sampling errors in the 1949 data. But for 1972 and 1973, our results receive some support from the Survey of Consumer Finances-based findings of Love and Wolfson (1976, Table 3). For all income-recipients, their estimates of Gini coefficient are .4843, .4681 and .4682 for 1971, 1972 and 1973; respectively; our estimates of σ^2 for these three years are .93107, .91914 and .91298. The reader can see for himself that both sets of numbers imply a similar pattern of income inequality over the same time-span.

3.3 The Pattern of Inequality in the Distribution of Pre-Tax Personal Income in Canada---An Overview

Our general conclusion based on the σ^2 values for 1946 to 1976 may be put in these words: In Canada, the inequality in the distribution of pre-tax personal income (in current dollars) has been increasing over time, except for a few years after the World War II.

We cannot afford to go into details, within the terms of reference of this thesis, to rationalize this finding. But it may be mentioned that this result is not unique. Similar conclusions have been reported for other countries as well; one can see, for example, Schultze (1971) for the U.S. In our opinion, the factors producing this general trend are likely to be demographic, rather than economic. The combined total of wages, salaries and current government transfers, as a percentage of personal income (in current dollars), has steadily gone up from 68.6% in 1947 to 81.9% in 1976. Over the same period, interest, dividends and miscellaneous investment income also registered a steady increase, but from 6.6% to 9.3% of the personal income (in current dollars). These two pieces of information enhance the view that _ceteris paribus_ income inequality should have declined or, at best, remained fairly stable within narrow bounds. This _ceteris paribus_ includes the demographic factors, as hinted above. They have not remained the same. Our data period is characterized by the increased labour-force participation of females (especially wives) and young males. This list also includes the gradual shift in the population structure, particularly over the second half of our data period, toward a more elderly population. All these groups are characterized by low-income earners. We do not hold these as the only factors for increasing inequality in the distribution of pre-tax personal income. One may, perhaps, find economic and structural shifts in the Canadian economy also contributing to this pattern. However, we think that these are the key factors.

FOOTNOTES TO CHAPTER 3

1. See Aitchison and Brown (1957, section 5.2), and McDonald and Ransom (1979).

2. This can be readily determined from information given in McDonald and Ransom (1979, Table II). Their estimates (with the two methods) also follow different time paths.

3. See Fisher (1922).

4. We set up our own program for the minimization of X^2; Dr. Harrison has been very helpful in this matter. The IMSL library routine DCADRE is used for numerical integration. It follows the cautious Rhomberg extrapolation technique. The starting income value assumed was "$1.00". The parameter estimates are refined up to five decimal digits; this effort ensures that, in our case, the X^2 function does not register a futher decline on the integer side. A copy of the program is available upon request.

5. McDonald and Ransom (1979,footnote 5) and Harrison (1981,forthcomming) also note this point.

6. This would happen if the income groups below $500 are modified to suit those in 1976, and those in the $3000+ range are redefined to follow the 1948 setting.

7. The income groups in the Taxation Statistics for 1958, 1959, 1960, 1961, 1962 and 1963 are reported in a discrete format. For example, the first two groups for 1958, 1959 and 1960 are Under $500 and $500-$999 --- in our text we refer to these as Under $500 and $500-$1000 . However, in actual estimations, these and all the other discrete groups (for the six years listed above) have been converted into their continuous equivalents. Thus, for example, Under $499.5 and $499.5-$999.5 replace the two above-mentioned groups.

8. By this form-ula, we have

$$AX^2 = \frac{X^2}{N.X^2_{(p,q)}};$$

where

N = the sample size, and

$X^2_{(p,q)}$ = the theoretical value of X^2 for q degrees of freedom, at an a priori fixed level of significance (p); we set p=5%.

The normalization of X^2 by N takes care of the sample size differences among different years. However, to compensate for the income group related differences, neither the number of income groups nor the degrees of freedom is an appropriate normalization factor, since X^2 is nonlinearly related to both of them. $X^2_{(p,q)}$ is suitable for this latter purpose, provided that p is kept the same for all the years.

9. The 1951 data also has but one group beyond the $50,000 income level. In re-structuring the 1952 data, the 1953 grouping is chosen over that for 1951 after comparing the AX^2's for 1953 to 1955 with those for 1946 to 1951---see column 7, Table 3.2.

10. The mean income (in current dollars) of all returns is $2789.26 for 1953, and $2785.33 for 1954; the mean personal income figures (also in current dollars) for all the Canadians in the age-bracket "15 years and above" are $1966.09, $1937.52 and $2049.43 for 1953, 1954 and 1955, respectively.

11. Figures available upon request.

Table 3.1 Evidence Supporting the Need for a Single Group in the "Under $1000" Income Category

Year	μ	σ^2	χ^2	Number of Groups
Part A	10 Sub-groups in the "Under $1000" Category			
1946	7.13716 (7.25681)	.84655 (.76083)	81044.05 (44332.95)	45 (36)
1947	7.24440 (7.38234)	.85483 (.74771)	117821.35 (61348.72)	45 (36)
1948	7.35513 (7.48923)	.84152 (.68553)	92749.19 (37237.14)	44 (35)
Part B	2 Sub-groups in the "Under $1000" Category			
1958	7.82330 (7.87690)	.81123 (.74022)	54801.52 (32651.60)	54 (53)
1959	7.84585 (7.89625)	.82245 (.74807)	55450.25 (32456.24)	54 (53)
1960	7.87238 (7.92601)	.84098 (.77254)	57783.22 (35407.38)	55 (54)

Note: The figures in parentheses are obtained with a single group in the under-$1000 income range. The grouping at the upper end of the distribution is retained as in Taxation Statistics.

Table 3.2 The Distribution of Pre-Tax Personal Income in Canada---Estimates Under the Lognormality Hypothesis

	The Lognormal Parameter Estimates		Supplememtary Information			
Year	μ	σ^2	Number of Groups	X^2	Sample Size (N)	AX^2
(1)	(2)	(3)	(4)	(5)	(6)	(7)
1946	7.25211	.71543	35	33940.77	402,501	.001803
1947	7.37313	.70034	35	46871.15	488,420	.002052
1948	7.48923	.68553	35	37237.14	354,503	.002246
1949	7.56833	.67609	35	39114.15	366,356	.002283
1950	7.59294	.68390	35	40341.08	376,616	.002290
1951	7.67721	.68041	35	40040.86	397,717	.002153
1952	7.71976	.67778	51	40298.29	424,571	.001448
1953	7.73824	.68899	51	42461.54	458,242	.001414
1954	7.73638	.69747	51	43479.46	470,341	.001410
1955	7.76002	.69970	51	42487.95	477,370	.001358
1956	7.81941	.73310	52	32028.66	327,311	.001470
1957	7.83857	.73833	52	30984.80	345,449	.001348
1958	7.87690	.74022	53	32651.60	352,329	.001373
1959	7.89625	.74807	53	32456.24	362,295	.001327
1960	7.92601	.77254	54	35407.38	372,684	.001380
1961	7.94517	.77811	58	35591.55	383,808	.001256
1962	7.98289	.77347	58	34045.54	370,814	.001244
1963	7.98952	.78642	58	29343.76	311,978	.001274
1964	8.02444	.81089	58	39026.48	399,345	.001324
1965	8.05803	.82010	58	44830.51	425,707	.001427
1966	8.09294	.83673	58	50071.14	459,580	.001476
1967	8.14111	.84308	60	57449.64	510,000	.001482
1968	8.17991	.88806	60	66994.17	553,000	.001593
1969	8.24827	.89353	60	67269.40	616,000	.001436
1970	8.29842	.91337	66	57408.05	519,000	.001334
1971	8.35214	.93107	66	61105.97	550,000	.001340
1972	8.43763	.91914	64	61238.24	598,204	.001273
1973	8.53811	.91298	64	42123.99	452,003	.001159
1974	8.67078	.94150	64	55090.60	525,664	.001304
1975	8.79609	.94287	64	63946.31	581,641	.001368
1976	8.91253	.95000	64	46174.53	381,622	.001505

Note: See section 3.1 of the text and Appendix A for matters relating to data.

CHAPTER 4

TESTING FOR THE DISTRIBUTION OF REAL PERSONAL DISPOSABLE INCOME EFFECTS IN THE CURRENT INCOME-CURRENT EXPENDITURE FRAMEWORK

In this chapeter we try to answer Question A, so labelled in Chapter 1: Does the distribution of real personal disposable income matter in the current income-current expenditure framework?

The current income-current expenditure approach to consumption is a short-run theory of the aggregate consumption function. It is attributed to Keynes (1936); probably this is the reason for the more familiar label "The Keynesian Consumption Function" in the literature.[1] Generally speaking, it is regarded as a linear specification (with an additive constant term) between current personal expenditure on consumer goods and services and current personal disposable income (both in real terms), at the aggregate level.[2] As mentioned in Chapter 1, there is no well-defined micro theory behind this aggregative specification. Therefore, we have to start at a very elementary level. This involves the specification of a micro relation between consumer expenditure and personal disposable income (both in real terms), and a set of intuitively plausible assumptions about stochastic elements. The micro unit in our study is an individual income-recipient, and the working hypothesis about the distribution of real personal disposable income among the income-recipients is lognormality. The argument in this chapter runs as follows.

In section 4.1 the relation between real consumer expenditure and real disposable income for a typical consumer unit is specified, and the problem defined. In section 4.2 four alternative assumptions about stochastic factors are made. Given these ingredients and the lognormality assumption about the distribution of real personal disposable income, three aggregative regression models are developed corresponding to the micro model of section 4.1. In section 4.3 each variable is identified with the data to be used. In section 4.4 the results of our estimation (for two of the three models developed in section 4.2) and tests (for the "no distribution effects" hypothesis) are reported. In a time-series analysis the problem of autocorrelation is almost unavoidable. This study is no exception. Accordingly, a stationary first-order autocorrelation process is postulated. An equivalent of the Full Maximum Likelihood (FML) procedure, discussed in Beach and MacKinnon (1978), is used for estimation. The main finding is that there exist distribution effects that conform to the conventional wisdom, but their statistical significance is open to argument. In section 4.5 some sensitivity tests are also performed for the sake of double-checking. The afore-mentioned result is, however, re-affirmed.

Before turning to details we may remind the reader that the policy implications of the question addressed in this chapter are very important. If the distribution of real personal disposable income can be shown to matter, it would mean that there are not one but two instruments of macro policy. One which also happens to be familiar in the current policy debates, is the level of government expenditure and taxes. The other would be the distribution of either (government expenditure or taxes) in so far as it alters the distribution of real disposable income and, thereby, affects aggregate demand.

4.1 The Problem Defined

Let c_{it} and y_{it} respectively be the consumer expenditure and disposable income (both in real terms) of the ith income-recipient, the consumer unit in our study, at time t. We postulate

$$(4.1) \qquad c_{it} = Ay_{it}^{B}; \qquad \text{A and B are parameters.}$$

This relation is the same as our model (2.3); the variables are, however, redefined. Equation (4.1) is not supposed to follow from a particular choice-theoretic framework. It simply reflects the proposition that there is no money illusion, and that at the micro level real consumer expenditure is a function of real disposable income alone. One may argue that there should be an additive constant term in this model to project the idea of subsistence consumption. But the specification of the constant term is not a theoretical necessity in order to test for distribution effects. On the other hand, its exclusion has been found to make estimation much easier. We will continue to treat equation (4.1) as if it holds on the average. In section 4.5.2, however, this matter will be re-opened.

Normally one would expect both A and B to be positive, but nothing categoric can be said about their magnitude. The parameter of interest for the purpose of distribution effects is B. If it were unity, this would imply that real consumer expenditure is a linear function of real disposable income, so that distribution would not matter. As seen in Chapter 2, a value of B other than unity would establish both the variability of the marginal propensity to consume (mpc) with respect to the income level, and the existence of distribution effects. Moreover, $B<1$ will imply that in the context of the current income-current expenditure framework, distribution effects conform to the conventional wisdom. On the other hand, $B>1$

will suggest the existence of distribution effects opposite to the afore-mentioned ones. We leave it for the data to decide which happens to be the case. With these points in mind we specify the following hypotheses:

(4.2) $H_0: \; B = 1;$
and
(4.3) $H_a: \; B \neq 1.$

The implementation of the tests of these hypotheses from aggregative time-series data requires two things:

(1) the knowledge of the distribution of real personal disposable income at every point in time; and

(2) the role of the stochastic factors affecting the consumption decisions of the income-recipients, both at a point in time and across time.

Our working hypothesis about the functional form of the distribution of real personal disposable income, $f(y_{it})$, is that it follows the lognormal pattern in every year. That is,

$$(4.4) \quad f(y_{it}) = \frac{1}{y_{it}\sigma_{dt}\sqrt{2\pi}} e^{(-0.5/\sigma_{dt}^2)(\ln y_{it} - \mu_{dt})^2}, \; y_{it}>0;$$
$$= 0, \text{ otherwise};$$

where μ_{dt} and σ_{dt}^2 are two parameters of $f(.)$ in year t. In the construction of the aggregate model, we will also make use of the following relation which is true for the lognormal distribution:

$$\bar{y}_t = e^{\mu_{dt} + 0.5\sigma_{dt}^2}; \tag{4.5}$$

$\bar{y}_t$ is the arithmetic mean income in year t.

Equation (4.1) when aggregated in the light of function (4.4) and relation (4.5), yields the following deterministic aggregate model:[3]

$$E(c_{it}) = Ae^{B\ln \bar{y}_t + 0.5B(B-1)\sigma_{dt}^2}. \tag{4.6}$$

For the application of regression procedures we need to give this a stochastic touch. This is where the item (2) enters the picture. This stochastic side is to take care of the factors other than income, affecting consumers' expenditure. Again, as far as possible, we try to maintain the distinction between the micro and macro aspects of this task. These issues are discussed in the next section.

4.2 Complete Specification of the Regression Model

We do not intend to develop any special theory here; rather, different assumptions are made to bring estimation issues in line with the existing econometric approaches. We presume that a model error and/or a measurement error (in the dependent variable) are present. The model error accounts for the role of the explanatory variables excluded from equation (4.1), our micro model. Given the assumptions about the model error, proper aggregation helps one to define the stochastic equivalent of relation (4.6), that is, the theoretical quantity $E(c_{it})$ with due allowance for the stochastic elements. Normally this is also treated as $\bar{c}_t$, the data-based average consumer expenditure. But how accurately can $\bar{c}_t$ be equated with $E(c_{it})$? Here lies the role of the measurement error in the dependent variable.

We present four simple alternative specifications and related assumptions. This is followed by a rationalization of each, and a look at the associated aggregative model.

Specification	Model Error	Measurement Error in the Dependent Variable
S.1	None	$\bar{c}_t = E(c_{it})+u_t$
S.2	$c_{it} = Ay_{it}^{B}e^{w_t}$	None
S.3	$c_{it} = Ay_{it}^{B}e^{w_t+v_{it}}$	None
S.4	Same as in S.3	Same as in S.1

In these specifications U, V, W and Y are random variables; the corresponding small letters refer to the values of these random variables. The presence of the subscript "i" indicates that the associated variable is to be viewed as distributed across the consumer units as well; otherwise, only a time dimension is to be perceived with reference to the subscript "t". Complementing each specification, depending on the variables involved, are the following assumptions.

1. As mentioned in section 4.1, $f(y_{it})$, the marginal density function of Y_{it}, is lognormal.
2. $h(v_{it})$, the marginal density function of V_{it}, is continuous with mean 0 and a finite, though unknown, variance. Moreover, $E(\exp(v_{it}))$ exists and equals a finite constant A'.
3. W_t and V_{it} are mutually independent both at each point in time and across time.
4. In S.1 u_t is $N(0,\sigma_u^2)$, in S.2 and S.3 w_t is $N(0,\sigma_w^2)$, and in S.4 u_t and w_t have a bivariate normal distribution.

According to the specification S.1 we discount the possibility of leaving out any explanatory variables while specifying equation (4.1). Thus the aggregative model is deterministic as in equation (4.6). However, an allowance for potential measuement errors in the dependent variable, yields the following regression model:

$$\bar{c}_t = Ae^{B\ln \bar{y}_t + 0.5B(B-1)\sigma_{dt}^2} + u_t. \tag{4.7}$$

The specification S.2 implies that income is not the sole factor affecting the consumption decisions of our micro units. The effect of the unspecified factors is, however, embodied in w_t, and is viewed as common to every consumer unit. As an example, one may consider the 1930 depression or the recent gas price increases. The aggregate model would be defined as

$$E(c_{it}) = \int_{y_{it}} c_{it} f(y_{it}) dy_{it},$$

$$= A \int_{y_{it}} y_{it}^{B} f(y_{it}) dy_{it} e^{w_t} .$$

In the absence of any measurement errors in the dependent variable, we write the regression model for this case as

$$(4.8) \qquad \overline{c}_t = Ae^{B\ln \overline{y}_t + 0.5B(B-1)\sigma_{\overline{dt}}^{2} + w_t} .$$

The specification S.3 is more realistic than S.2. In this case only part of the effect of the left-out variables is treated as common to every consumer unit (as embodied in w_t). The total effect of these variables also includes a component v_{it} which may not be the same for all the consumer units. But as mentioned earlier, we treat V and Y as two random variables. Therefore, the aggregative model would be defined as below.

$$E(c_{it}) = \int_{y_{it}} \int_{v_{it}} A y_{it}^{B} e^{w_t + v_{it}} f(y_{it}) h(v_{it}) dy_{it} dv_{it};$$

$$= A \int_{v_{it}} e^{v_{it}} h(v_{it}) dv_{it} \int_{y_{it}} y_{it}^{B} f(y_{it}) dy_{it} e^{w_t};$$

$$= AE(e^{v_{it}}) e^{B \ln \overline{y}_t + 0.5B(B-1)\sigma_{\overline{dt}}^{2} + w_t};$$

$$= AA' e^{B \ln \overline{y}_t + 0.5B(B-1)\sigma_{\overline{dt}}^{2} + w_t};$$

$$= A'' e^{B \ln \overline{y}_t + 0.5B(B-1)\sigma_{\overline{dt}}^{2} + w_t};$$

where $A''=AA'$. Since $\overline{c}_t = E(c_{it})$, according to our assumption about the measurement error in the dependent variable, we get the following regression model:

$$(4.9) \qquad \overline{c}_t = A'' e^{B \ln \overline{y}_t + 0.5B(B-1)\sigma_{\overline{dt}}^{2} + w_t}.$$

Equation (4.9) is no different from equation (4.8) from an estimation point of view. It only raises the possibility of attaching a broader interpretation to the coefficient A in our micro model (4.1). This would be in line with Friedman's interpretation of the constant of proportionality (in his model) which depends on several factors such as the rate of interest and the ratio of nonhuman to human wealth, but not on the income variable itself.[4]

The specification S.4 offers no innovation, compared to S.3, as far as the model error is concerned. However, now the possibility of measurement error in the dependent variable is not ruled out. The construction of the aggregative regression model requires the same steps to be taken as in the case of S.3; only in the last stage we need to replace $E(c_{it})$ by $\bar{c}_t - u_t$. Therefore, the aggregative regression model for S.4 would be:

$$(4.10) \qquad \bar{c}_t = A'' e^{B \ln \bar{y}_t + 0.5B(B-1)\sigma_{\overline{dt}}^2 + w_t} + u_t.$$

We may also mention that one could assume $w_t = 0$ in the specification S.3, and allow for measurement errors in the dependent variable. In that event the regression model would look like equation (4.7) with A" as the coefficient of the exponential term, instead of A. Therefore an extended interpretation can be attached to the estimate of A in equation (4.7).[5]

To recapitulate, the preceding discussion suggests that at the estimation stage the model (4.6) can assume one of the <u>three</u> forms represented by equations (4.7), (4.8) and (4.10). Equation (4.7) requires that we apply a logarithmic transformation to the regression equation before proceeding with estimation. Equation (4.8) precludes such a step, and the dependent variable remains $\bar{c}_t$. Equation (4.10) lies in between the two extremes represented by equatons (4.7) and (4.8). Goldfeld and Quandt (1970) offer some advice on the maximum likelihood estimation of equation

(4.10); but such an exercise appears not to be justified on cost-benefit considerations. We, therefore, do not deal with equation (4.10) here, and trust that the results of the other two cases are similar enough to give us some confidence in the estimates.

4.3 Data

The data period is 1947-1976.[6]

The consumer unit in our analysis is an individual income-recipient. We use the total population in the age bracket 15 years and above (POP15) as a proxy for the total number of income-recipients.[7]

In an aggregative framework, such as ours, the implicit assumption is that there is only one good which serves both as the wage-good and the commodity consumed in the economy. This implies that the unit of account for both the income and expenditure variables ought to be the same. To meet this consistency requirement, we proceed in the following manner. First the current dollar data are drawn from the National Accounts for the monetary aggregates. Then these data are converted into their real equivalents using the CPI, Consumer Price Index---the CPI acts as our proxy for the price level in the system.[8] Accordingly, we take CEXP as the total personal expenditure on consumer goods and services corrected for government grants to hospitals, in current dollars. Then $\bar{c}_t$ is defined as "(CEXP/CPI)/POP15" in year t. Similarly, $\bar{y}_t$ equals "(YD/CPI)/POP15" in year t, where YD is the current dollar data for total personal disposable income corrected for government grants to hospitals (also in current dollars).[9]

The question of data for $\sigma_{\overline{dt}}^2$ is somewhat complicated. In principle, $\sigma_{\overline{dt}}^2$ should be based on the distribution of *aggregate real personal disposable income (YD/CPI) among the total number of income-recipients (POP15);* and it should be *calculated subject to the lognormality hypothesis* about the distribution. But there is practically no statistical information on the distribution of real personal disposable income. We are obliged to use $\sigma_{\overline{\ \ }}^2$ at time t (hereafter $\sigma_{\overline{t}}^2$) in place of $\sigma_{\overline{dt}}^2$, where $\sigma_{\overline{t}}^2$ is the parameter estimated from taxation data in the previous chapter. The $\sigma_{\overline{t}}^2$'s are based on the distribution of pre-tax personal income (in current dollars). This choice may be questioned in two ways. *One*, is it justified to use the variance of logs constructed from nominal data ($\sigma_{\overline{t}}^2$) in place of the same for real data? *Two*, is it reasonable to use $\sigma_{\overline{t}}^2$ which is based on the distribution of pre-tax personal income in place of the same for personal disposable income?

As for the first question, we see no problem. Under the pretext of a one-commodity model, it could be argued that fluctuations in the price level are going to affect every income-recipient alike. Thus inflation will only shift the mean of the distribution at a point in time; but the variance of logs of real disposable income will remain unchanged, since it is a scale-free measure. Theoretically speaking, therefore, $\sigma_{\overline{dt}}^2$ can be approximated by $\sigma_{\overline{t}}^2$ on the first count. As for the second issue, there is really no simple answer. We recall that the $\sigma_{\overline{t}}^2$'s exhibit a long-term trend toward greater inequality in the distribution of pre-tax personal income. We see no reason why the same should not be true for the distribution of real disposable income. Thus by using $\sigma_{\overline{t}}^2$ we can, at least, hope that the signs of the parameters will not be affected. This usage can also be evaluated with the help of sensitivity analysis. Accordingly, we will address this problem again in section 4.5. In our general experiments, however, we do not complicate the estimation procedure by invoking the errors-in-variables argument for using $\sigma_{\overline{t}}^2$; the reason again lies in cost-benefit considerations.

4.4 Tests for Distribution Effects: Results

In the following discussion the phrase "the additive-errors specification (or, version) of the model" refers to equation (4.7) which is

$$\overline{c}_t = Ae^{B\ln \overline{y}_t + 0.5B(B-1)\sigma_{\overline{dt}}^2} + u_t. \tag{4.7}$$

On the other hand, the terms "the multiplicative-errors specification" or "the logarithmic version (of the model)" apply to

$$(4.11) \qquad \ln \overline{c}_t = A^* + B\ln \overline{y}_t + 0.5B(B-1)\sigma_{\overline{dt}}{}^2 + w_t,$$

where $A^*=\ln A$; this is the logarithmic equivalent of equation (4.8).

Test-runs for these two models revealed the presence of significant autocorrelation. We, therefore, append both models with additional assumptions about the stochastic terms. These involve the specification of a stationary first-order autocorrelation scheme for both u_t of equation (4.7) and w_t of equation (4.11).[10] As mentioned earlier, our estimation method is equivalent to the FML procedure discussed in Beach and MacKinnon (1978). According to this procedure, the constraint on the autocorrelation coefficient (ρ), that it should lie in the open interval (-1,+1), becomes a part of the estimation process.[11] We report results both before and after the correction for autocorrelation. The maximum-likelihood parameter estimates for our models are given in Table 4.1a. The estimates under H_0:B=1 are reported in order to give a rough idea of the performance of our models.[12] Their primary purpose, however, is to provide input for the likelihood-ratio tests of the null hypothesis. A statistical summary of these tests, and those based on the asymptotic t-values, is given in Table 4.1b. Before discussing the results regarding B directly, we first comment on the validity of the models.

4.4.1 On the Validity of the Regression Models

Under this heading, the following issue is addressed: do the models (4.7) and (4.11) form valid bases for tests of the "no distribution effects" hypothesis, that is, B=1?

Let us take up the case of (4.7) first. In this model, A is the coefficient of the exponential term which contains our income and distribution variables. If the parameter A is not different from zero, both $\bar{c}_t$ and C_t will be zero for all levels of income and its distribution. In that event, this model cannot be used for any further discussions, including those of distribution effects. Thus the validity of this model hinges on the contradiction of the hypothesis A=0. Our estimates for equation (4.7) show that such an inference is warranted. Both before and after the correction for autocorrelation, the estimated A is no less than three times its asymptotic standard error.[13] This obervation, therefore, affirms our confidence in the use of the model (4.7) to test for the existence of distribution effects.

In the case of model (4.11), in principle we need to ask a similar question about A. But now A=0 if and only if

$A^*=-\infty$, because $A=\exp(A^*)$. There is no straight-forward way to cast the hypothesis A=0, and to test it in the logarithmic set-up of equation (4.11). However, the appropriateness of (4.11) for distribution effects testing can be justified in an indirect manner. We propose to do so with the help of the statistic RMSE (given in column 5 of Table 4.1a). It is the root mean-square error in predicting aggregate real consumer expenditure (CEXP/CPI)---all the RMSE's are in terms of millions of constant dollars. Before the correction for autocorrelation, RMSE is 1066 for equation (4.11) compared to 898.2 for equation (4.7); after the correction these figures are 508.6 and 474.7 for (4.11) and (4.7), respectively. Numerically, the RMSE's are slightly larger with equation (4.11) than with equation (4.7) in both the cases. However, these differences become trivial once we recall that the actual real consumer expenditure figures lie between 18953 and 74375 million dollars for the years 1947 and 1976, respectively. This means that the statistical fit of (4.11) can be regarded almost as good as that of (4.7). On this basis, we will treat equation (4.11) at par with equation (4.7) while testing for distribution effects in the current income-current expenditure framework.

Our discussion of tests for distribution effects is divided into two parts. First, in section 4.4.2 the magnitude of B and the implied relation between income inequality and aggregate real consumers' expenditure are discussed. Second, in section 4.4.3 the statistical significance of the difference of the estimated B from unity, is studied.

4.4.2 The Magnitude of B

According to our estimates in column 4 of Table 4.1a, B is always less than unity. Its estimated magnitude is:

0.9319 in (4.7), before the correction for autocorrelation;
0.8962 " " , after " " " " " " " ;
0.9627 " (4.11), before " " " " " " " ; and
0.9346 " " , after " " " " " " " .

This set of results has two implications. First, for a typical consumer unit the mpc (out of real disposable income) declines as the level of real disposable income rises. Second, +0.5B(B-1), the coefficient of our distribution variable $\sigma_{\overline{dt}}{}^2$, is negative. This, in turn, yields a conclusion that conforms to the conventional wisdom:

$$(\partial \bar{c}_t / \partial \sigma_{\overline{dt}}{}^2) < 0;$$

and, for a given number of income-recipients:

$$(\partial C_t / \partial \sigma_{\overline{dt}}{}^2) < 0.$$

That is, given the mean income $\overline{y}_t$, as the distribution of aggregate real personal disposable income becomes more unequal, the level of aggregate real consumers' expenditure would decline, and *vice versa* .

4.4.3 The Statistical Significance of B Being Different from Unity

As proposed in section 4.1, H_0:B=1 is tested against H_a:B$\neq$1. The likelihood-ratio test statistic -2ln λ is defined as the negative of twice the difference between the log-likelihood for the constrained model (with B=1) and that for the unconstrained one. For each of the additive- and multiplicative-errors versions of our model, two such statistics are calculated---one before the correction for autocorrelation, and the other after such correction. In large samples, -2.ln λ has a chi-square distribution with degrees of freedom equal to the number of constrained parameters. In our case this would be a chi-square distribution with one degree of freedom. Based on this theoretical distribution, we report critical values for -2.ln λ both at the 5% and 10% levels of statistical significance in Table 4.1b (column 4, against -2.ln λ). Moreover, following the example of Spitzer (1977), the asymptotic standard errors are used to define asymptotic t-values under the null hypothesis B=1 for models (4.7) and (4.11), both before and after the correction for autocorrelation. The question of degrees of freedom for the t-statistic, when models are nonlinear in parameters, is not entirely clear; however, to be conservative, the rule of thumb used is the number of observations minus the number of parameters.[14] We report the critical t-values in Table 4.1b (column 4, against $t_{(B=1)}$) for a two-tailed test at the 5% level of statistical significance. The absolute magnitude of the theoretical t at the 10% level of significance is smaller than the reported ones, and is omitted.

Based on Part I of Table 4.1b, our general finding is the following. If the correction for autocorrelation is ignored, the null hypothesis B=1 is always rejected. This is so whether we use the likelihood-ratio test or the asymptotic t-values, both at the 5% and 10% levels of statistical significance.

Part II of Table 4.1b contains results after the autocorrelation correction. The correction for serial correlation means replacement of the dependent variable by a transformed one; and the same applies to the independent variables (or functions of independent variables, as in the additive-errors case). This is accompanied by an improved fit for the transformed models, as gauged by the sum of squared-residuals, and an increase in the associated log-likelihood. However, if the autocorrelation phenomenon is more pronounced in the constrained than in the unconstrained case, one will have to reconsider some of the conclusions reached <u>before</u> the correction for serial correlation. This is exactly what happens in our case. For the additive-errors version, the new $-2.\ln \lambda$ is 2.8238 compared to 14.7076 before the correction for autocorrelation; and for the multiplicative-errors specification it reduces to 3.4676 from 4.7834 before the correction. The critical values for $-2.\ln \lambda$ are still 3.8415 and 2.7055 at the 5% and 10% levels of significance, respectively. Consequently, B is indiscernable from unity at the 5% significance level. But if one chooses this probability to be 0.10 (or 10%), the null hypothesis does not hold ground; that is, the existence of distribution effects becomes credible.

The use of asymptotic t-values produces further ambiguities---again after the correction for autocorrelation. For the multiplicative-errors specification, the result corresponds to the likelihood-ratio based one. However, for the additive-errors case, this correspondence does not hold; the asymptotic t-value calls for a rejection of H_0:B=1 at the 5% level of significance! We have no explanation for this anomaly. It can only be said that the likelihood-ratio test has a strong theoretical basis compared to the test with the asymptotic t-value. So we would recommend in favour of the results based on the likelihood-ratio method. In reporting further results, we will confine ourselves to the likelihood-ratio test; but on most occasions the asymptotic standard errors will also be reported for the interested researchers.

Next, we proceed to the sensitivity analysis of these results.

4.5 Sensitivity Analysis

Here we examine three issues. The first relates to the data on $\sigma_{\overline{dt}}^2$. The second is about a technicality. An additive constant term is a prominent feature of the aggregative consumption models in the current income-current expenditure framework. Our micro and, hence, aggregate models do not have one. As we will see in section 4.5.2, its existence has a direct bearing on the estimates of B and, therefore, the existence of distribution effects. So we check how far this neglect is responsible for our results. The third point deals with the sensitivity of our results to the presence of simultaneity bias in our data on the income and expenditure variables. Each of these three points is addressed separately, starting with the first in section 4.5.1.

4.5.1 Sensitivity of the Results to Data on the Variance of Logs of Real Personal Disposable Income

The reader may recall, from our discussion in section 4.3, that we use the data on the variance of logs of pre-tax personal income in period t for $\sigma_{\overline{dt}}^2$ in our experiments. Now we follow an approach of Blinder (1975), and use a one-period lagged value of the same (that is, $\sigma_{\overline{t-1}}^2$) for the true $\sigma_{\overline{dt}}^2$.[15] Re-estimation of equations (4.7) and (4.11) results in a slight increase in the asymptotic standard errors of the parameter estimates; the likelihood-ratio statistics also follow this trend. Generally the results in both magnitude and conclusions are the same as before, however, with one exception. This exception arises in the event of correction for serial correlation. The autocorrelation-corrected estimates for both the models (4.7) and (4.11), along with tests of H_0:B=1, are given in Table 4.2.[16] Unlike before, we notice that the no-distribution effects proposition is rejected now even at the 5% level of significance for model (4.11). This is puzzling in view of the fact that the asymptotic standard error of B actually registers an increase with $\sigma_{\overline{t-1}}^2$ used instead of $\sigma_{\overline{t}}^2$. However, it is too early to comment on the importance of this single piece of evidence. A meaningful discussion of it should await the clarification of other issues under consideration. In particular, we must re-examine this when the simultaneity bias issues are addressed in section 4.5.3. For now we will act on the premise that the problem of finding the appropriate proxy for $\sigma_{\overline{dt}}^2$ is not important. With this in mind, in the remaining experiments we will approximate $\sigma_{\overline{dt}}^2$ by $\sigma_{\overline{t}}^2$, except when mentioned otherwise.

4.5.2 Test for an Additive Constant Term in the Micro Model (4.1) and Its Associated Aggregative Models

The importance of the exclusion of an additive constant term when it should have been a part of equation (4.1) (and its associated aggregative models) is best explained by a reference to Figures 4.1a and 4.1b (found on page 102). Both these figures are based on hypothetical data on average (real) consumer expenditure $\bar{c}_t$ and average (real) disposable income $\bar{y}_t$. According to Figure 4.1a, the data dictate the relation OAB. But if we suppress the constant term, a nonlinear relation OD may find support from the data. OD is the case of declining mpc. For $\overline{\text{OAB}}$, however, it is constant at all levels of income (of course, except $\bar{y}_t=0$ at which point it is undefined). Clearly, for OAB the distribution would not matter. On the other hand, Figure 4.1b shows a case when the mpc increases with income, along the true relation OAB. This case would yield B>1, and a positive coefficient for the distribution variable. But again, constraining the relation through the origin may produce quite the opposite result, that is, the relation OD which has B less than unity.

To test for the existence of an additive constant term implied by our data, we postulate

(4.12) $$c_{it} = A_s + Ay_{it}^{B},$$

and test for

(4.13) $\overline{H}_0$: $A_s = 0$,

against

(4.14) $\overline{H}_a$: $A_s \neq 0$.

For obvious reasons, a more appropriate alternate hypothesis is $A_s>0$. But a straight-forward implementation of the likelihood-ratio test for such a constraint is not possible; so we retain (4.14). Under the error specification S.1 and the lognormality hypothesis about y_{it}'s, the aggregative regression model for (4.12) is

$$(4.15) \qquad \overline{c}_t = A_s + Ae^{B\ln \overline{y}_t + 0.5B(B-1)\sigma_{\overline{dt}}^2} + u_t;$$

and with specification S.2, it becomes

$$(4.16) \qquad \ln (\overline{c}_t - A_s) = A^* + B\ln \overline{y}_t + 0.5B(B-1)\sigma_{\overline{dt}}^2 + w_t.$$

It should be readily apparent that equation (4.16) is more difficult to handle than equation (4.15), especially if we want to search for both the autocorrelation coefficient ρ and A_s. At this stage we recall the broad similarities in our earlier results between equations (4.7) and (4.11), and confine our attention now to (4.15).

The regression model under $\overline{H}_0$:A_s=0 is identical to equation (4.7). The estimates for this model may be found in Table 4.1a---row 2, for estimates before the correction for autocorrelation, and row 6, for the same after the correction. The estimates of equation (4.15), both before and after the correction for autocorrelation, are given in Table 4.3. We report both the uncorrected and corrected versions because of the interesting contrasts they offer regarding the magnitudes of A and B; the reader is referred to these in Table 4.3. In either case the asymptotic standard error of A is greater than the magnitude of the parameter itself. This difference is especially great in the case of the autocorrelation-corrected estimates where the asymptotic standard error of A is nearly 250 percent of the size of the coefficient. This obviously casts doubt on the role of income in the first place! Moreover, the estimates after the correction for autocorrelation are the ones that really matter. According to these, the hypothesis $A_s=0$ has strong support from our data. The calculated value of $-2\ln\lambda$ (=1.042) is far below 3.8415 (or 2.7055), the critical value based on the chi-square distribution with one degree of freedom at the 5% (or the 10%) level of significance.

In the light of the evidence presented, the following is our assessment: the exclusion of a constant term from the micro model (4.1) and its associated aggregate models has no bearing on our finding of B<1.

4.5.3 The Impact of Simultaneity Bias

In our tests in section 4.4, current consumer expenditure and current personal disposable income, respectively, have been used to define our dependent variable and one of the independent variables. It may be argued that the causation is not strictly from income to consumer expenditure. From the expenditure side of the National Accounts, it is obvious that personal expenditure on consumer goods and services is one of the determinants of income, in fact, the major one. This led Friedman and Becker, in their criticism of the current income-current expenditure consumption functions, to remark: "Consumption is, as it were, being correlated with itself."[17] From the point of view of the sum of squared-residuals and the correlation between the dependent variable and regressors, the odds will be in favour of finding a strong relation. But the presence of this simultaneity is likely to affect the coefficient estimates also. Haavelmo (1947) is the first one to recognise this possibility. He uses a two-equation model consisting of a linear regression relation between consumer expenditure and income, and an income-expenditure identity. In this context Haavelmo shows that the direct regression of consumer expenditure on income, according to the least-squares principle, will not yield consistent estimates. More specifically, he concludes that the constant term in the (linear) consumption function will be under-estimated and the mpc over-estimated, in the presence of the simultaneity bias. To remedy this problem, he suggests the use of indirect least-squares for his two-equation model.

In our nonlinear consumption model, we have similarly followed a two-stage least-squares type (2SLS-type) procedure to study the impact of simultaneity bias. The first stage involves the search for an appropriate proxy for real personal disposable income in the context of a theoretical model to be specified below. This proxy is used at the second stage in place of YDR (=YD/CPI) in estimating the multiplicative-errors version of our model, that is, equation (4.11). The same exercise could also be repeated for equation (4.7), the additive-errors vintage of our model. Part of the reason for our going ahead with (4.11) lies in the similarities in the results reported in Tables 4.1a and 4.1b. From this point of view, an exercise with only one version of the model should suffice. However, working with equation (4.11) will also enable us to re-examine a potential anomaly reported in section 4.5.1 regarding the use of the lagged σ_{-}^{2} in place of $\sigma_{\overline{dt}}^{2}$. We will come to this point later in this section.

Mathematically, the causal process linking consumer expenditure and personal disposable income may be specified in the following manner:[18]

(4.17a) $$CEXPR = r(YDR, N, \sigma_{\overline{dt}}^{2});$$

(4.17b) $$YDR \equiv YPR - TR;$$

(4.17c) $$YPR = s(GNPR); \quad \text{and}$$

(4.17d) $$GNPR \equiv GNER \equiv CEXPR + IR + GR + XR - MR.$$

The letter "R" at the end of a variable's name indicates that it has been converted into real terms. Apart from this, CEXP and YD, respectively, are personal expenditure on consumer goods and services and personal disposable income (both in current dollars) as before. Among the other monetary aggregates, YP is personal income; I, G, T, X and M, in this order, stand for investment, government expenditure, taxes, exports and imports; and GNP and GNE are Gross National Product and Gross National Expenditure, respectively (all in current dollars). Of the remaining variables, N=POP15 and σ_{dt}^2 is the variance of logs, as defined in section 4.3. The time subscript is omitted from names of almost all variables, but the values in every case belong to the same time period. This setting follows the National Income Accounts, except that equations (4.17a) and (4.17c) are behavioural relations.

Equation (4.17a) is our aggregate consumption function written in a general form; N and σ_{dt}^2 are exogenous variables, and YDR is related to CEXPR via (4.17b), (4.17c) and (4.17d). (4.17b) and (4.17d) are identities based on the National Accounts, whereas equation (4.17c) implies that YPR is some function "s" of current (real) GNP. In this framework, the causation from real personal disposable income to real consumer expenditure, and back, works in the manner shown in Figure 4.2.

The ideal procedure to tackle the problem of simultaneity bias would be to write down the complete model, that is, to specify each element in the system (4.17a)-(4.17d) precisely, and to follow an appropriate simultaneous equation estimation technique. For our purpose, however, we look only at the main components of such a model, and supplement the four-equation model by the following relations.

(4.17e) $TR = TRATE.YPR$, where $TRATE \equiv (TR/YPR)$;

(4.17f) $IR = I(GNPR, GNPR_{-1}, GNPR_{-2})$;

(4.17g) $GR = \overline{GR}$;

(4.17h) $XR = \overline{XR}$; and

(4.17i) $MR = m(CEXPR, IR, \overline{GR}, \overline{XR})$.

$GNPR_{-1}$ and $GNPR_{-2}$ are one- and two-period lagged values of GNPR, respectively. A bar on GR and XR indicates that these are treated as exogenous variables. (4.17e) is based on the definition of income-tax rate (TRATE); (4.17f) is a simple investment function; and (4.17i) is an import function.

Combining (4.17a)-(4.17i) yields a solution for YDR to this effect:

$$(4.18) \qquad YDR = z(GNPR_{-1}, GNPR_{-2}, \overline{GR}, \overline{XR}, TRATE, N, \sigma_{\overline{dt}}{}^{2}).$$

Of course, we are assuming that there are no nonlinearities in terms of variables, at any step leading to the solution (4.18). Next we estimate the following linear approximation to function (4.18):[19]

$$(4.19) \qquad YDR = a_0 + a_1 GNPR_{-1} + a_2 GNPR_{-2} + a_3\overline{GR} + a_4\overline{XR} + a_5 TRATE + a_6 N + a_7\sigma_{\overline{dt}}{}^{2} + \text{error term}.$$

Again as in section 4.3, all monetary aggregates have been deflated by the CPI in order to arrive at their real values.[20] Least-squares estimation of equation (4.19) provides us with an estimate for $\widehat{YDR}$, the predicted value of YDR. Statistical results for this exercise are reported in Part A of Table 4.4. This completes the first stage of our 2SLS-type procedure. At the second stage, $\widehat{YDR}$ generated in the afore-mentioned fashion is used in place of YDR. Keeping the other data the same <u>as explained in section 4.3</u>, we re-estimate equation (4.11) both under H_0:B=1 and H_a:B$\neq$1. The maximum likelihood estimates, both before and after the correction for autocorrelation, make up Part B of Table 4.4.

The use of $\widehat{YDR}$, instead of YDR which is an endogenous variable, ensures that none of the regressors is correlated with the stochastic term in the model. This should improve the consistency aspect of the parameter estimates. As mentioned above, Haavelmo also predicts a decline in the mpc after accounting for simultaneity bias. There is no strong reason to expect a strict analogue of this result to hold in nonlinear models. Our results, in fact, show almost no change in the magnitude of B, the income elasticity, before the autocorrelation correction. However, after the correction for serial correlation B registers a small decline, in line with Haavelmo's conclusion. The important thing to look for is the value of the log-likelihood function. After the adjustment for simultaneity bias, it declines in every case; the absolute magnitude of the decline is greater for the unrestricted model than for the restricted one, both before and after the correction for autocorrelation. As for the impact on the conclusion about distribution effects, we experience only one reversal. Now the null hypothesis B=1 is not rejected at the 5% level of significance even before the autocorrelation correction; otherwise, the results are the same as reported in section 4.4. In view of this finding, we ask whether it is possible to explain the odd result noted in section 4.5.1.

Just to recall the issue, the equation (4.11)-based results with $\sigma_{\overline{t-1}}^2$ (as a proxy for $\sigma_{\overline{dt}}^2$) contradicted H_0:B=1 at the 5% level of significance after the correction for serial correlation. At that stage we speculated that this result might not hold if one took into account the simultaneity bias also. The present querry is an attempt to verify that conjecture. For this purpose we replace data on YDR by $\widehat{YDR}$ which is used to determine the other results in this section. But we then re-estimate equation (4.11) with $\sigma_{\overline{t-1}}^2$, as in section 4.5.1. Notice that $\widehat{YDR}$ itself is based on $\sigma_{\overline{t}}^2$. It may be recommended that $\widehat{YDR}$ must be reconstructed, in the light of equation (4.19), with $\sigma_{\overline{t-1}}^2$ (as the proxy for $\sigma_{\overline{dt}}^2$). But this is not necessary. $\widehat{YDR}$ is an exogenous quantity as far as the second stage of our correction for simultaneity bias is concerned. Moreover, the advantage of using $\widehat{YDR}$ (in its present form) is that our results will also relate to the others determined in this section. Briefly, re-estimation of equation (4.11), with $\widehat{YDR}$ and $\sigma_{\overline{t-1}}^2$, yields the following results after correction for serial correlation.

	ρ	A^*	B
Estimated coefficient: (asymptotic standard error)	.693	.5459 (.3310)	.9264 (.0392)
DW=1.83	ln L=72.5508		$-2\ln\lambda$ =3.4028

The new value of the likelihood-ratio test statistic (=3.4028) is less than the critical value (3.8415) at the 5% level of significance. This, therefore, confirms our earlier reservations about the contradiction of H_0:B=1 at the 5% level of significance after the correction for serial correlation, with equation (4.11) and $\sigma_{\overline{t-1}}^2$. This clarifies the only result that was not in line with our other findings.

To sum up then, the problem of simultaneity bias may have important bearings on the quality and, perhaps, the magnitude of coefficient estimates. In the case of our work, the magnitude of B registers no significant change. However, three things regarding the null hypothesis B=1 come to light, after an adjustment for the simultaneity bias. One, unlike our initial finding in section 4.4, at the 5% significance level the support for H_0:B=1, before the correction for autocorrelation, is doubtful. Two, the H_0:B=1 definitely cannot be rejected at the 5% level of significance after the correction for serial correlation. Three, at the 10% significance level there remains no doubt about the contradiction of H_0:B=1, both before and after the correction for serial correlation.

4.6 Summary and Conclusions

In this chapter the analysis is carried out with aggregative (time-series) data, but it is based on micro foundations. The maintained hypothesis is that the distribution of real disposable income follows the lognormal pattern, in every year. The process of going from the micro to the macro consumption model determines two things: (1) the choice of the distribution variable σ_{dt}^2 (the variance of logs of real disposable income); and (2) the constraint that should be placed on the coefficient of σ_{dt}^2. The test is simultaneously a test of the no-distribution effects hypothesis and the nonvariability of the mpc with respect to income at the micro level---in fact, one implies the other. _Ex post_ B, the income elasticity of consumption in our model, turns out to be less than unity. This result implies that the mpc for a typical consumer unit declines with its income level, and that σ_{dt}^2 has a negative coefficient in the aggregative model. Thus the conventional wisdom is established in the context of the current income-current expenditure theory of consumption. _Before the correction for autocorrelation_, the hypothesis of unitary B is generally rejected both at the 5% and 10% significance levels. However, the hypothesis is not rejected at the 5% level if one also takes into account the simultaneity bias. _After the correction for autocorrelation_, the result of a nonunitary B is supported, though only at the 10% level of significance. Of course, all these results are conditional upon the functional form of the distribution of real disposable income (among the income-recipients) being lognormal. They are not, however, sensitive to the following: (1) use of the particular data on the variance of logs; and (2) exclusion of an additive constant term from the analysis.

FOOTNOTES TO CHAPTER 4

1. See, for example, Branson (1979,p.184).

2. No particular functional form can be traced to Keynes (1936), however.

3. See Chapter 2,pp.22-23, for its derivation.

4. We will explain Friedman's model in Chapter 5.

5. In our analysis of stochastic errors, we try all but one of the sources through which they may be introduced to obtain a regression model. The neglected one is the possibility of measurement errors in our regressors, that is, $\bar{y}_t$ and σ_{dt}^2. An attempt in this direction carries theoretical complications which have no straightforward answer in the nonlinear estimation framework, at least, in the present state of the art. So we avoid the issue.

6. Definitionally speaking, the consumer expenditure figure in 1946 is not consistent with those in the other years. So the year 1946 is dropped. In section 4.5.1, and also once in section 4.5.3, we use the variance of logs figure for 1946. Similarly, while correcting for the simultaneity bias in section 4.5.3, we also use the Gross National Product data for 1945 and 1946. In all these cases, however, the first observation on consumers' expenditure belongs to 1947. Thus each equation is estimated with 30 observations.

7. Keeping in view the argument of Appendix A, an alternative proxy for the number of income-recipients is RET, the total number of all returns (both taxable and non-taxable). On the criterion of the maximum log-likelihood <u>standardized</u> by the mean of the dependent variable, however, the estimates are inferior with RET. So it has been dropped in favour of the POP15 series. The annual figures for POP15, for 1947-1976, are (in millions):

8.751*(in 1947)	8.881*	9.220*	9.342*	9.452
9.673 (in 1952)	9.879	10.099	10.298	10.501
10.803 (in 1957)	11.053	11.256	11.467	11.672
11.880 (in 1962)	12.115	12.382	12.682	13.017

13.403 (in 1967)	13.780	14.140	14.504	14.864
15.214 (in 1972)	15.586	16.018	16.470	16.873.

The starred figures are our estimates. The population age 15 years and over is published only for the latter years of the period. The series for population age 15 and over back to 1951 was constructed at McMaster University by Professors F.T. Denton, A.L. Robb and B.G. Spencer.

8. In 1971, CPI=1.0. The annual data for the period 1947-1976 are given below.

.492 (in 1947)	.562	.580	.597	.660	.676
.670 (in 1953)	.674	.675	.685	.707	.726
.734 (in 1959)	.743	.750	.759	.772	.786
.805 (in 1965)	.835	.865	.900	.941	.972
1.000 (in 1971)	1.048	1.127	1.250	1.385	1.489.

The CPI data for 1945 and 1946 (used only in the work in section 4.5.3) are .435 and .449, respectively.

For the more recent years these figures are comparable to those in the Canada Year Book (Special Edition, 1976-77,p.1022). For the remaining years the Canada Year Book (1972,p.1050) figures are converted to the 1971 base.

9. In 1961, non-profit hospitals were transferred from the personal sector to government, so government grants to hospitals have been discontinued as of 1961. Prior to 1961, the disposable income and consumer expenditure data includes this item. For consistency, we subtract this element from all the data.

The sources of all monetary data used in this thesis are the following volumes of Statistics Canada (National Income and Expenditure Accounts):

1) Catalogue 13-551 Occasional, and
2) Catalogue 13-201 Annual (the November-1978 issue).

For each series we used the latest data in these volumes.

10. The autocorrelation scheme for u_t:

$$u_t = \rho u_{t-1} + \epsilon_t \quad (t=1,2,\ldots,n);$$

where ϵ_t is distributed as $N(0,\sigma_\epsilon^2)$, for all t;

$Cov(\epsilon_t, \epsilon_s) = 0$, for all $t,s=1,\ldots,n$, except $s=t$;

$$Cov(\epsilon_t, u_{t-1}) = 0, \text{ for all } t;$$

and u_0 is distributed as $N(0, \sigma_\epsilon^2/(1-\rho^2))$.

The autocorrelation scheme for w_t is also the same, however, with "w" in place of "u".

11. Suppose the model at hand can be re-written as

(i) $$y_t = g(X_t, b) + u_t, \qquad (t=1,2,3...n)$$

where y stands for the dependent variable, X and b symbolize parameters, respectively, and u follows the stationary first-order autocorrelation process given in footnote 10. For a given ρ, the model (i) can be transformed as follows:

(ii) $$\sqrt{(1-\rho^2)}y_1 = \sqrt{(1-\rho^2)}g(X_1,b) + \epsilon_1, \qquad \text{for } t=1, \text{ and}$$
$$y_t - \rho y_{t-1} = g(X_t,b) - \rho g(X_{t-1},b) + \epsilon_t, \quad \text{for } t=2,3...n.$$

Under the assumptions about ϵ_t, given in footnote 10, the concentrated log-likelihood function (ln L) for this model can be shown to equal:

(iii) $$\ln L = (-n/2)\{1+\ln(2\pi/n)\} + (1/2)\{\ln(1-\rho^2)\} - (n/2)\{\ln(S(b))\},$$

where

(iv) $$S(b) = \{\sqrt{(1-\rho^2)}y_1 - \sqrt{(1-\rho^2)}g(X_1,b)\}^2 + \sum\{y_t - \rho y_{t-1} - g(X_t,b) + \rho g(X_{t-1},b)\}^2,$$

the summation is from t=2 to t=n. This log-likelihood function forms the basis of our FML procedure.

For a given ρ, maximizing ln L is the same as minimizing S(b) with respect to b. We use the nonlinear least-squares routine LSQ in the computer package TSP, Version 2.5, to minimize S(b) with ρ <u>a priori</u> specified. In the process different initial values of other parameters are tried to ensure the uniqueness of the solution. Then for the given ρ and the parameter values so determined, ln L is calculated. This exercise is repeated until ln L is maximized---we continue our search until a three-digit estimate of ρ is established. The asymptotic standard errors reported in our tables correspond to this maximum of ln L, and not the global minimum of S(b).

12. Under the hypothesis B=1, (4.7) becomes

$$\bar{c}_t = A\bar{y}_t + u_t,$$

and (4.11) reduces to

$$\ln \bar{c}_t = A^* + \ln \bar{y}_t + w_t.$$

13. The parameter estimates for (4.7) show that A is 1.6495 before the the correction for autocorrelation, and 2.2222 after the correction; the respective asymptotic standard errors are 0.2272 and 0.7061.

14. ρ is treated as a parameter in the event of correction for autocorrelation.

15. Blinder attributes this approach to a suggestion of Lubell. See Blinder (1975,p.465).

16. Equation estimates before the correction for serial correlation are available upon request.

17. Friedman and Becker (1957,p.70).

18. The author owes much of the following argument to the insight of Professor F.T. Denton.

19. This linear approximation was found to yield better results (in terms of predicting YDR) than an alternative dougle-logarithmic approximation.

20. See footnote 9 for the data sources. The data on Personal Income (used in the definition of TRATE) and GNP are also adjusted for government grants to hospitals.

Table 4.1a Maximum-Likelihood Estimates of Equations (4.7) and (4.11), Before and After the Correction for Autocorrelation

Row	Regression Equation, Hypothesis (1)	Parameter Estimates (a.s.e) ρ (2)	CONST (3)	B (4)	Some Relevant Statistics RMSE (5)	DW (6)	ln L (7)
	Part I	**Before the Correction for Autocorrelation**					
1	4.7, H_0	.000	.9197 (.0050)	1.0	1341.	.28	-178.3317
2	" H_a	.000	1.6495 (.2272)	.9319 (.0161)	898.2	.39	-170.9779
3	4.11, H_0	.000	-.0773 (.0050)	1.0	1430.	.47	66.1163
4	" H_a	.000	.2389 (.1444)	.9627 (.0171)	1066.	.53	68.5080
	Part II	**After the Correction for Autocorrelation**					
5	4.7, H_0	.998	.7391 (.0595)	1.0	868.5	2.08	-157.0124
6	" H_a	.838	2.2222 (.7061)	.8962 (.0369)	474.7	1.73	-155.6005
7	4.11, H_0	.779	-.0811 (.0134)	1.0	644.4	1.86	77.8676
8	" H_a	.747	.4735 (.2955)	.9346 (.0350)	508.6	1.80	79.6014

Note: a.s.e. = asymptotic standard error;
CONST = A for equation (4.7), and A^* for (4.11);
RMSE = the Root Mean-Square Error in predicting aggregate real consumer expenditure;
DW = the Durbin-Watson statistic; and
ln L = log-likelihood.

Table 4.1b Tests for Distribution Effects Based on Table 4.1a (H_0:B=1; H_a:B$\neq$1)

Row	Regression Equation (Error Type) (1)	Test-Statistic (2)	Calculated Value (3)	Critical Value (4)	Conclusion for Null Hypothesis (5)
	Part I Before the Correction for Autocorrelation				
1	4.7 (additive)	$-2\ln\lambda$	14.7076	3.8415 (2.7055)	Rejected (Same as above)
2		$t_{(B=1)}$	-4.2298	2.048	Same as above
3	4.11 (multiplicative)	$-2\ln\lambda$	4.7834	3.8415 (2.7055)	Same as above (Same as above)
4		$t_{(B=1)}$	-2.1813	2.048	Same as above
	Part II After the Correction for Autocorrelation				
5	4.7 (additive)	$-2\ln\lambda$	2.8238	3.8415 (2.7055)	Not Rejected (Rejected)
6		$t_{(B=1)}$	-2.8310	2.052	Same as above
7	4.11 (multiplicative)	$-2\ln\lambda$	3.4676	3.8415 (2.7055)	Not Rejected (Rejected)
8		$t_{(B=1)}$	-1.8657	2.052	Not Rejected

Note on Critical Values and Conclusions: The statistics and comments in parentheses relate to the 10% level of significance.

Table 4.2 Sensitivity of the Results to Data on σ_{dt}^2

Estimates and Results After the Correction for Autocorrelation[1]

Part I Equation (4.7) with Lagged σ^2

ρ	A	B	RMSE	DW	ln L
.846	2.3335 (.7554)	.8905 (.0376)	468.2	1.73	-155.2635

$-2\ln\lambda = 3.4978$ Critical Value = 3.8415 (2.7055)

Conclusion for H_0:B=1 :- Not Rejected (Rejected)

Part II Equation (4.11) with Lagged σ^2

ρ	A^*	B	RMSE	DW	ln L
.754	.5140 (.3013)	.9297 (.0357)	497.8	1.81	79.8048

$-2\ln\lambda = 3.8744$ Critical Value = 3.8415 (2.7055)

Conclusion for H_0:B=1 :- Rejected (Same as above)

1. For Part A, comparable estimates under H_0:B=1 are given in row 5 of Table 4.1a, while for Part B the same are reported in row 7 of that table.

Note: For notation and outlay of this table, see footnotes to Tables 4.1a and 4.1b.

Table 4.3 Tests for the Constant Term (A_s) in the Micro Model (4.1)---$\bar{H}_0$:A_s=0, $\bar{H}_a$:$A_s \neq 0$

Parameter Estimates	Some Relevant Statistics, and Results
Regression Equation (4.15), Not Corrected For Autocorrelation[1]	
A_s = -3961.76 (1930.43)	RMSE = 668.9
	DW = .95
A = 187.212 (222.571)	ln L = -162.3684
B = .4603 (.1133)	-2ln λ = 17.2190
	CONCLUSION: H_0 Rejected[2]
Regression Equation (4.15), Corrected for Autocorrelation (ρ=.937)[3]	
A_s = 884.97 (613.51)	RMSE = 482.6
	DW = 1.69
A = .187 (.471)	ln L = -155.0795
B = 1.1442 (.2577)	-2ln λ = 1.0420
	CONCLUSION: H_0 Upheld[2]

1. For comparable estimates under $\bar{H}_0$: A_s=0, see row 2 of Table 4.1a.

2. This result holds for both the 5% and 10% levels of significance, the respective critical values being 3.8415 and 2.7055 (based on a chi-square distribution with one degree of freedom).

3. For comparable estimates under $\bar{H}_0$: A_s=0, see row 6 of Table 4.1a.

Table 4.4 Correction for Simultaneity Bias, and Tests for Distribution Effects

---

Part I **Equation for the Predicted YDR**

$$YDR = 3182.24 + 0.29GNPR_{-1} + 0.16GNPR_{-2} + 0.54GR + 0.52XR - 94660.30TRATE + 1237.66N - 5315.34\sigma_{dt}^{2}$$

$R^2 = .9989$ DW = 1.84 F-Statistic(7,22) = 2883.57

Correlation Coefficient between YDR and YDR = .9995

Part II **The Maximum Likelihood Estimates for Equation (4.11), and Conclusions about H_0: B=1** [1,2,3]

ρ	A^*	B	ln L	$-2\ln\lambda$	CONCLUSION
(1)	(2)	(3)	(4)	(5)	(6)
		Before the Correction for Autocorrelation			
.000	-.0773 (.0773)	1.0	62.0830 (66.1163)		
				3.7062 (4.7834)	Not Rejected (Rejected)
.000	.2442 (.2389)	.9620 (.9627)	63.9361 (68.5080)		
		After the Correction for Autocorrelation			
.726 (.779)	-.0769 (-.0811)	1.0	70.8494 (77.8676)		
				3.1794 (3.4676)	Not Rejected (Rejected)
.689 (.747)	.5187 (.4735)	.9296 (.9347)	72.4391 (79.6014)		

(Continued on the next page)

Table 4.4 (Continued)

1. In the cases where B is indicated as being precisely 1.0, the etimates are constructed under the assumption that the null hypothesis is true.

2. In columns 1 to 5, the estimates in parentheses are not the asyptotic standard errors, but the estimates of the same parameters as reported in Table 4.1a. These are provided for comparison purposes only, and are not involved in the present tests (reported in column 6).

3. In column 6, the comments in parentheses relate to the 10% level of significance, while those immediately above them are arrived at for the 5% level of significance. The critical values for the test-statistic $-2\ln\lambda$ are given in Table 4.1b.

Figure 4.1 Implications of Ignoring an additive Constant term in the Consumption Function

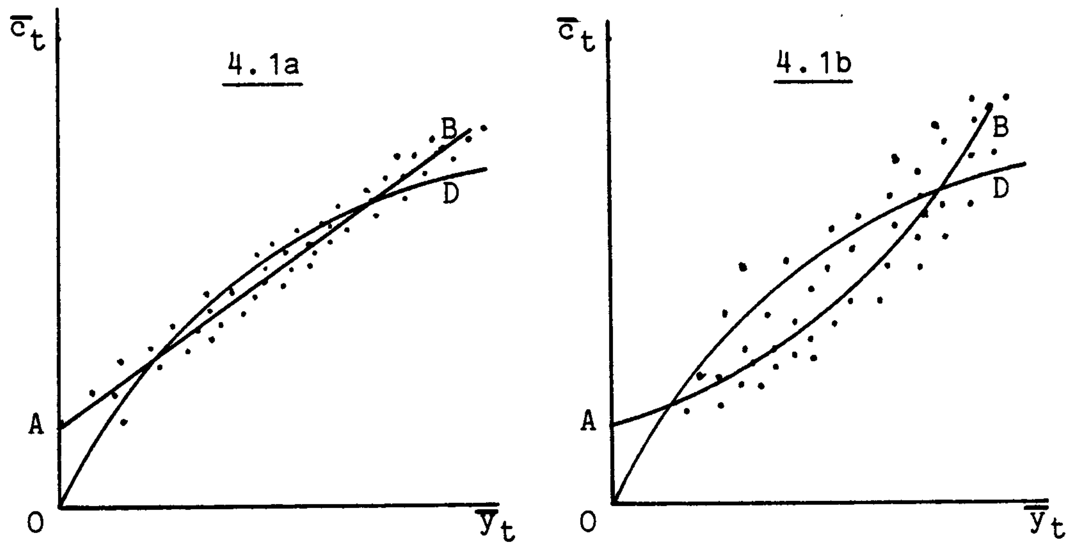

NOTES: 1) Dots represent hypothetical data.
2) OAB is the true relation in either case.
3) OD is the implied relation if the constant OA is supressed in estimation.

Figure 4.2 Causal Link between Consumer Expenditure and Disposable Income

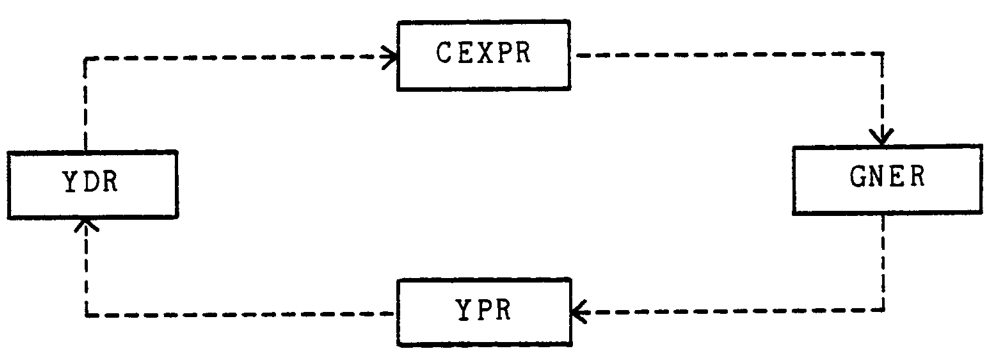

CHAPTER 5

TESTING FOR THE DISTRIBUTION OF REAL PERMANENT INCOME EFFECTS IN THE FRAMEWORK OF THE PERMANENT INCOME HYPOTHESIS

In this chapter we attempt to answer Question B of Chapter 1: Does the distribution of real permanent income matter in the Permanent Income Hypothesis (PIH) framework?

The PIH of Friedman (1957) offers a proposition about the long-term propensity to consume; it has come to be known as "the proportionality thesis". This proposition, together with the other postulates of the PIH, implies that aggregate consumption depends only on the first moment of the distribution of real permanent income, which is the mean permanent income. That is, the dispersion and skewness parameters of this distribution play no role in the aggregate consumption function. A logical implication of this view is that government demand-management policies involving only a redistribution of real permanent income will be neutral with respect to the level of aggregate consumption. This aspect of Friedman's PIH can be analyzed effectively using the approach suggested in Chapter 2, and employed in Chapter 4. Again, the micro unit in our analysis is an "individual income-recipient". We keep all the postulates of the PIH; yet a few additional assumptions are needed. In particular, _a priori_ knowledge of the functional form of the distribution of real permanent income is indispensible. We choose the lognormality hypothesis. The plan of this chapter is as follows.

The PIH is restated in section 5.1. This effort defines the dimensions of our argument. The problem is formally stated in section 5.2. The additional assumptions appended (by us) to Friedman's model are discussed in section 5.3. Matters relating to the data occupy section 5.4. A complete account of the estimation and tests is given in section 5.5. The null hypothesis addressed is the unitary elasticity of consumption out of real permanent income. This is synonymous with the "no distribution (of real permanent income) effects" hypothesis in our study. It is shown to have empirical support conditional, of course, upon the lognormality proposition about the distribution of real permanent income. This general finding, however, is robust with respect to _one,_ the presence of serial correlation, _two,_ the possibility of a constant term embedded in the data, and _three,_ an asset component in the consumption variable; these points are explained in section 5.6.

5.1 The Permanent Income Hypothesis Restated

In this restatement we draw upon Chapters II, III and V of "A Theory of the Consumption Function" (Friedman, 1957), and Friedman and Becker (1957).[1]

Friedman's macro theory is based on well-defined micro foundations. It focuses on "consumption", in the sense of the value of the services actually consumed. Thus the purchase of durables such as cars, fridges and the like, is considered to be investment of consumers; only the use value of these items, during the given accounting period, is treated as part of consumption. Friedman builds his theory around permanent (or, planned) consumption and permanent (or, expected) income. His basic argument has a long-term perspective. The complete PIH is specified in terms of the following equations and assumptions:

The micro model:

$$c_{Pt} = k(i_t, w_t, u_t)\, y_{Pt}; \tag{5.1}$$

the definitional equations:

$$y_t = y_{Pt} + y_{Tt}; \tag{5.2}$$

$$c_t = c_{Pt} + c_{Tt}; \tag{5.3}$$

the correlation assumptions:

$$(5.4) \qquad r_{(y_{Tt}, y_{Pt})} = r_{(c_{Tt}, c_{Pt})} = r_{(y_{Tt}, c_{Tt})} = 0;$$

and, the aggregation assumption:

$$(5.5) \qquad f(i_t, w_t, u_t, y_{Pt}, c_{Tt}) = g(i_t, w_t, u_t)h(y_{Pt})l(c_{Tt}).$$

In our listing, subscripts "P" and "T" convey the Friedmanian distinction between "permanent" and "transitory", and "t" stands for the time index. y_{Pt} and y_{Tt}, respectively, are the permanent and transitory components of the consumer unit's observed (or measured) income in period t, y_t. Similarly, c_{Pt} and c_{Tt} are the two components of c_t, observed (or measured) consumption of the consumer unit during period t. The symbol "r" stands for the correlation coefficient for the variates specified as its subscripts.[2] f(.), g(.), h(.) and l(.) are the theoretical relative-frequency distributions, across the consumer units, of the variables listed in their corresponding parentheses---for simplicity, each of these probability density functions will be treated as a continuous one. We will explain the meanings accorded to k, i_t, w_t and u_t as part of the explanation of equation (5.1).[3]

> Equation (5.1) asserts that planned or permanent consumption (c_{Pt}) is a fraction (k) of planned or permanent income (y_{Pt}) that does not depend on the size of permanent income but does depend on other variables, in particular, the interest rate (i_t), the ratio of nonhuman wealth to income (w_t), and other factors affecting the consumer unit's tastes for current consumption versus accumulation of assets (u_t), 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 behaviour as presented in Chapter II.
>
> (Friedman,1957,p.222)[4]

Chapter II of Friedman (1957), referred to in this quotation, provides the choice-theoretic basis of the micro relation (5.1). Friedman employs the two-period Fisherian apparatus to explain his point. In this framework the consumer unit is assumed to be an inter-temporal utility-maximizer and the constraint on his choice of present versus future planned consumption is his wealth (W_t). The key assumption made by Friedman is homotheticity of the inter-temporal utility function.[5] This assumption together with the transformation $y_{Pt}=i_t W_t$, where $i_t \neq 0$, implies that the relation (5.1) between permanent consumption and permanent income is that of proportionality; the exact fraction k, however, depends on factors embodied in i_t and u_t. It should be noted that i_t is the subjective rate of discount used by a consumer unit in defining its wealth at time t.[6] The rest of the argument in Friedman's Chapter II involves dropping the restrictive assumptions one by one, and noting their implications for the proportionality thesis. The important claim is that, with the assumption of homogeneity of the consumer unit over time aside and the uncertainty about future incomes and consumptions allowed for, there is no need to modify the postulate of proportionality. All that one needs is to attach an extended interpretation to the coefficient of proportionality, like the one given in the above-listed quotation.

Equation (5.2) is assumed to reflect standard income accounting practices; the observed income of a consumer unit equals its expected or permanent income plus (minus) any unexpected gain (loss) in the given income-accounting period. In its nature, the transitory component of income is analogous to an unexpected gain, if positive, or a loss, if negative. Friedman also relegates the chance errors of measurement to this category; but in our general discussion we shall avoid their mention to maintain clarity in the argument. Equation (5.3) conveys similar ideas with respect to consumption; a consumer unit's observed consumption is treated as the sum of a planned or permanent and a random or transitory component. Some of the factors producing the transitory components may be specific to particular consumer units, while others may affect all the consumer units alike. As an example of the former, in the context of consumption, Friedman cites unusual sickness and a special opportunity to purchase; and for the latter he mentions special circumstances such as a bountiful harvest or an unusually cold spell. The effect of the former tends to average out, whereas the latter may result in a positive or a negative mean transitory component for all the consumer units.[7]

The correlation assumptions (5.4) specify "some of the characteristics of the probability distributions of the transitory components."[8] These are invoked on the basis of intuitive plausibility. Their primary function is to supply Friedman with logical reasons to justify the compatibility of the observed cross-section relations between c_t and y_t with his basic consumption relation (5.1).[9] Since they are not relevant for our argument at any stage, we shall omit their rationalization; however, the interested reader may refer to Friedman (1957,pp.26-29).

Normally the condition (5.5) is not mentioned as a part of the PIH.[10] But it is crucial for Friedman's aggregative model which emphasizes that aggregate consumption per consumer unit is proportional to aggregate permanent income per consumer unit. In (5.5), f(.) is the joint probability density function of the characteristics i_t, w_t, u_t, y_{Pt} and c_{Tt}. g(.) happens to be the marginal density function of i_t, w_t and u_t; similarly h(.) and l(.), respectively, give the marginal densities of y_{Pt} and c_{Tt}. Underlying (5.5) are two assumptions at every point in time. First, the distribution of consumer units by i, w and u is independent of their distributions according to y_P and c_T; second, their distributions by y_P and c_T are also mutually independent. The first one is a rather strong assumption. As an example to substantiate this claim, note that u_t includes such factors as age and education which are likely to be directly related to y_{Pt}. But as mentioned earlier, this is an empirical necessity.[11]

Equations (5.1) and (5.3), in combination with the aggregation assumption (5.5), yield the aggregative, long-term model of consumption.[12] As we shall see in section 5.6.3, equation (5.2) helps to approximate aggregate transitory income. The assumptions (5.4) are not relevant for time-series work. Friedman also mentions the possible need for extra assumptions to facilitate estimation.[13] But he does not elaborate on this, except for assuming that the mean transitory component of consumption (for the population as a whole) equals zero (Friedman,1957,p.144). We do add some extra assumptions to which we shall return after defining our problem.

5.2 The Object of the Present Study: the Elasticity of Consumption out of Real Permanent Income

In our restatement of the PIH, it is seen that the proportionality thesis is rooted in the assumption of a homothetic inter-temporal utility function for a typical consumer unit. This implies that the elasticity of consumption out of real permanent income is unity. In other words, an increase (or, a decrease) in the permanent income of a consumer unit will result in an equi-proportionate increase (or decrease) in its consumption. Therefore, if a mean-preserving redistribution of real permanent income occurs, there will be offsetting changes in the consumptions of gainers and losers, leaving the aggregate consumption and, hence, aggregate savings of the economy unchanged. From a neo-classical point of view, this scenario has important policy implications. If the consumer units behave according to the logic of the PIH, government intervention in the market place involving only the redistribution of real permanent income would have no economic justification. The reason is that such a redistribution would serve no useful purpose in stimulating the economy by influencing capital accumulation.[14] We examine these possibilities by studying the elasticity of consumption out of real permanent income. The approach followed is essentially the same as developed in the previous chapters.

We stipulate the following augmented relation

$$c_{Pt} = k(i_t, w_t, u_t)\, y_{Pt}^{M}, \tag{5.6}$$

instead of (5.1), as the model for a typical consumer unit, and propose to test

$$H_0\colon M = 1, \tag{5.7}$$

against the alternative

$$H_a\colon M \neq 1. \tag{5.8}$$

M is the elasticity of consumption out of real permanent income. A non-unitary M also means a variable marginal propensity to consume (out of real permanent income). Thus the result should have a direct bearing on the no-distribution effects proposition.

The existing argument of the PIH, as noted in section 5.1, is not sufficient to test this hypothesis. This requires some extra assumptions which we now consider.

5.3 Additional Assumptions in the Present Study

Substituting (5.6) into (5.3), and aggregating over all the consumer units in the economy, with respect to (5.5), we get[15]

$$(5.9) \qquad \bar{c}_t = k^* \int_{y_{Pt}} y_{Pt}^M h(y_{Pt}) dy_{Pt} + u_{3t}.$$

k^* is an aggregative constant which, conceptually, depends on various moments of the joint density $g(i_t, w_t, u_t)$, but is independent of the level of permanent income. u_{3t} is the mean of the marginal density function $l(c_{Tt})$. It also plays the role of the stochastic term in our regression model. Now it should be obvious that we need some a priori knowledge about the properties of u_{3t} and $h(y_{Pt})$ for the estimation of (5.9), and the subsequent testing of our hypotheses about M.

Friedman gives good reasons as to the possibility of a non-zero, positive or negative, mean transitory consumption component for all consumer units.[16] We carry his point a little further in the following manner. u_{3t}, the mean transitory consumption component at a point in time, may assume any value from the set of real numbers defined over $(-\infty,+\infty)$. The probability of each possible value is given according to a normal distribution with a zero mean and a constant, though unknown, variance. When correcting for serial correlation, we will also assume that

$$u_{3t} = \rho u_{3t-1} + \epsilon_t,$$

where ρ, the autocorrelation coefficient, belongs to the open-internal $(-1,+1)$, and ϵ_t, the true stochastic component, satisfies the Gauss-Markov assumptions. These assumptions allow us to use the maximum likelihood procedures for estimation, as in the previous chapter, and the likelihood-ratio method to test our hypotheses about M.

As for $h(y_{Pt})$, our interest is in two things: <u>first</u> its functional form, and <u>second</u> the parameters of the distribution. The discussion of the parameter values will be postponed until the data section. For the moment we restrict ourselves to the first item.

$h(y_{Pt})$ is not the (relative-frequency) distribution of incomes over an accounting period larger than one year, say two, three, or perhaps more years. It is also not to be confused with the (relative-frequency) distribution of life-time incomes. In fact, at each point in time every consumer unit has some notion about its real permanent income in view of what it owns at present, and expects to have over a number of future periods (years in our study). $h(y_{Pt})$ refers to the relative-frequency distribution of these perceived real permanent incomes.

The concept of aggregate real permanent income is still an unsettled issue; the nature of its distribution is even more vague. Friedman is not explicit on the matter either. Only once in "A Theory of the Consumption Function" we find an indirect mention of it. On this occasion Friedman explains his reasons for the compatibility of the observed consumption relation

$$c_t = a_0 + a_1 y_t, \tag{5.10}$$

often claimed to have support from cross-section data, with his theoretical relation between c_{Pt} and y_{Pt}, that is equation (5.1). He hints at the following mathematical justification in a footnote:

> On our hypothesis, the relation between the mean value of c_t and y_t will be linear only under special conditions. For example, it will be if y_{Pt}, y_{Tt}, and c_{Tt} are distributed according to a trivariate normal distribution....
>
> (Friedman, 1957, p.31)[17]

Since Friedman treats equation (5.10) as the truth for observed consumption and observed income, compatible with equation (5.1), it seems reasonable to infer that he implicitly assumes the joint distribution of y_{Pt}, y_{Tt} and c_{Tt}, across the consumer units, to be trivariate normal. We may also recall that the independence of the marginal distributions of c_{Tt} and y_{Pt} is implicit in Friedman's work. If we further attach the strong interpretation of "independence" to his "no correlation" assumptions (5.4), it would follow that $h(y_{Pt})$ is also normal in its shape. This may be a correct deduction, but it is difficult to rationalize. A normal density is symmetric around its mean, and its domain requires the existence of negative values too. Negative permanent incomes are not possible in real life situations. This alone is a strong reason to cast aside the normality hypothesis about $h(y_{Pt})$.

The existing literature (other than Friedman (1957, 1963)) does not offer any definitive clues as to the shape of $h(y_{Pt})$ either. Perhaps, one could begin with a stochastic process whose tenets do not clash with any of the existing pillars of the PIH. From here, it may be possible to work one's way toward a precise form for $h(y_{Pt})$. But the very lack of a means to verify such a scheme casts doubt on the wisdom of the whole exercise. Accordingly we take a more direct step, and assume: $h(y_{Pt})$, in every year, follows the pattern of a two-parameter lognormal distribution with the logarithm of geometric mean (real) permanent income μ_{Pt}, and the variance of logs of (real) permanent income σ_{Pt}^2.

There are two considerations behind this assumption. The first follows from the observation that the distribution of variates such as wealth and measured income all have positive skewness; there is no reason why the same should not be true for the distribution of real permanent income. The second stems from data considerations though these are a less binding constraint for the lognormality hypothesis. In the end, however, the results are to be given a conditional interpretation, as done in the previous chapter.[18]

In closing the discussion of our additional assumptions, we would like to remind the reader that the consumer unit in this study is an individual, who happens to be the primary income-receiving unit. Friedman's theoretical discussions revolve around the "household". At the theoretical level, our usage of the individual income-recipient should not create any complications for the PIH. One only needs to replace the word "household" by "income-recipient", and to qualify the relevant statements to suit the latter. On the empirical side, the advantages are enormous. For example, one need not bother about the income- and consumption-related characteristics of the households, such as their distribution according to the number of income-recipients and household size.[19] With this reminder, we move to more practical matters.

5.4 Data

In the light of equation (5.9), data for $\bar{c}_t$, μ_{Pt} and σ_{Pt}^2 are needed, preferably with μ_{Pt} and σ_{Pt}^2 coming from the same source. Permanent income is a theoretical concept; no readily available statistical information (on it) is to be expected. Moreover, the consumption figures (needed to define $\bar{c}_t$) are also not formally compiled. The general approach in a situation such as this, is to identify some of the required series *a priori*, and to search for the rest under one's null (or, alternative) hypothesis. Thus the task of complete data specification becomes a part of the estimation and testing. For example, in Friedman's empirical work with time-series data, the consumption and consumer unit series are taken as given. The per capita permanent income ($\bar{y}_{Pt}$) is searched using the data on personal disposable income and the aggregative equivalent of (5.1), that is, $\bar{c}_{Pt}=k^*\bar{y}_{Pt}$. Of course, in the process, $\bar{y}_{Pt}$ is defined as a weighted average of past per capita disposable incomes, with the weights declining into the past. In our work, a rather flexible approach is followed. According to our procedure, only the consumer unit index and σ_{Pt}^2 are assumed *a priori*. Of the rest, search is conducted, not for $\bar{c}_t$ and μ_{Pt} but, for $\bar{c}_t$ and $\bar{y}_{Pt}$. Further use is made, in the search, of the relation $\bar{y}_{Pt}=\exp(\mu_{Pt}+0.5\sigma_{Pt}^2)$ which is true for the lognormal distribution. The procedure is flexible in the sense that we use not one but two proxies each for consumption and income (the latter to act as an input in the search for permanent income). In addition, both the null and alternative hypotheses, in turn, are treated as truth in the search for the data on permanent income. We return to the technical details in the next section. The remainder of this section is devoted to the examination of our data choices.

The data period in our study is 1947-1976. Some features of the previous chapter are carried over to the present one. Firstly, we try to preserve consistency in the analysis by maintaining the assumption of one good and, hence, one price level. As before, the Consumer Price Index (CPI) is used as the proxy for the price level. Thus all the original monetary data are converted into real terms by using the CPI as the price deflator. Secondly, total population in the age-bracket 15 years and above (POP15) is used to approximate the total number of consumer units who are income-recipients in our study.[20] Finally, σ^2 at time t (hereafter σ_t^2) is used as the proxy for σ_{Pt}^2, where σ_t^2 is the variance of logs estimated from the taxation data in Chapter 3. As in the case of the current income-current expenditure experiments, this choice is made out of necessity. However, we do look at the ramifications of this undertaking, and note its potential implications for the results.

Our choice of data on $\sigma_{\overline{Pt}}^2$ can be defended as far as nominal versus real considerations are concerned.[21] The variance of logs is a scale-free measure. Therefore, in our one-good framework, the use of $\sigma_{\overline{t}}^2$ raises no problem on this score. Next, we notice that the definitional equation (5.2) implies:

$$Var(y_{Pt}) = Var(y_t) - Var(y_{Tt}) - 2Covar(y_{Pt}, y_{Tt}).$$

Here, "Var" and "Covar", respectively, stand for the variance and the covariance. $Var(y_{Tt})$ is likely to be a non-zero (but, of course, a positive) quantity. If $Covar(y_{Pt}, y_{Tt})$ can be assumed negligible, it would follow that at any point in time the distribution of permanent incomes has less dispersion than the distribution of measured incomes. This general conclusion is also supported by the following observation of Friedman:

> The existence of large negative savings is a symptom that the observed inequality of measured income overstates substantially the inequality of permanent income.
>
> (Friedman, 1957, p.40)

This discussion leads us to believe that by the use of $\sigma_{\overline{t}}^2$ we shall be over-estimating $\sigma_{\overline{Pt}}^2$ at every point in time. The $\sigma_{\overline{t}}^2$'s are in the neighbourhood of unity. Any fractional adjustment in these values should make little difference to our estimates---see the way $\sigma_{\overline{Pt}}{}^2$ enters our regression model (5.13 below). However, the pattern of the $\sigma_{\overline{Pt}}{}^2$ values over time may differ from the inter-temporal pattern of $\sigma_{\overline{t}}{}^2$ values at hand, and this may be important for correct regression analysis. We shall maintain that this is not a significant problem, until evidence to the contrary comes up.[22]

Thus far we have looked at statistical features that are common to Chapter 4 and the present one. The current study takes a different course when it comes to data on the expenditure and income variables. This happens because it is no longer feasible to have only one proxy each for both of these variables. In theory, the expenditure variable is assumed to reflect the value of services that the consumer units enjoy, rather than the value of the means that generate these services. On this count, only the monetary equivalent of depreciation---not the expenditure on purchases---of consumer durables forms a part of consumption during a given accounting period. A similar approach is to be followed in connection with the semi-durables and services which last beyond the current accounting period. Similarly, there is a case for treating expenditures on education as an investment in human capital, not consumption. Moreover, as is pointed out by Mayer

(1972,p.14), part of the expenditure on children (or elderly parents) is in fact saving (or payment of debts, that is, negative saving). Appropriate adjustments on all these scores would be necessary to bring us close to the theoretically correct concept of consumption. However, conventional data sources, such as the National Accounts, do not permit access to such information. In the past, two expenditure series have been prominent in the tests of both the adherents and the challengers of the PIH---see, Darby (1974) and Mayer (1972). These are:

CEXP = Personal Expenditure on Consumer Goods and Services, as in the National Accounts; and

CXP1 = CEXP minus the personal expenditure on purchases of Durable Goods.[23]

Neither of the two series adequately represents the required variable. Both are potential candidates for use in this study, and rather than choosing between them *a priori*, we consider both alternatives. Thus our results can be compared to most earlier studies on the PIH.

As mentioned earlier, permanent income is simply a theoretical construct. The general practice, in the empirical work on the PIH, is to determine it as a part of the actual experiments, and we follow this procedure. The primary inputs of such an exercise are a variable which forms the basis of consumer units' projections of their permanent incomes, and a scheme generating permanent income (to be outlined in the next section). Friedman's regression with aggregative time-series data (reported in Chapter V of Friedman (1957), and also in Friedman and Becker (1957)) employs the personal disposable income concept. This amounts to assuming that every consumer unit's guess of its permanent income is based on its take-home pay. But Mayer (1972) recommends the broadening of this concept to include undistributed corporate profits. This is, in Mayer's opinion, consistent with the "inherent logic of the permanent income theory"; he argues thus:

> If a household exercises foresight and rationality one would expect it to be rational enough to count corporate saving as part of its income.
>
> (Mayer,1972,p.365)[24]

In addition to the two above-mentioned series, Darby (1974) also uses personal income (gross of taxes) in his search for permanent income. But it is hard to see why the consumer units, though fully aware of the tax-bite, should base their projections of permanent income on the pre-tax personal income.[25] We, therefore, consider only the following two income series in our experiments:

YD = Personal Disposable Income, corrected for government grants to hospitals; and

PVY = YD plus Undistributed Corporate Profits.[26]

The statistics on consumption, income, $\sigma_{\overline{Pt}}^{2}$, the price level and the consumer units' index provide us, in effect, with four data sets:

$$\underline{S1} = \{CXP1, YD, \sigma_{-}^{2}, CPI, POP15\},$$
$$\underline{S2} = \{CXP1, PVY, \sigma_{-}^{2}, CPI, POP15\},$$
$$\underline{S3} = \{CEXP, YD, \sigma_{-}^{2}, CPI, POP15\},$$
and
$$\underline{S4} = \{CEXP, PVY, \sigma_{-}^{2}, CPI, POP15\}.$$

As indicated earlier, these will be used to generate aggregate permanent income, and to calculate M, the permanent income elasticity of consumption. We turn to these matters in the following section.

5.5 The Parameter Estimates and Tests for the Elasticity of Consumption out of (Real) Permanent Income

The proposition in section 5.3 that $h(y_{Pt})$ is lognormal with parameters μ_{Pt} and $\sigma_{\overline{Pt}}^{2}$, implies that

$$(5.11)\quad h(y_{Pt}) = \frac{1}{y_{Pt}\sigma_{\overline{Pt}}\sqrt{2\pi}} e^{(-0.5/\sigma_{\overline{Pt}}^{2})(\ln y_{Pt} - \mu_{Pt})^{2}}, \quad y_{Pt}>0;$$
$$= 0, \text{ otherwise.}$$

Substituting this into equation (5.9), and evaluating the definite integral therein, we get

$$(5.12)\qquad \bar{c}_{t} = k^{*} e^{M\mu_{Pt} + 0.5M^{2}\sigma_{\overline{Pt}}^{2}} + u_{3t}.$$

Upon using the relation

$$\overline{y}_{Pt} = e^{\mu_{Pt} + 0.5\sigma_{\overline{Pt}}^2},$$

equation (5.12) further simplifies to

$$(5.13) \qquad \overline{c}_t = k^* e^{M \ln \overline{y}_{Pt} + 0.5M(M-1)\sigma_{\overline{Pt}}^2} + u_{3t}.$$

Equation (5.13) is referred to as the "unrestricted model" in the following discussion. It will be the regression model when the null hypothesis M=1 is not true. On the other hand, under the hypothesis M=1, equation (5.13) reduces to

$$(5.14) \qquad \overline{c}_t = k^* \overline{y}_{Pt} + u_{3t}.$$

This will be referred to as the "restricted model". It is the same as Friedman's regression model, reported in Friedman and Becker (1957).

As for the data, two definitions of consumption per income-recipient at time t ($\overline{c}_t$), based alternatively on CXP1 and CEXP, are available to us. The specification of data on $\sigma_{\overline{Pt}}^2$ has also been taken care of. For estimation and testing, there is a further need to define aggregate real permanent income which, together with POP15, will yield estimates of $\overline{y}_{Pt}$. Following Darby (1974), we define this variable as

$$(5.15) \qquad Y_{Pt} = bY_t + (1-b)(1+s)Y_{Pt-1}.$$

Again, t is the time-subscript. Y is a given aggregate measured income series (in real terms), and Y_P is its permanent counterpart. b can be thought of as an adjustment coefficient in an adaptive expectations model, and s is the trend rate of growth. The income-trend regression

$$\ln Y_t = g_1 + g_2 t + v_t, \tag{5.16}$$

where v_t is a stochastic term, supplies an estimate of s $(=\hat{g}_2)$ and the initial value of aggregate permanent income $(Y_{P0}=\exp(\hat{g}_1))$---$\hat{g}_1$ and $\hat{g}_2$ are the least-squares estimates of g_1 and g_2, respectively. The assumptions here are that the long-term growths in Y and Y_P are the same, and that at time t=0 the expectation of Y from the regression equation (5.16) is equal to Y_{P0}. According to this approach, if the economy is allowed to follow its long-term growth path, aggregate permanent income in period t should equal $(1+s)Y_{Pt-1}$. However, the actual experiences embodied in Y_t, if different than $(1+s)Y_{Pt-1}$, force the participants to revise the estimate for Y_{Pt} in accordance with equation (5.15); b and 1-b are the weights attached to Y_t and $(1+s)Y_{Pt-1}$, respectively. In the context of equation (5.15), the task of identifying the aggregate permanent income series becomes one of determining the appropriate income series on Y_t and the adjustment coefficient b. This parameter is supposed to take one of the values from the following set:

$$\{0,\ x,\ 2x,\ 3x, \ldots\ldots, 1;\ x=.025\}$$

Of course, specification of the exact value of b, which amounts to the identification of an aggregate permanent income series (for a given measured income series), takes place as part of our estimation, to which we turn next.

One approach to estimation and testing is as follows:

(1) Pick one of the data sets <u>S1</u>, <u>S2</u>, <u>S3</u> and <u>S4</u> ;

(2) Treat H_0:M=1 as the truth;

(3) In light of step (2), use the restricted model (5.14) to search for aggregate real permanent income; and

(4) With permanent income as specified in step (3) and the rest of the data as chosen in step (1), estimate the unrestricted model (5.13), and test for M=1.

Repeating steps (2)-(4) for each of the data sets <u>S1</u>, <u>S2</u>, <u>S3</u> and <u>S4</u> allows one to gauge the robustness of the conclusions. to the data set chosen. However, this still would not be sufficient, since there is no strong <u>a priori</u> reason for undertaking step (2). Moreover, treating M=1 (while searching for permanent income) may be claimed to bias the tests in the direction of the hypothesis of unitary elasticity. This situation can be rectified by replacing steps (2), (3) and (4) by the following ones, for each data set.

(2') To begin with, treat H_a:M$\neq$1 as the truth;

(3') Use the unconstrained model (5.13) in the search for aggregate real permanent income; and

(4') Finally, estimate the restricted model (5.14) with the data as specified in the preceding steps, and test for H_0:M=1.

In the end if both the steps (2)-(4) and (2')-(4') are alternatively undertaken for each of the four data sets at hand, we should be in a good position to judge the issue. Accordingly, this dual approach is followed here. First we search for permanent income under the assumption that the hypothesis M=1 is in fact true. The parameter estimates and results of the tests for this exercise are reported in Table 5.1.

We start with the data set S1. Given that u_{3t} has a normal distribution, the maximum-likelihood (or, equivalently, the least-squares) method can be used to estimate equation (5.14), and the likelihood-ratio test procedure to verify H_0:M=1. This equation is estimated for the different values of b mentioned earlier. Of course, in the process Y_{Pt} is defined in the light of equations (5.15) and (5.16). b=.150 maximizes the log-likelihood function (and minimizes the sum of squared-residuals). Parameter estimates for equation (5.14), with S1 and b=.150, are reported in row 1 of Table 5.1. Next using S1 and b=.150, we estimate the unrestricted model, (5.13). The estimates for this model (given in row 2 of Table 5.1) show that M is marginally higher than unity.

The difference (from unity) is less than one percent; in fact the estimated M is 1.0072. The log-likelihood function (values reported in column 8 of Table 5.1) registers a very small improvement under $H_a{:}M \neq 1$; but again the change is trivial. The value of the likelihood-ratio test-statistic, $-2\ln\lambda$, is .310 (given in Part B of Table 5.1). This lends strong support to $H_0{:}M=1$ at all the conventional levels of statistical significance. The slight difference (from unity) in the value of M can, therefore, be attributed to noise in the data.

S2 has the same aggregate consumption series as S1, that is, CXP1/CPI. But data input in the construction of permanent income is now PVY/CPI, instead of YD/CPI. However, this difference does not translate into a different outcome for this experiment. Aggregate real permanent income is defined as the combination of PVY/CPI and b=.125 (row 3 and column 2, Table 5.1), in the light of equations (5.15) and (5.16). Other than this, M (=1.0036) is more or less the same as before, and the likelihood-ratio test result (Part B, Table 5.1) reaffirms $H_0{:}M=1$.

Experiments with the data sets S3 and S4 are parallel to those with S1 and S2, with respect to the measured income proxies. However, these two new experiments are based on a broader concept of consumption, that is, CEXP/CPI. Nonetheless these experiments also carry no surprises. In the case of S3, b=.200, while for S4, b=.150. Estimates in rows 6 and 8 of Table 5.1 show that M is more or less the same in both these cases---.9945 for S3 and .9943 for S4. As in the two previous cases, the likelihood-ratio test results favour H_0:M=1 for both these data sets.

These experiments have another interesting feature which can be observed from the numbers in columns 7 and 8 of Table 5.1. C.V. is the standard error of the regression divided by the arithmetic mean of the dependent variable $\bar{c}_t$ used in that regression. It gives a rough idea of the relative performance of both the restricted and unrestricted models for each of the four data sets S1, S2, S3 and S4. ln L is a probability indicator of the fit of each model to these data sets. According to the figures in column 8, the value of ln L improves as one goes from the restricted to the unrestricted model. But the estimates in column 7 show that the addition of another parameter like M to the micro model (5.1), which is the basis of equation (5.14), in fact leads to a deterioration in the empirical fit. The direction of the change is uniform for all the data sets, regardless of the differences in the value of b. Of course, as mentioned earlier, the need for an extra parameter is clearly set-aside by the likelihood-ratio tests. At this point, we may also inform the reader that all these findings show up again when $H_a:M\neq 1$ is considered as the truth in searching aggregate real permanent income.

Table 5.2 contains the parameter estimates and results for this second set of experiments. The odd-numbered rows in Part A give the parameter values under the assumption that $H_a:M\neq 1$ is true; the even-numbered ones have comparable estimates for the restricted model. For each data set, b's are determined such that the log-likelihood function reaches its maximum for the unrestricted model, (5.13). The new b values are .125, .100, .200 and .175 for S1, S2, S3 and S4, respectively (as compared with .150, .125, .200 and .150, in this order, determined under $H_0:M=1$). For each of the four cases, the model is re-estimated with the restriction M=1, and the likelihood-ratio test-statistic (reported in Part B of Table 5.2) is formed. As mentioned above, the parameter estimates and results are substantially the same as determined in the first set of experiments (reported in Table 5.1). The only noticable thing is that now the magnitude of M is smaller with the data set S4 than with S3; but it is still insignificantly different from unity. So there is no need to retract our finding of the unitary elasticity of consumption as the result of this second experiment.

To conclude the matter as it stands now, one can say that the result of unitary elasticity is symmetric. It is unaffected by our treatment of either the restricted or the unrestricted model as the truth in the search for the data on aggregate real permanent income. Moreover, the result does not depend on the choice of a narrow (as with CXP1) versus a broad (as in the case of CEXP) definition of consumption. But before these may be called our final results, some additional matters need clarification. This is what we intend to do in the next section.

5.6 Further Analysis of the Results

In this section we shall address three issues. The first of these three is a mere technicality; it relates to the problem of serial correlation. The other two are about matters of substantial interest. The first concerns the question surrounding the constant term debate; and the second relates to the presence of the consumer asset component in the consumption data and the role of transitory income. We shall address the three issues in turn.

5.6.1 Correction for Serial Correlation and the Results

The maximum likelihood estimates reported in Tables 5.1 and 5.2 are also least-squares estimates. So the Durbin-Watson statistic DW (given in the column 6 of both tables) can be taken to imply a high degree of positive autocorrelation. The correction for serial correlation would increase the sum of squared-residuals, and thereby affect the values of the log-likelihood functions for both these models. Conclusions favouring the null hypothesis, such as ours, generally do not change after the correction for serial correlation. Our checks with two data choices S1 and b=.150 (determined under H_0) and S1 and b=.125 (determined under H_a) confirm this point. In each case our working hypothesis has been a first-order autocorrelation scheme, as metioned in section 5.3, and the FML procedure of Chapter 4 was used in the estimation. For S1 and b=.150, the autocorrelation coefficient turns out to be .705 and .711 for the restricted and unrestricted models, respectively; M is .9928 (down from 1.0072); and $-2\ln\lambda$ falls to .108 (from .310). Similarly for S1 and b=.125, the autocorrelation coefficient is .709 for the restricted model and .711 for the unrestricted one; M turns out to be .9985 (down from 1.0010); and $-2\ln\lambda$ is.010 (down from .610). These estimates indicate even more dramatically the closeness of M to unity, our principal result so far. In the face of this finding, the correction for serial correlation for every case seems unnessary. So we move to the next item on our agenda.

5.6.2 An Additive Constant Term and the Results

Since 1957 there has been a heated empirical debate on the existence of an additive constant in the relation (5.1). By taking this route, the critics have hoped to destroy the proportionality thesis, and to show that the income elasticity is not unity. However, its implications go beyond those relating to the magnitude of the income elasticity. The presence of an additive constant in the micro model (5.1) means that permanent consumption would be nonzero even if permanent income were zero. A zero permanent income means that wealth (which generates the income flow) is zero. A consumer unit's wealth includes both the physical and monetary assets at hand and the discounted present value of the (expected) future income stream. Thus a zero wealth must mean that the consumer unit is penniless and it also does not expect anything in the future. In such circumstances, the idea of "permanent" consumption is beyond comprehension. Looked at from a different angle, if one considers permanent consumption to be nonzero, this must mean either that the consumer unit has something currently at its disposal (carried over from the past) or that it expects something in the future. In both cases permanent income would be nonzero. These twists of logic suggest that the idea of permanent consumption is consistent with nonzero permanent income only. Thus if the data yields a significant constant term, it calls into question the very foundations of the PIH. But the available data and analytic techniques do not permit an effective challenge to the theory on this score. Moreover, one's estimation may yield a constant term for purely statistical reasons. This point is explained with the help of Figure 5.1 (found on page 152).

The basket of goods underlying the consumption variable is different from that underlying the income variable; this is also true of the available price indices for these variables. It is possible that the true relation between real consumption and real income is that of proportionality, such as OA' in Figure 5.1. Now imagine what would happen if one started out with nominal data on these variables, and converted them into real terms using the price index relevant for, say, the income variable.[27] If this price index were consistently lower than that for the price index were consistently lower than that for the consumption variable, resulting real consumption would be inflated for all the levels of real income. In that event we would observe BB' instead of OA'. Such a finding will not represent the facts. The message from this example is clear enough to warrant any further comment on our part.[28] However, the existence of the constant term has important bearings on the estimates of our elasticity parameter; the argument is analogous to that presented in Chapter 4 (section 4.5.2). Thus we thought it prudent to check for the presence of an additive constant in the data choices underlying our parameter estimates and tests.

Our choice of S1 and b=.150 is based on the assumption M=1; the same is also true of S2 and b=.125, S3 and b=.200, and S4 and b=.150. We maintain this proposition, and modify the micro relation (5.1) to include a constant term, k_0. Under the aggregation restriction (5.5) and the assumptions of section 5.3, the aggregative model becomes

$$\bar{c}_t = k_0 + k^* \bar{y}_{Pt} + u_{3t}. \qquad (5.17)$$

This model is estimated for all of the afore-mentioned data choices, and the significance of k_0 is studied. The parameter estimates and results of the t-tests for $k_0=0$ are given in Table 5.3, Part B. These results leave no doubt as to the nonexistence of the constant term in all the four data choices under consideration.

The data choices S1 and b=.125, S2 and b=.100, S3 and b=.200, and S4 and b=.175 arose from our treatment of H_a:M$\neq$1 as the truth in searching for permanent income. Now, this proposition is assumed, and the micro model (5.6) is modified thus:

$$c_{Pt} = k_0 + k^* y_{Pt}^M. \qquad (5.6')$$

Aggregating this in the light of the aggregation condition (5.5) and other assumptions made in section 5.3, we get

$$\bar{c}_t = k_0 + k^* e^{M \ln \bar{y}_{Pt} + 0.5M(M-1)\sigma_{\overline{Pt}}^2} + u_{3t}. \qquad (5.18)$$

This model is estimated for all of the four data choices determined in light of H_0:M$\neq$1, and a likelihood-ratio test is performed for the hypothesis $k_0=0$. The statistical details for this endeavour may be found in Table 5.4. These results also confirm the proposition $k_0=0$. Thus our elasticity estimates may be treated as being free from any potential errors associated with the existence of a constant term. Can the same be said about the effect of the asset component that might be present in our consumption data? We address this issue next.

5.6.3 The Asset Component in the Consumption Variables and the Results

Another of the controversies, surrounding tests of the PIH, relates to the presence of an asset component in the consumption variable. The theory requires the data to be free from any consumers' expenditure which does not translate into the satisfaction of their consumption needs within the given accounting period. But at the empirical level, the data invariably contain many questionable items. This situation cannot be corrected either, due to the lack of proper information. Friedman (1957,p.28) suggests that consumer asset formation is related to transitory, rather than permanent, income. Thus use of consumption data that may contain asset purchases to calculate the permanent income elasticity, is questionable. This recommends caution in the interpretation of our results.

One way to solve the problem would be to model transitory income in the micro relation (5.6). But this is not feasible, because exact specification of the role of transitory income is not possible in the present state of the art. Moreover, these specification errors may find their way into the estimates of the permanent income elasticity of consumption. Further, we would also need to know the functional form of the joint distribution of permanent and transitory incomes, in order to aggregate the modified version of the relation (5.6) with respect to the permanent and transitory incomes. These difficulties rule out a direct approach to the problem. However, we can still make some progress by looking at the relation between the variation in aggregate consumption that remains unaccounted for by our model and aggregate transitory income. If this relationship turns out to be insignificant, it may be said that the problem is not serious in the case of our experiments. This is essentially what we do. The statistical details pertaining to this exercise are given in Table 5.5.

We take as the maintained hypothesis the constrained model, that is, (5.14), which has been supported by the data so far. The parameter estimates for this model for each of our eight data choices may be found in Tables 5.1 and 5.2. These are used to generate the predicted values of $\overline{c}_t$ for each of the eight data choices. The difference between actual $\overline{c}_t$ associated with each data choice and the corresponding predicted value is then inflated by POP15. This gives an estimate of the unexplained variation in aggregate real consumption for each data choice. Next, for every data choice, aggregate real permanent income is redefined in the light of (5.15) and (5.16), and the corresponding aggregate real transitory income Y_{Tt} is calculated as follows.

$$(5.19) \qquad Y_{Tt} = Y_t - Y_{Pt}.$$

Finally we study the correlation between Y_{Tt} and the unexplained variation in aggregate consumption (both for the same data choice); the product moment correlation coefficient r is used toward this end.[29] As the estimates in column 2 (Part B of Table 5.5) reveal, these correlations are very weak; r ranges between -.0068 for S1 and b=.150, and .1954 for S4 and b=.150. Upon subjecting these correlation coefficients to a t-test, it becomes evident that r is insignificantly different from zero in each case. We hope that this explanation establishes, *ex post*, that there is no cause for alarm on account of the presence of the asset component in our data. Even if it is there, its magnitude appears to be relatively minor. So the estimates of M can be safely assumed as being free from bias on this score.

With the conclusion of the preceding discussion, our analysis of the results of section 5.5 is now complete. The results emerge unchanged. Now is the time to recapitulate our main finding, and to note its implications for the role of the distribution of real permanent income in the PIH framework.

5.7 A Summing-up: The main finding and Its implications for the Distribution of Real Permanent Income Effects

The PIH, as presented in Friedman (1957) and Friedman and Becker (1957), has supplied the terms of reference for this chapter. It focusses on Friedman's maintained hypothesis of the unitary elasticity of consumption out of real permanent income (M=1). This particular proposition, which is also known as "the proportionality thesis", rules out any role that might otherwise be attributable to the distribution of real permanent income (in determining the level of aggregate real consumption). Our study presumes the functional form of this distribution to be lognormal. The data is found to give a strong support to the hypothesis M=1.

In the framework of the PIH, M=1 affirms that the marginal propensity to consume is constant for all levels of real permanent income, for a typical consumer unit who is an income-recipient in this study. Another consequence of this result is that the coefficient of the distribution variable $\sigma_{\overline{Pt}}{}^2$ in the aggregative model (5.13) is zero. This implies that <u>the distribution of real permanent income does not</u>

matter in the PIH framework. This is our answer to the question posed at the beginning of this chapter. Its credibility is, of course, conditional upon the reasonableness of the lognormality hypothesis about the distribution of real permanent income.

We notice that the same result does not follow from the existing studies on the PIH, which try to establish the existence of the constant term in the PIH model. The presence of an additive constant, if established, raises the possibility of logical flaws in the PIH. Even if this were not the case, the existence of an additive constant does not necessarily imply the existence of distribution effects---though it implies a non-unitary income elasticity. All that it amounts to is this: the average (not the marginal) propensity to consume is a function of real permanent income. As shown in Chapter 2, this is neither necessary nor sufficient to verify the existence of distribution (of real permanent income) effects.

Most analysts of the PIH also try to determine as part of their empirical studies, an estimate of the real rate of interest, and sometimes also the appropriate measures of consumption and of permanent income.[30] No opinion on these matters has been given in the preceding sections of this chapter, as this has not been the intent. To say anything meaningful of these subjects would require us to establish criteria to choose between the data sets examined here. That study is left for another occassion.

FOOTNOTES TO CHAPTER 5

1. Friedman (1963) is not important for the basic argument.

2. It reflects the properties of the probability distributions of the variables involved.

3. For the record, it may be mentioned that (5.1), (5.2), (5.3) and (5.4) are mentioned under the heading of "A Formal Statement of the Permanent Income Hypothesis" in Friedman (1957 Chapter III, section 2). We infer (5.5) from the relation (2.9) in Friedman (1957, Chapter II) and the aggregative model reported in Friedman and Becker (1957).

4. The equation number is altered, and a time subscript added to the variable names in this quotation.

5. Friedman uses the word "homogeneous"---see Friedman (1957,p.13). But the idea implied is that of homotheticity.

6. In perfect capital markets, i is the same for all the consumer units.

7. Of course, in special circumstances these will be restricted to specific groups. For example, excessive rain affecting crops will produce negative transitory (income) component for farmers. A rather comprehensive discussion of transitory factors, both in income and consumption, is given in Friedman (1957,pp.21-23).

8. Friedman (1957,p.26).

9. His argument may be found in Friedman (1957, Chapter III, section 3).

10. See, for example, Mayer (1972,p.38).

11. Otherwise the aggregate consumption function will contain the first- and higher-order moments of the joint density function f(.), and the aggregative theory would not be as simple as Friedman proposes.

12. This is not how Friedman builds his aggregate model. He uses (5.1) and (5.5) in the first step; next he assumes $\bar{c}_{Tt}=0$ or, in other words, $\bar{c}_t=\bar{c}_{Pt}$---see Friedman (1957, p.144).

13. On page 30 of his 1957 book, where Friedman first mentions the need for extra assumptions, he does introduce something controversial. It is the so-called logarithmic variant of the PIH. According to it, the micro model (5.1) would be replaced by its logarithmic equivalent, and equations (5.2) and (5.3) by

$$(5.2') \qquad \ln y_t = \ln y_{Pt} + \ln y_{Tt},$$

and

$$(5.3') \qquad \ln c_t = \ln c_{Pt} + \ln c_{Tt}.$$

Of course, the rest of the model (5.1)-(5.5) also needs suitable alterations to fit into this logarithmic setting. Friedman considers this logarithmic structure as <u>an alternative statement</u> of the PIH. His reason for doing so is: "Its (i.e., the logarithmic variant's) implications are essentially the same as those of the arithmetic variant (the model (5.1)-(5.5)), since one can be regarded as a first order approximation to the other, and (therefore) most verbal statements of the implications apply equally to both..." (Friedman, 1957, p.223; the bracketed insertions are ours). This assertion is questionable; it is not possible to apply a linear transformation to (5.2) and (5.3), and get exactly (5.2') and (5.3'), or vice versa. So the two set-ups cannot be treated as equals for any practical purposes. Accordingly, one has to base one's argument on either the model (5.1)-(5.5) or the logarithmic set-up proposed by Friedman, but certainly not on both. We give up the logarithmic variant.

The reader should also note that the ideas of negative transitory incomes and consumptions, an important feature of the arithmetic version, cannot be conveyed simply by this logarithmic variant. The logarithms of negative numbers are undefined. Moreover, for empirical work one needs data on the logarithms of geometric mean consumption; Friedman's results using the logarithms of the arithmetic means appear to need some re-thinking.

14. There may be some immediate reaction to the government measures; but, according to our inference it will fizzle out in the long-run.

15. $$c_t = c_{Pt} + c_{Tt}.$$

$$= k(i_t, w_t, u_t) y^M_{Pt} + c_{Tt}.$$

Under our assumptions

$$\bar{c}_t = E(c_t) = E(k(i_T, w_t, u_t))E(y^M_{Pt}) + E(c_{Tt}).$$

This is the basis of (5.9).

16. See footnote 7 for the appropriate reference.

17. The time-subscript "t" is added to each variable name, in line with the notation followed here.

18. In passing we may note that lognormal is found unsatisfactory by Carlton and Hall (1978). This study is based on the logarithmic version of the PIH. We noted earlier (see footnote 13) that this is, in fact, a different theory and not a simple approximation to the arithmetic version considered here.

19. Friedman, in fact, uses per capita magnitudes, not per household ones, in his empirical work with time-series data---see Friedman (1957, Chapter V).

20. This choice is, empirically determined. As an alternative to POP15, we also tried the total number of all returns filed for a given taxation year. But in every instance, this index yielded inferior fits than those with POP15.

21. See section 4.3 in Chapter 4 of this thesis.

22. The PIH offers no predictions about the pattern of inequality in the distribution of real permanent income over time. It may also be undergoing a change in the direction of greater inequality.

23. All the expenditure and income data, used in this study, come from the Canadian National Accounts. As in Chapter 4, we subtracted the government grants to hospitals from both CEXP and YD. See also footnote 9 to Chapter 4.

24. A household is treated as the consumer unit in Mayer's discussion.

25. Darby's use of the personal income series would be justified, however, if one presumes that people expect to consume the tax-equivalent of government services. In that event, the consumption proxy should also be assumed to reflect the govrnment services consumption.

26. See footnote 23. Undistributed corporate profits are also taken from the same data source.

27. Aggregative analyses, both theoretical and empirical, are based on the assumption of one good in the economy. The only way to accomodate this in the empirical work, on the consumption function, is to apply the same price index to both the income and expenditure variables. In this example, it is the price index of the income variable. In the actual empirical work in this chapter the CPI is used to deflate both series.

28. There can be many other instances in which a constant, positive or negative, may show up in estimation. One potential source is the scheme used in generating permanent income. In our case if Y_{P0} were substantially underestimated, it would result in an understatement of the permanent incomes for all the years, because the effect is cumulative. With a consumption variable properly identified, the resultant consumption function would have a negative constant term.

29. This r is not to be confused with that in the correlation assumptions (5.4) in our text.

30. See, for example, Darby (1974).

Table 5.1 Tests for the Unitary Elasticity of Consumption, with Permanent Income Searched under H_0

Part A Parameter Estimates and Other Statistics [1,2,3]

Row	Data Set (1)	b (2)	k^* (3)	M (4)	R^2 (5)	DW (6)	C.V. (7)	ln L (8)
1	S1	.150	.6689 (.0021)	1.0	.9937	0.59	.01787	-152.314
2	""	" "	.6291 (.0691)	1.0072 (.0128)	.9937	0.59	.01809	-152.159
3	S2	.125	.6300 (.0020)	1.0	.9936	0.61	.01790	-152.377
4	""	" "	.6110 (.0678)	1.0036 (.0128)	.9937	0.61	.01820	-152.340
5	S3	.200	.9323 (.0027)	1.0	.9945	1.05	.01651	-159.961
6	""	" "	.9770 (.0987)	.9945 (.0118)	.9945	1.05	.01674	-159.848
7	S4	.150	.8786 (.0026)	1.0	.9945	1.07	.01648	-159.911
8	""	" "	.9227 (.0938)	.9943 (.0118)	.9946	1.08	.01671	-159.790

Part B Conclusions For the Hypothesis of Unitary Elasticity (H_0)[4]

Data Set	$-2\ln\lambda$	CONCLUSION
S1	.310	H_0 Upheld
S2	.074	Same as Above
S3	.113	Same as Above
S4	.242	Same as Above

(Continued on the next page)

Table 5.1 (Continued)

1. R^2 = the coefficient of determination;
   DW = the Durbin-Watson statistic.
   C.V. = the standard error of the regression divided by the mean of the dependent variable at hand---it is the coefficient of variation for the regression model; and
   ln L = log of the likelihood function.

2. The standard errors are reported in parentheses under each estimated coefficient---for the nonlinear models, these standard errors are asymptotic.

3. If M=1.0, the estimates belong to equation (5.14); otherwise, they are for equation (5.13).

4. The likelihood-ratio ($-2\ln\lambda$) method is used to test the hypothesis at hand. $-2\ln\lambda$ is defined as twice the difference between ln L for the constrained model and that for the unconstrained one. Its critical value, used in the tests, is 3.8415; it is based on a chi-square distribution with one degree of freedom, at the 5% level of significance.

Reminder: $\underline{S1}$ = {CXP1, YD, σ^2, CPI, POP15}

$\underline{S2}$ = {CXP1, PVY, σ^2, CPI, POP15}

$\underline{S3}$ = {CEXP, YD, σ^2, CPI, POP15}

$\underline{S4}$ = {CEXP, PVY, σ^2, CPI, POP15}

Table 5.2 Tests for the Unitary Elasticity of Consumption, with Permanent Income Searched under H_a

Part A Parameter Estimates and Other Statistics

Row	Data Set (1)	b (2)	k^* (3)	M (4)	R^2 (5)	DW (6)	C.V. (7)	ln L (8)
1	S1	.125	.6144 (.0676)	1.0100 (.0128)	.9938	0.59	.01805	-152.095
2	""	" "	.6694 (.0021)	1.0	.9936	0.58	.01792	-152.400
3	S2	.100	.5975 (.0665)	1.0062 (.0129)	.9937	0.61	.01816	-152.279
4	""	" "	.6305 (.0020)	1.0	.9936	0.60	.01791	-152.395
5	S3	.200	.9770 (.0987)	.9945 (.0118)	.9945	1.05	.01674	-159.848
6	""	" "	.9323 (.0027)	1.0	.9945	1.05	.01651	-159.961
7	S4	.175	.9536 (.0964)	.9904 (.0117)	.9946	1.07	.01669	-159.770
8	""	" "	.8779 (.0026)	1.0	.9944	1.05	.01660	-160.117

Part B Conclusions For the Hypothesis of Unitary Elasticity (H_0)

Data Set	$-2\ln\lambda$	CONCLUSION
S1	.610	H_0 Upheld
S2	.232	Same as Above
S3	.226	Same as Above
S4	.694	Same as Above

Note: See footnotes to Table 5.1.

Table 5.3 Tests for an Additive Constant Term (I)

Part A DATA

Data Base: $\bar{c}_t = k^* \bar{y}_{Pt} + u_{3t}$. (5.14)

Data Choices:[1] (i) S1 and b=.150;

(ii) S2 and b=.125;

(iii) S3 and b=.200;

(iv) S4 and b=.150.

Part B Estimates for Equation (5.17), and Tests[2,3]

Data (1)	k_0 (2)	k^* (3)	t-value for $k_0=0$ (4)	Is k_0 significantly different from zero? (5)
(i)	-24.872 (34.165)	.6761 (.0101)	-.728	No
(ii)	-15.139 (34.289)	.6341 (.0096)	-.442	Same as Above
(iii)	14.655 (44.119)	.9280 (.0130)	.332	Same as Above
(iv)	17.234 (43.976)	.8739 (.0122)	.392	Same as Above

1. See 'Reminder' to Table 5.1.

2. See footnote 2 to Table 5.1.

3. A two-tailed t-test is used to test the hypothesis at hand. The critical values of t (for 28 degees of freedom, at the 5% level of significance) are ±2.048.

Table 5.4 Tests for an Additive Constant Term (II)

Part A DATA

Data Base: $\bar{c}_t = k^* e^{M \ln \bar{y}_{Pt} + 0.5M(M-1)\sigma_{\overline{Pt}}^2} + u_{3t}$. (5.13)

Data Choices:[1] (i)* S1 and b=.125;

(ii)* S2 and b=.100;

(iii)* S3 and b=.200;

(iv)* S4 and b=.175.

Part B Estimates for Equation (5.18), and Tests[2,3]

Data (1)	k_0 (2)	k^* (3)	M (4)	$-2\ln\lambda$ for $k_0=0$ (5)	Is k_0 significantly different from zero? (6)
(i)*	-245.388 (309.007)	1.3921 (1.3229)	.9262 (.0974)	.846	No
(ii)*	-218.245 (305.868)	1.2493 (1.2022)	.9311 (.0981)	.670	Same as above
(iii)*	-370.037 (419.596)	2.3165 (2.0784)	.9062 (.0919)	1.062	Same as above
(iv)*	-442.597 (432.142)	2.6370 (2.3594)	.8870 (.0910)	1.484	Same as above

1. See 'Reminder' to Table 5.1.

2. See footnotes 2 and 4 to Table 5.1.

3. The figures in column 5 are defined as twice the difference between ln L for equation (5.13) (reported in Table 5.2) and that for equation (5.18) (not reported here).

Table 5.5 Relation Between the Unexplained Variation in Aggregate Consumption and Aggregate Transitory Income[1]

Part A. DATA CHOICES[2]

Based on the Proposition M=1	Based on the Proposition M≠1
(i) S1 and b=.150	(i)* S1 and b=.125
(ii) S2 and b=.125	(ii)* S2 and b=.100
(iii) S3 and b=.200	(iii)* S3 and b=.200
(iv) S4 and b=.150	(iv)* S4 and b=.175

Part B. Tests for the Said Relation[3]

Data Choice	r	Calculated t-value	Critical t-values	Conclusion for the Hypothesis r=0
(1)	(2)	(3)	(4)	(5)
(i)	-.0068	-.036	±2.048	Not rejected
(ii)	-.0695	-.369	" "	Same as above
(iii)	.0945	.502	" "	Same as above
(iv)	.1954	1.054	" "	Same as above
(i)*	.1493	.638	" "	Same as above
(ii)*	.1068	.323	" "	Same as above
(iii)*	----------Same as those for (iii)----------			
(iv)*	.0273	.013	" "	Same as above

1. The constrained model, which has been supported by the data so far, is taken as the maintained hypothesis here. The relation examined is between the unexplained aggregate consumption and aggregate transitory income. The correlation coefficient (r) is used in the analysis. For further details, see section 5.6.3 of the text.

2. For S1, S2, S3 and S4, see 'Reminder' to Table 5.1.

3. The critical t-values are for 28 degrees of freedom, at the 5% level of significance.

Figure 5.1 Potential Shift in the Consumption Function due to Data Construction

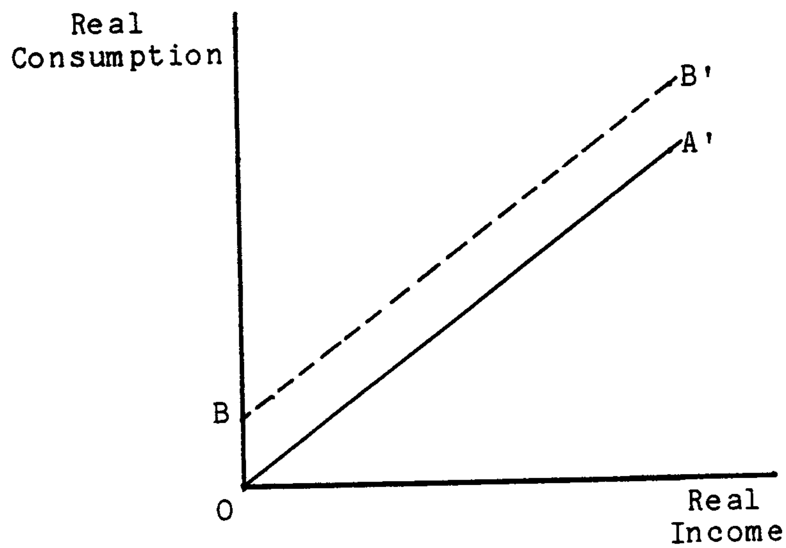

CHAPTER 6

EPILOGUE

In the consumption function literature, usage of the term "distribution effects" is quite imprecise. It only gains meaning and precision when put into a particular context. For example, within the context of the current income-current expenditure framework of the Keynesians, it refers to the effect of a redistribution of real disposable income on real consumers' expenditure. When the Permanent Income Hypothesis (PIH) is the reference framework, it can imply the effect of a redistribution of <u>either</u> real disposable income <u>or</u> real permanent income, on real consumption. Friedman (1957,p.17) suggests that a redistribution of real disposable income enters the PIH picture via its effect on uncertainity. On the other hand, a redistribution of real permanent income plays a similar role with respect to real consumption <u>as</u> a redistribution of real disposable income does with respect to real consumers' expenditure. These last two cases are identical from a methodological point of view; one only needs to read "consumption", "permanent income" and "redistribution of the permanent income" in place of "consumers' expenditure", "disposable income" and "redistribution of the disposable income". Both cases have been empirically studied in the past; we also concentrate on them in this thesis.

In testing for distribution effects, one ought to realize that the size distribution of income is (1) a relative-frequency distribution of income; and, therefore, (2) the theoretical link between one's micro and macro consumption relations. From this, it follows that the problem of testing for distribution effects has to be formulated at the micro level; moreover, a knowledge of the functional form of the income distribution is indispensable in the analysis.

Two points vital to any test of distribution effects may be noted:

(i) The variability of the marginal (not the average) propensity to consume with respect to income, at the micro level, is the key to the existence of distribution effects.

(ii) Empirically, the distribution variable (which measures inequality in the income distribution) in the consumption function will not enter independently of the income variable. A constraint will link the coefficients of the distribution and income variables in the aggregate consumption function.

Regression procedures that have not taken account of these points need reconsideration. With due allowance for these points, we demonstrate two things:

I. In the current income-current expenditure framework, the distribution effects appear to be consistent with the conventional wisdom. The empirical evidence supports this view at a 10% significance level, although the results are mixed at the 5% level of significance.

II. The distribution of real permanent income does not matter in the framework of Friedman's Permanent Income Hypothesis. The data unequivocally support the hypothesis of unitary elasticity of consumption out of real permanent income.

I and II are our answers to Questions A and B, respectively, addressed in this thesis. These results are conditional upon the lognormality hypothesis about the distribution of real disposable income (in the case of I) and the distr oution of real permanent income (in the case of II). In e.ther case, the consumer unit is taken to be an individual income-recipient.

Unfortunately it is not possible to compare these findings with those of the six works surveyed in Chapter 2: Staehle (1937), Polak (1939), Metcalf (1972), Blinder

(1975), Van Doorn (1975) and Della Valle and Oguchi (1976). The reason lies in the methodological problems in these studies, as assessed in Chapter 2. Only the approach in the first experiment of Blinder (1975) may be considered comparable to our method. However, his process of testing for different marginal propensities to consume (across the income quintiles) is not ideal; correlation analysis built on these lines may obscure the pattern of the marginal propensity to consume for a typical consumer unit, and consequently the distribution effects.

We want to assure the reader that the choice of consumer unit and the data have in no way biased our results. Throughout this thesis, serious effort has been made to preserve the norms of consistency. The extra effort in Chapter 3 to compile evidence on the distribution of pre-tax personal income (in current dollars) in Canada, under the lognormality hypothesis, is a case in point. In addition, we can cite as evidence the extensive sensitivity analyses that have been performed in Chapters 4 and 5.

The emphasis in this thesis has been on methodological issues. But the implications go beyond such matters. One area of relevance is the eventual comparison of the Keynesians' and Friedman's theories of consumption---of course, with due allowance for the different consumption variables. An across-theory comparison makes sense only after all the elements in each theory are in place, properly identified, and empirically verified. The income distribution is regarded in Keynes (1936) to be one such factor; but so far it has not been a part of the empirical implementations of the theory by the Keynesians. Similarly, there is a need to establish the maintained hypotheses in Friedman (1957) before the final step of inter-theory comparison may be taken. Chapter 4 of this thesis helps this cause by incorporating the distribution of real disposable income in the current income-current expenditure framework. The work in Chapter 5 fills the vacuum on the side of the PIH by confirming the role assigned to to the distribution of real permanent income. These efforts bring us a step closer to the said objective. However, one still needs to identify the role of the distribution of real disposable income in the PIH-framework, and verify the aggregation assumption of Friedman, before passing a judgment on the superiority of either theory.

The emphasis in this thesis has been on methodological issues; policy matters were not our primary concern. Nevertheless, some conditional inferences can be drawn from our work in Chapters 4 and 5. Within the framework of text-book Keynesian models, the government has not one, but two instruments of policy: one is the level of government expenditures and taxes, and the other is the distribution of either (with respect to income-classes). In a neo-classical framework, however, the redistribution of real disposable income will not be useful, if it only affects the distribution of a given real permanent income. However, the redistribution of real disposable income may still matter, if it affects the level of real permanent income---this is a possibility if the said redistribution alters people's expectations. This suggests a promising, as well as challenging, area for future research.

Of course, our thesis also draws attention to other avenues for further research. For example, one needs to define the type of questions which can be asked in the Life-Cycle Hypothesis framework of the consumption theory. Just to drop a hint, it may be mentioned that in this context aggregation across age-groups, at any point in time, is also necessary to arrive at an aggregate consumption function. Most probably, one shall need longitudinal data---income information for the same consumer units for more than one year, to answer any distribution question within the framework of this theory. In the context of our own work, serious effort is needed to compile information on the distribution of real disposable income, and also on the distribution of real permanent income---along with data on real permanent income itself. Moreover, there is also a genuine need to explore alternative functional forms to characterise the income distribution.

We have used the lognormal distribution in our illustrations. Its advantage is that distribution notions can be identified with only one of the parameters of the income density (σ^{2} in our case). The same also holds for the gamma distribution (Salem and Mount, 1974). However, three- or four-parameter functions are very unlikely to offer such flexibility. For example, for the beta distribution inequality notions relate to more than one parameter; this follows the Gini coefficient formula for this distribution in McDonald and Ransom (1978). In such cases, one still has to establish the nonlinearity of the micro relation in order to demonstrate the existence of distribution effects. However, even if this is done, the distribution effects would not be as precisely defined as in the case of the lognormal (or gamma). The researcher would have to resort to policy simulations, by _a priori_ specifying changes in the parameter values.

APPENDIX A

'TAXATION STATISTICS' AS A SOURCE OF INFORMATION ON THE DISTRIBUTION OF PRE-TAX PERSONAL INCOME IN CANADA

In Canada, the more well-known source of information on personal income distribution is the Survey of Consumer Finances of Statistics Canada. However, it is available on an annual basis only since 1971. For the purpose of this thesis, a much longer annual time-series was necessary, so we have drawn on Taxation Statistics of Revenue Canada.

We use data on the distribution of all returns, both taxable and nontaxable, by income groups. For the more recent years, these data can be found in Table 2 under the classification "Basic Tables of Individual Statistics" in Taxation Statistics. Comparable information for the earlier years is reported in either Table B or Table 2 of the sections (of this publication) relevant for individual income statistics. In this appendix we give our assessment of two things, namely the data base and the income concept. Following this assessment, we describe some of our adjustments to the reported data. Before getting into the specifics of these items, we may remind the reader that we use data on income ranges and the number of all returns therein. The information on total gross income of these returns is not necessary for our method of estimating the lognormal function, and is, therefore, not used.

The micro unit in our analysis is an individual income-recipient. Taxation Statistics are based on samples drawn from the universe of all those who file income-tax returns---whether they are taxpayers or not is a secondary matter. According to taxation laws, a tax-filing unit may consist of both spouses who pool their incomes, and, thereby, choose to be taxed together. The existence of such joint returns raises problems regarding our assumption of one income-recipient per tax-return. However, we notice that for the entire length of our data period, except 1946, the tax rate structure is independent of marital status. Moreover, beyond a certain minimum income level, each spouse has had to file a tax return separately. This minimum level equalled the Basic Personal Exemption upto 1971, and it has been slightly more than twice this exemption level since 1972. This basic exemption was $750 prior to 1949, $1000 from 1940 to 1971, and it has been adjusted for the inflation rate annually while starting at $1500 in 1972. Thus the phenomenon of joint returns is restricted to cases in which one spouse is in the lower income brackets. The impact of this may be a slight understatement of inequality in the distribution of income among the individual income-recipients. Given that the tax rate structure itself is independent of marital status, it would not be unreasonable to conclude that this impact is negligible. The existence of joint-returns has another aspect, namely its effect on the sample base for income-recipients. This takes us to questions related to coverage of tax data for purposes of the distribution (of income) among income-recipients.

Joint returns are one reason why one may question the use of tax data for our stated objective. One can think of other instances in which income-recipients may not be filing tax returns, thereby eroding the sampling universe for the distribution of income among the income-recipients. For example, at very low levels of income, people may choose not to file a tax return, or recipients of illegal incomes may avoid filing altogether. But these problems do not deflate the usefulness of the taxation data. For one thing, we note that the Basic Personal Exemption limit has been fairly low, and it is a criminal offence not to file a tax-return if one is in the taxable-income range. Thus over and above these exemption limits, the filing of tax-returns should be fairly common. Recently, there has been the added incentive of tax-credits which should have boosted the tax-return filing. We also note that the number of all returns (as a fraction of total popultion aged 15 years or more) has been over 50% since 1957; it is over 73% in 1976.

As well, we should also remember that not everybody is an income-recipient. In our opinion, the data for the more recent years is fairly representative of the true picture regarding income-recipients. The stability of the parameter values obtained in Chapter 3 suggests that situation may not be as bad in the earlier years as one might first think. Thus we have few reservation in concluding that the data should be an adequate representation of the income distribution among the individual income-recipients.

Now we shall briefly discuss matters related to the income concept. Pre-tax personal income in the taxation data is gross income before any allowable deductions are made or taxes paid. Since 1972, unemployment insurance, military pay and allowances, and capital gains and losses are also recorded as income. The first two items, taken together, are likely to have a positive impact on the degree of equality in the income distribution. On the other hand, capital gains and losses may be viewed as items adding to the magnitude of income inequality. On the whole, the inclusion of these three items in income may as well have offsetting impacts for the inequality picture. The stability of the parameter values, reported in Chapter 3, also enhance this view. So we would presume that the definitional change in the income concept in 1972 does not render the data inconsistent for analytic purposes.

Now we look at the compatibility of the gross income concept in Taxation Statistics to that of personal income (gross of taxes) in the National Accounts. Unlike the National Accounts, income in taxation data excludes imputed items such as the rental value of owner-occupied dwellings and the value of produce consumed on the farm. Of course, since 1972 capital gains and losses are an extra item which are excluded from income in the National Accounts. But how far are these differences important so as to affect the pattern of the distribution of pre-tax personal income? We suggest that to answer this one should discount these marginal differences in income concepts against all the things that allegedly escape the tax man. Intutively speaking, the inclusion of the imputed items in the tax data might push many of low income earners toward the median income, thereby increasing income equality. On the other hand, a proper allowance for high incomes, generally believed to be under-reported, would increase income inequality. Thus, on an average basis, the existing tax data may be representative of the inequality picture in the distribution of pre-tax personal income. The empirical evidence may go either in favour of or against this inference. However, pending such an evidence, we trust that the tax data can be used to capture the pattern of the distribution of personal income (gross of taxes).

As mentioned in the begining, now we shall briefly explain some of our adjustments to the tax data. The first of these relate to the allocation of composite figures of non-taxable returns to income groups (where they actually belong). For a few years, we have complete distribution of taxable returns by income-classes. But for selected income ranges, either in the middle or along the upper end of the distribution, only an aggregate figure of non-taxable returns is given in Taxation Statistics. We assume these non-taxable returns to be distributed in the same manner as the comparable taxable ones. The second of our adjustments relate to "loss figure" in the tax data. This category includes returns which report negative incomes. A careful look at the data revealed that for the years 1949-1955, the entire loss figure belonged to the nontaxable category of "Under $1000" group. So we put the number of returns in the loss category back into this group (although lowest income in our estimations is still $1.00). Of the remaining years, the tax data for 1946-1948 and 1969-1976 contain numbers in this loss category. As these could not be traced to any specific group (regarding the income before adjustment for losses), so we made no adjustment to the data for these years.

BIBLIOGRAPHY

Aitchison, A.J. and J.A.C. Brown (1957), The Lognormal Distribution, Cambridge University Press.

Alamgir, M. (1974), "Some Analysis of Distribution of Income, Consumption, Saving and Poverty in Bangladesh", The Bangladesh Economic Review, 2, 737-818.

Beach, C.M. and J.G. MacKinnon (1978), "A Maximum Likelihood Procedure for Regressions with Autocorrelated Errors", Econometrica, 46, 51-58.

Blinder, A.S. (1975) "Distribution Effects and the Aggregate Consumption Function", Journal of Political Economy, 83, 447-75.

Brady, D.S. and R.D. Friedman (1947), "Savings and the Income Distribution", Studies in Income and Wealth, Vol. X, National Bureau of Economic Research.

Branson, W.H. (1979), Macroeconomic Theory and Policy 2nd Edition, Harper & Row.

Carlton, D.W. and R.E. Hall (1978), "The Distribution of Permanent Income", in Z. Griliches, W. Krelle, H. Krupp and O. Kyn (eds.) Income Distribution and Economic Inequality, Halsted Press (John Wiley & Sons).

Christ, C.F. (1967), Econometric Models and Methods, John Wiley & Sons.

Darby, M.R. (1974), "The Permanent Income Theory of Consumption---A Restatement", Quarterly Journal of Economics, 88, 228-50.

Della Valle, P.A. and N. Oguchi (1976), "Distribution, the Aggregate Consumption Function, and the Level of Economic Development: Some Cross-Country Results", Journal of Political Economy, 84, 1325-34.

Fisher, R.A. (1922), "On the Mathematical Foundations of Theoretical Statistics", Philosophical Transactions of the Royal Society of London, A, 222, 309-368 (Reprited in Collected Papers of R.A.Fisher, The University of Adelaide, 1971).

Friedman, M. (1957), A Theory of the Consumption Function, Princeton University Press.

------------ (1963), "Windfalls, the "Horizon," and Related Concepts in the Permanent Income Hypothesis", in Christ, C.F. and Others (eds.) Measurement in Economics, Stanford University Press.

------------ and G.S. Becker (1957), "A Statistical Illusion in Judging Keynesian Models", Journal of Political Economy, 65, 64-75.

Goldfeld, S.M. and R.E. Quandt (1972), Nonlinear Methods in Economics, North-Holland.

Harrison, A.J. (1981), "Earnings by Size: A Tale of Two Distributions", forthcoming, Review of Economic Studies.

Haavelmo, T. (1947), "Methods of Measuring the Marginal Propensity to Consume", Journal of the American Statistical Association, 42, 105-22.

Husby, R.D. (1971), "A Nonlinear Consumption Function Estimated from Time-Series and Cross-Section data", Review of Economics and Statistics, 53, 76-79.

----------- (1974), "The Husby Consumption Analysis: A Reply", Review of Economics and Statistics, 56, 403.

Johnson, H.G. (1952), "The Effects of Income-Redistribution on Aggregate Consumption with Interdependence of Consumers' Preferences", Economica, 19, 131-47.

Keynes, J.M. (1936), The General Theory of Employment Interest and Money, Harcourt, Brace and Company.

Love, R. and M.C. Wolfson (1976), Income Inequality: Statistical Methodology and Canadian Illustrations, Statistics Canada (Catalogue 13-559 Occasional).

Lubell, H. (1947), "Effects of Income Redistribution on Consumers Expenditures", American Economic Review, 37, 157-70.

Lydall, H.F. (1968), The Structure of Earings, Oxford University Press.

Mayer, T. (1972), Permanent Income, Wealth and Consumption, University of California Press.

McDonald, J.B. and M.R. Ransom (1979), "Functional Forms, Estimation Techniques and the Distribution of Income", Econometrica, 47, 1513-26.

Metcalf, C.E. (1972), An Econometric Model of Income Distribution, Markham.

Musgrove, P. (1980), "Income Distribution and the Aggregate Consumption Function", Journal of Political Economy, 88, 504-25.

Polak, J.J. (1939), "Fluctuations in United States Consumption, 1919-1932", Review of Economics and Statistics, 21, 1-12; "A Correction", 21, 88.

Salem, A.B.Z. and T.D. Mount (1974), "A Convenient Descriptive Model of Income Distribution: The Gamma Density", Econometrica, 42, 1115-27.

Schultz, T.W. (1971), Long Term Change in Personal Income Distribution: Theoretical Approaches, Evidence and Explanations, Rand Corporation.

Spitzer, J.J. (1977), "A Simultaneous Equations System of Money Demand and Supply Using Generalized Functional Forms", Journal of Econometrics. 5, 117-28.

Staehle, H. (1937), "Short Period Variations in the Distribution of Income", Review of Economics and Statistics, 19, 133-43.

Van Doorn, J. (1975), "Aggregate Consumption and the Distribution of Incomes", European Economic Review, 6, 417-23.