Federal Reserve Bank
of Minneapolis

Fall 1987

Quarterly Review

Why Is Consumption Less Volatile Than Income? (p. 2)

Lawrence J. Christiano

Federal Reserve Bank of Minneapolis

Quarterly Review

Vol. 11, No. 4 ISSN 0271-5287

This publication primarily presents economic research aimed at improving policymaking by the Federal Reserve System and other governmental authorities.

Produced in the Research Department. Edited by Preston J. Miller, Kathleen S. Rolfe, and Inga Velde. Graphic design by Phil Swenson and typesetting by Terri Desormey, Public Affairs Department.

Address questions to the Research Department, Federal Reserve Bank, Minneapolis, Minnesota 55480 (telephone 612-340-2341).

Federal Reserve Bank of Minneapolis
Quarterly Review Fall 1987

Why Is Consumption Less Volatile Than Income?

Lawrence J. Christiano
Senior Economist
Research Department
Federal Reserve Bank of Minneapolis

Many of today's public policy issues—such as the potential economic impacts of the federal budget deficit and proposed changes in taxes—require an understanding of how households make consumption spending decisions. Indeed, some issues hinge on it. For example, recently in this *Quarterly Review,* Aiyagari (1987a) described two views about how consumption decisions are made which have strikingly different implications for large government deficits. One says such deficits have no effect on interest rates, investment, or saving; the other, that they drive up interest rates, reduce investment, and impoverish the next generation.

Despite the substantial policy issues at stake, economists are still struggling to understand consumption decisionmaking at even the most basic level. In fact, recently their uncertainty has seemed to increase.

For a long time, economists thought they at least had generally agreed on an explanation for what is perhaps the most basic fact about these decisions, the fact that over time the total amount people spend changes much less than the total amount they earn—or as an economist would say, aggregate consumption is much less volatile than aggregate income. The traditional explanation (due to Milton Friedman) is known as the *permanent income hypothesis.* It assumes, roughly, that people make their spending decisions based on what they expect their income to be in the long run, not just the short run. Their spending doesn't necessarily change, therefore, whenever their income changes. Many of those income changes are expected and so already built into people's spending plans. The rest are surprise income changes, but many of them may not change spending, either, if people don't think they are going to last. Spending should be considerably affected by surprise income changes that seem *permanent,* that is, but only minimally by those that seem *temporary.* The traditional assumption has been, moreover, that most surprise income changes are temporary; hence, consumption should be less volatile than income, which it is.

Recently, though, some economists have lost confidence in this explanation for the basic consumption/income relationship. Based on their analysis of aggregate income data, they have concluded that the traditional view, according to which surprise changes in income are temporary, is implausible. Instead, they are convinced that it is more plausible for households to respond to a surprise increase in income of, say, \$1 by raising their outlook for income forever. Specifically, the long-run income outlook should be raised about \$1.60, according to Deaton (1986). Deaton pointed out the dramatic implications of these developments for economists' traditional way of thinking about consumption. He showed that when the permanent income hypothesis is combined with this new view about the nature of surprise income changes, then the implications of the traditional argument are stood on their head. The prediction now is that consumption ought to be more volatile than income, not less. Deaton's paradox (as it has come to be called) suggests to many that something is seriously wrong with the traditional explanation for a basic economic fact.

The situation is actually not so grave. Despite Deaton's challenge, the traditional explanation for the consumption/income relationship might still be right. Several researchers have argued that estimates of the long-run impact of a surprise change in income based on aggregate postwar U.S. income data are too imprecise to support any conclusion, including Deaton's. This should not be surprising, since—as I show here—the relevant long run for this issue is from 10 to 15 years and U.S. postwar data offer only three or four nonoverlapping intervals of this length. Efforts are currently under way to increase the precision of estimates of the impact of surprise income changes by, for example, bringing income data from many countries into the analysis. Still, controversy on the magnitude of the impact continues, and in the end, Deaton's estimate may prove to be the best. However, even if that were true, the traditional explanation for the consumption/income relationship may not be too wrong. Results in Christiano 1987c—summarized here—suggest that even if Deaton's view of the impact of a surprise change in income is accepted, a small change in another of the permanent income model's assumptions is enough to pull the theory's implications back into line with the basic economic fact.

Essentially, the model adjustment involves changing the assumed source of the surprise income changes. The permanent income model, by assuming that the return to investing in capital is fixed, implicitly posits that the return is unrelated to the factors that produce output movements. The adjusted model I describe embodies the *real business cycle* perspective (associated with Kydland and Prescott 1982, Long and Plosser 1983, and Prescott 1986).[1] In that type of model, income changes stem from productivity shocks that simultaneously change the reward to investing. Therefore, the response of consumption to an unexpected, permanent jump in income is relatively restrained, as households simultaneously increase saving in order to take advantage of the increased return to investing. Evidently, the key ingredient of this proposed way to resolve the Deaton paradox is a presumed positive correlation between income and the anticipated return to investing in capital. Determining whether this proposal actually resolves the paradox, then, requires determining whether the required correlation is empirically plausible. This is the subject of ongoing research.

Thus, there are at least two possible ways to resolve the Deaton paradox. One is to revert to the traditional view that the effect of a surprise change in income is temporary. The other is to modify the permanent income model's assumed source of income fluctuation. Neither of these two proposed ways to resolve the paradox does so decisively yet; each requires further research to establish its plausibility. That is why the title of this paper is a question.

My objective here is to describe the Deaton paradox and the proposals to resolve it, at a fairly basic level. I thus describe the basic building blocks: the permanent income model, the real business cycle model, and several mathematical models about the persistence of surprise changes in income. Since the second proposed resolution to the paradox requires comparing the permanent income and real business cycle models, they need to be placed on a comparable basis. This is done by showing that they both belong to a general family of models called the *equilibrium growth* (or simply *growth*) model, which is increasingly becoming the standard framework for analysis in macroeconomics.[2] (The observation that the permanent income model is a special case of the growth model—something not generally known—is due to Sargent 1986 and L. Hansen 1985.)

The Equilibrium Growth Model

Here I lay the groundwork for describing the permanent income and real business cycle models by describing the simplest equilibrium growth model which includes them as special cases. First I describe the growth model in general, informal terms. This description follows recent developments which show that the growth model can be thought of as an economy populated by numerous heterogeneous, mortal people. I also describe an alternative interpretation of that model economy according to which it is populated by just one fictitious agent who lives forever. Because of this agent's resemblance to Defoe's fictional character, I call the agent *Robinson Crusoe.* This interpretation of the growth model is conventional and very convenient for my formal discussion of the growth model.

An Informal Look at the Multiagent Economy

The growth model is an abstract economy populated by a large number of households and firms. As time

[1] In these models, only *real* things, like unexpected productivity changes (or *shocks*), cause recurring fluctuations in general activity. *Nominal* things, like money, play no role.

[2] Since the permanent income and real business cycle models do not include money, neither does the growth model I describe here. For a growth model that incorporates money, see Marshall 1987.

evolves in the growth economy, some people are born into households and others die. During their lifetimes, people choose how much of the economy's good to consume and how to divide their time between *labor,* defined as work in the marketplace, and other activities, which are conventionally (and perhaps inappropriately) called *leisure.* If their income exceeds current consumption, they lend (or save) the difference; otherwise, they borrow (save a negative amount). People are purposeful in that their choices reflect their *preferences,* their attitudes about consumption today versus consumption tomorrow and about work versus leisure. The mathematical representation of preferences is the *utility function.*

Different people in the growth model have different preferences and labor productivities, reflecting in part their different ages. Therefore, a given person's saving might be positive at some times and negative at others. Moreover, at any given time, some people's saving is positive and others' is negative. Total saving, if positive, results in the accumulation of capital—things like factories, office buildings, and airplanes. If total saving is negative, the capital stock wears out. The fact that capital can accumulate means that production in this economy can increase over time; this is why it is called a *growth* model.

Firms possess the economy's *technology,* the knowledge about how to convert capital and labor effort into output. The mathematical representation of technology is the *production function.*

The growth model economy is simpler than an actual economy in several respects. In an actual economy, output is composed of many different goods (for example, cars, houses, food, health services, transportation). The growth model abstracts from this diversity by consolidating everything into one homogeneous good. Similarly, it abstracts from the many different actual types of capital by assuming that the capital stock is homogeneous and reflects past accumulation of the single produced good. The relationship of capital and output in the growth model is much like that of clay and putty: one is a hardened, congealed version of the other.

The multitude of people and firms in the growth model interact anonymously in markets.[3] These are *competitive* markets in the sense that prices are not set by any individual households or firms. Given the prices of labor and goods, people determine how much labor they want to supply and how many goods they want to buy. In addition, through their saving behavior, households acquire ownership of the stock of capital, which earns a competitive return (a rental rate) from firms. Given market prices, firms demand labor and the services of capital and they supply goods. In a competitive equilibrium, then, market prices are such that labor, goods, and capital markets clear; that is, they are *in equilibrium.* This is why the growth model is called an *equilibrium* model.

I must emphasize, however, that markets being in equilibrium in this model does not imply that everything in it is fixed over time. Quite the contrary: the equilibrium values of all variables in the growth model economy can change constantly over time in ways that are only imperfectly predictable. This reflects the constant, difficult to predict changes in the factors that determine the productivity of capital and labor (including weather, strikes, inventions, managerial skills, and the educational level and technical skills of workers). These factors, called *fundamentals* or *technology shocks,* account in part for the unpredictability inherent in actual economies.

A Formal Look at the Robinson Crusoe Economy

A *solution* to the growth model is a detailed specification of everyone's consumption, hours worked, and savings activity, along with prices such as wages and interest rates on loans. Here, though, I am only concerned with economywide averages of variables, particularly average consumption and income. And I only need some prices, such as the wage rate and the rate of interest. Fortunately, economywide average values of a growth model's equilibrium variables can be computed without first calculating each person's consumption, hours worked, and saving. These average values can instead be viewed as reflecting the choices of a single, fictitious agent who lives forever—who represents an average across the multitude of diverse, finite-lived people living in the growth economy. This representative Robinson Crusoe–like agent has preferences over economywide average consumption and work and chooses these variables subject to the available technology for converting capital and work into output, so as to maximize utility over an infinite-horizon lifetime.[4]

---

[3]For a detailed discussion of the competitive markets in which households and firms are assumed to interact, see Prescott 1986 and G. Hansen 1985. They describe two different—though equivalent—market settings. Prescott considers a *sequence of markets* equilibrium in which households and firms meet and trade every quarter. Hansen describes a *date zero* market equilibrium in which households and firms meet just once, at the beginning of time, to sign contracts. Hansen's discussion takes place explicitly in the context of the permanent income model.

[4]It is beyond the scope of this paper to formally present the growth model economy at the level of the finite-lived individual agents who are assumed to

□ *Technology*
In deciding what to do with output in the growth economy, Robinson Crusoe faces a *resource constraint:* uses of final output cannot exceed the amount available. The growth model assumes that gross output, which I call y_t, is allocated to only two uses: consumption, c_t, and gross capital investment, dk_t. (For a discussion of the empirical measures of these and other variables used later, see the Appendix.) This, then, is the resource constraint:

(1) $c_t + dk_t = y_t$.

The production function specifies how output is related to the factors of production. The permanent income and real business cycle models assume only three factors of production: capital, hours worked, and a term representing the fundamentals. These are here denoted by k_t, h_t, and z_t, respectively. An abstract representation of the production function, which includes those assumed by the permanent income and real business cycle models as special cases, is

(2) $y_t = f(k_t, z_t h_t)$.

The production function, f, indicates that the greater the value of the fundamentals, z_t, the more productive is an average hour's work.

Gross capital investment, dk_t, and the stock of capital, k_t, are linked in this way:

(3) $k_{t+1} = (1-\delta)k_t + dk_t$.

Here δ is the quarterly rate of depreciation of the stock of capital.[5] Thus, if the stock of capital is k_t at the beginning of quarter t, then the undepreciated part of that stock which remains at the end of the quarter is $(1-\delta)k_t$. The stock of capital available at the beginning of next quarter, k_{t+1}, is composed of this plus the part of y_t devoted to investment, dk_t.

Taken together, equations (1)–(3) describe the basic tradeoffs faced by Robinson Crusoe. In particular, the more leisure, l, is taken now (that is, the lower is h_t), the less current output is available for consumption and investment. High current consumption reduces the rate of capital accumulation, which reduces the future stock of capital. So high current consumption comes at the expense of reduced future output. In principle, even the determination of the fundamentals, z_t, involves tradeoffs. For example, time devoted to education must be taken from time in the home or at work. The permanent income and real business cycle models abstract from the factors that determine z_t by simply assuming it evolves exogenously, or outside the model. (See Lucas 1985 and Romer 1986 for models in which the economic decisions that determine z_t are modeled explicitly.)

□ *Preferences*
How Robinson Crusoe resolves the basic tradeoffs is determined by this agent's preferences. As of date t, Crusoe is assumed to value alternative uncertain paths of consumption and leisure, $\{c_{t+j}, l_{t+j}; j=0,1,2,\ldots\}$, according to the following expected, discounted utility function:

(4) $$E_t\{u(c_t,l_t) + \beta u(c_{t+1},l_{t+1}) + \beta^2 u(c_{t+2},l_{t+2}) + \ldots\} = E_t \textstyle\sum_{j=0}^{\infty} \beta^j u(c_{t+j}, l_{t+j})$$

where $0 < \beta < 1$. Here leisure and hours worked are linked by the restriction $l_t + h_t = T$, the total number of hours available in a quarter; u is the period utility function; and $E_t x$ is the expected value of x, conditional on information available at the beginning of quarter t.

According to (4), Crusoe cares not only about current consumption and leisure, but also about their future values. Therefore, at date t all of y_t is not necessarily consumed; some of it is saved (that is, $dk_t > 0$) for future consumption.

□ *Model Solution*
A solution to this model is a mathematical equation relating Robinson Crusoe's decisions at a given date to information available to Crusoe at that date. Such an equation is called a *contingency plan.* In particular, at date t Crusoe chooses contingency plans for setting c_{t+j}, h_{t+j}, and k_{t+1+j} for $j = 0, 1, 2, \ldots$ as a function

---

populate it. (For this, see Aiyagari 1987b.) Instead, I start the formal part of my discussion at the level of the Robinson Crusoe interpretation. However, it is important to remember that this particular interpretation of the growth model has little economic interest in itself and is intended only to facilitate solving for the economywide average values of the model's variables. The economically interesting interpretation of the growth model is the one in which numerous heterogeneous agents with finite lives interact in competitive markets. In practice, once Robinson Crusoe's preferences and the technology available to this agent are specified, the degree of heterogeneity in the underlying multiagent economy is severely restricted.

[5]Equation (3) implicitly abstracts from population growth. If the gross rate of population growth were a constant, n, then (3) would have to be written $k_{t+1} = [(1-\delta)/n]k_t + dk_t$. Although I implicitly set $n = 1$ throughout this paper, in the computation of the solution to the real business cycle model, I set $n = 1.00325$, which corresponds to an annual population growth rate of roughly 1.2 percent. For a formal treatment of that model which exhibits where and why n enters into the mathematical equations that describe the real business cycle model, see Christiano 1987c, forthcoming.

of information available contemporaneously, $\Omega_{t+j} = \{z_{t+j-s}, c_{t+j-1-s}, h_{t+j-1-s}, k_{t+j-s};\ s = 0, 1, 2, 3, \ldots\}$, to maximize (4) subject to (1)–(3) and a specification of the statistical properties of z_t.[6] These contingency plans are functions c, h, and k:

$$
(5) \qquad \begin{aligned} c_{t+j} &= c(\Omega_{t+j}),\ h_{t+j} = h(\Omega_{t+j}), \\ k_{t+1+j} &= k(\Omega_{t+j}). \end{aligned}
$$

Crusoe's contingency plans for dk_{t+j} and l_{t+j} can be derived from k and h using (3) and the fact that $l_t = T - h_t$.

The Robinson Crusoe perspective on the growth model not only makes solving the model for equilibrium quantities easier, but also can be used to compute equilibrium prices. I will use this representative agent approach to get formulas for the prices I will need later. One is the market-clearing *wage rate*, w_t, which is the marginal product of labor: $\partial f(k_t, z_t h_t)/\partial h_t$. Therefore, income attributable to labor effort, y_{ht}, is $w_t h_t$ or

$$
(6) \qquad y_{ht} = [\partial f(k_t, z_t h_t)/\partial h_t] h_t.
$$

Note that w_t and y_{ht} can be computed using the equilibrium values of k_t and h_t. Another market price I need is the *risk-free rate of interest*, r_t. This is the yield on a bond which, at a cost of one unit of the time t consumption good, entitles the holder to $1 + r_t$ units of the consumption good in period $t + 1$ with certainty. From the Robinson Crusoe perspective, $1 + r_t$ is the number of date $t + 1$ goods Crusoe requires to be compensated for giving up one unit of the good in period t. In utility terms, Crusoe's cost of giving up that unit is approximately $u'(c_t) \equiv \partial u(c_t)/\partial c_t$. From the perspective of date t, Crusoe's benefit of $1 + r_t$ goods in period $t + 1$ is $\beta E_t(1+r_t)u'(c_{t+1}) = \beta(1+r_t)E_t u'(c_{t+1})$. The risk-free rate of interest is the value of r_t that equates benefits and costs:

$$
(7) \qquad 1 + r_t = u'(c_t)/[\beta E_t u'(c_{t+1})].
$$

The consumption contingency plan, c, can be used to express r_t as a function of Ω_t. According to (7), the risk-free rate is low if the marginal utility of consumption next period is expected to be high. This reflects the fact that if Crusoe values consumption next period highly, then Crusoe requires few goods next period to be compensated for giving up a relatively low-valued date t good. What this corresponds to in the multiagent market economy that Crusoe stands in for is this: When people value consumption next period highly, then the current supply of loans is large and a low interest rate is sufficient to clear the loan market.

The risk-free rate of interest, r_t, must be distinguished from the *return on investment in capital*, which is $1 + R_t = [\partial f(k_{t+1}, z_{t+1}h_{t+1})/\partial k_{t+1}] + (1-\delta)$. The bracketed part of this sum is the direct increment to period $t + 1$ output due to a one-unit increment in the capital stock, k_{t+1}. The other term $(1-\delta)$ is the amount of extra capital, per unit of capital invested, left over at the end of $t + 1$. Unlike the value of r_t, the value of R_t is uncertain at date t since it will be determined in part by z_{t+1}, which is unknown at date t. (Recall that k_{t+1} is chosen at date t, so it is known then.) A particular difference between r_t and R_t, then, is that the former is not and the latter is a random variable as of date t. In general, it is not even true that $E_t R_t = r_t$.

The difference between $E_t R$ and r_t is called the *risk premium*. A formula for it can be obtained by studying the costs and benefits Robinson Crusoe weighs when contemplating investment in an extra unit of capital. In particular, to increase k_{t+1} by one unit, Crusoe has to reduce c_t by one unit, with utility cost $u'(c_t)$. The payoff at date $t + 1$ of this investment is $1 + R_t$, and the utility of this from the perspective of date t is $\beta E_t u'(c_{t+1})(1+R_t)$.[7] Now, Crusoe invests up to the point where benefits equal costs, so that $u'(c_t) = \beta E_t u'(c_{t+1})(1+R_t)$. Any two random variables x and y are related as $\text{cov}(x,y) = Exy - ExEy$, where $\text{cov}(x,y)$ denotes the covariance between x and y. Thus, $u'(c_t)/\beta = \{\text{cov}_t[u'(c_{t+1}), 1+R_t] + E_t u'(c_{t+1})E_t(1+R_t)\}$. Dividing by $E_t u'(c_{t+1})$, using (7), and rearranging yields

$$
(8) \qquad E_t R_t - r_t = -\text{cov}_t[u'(c_{t+1}), 1+R_t]/E_t u'(c_{t+1}).
$$

Since $E_t u'(c_{t+1}) \geq 0$, the risk premium is negatively related to the indicated conditional covariance term. So if the payoff on investing in capital $(1+R_t)$ is negatively correlated with marginal utility, then the expected return on capital, $E_t R_t$, exceeds the risk-free return.[8]

---

[6]The maximization problem must also obey the nonnegativity constraints: $k_t, h_t, c_t, y_t, dk_t \geq 0$. According to the last one, capital investment is *irreversible;* once put in place it cannot be reduced at a rate faster than the depreciation rate. This is consistent with the putty/clay analogy mentioned above.

[7]In general, $E_t u'(c_{t+1})(1+R_t) \neq (1+R_t)E_t u'(c_{t+1})$ since R_t is not known at date t. (As pointed out below, the permanent income model R_t is an exception; it is a constant.) This is to be contrasted with the risk-free rate, r_t, which is known at date t, so that $E_t u'(c_{t+1})(1+r_t) = (1+r_t)E_t u'(c_{t+1})$.

[8]This illustrates an important result in the theory of finance, namely, that the risk premium on an asset reflects not its variance but the covariance of its payoff with consumption. For example, if the covariance term in (8) is positive, so that capital investment represents insurance against periods in which the marginal utility of consumption is high, then the risk premium is negative. In this case, the volatility of the payoff on capital investment makes it more desirable than the risk-free asset.

This is because investments in capital are relatively unattractive—they yield the highest return when returns are least valued—so that a high yield is needed to induce people to invest. Note that if f is linear in k_t—so that R_t is a constant—then the risk premium is zero and $r_t = R_t = r$, a constant. This is not surprising, since investment in capital is then risk free. (In the permanent income model, f is linear.)

The Permanent Income Model

Here I follow Sargent 1986 and L. Hansen 1985 in deriving the permanent income hypothesis from a growth model with a particular specification of preferences and technology. I call this model the *permanent income model.* I derive a formula for the ratio of the volatilities of consumption and labor income that it implies and show how the size of this ratio is determined by the properties of labor income over time (its *dynamic* properties). I present the result not well known before Deaton 1986: that when labor income takes a particular form, this model can imply that consumption is more volatile than labor income.

The Model and an Approximate Solution

The permanent income version of the growth model has the following restrictions:

$$(9)\quad \begin{aligned} &u(c_t,l_t) = -(c_t-b)^2/2,\ \ f(k_t,z_th_t) = qk_t + z_th_t,\\ &\beta(q+1-\delta) = 1 \end{aligned}$$

where $0<\beta<1$ and $b>0$. Note that the utility function u does not include l_t explicitly. Instead, to avoid complications of no concern here, I abstract from the determination of h_t (and, hence, l_t) by assuming it evolves exogenously over time.[9] This and the fact that the production function f is linear in z_th_t implies that labor income is exogenous, with $y_{ht} = z_th_t$. The fact that f is linear in k_t implies that the risk-free rate of interest in the economy is a constant r, equal to $R = q - \delta$. That is,

$$(10)\quad r = q - \delta.$$

This and the assumption on preferences imply, approximately, that households set consumption c_t at the highest possible level they think they can sustain indefinitely, given their expectations for labor income. This highest sustainable level of consumption is called *permanent income,* which accounts for the model's name.[10]

To compute permanent income, substitute for y_t and dk_t in (1) from (2), (3), and (9). Then use (10) and the fact that $z_th_t = y_{ht}$ to get

$$(11)\quad k_{t+1} = (1+r)k_t + y_{ht} - c_t.$$

Now consider a consumption decision made at date t, when k_t and y_{ht} are known to the household. Let c_t be the maximum consumption level households expect to be able to sustain indefinitely. It must satisfy not only (11), but also future versions of (11):

$$(12)\quad k_{t+1+i} = (1+r)k_{t+i} + E_ty_{ht+i} - c_t$$

for all $i = 1, 2, 3, \ldots$. It can be shown that the maximum value of c_t that satisfies (11), (12), and $k_t \geq 0$ is

$$(13)\quad c_t = rW_t$$

where

$$(14)\quad W_t = k_t + \left[(1+r)^{-1}\textstyle\sum_{i=0}^{\infty}(1+r)^{-i}E_ty_{ht+i}\right].$$

Equations (13) and (14) define the permanent income hypothesis. There, rW_t and W_t are per capita *permanent income* and *wealth,* respectively. Note that wealth has two parts: k_t, which is called *nonhuman wealth,* and the other, bracketed piece, *human wealth.* Nonhuman wealth is just the existing stock of capital. The term for human wealth is the present discounted value of expected future labor income. It is the value that the work force of the society would have if there were a stock market in workers.

Equation (13) expresses c_t as a function of information available at date t and so is a solution to the permanent income model. This can be seen by noting that (13) [with (14)] is in the same form as (5), expressing c_t as a function of date t information: k_t and the information in the date t conditional expectation. Equation (13) is not an explicit solution for c_t until the

---

[9]See Christiano, Eichenbaum, and Marshall 1987 for a version of the permanent income hypothesis which models the household hours decision explicitly. I avoid that approach here because it only complicates matters without altering the principal implications of concern here: those for the consumption/income relationship.

[10]The exact solution to the permanent income model is not analytically tractable (Chamberlain and Wilson 1984). It is a contingency plan for c_t and k_t which maximizes (4) subject to (1), (9), and $c_t, k_t, dk_t \geq 0$. The contingency plan for c_t that I describe—in which c_t is equated with permanent income—is the exact, unique solution to a modified version of the permanent income model in which the restrictions $c_t, k_t, dk_t \geq 0$ are not imposed and $E_t\sum_{j=0}^{\infty}\beta^jk_{t+j}^2 < \infty$ is imposed. (See L. Hansen 1985 for a formal derivation of this solution.) I hope that this solution is a good approximation to the solution of the original model. Research to investigate this issue would be worthwhile.

statistical properties of y_{ht} are specified, so the conditional expectation can be evaluated. As we will see, those statistical properties can make a great difference in the model's implications for the relative volatility of consumption and labor income.

A Formula for the Consumption/Income Relationship

The key implication of the permanent income model here is its implication for the relative volatility of consumption and income. To get a formula for this, I first need an expression for the change in consumption $c_t - c_{t-1}$.

Recall the permanent income model's implication that c_t is set at a level that households believe is sustainable indefinitely. This implies that c_t differs from c_{t-1} only when something happens to income at the start of quarter t that was not anticipated in quarter $t-1$, when c_{t-1} was set. Since the interest rate in the permanent income model is fixed by the linearity assumption on f, the only component of earnings in t that is uncertain as of $t-1$ is labor income. Thus, c_t differs from c_{t-1} in this model only if period t labor income differs from what was expected in quarter t.

But the magnitude of the difference, $c_t - c_{t-1}$, depends only in part on the exact magnitude of the unexpected component of period t labor income. It also depends on how much the unexpected part of y_{ht} induces households to revise their expectations about period $t+s$ labor income, for $s = 1, 2, 3, \ldots$. The permanent income model implies that $c_t - c_{t-1}$ is the *annuity value* of the revision to expectations about period $t+s$ labor income for $s = 0, 1, 2, 3, \ldots$.[11] This can be derived formally by subtracting (8) lagged one period from itself and using (11) and (14) to arrive at

$$(15)\quad c_t - c_{t-1} = r\big[(1+r)^{-1} \times \textstyle\sum_{i=0}^{\infty}(1+r)^{-i}(E_t y_{ht+i} - E_{t-1}y_{ht+i})\big].$$

The expression $E_t y_{ht+i} - E_{t-1}y_{ht+i}$ is the revision in households' expectation about future labor income y_{ht+i} due to the new information available in period t, but not in period $t-1$. One of these terms, $E_t y_{ht} - E_{t-1}y_{ht} = y_{ht} - E_{t-1}y_{ht}$, is called the *innovation* (or surprise change) in y_{ht} and is the difference between y_{ht} and what people expected it to be as of date $t-1$. Note that if all of these revised expectation terms are zero, then $c_t = c_{t-1}$: consumption does not change. The right side of (15) is the annuity value of the revisions to the outlook for current and future income. So a compact way to characterize (15) is this: according to the permanent income model, the change in consumption is the annuity value of revisions to the outlook for current and future income. Equation (15) is what I call the *fundamental equation* of the permanent income model.[12]

I make a proportionality assumption so that (15) can be substantially simplified:

$$(16)\quad E_t y_{ht+i} - E_{t-1}y_{ht+i} = \psi_i(y_{ht} - E_{t-1}y_{ht})$$

for $i = 1, 2, 3, \ldots$. Also, let $\psi_0 \equiv 1$. The parameter ψ_i is a multiplier which says how much the forecast of y_{ht+i} for $i > 1$ is revised as a result of an innovation in y_{ht}.[13]

The simplification of the fundamental equation, (15), is obtained by substituting (16) into it:

$$(17)\quad c_t - c_{t-1} = \Psi(y_{ht} - E_{t-1}y_{ht})$$

where

$$(18)\quad \Psi = r(1+r)^{-1}\textstyle\sum_{i=0}^{\infty}(1+r)^{-i}\psi_i.$$

For obvious reasons, we can call Ψ the *annuity value of a \$1 innovation in labor income.* It can also be thought of, though, as the ratio of the volatilities (or the *relative volatility*) of consumption and labor income. This interpretation derives from the fact that one measure of the volatility of a variable is the standard deviation of its innovation. The innovation in consumption, according to the permanent income model, is $\Psi(y_{ht} - E_{t-1}y_{ht})$. This has standard deviation $\Psi\sigma$, where σ is the standard deviation of the innovation in y_{ht}.

[11] The *annuity value* of a stream of possibly different payments—say, $x(1)$, $x(2)$, $x(3)$, ...—is the corresponding constant payment stream with equal present value. Suppose, for example, that the interest rate is $1 + r$. Then the present value PV of $x(s)$ for $s = 1, 2, 3, \ldots$ is $\sum_{s=1}^{\infty}(1+r)^{-s}x(s)$. The corresponding annuity value is $r \times PV$. As expected, if $x(s) = x$, a constant for all s, then $r \times PV = x$.

[12] An implication of a well-known property of conditional expectations is that $E_{t-1}[E_t y_{ht+i} - E_{t-1}y_{ht+i}] = 0$ for $i \geq 0$. Thus, according to equation (15), $E_{t-1}c_t = c_{t-1}$: consumption is a random walk. This implication of the permanent income model was pointed out in the famous paper by Hall (1978), who also derived equation (15).

[13] The proportionality assumption, (16), amounts to an assumption that only past y_{ht} is useful in forecasting future y_{ht}, that the addition of past c_t, for example, does not help (c_t does not Granger-cause y_{ht}). Also, this assumption implies that the regression of $c_t - c_{t-1}$ on current and past y_{ht} has a fitted disturbance identically equal to zero. This can be seen from (17) below, which has current and past y_{ht} on the right side of the equality and no error term. [Past y_{ht} appears only implicitly, in $E_{t-1}y_{ht}$, which is a function only of y_{ht} (because of the Granger-causality observation).] Christiano, Eichenbaum, and Marshall (1987) avoid this grossly counterfactual implication by working in an environment in which a proportionality relation like (16) does not hold. To follow their lead here would complicate matters without altering my conclusions, according to Campbell and Deaton 1987 and West 1986.

The Importance of the Dynamic Properties of Labor Income

Now I can show how the value of Ψ depends on the dynamic properties of y_{ht}, particularly on how long an innovation lasts.

□ Some Simple Examples

Recall that, according to the fundamental equation of the permanent income model, the only thing that prompts households in this economy to change consumption is something that induces a revision in the outlook for current or future labor income. Only $y_{ht} - E_{t-1}y_{ht} \neq 0$ can trigger a revision in the outlook for income at dates $t+j$ for $j>0$. There are two extreme possibilities. On the one hand, an innovation in y_{ht} could trigger a revision of the same size in the outlook for income at all future dates ($\psi_i = 1$, $i \geq 0$); then the innovation in y_{ht} is said to be *permanent.* On the other hand, an innovation in y_{ht} may trigger no revision in future income ($\psi_i = 0, i > 0$); then the innovation is said to be *temporary.* According to (15), the *persistence* of the income innovation (its degree of permanence) makes a great deal of difference in terms of the impact on consumption. Two simple examples show this.

EXAMPLE 1. *A Permanent Innovation*
Suppose the innovation in labor income today is \$100, and households are induced to also raise their estimate of future income \$100. That is, $E_t y_{ht+j} - E_{t-1} y_{ht+j} = \$100$ for $j = 0, 1, 2, \ldots$. Then, according to (15), households adjust consumption by the annuity value of this innovation. For any interest rate, r, the annuity value is

$$r[(1+r)^{-1}(\$100) + (1+r)^{-2}(\$100) + (1+r)^{-3}(\$100) + \ldots] = r(\$100/r) = \$100.$$

Thus, when the innovation in income is permanent, households increase consumption by the full amount of the innovation. Under these circumstances, that is the maximum increase in consumption they can sustain into the indefinite future.

EXAMPLE 2. *A Temporary Innovation*
Now suppose that a \$100 innovation in labor income is thought by households to be temporary, so it has no effect on their outlook for income in subsequent periods. That is, $E_t y_{ht} - E_{t-1} y_{ht} = \$100$ and $E_t y_{ht+j} - E_{t-1} y_{ht+j} = 0$ for $j > 0$. If the quarterly interest rate $r = 1$ percent, the annuity value of this innovation is

$$0.01[(1.01)^{-1}(\$100) + 0 + \ldots] \cong 0.01(\$100) = \$1.$$

Thus, when a \$100 innovation is viewed as temporary, (15) implies that households increase consumption only \$1 and invest the remaining \$99. With the rate of interest 1 percent and a \$100 temporary increase in income, \$1 is the maximum increase in consumption that can be sustained indefinitely.

The dependence of consumption's response on the persistence of an innovation is further illustrated by Example 3. This is more realistic than the first two because it explicitly models the uncertainty in y_{ht}. Doing so clarifies the nature of an innovation and makes precise the link between it and revisions to future forecasts. In addition, the form that the labor income process takes in Example 3 is a simple version of the one used later with actual U.S. data. Finally, Example 3 shows that there are many other possibilities in addition to the extreme ones in which an innovation is either permanent or temporary.

EXAMPLE 3. *Innovations With Explicit Uncertainty*
Suppose this is the labor income process:

$$y_{ht} = \mu + \phi y_{ht-1} + \epsilon_t.$$

This is called a *first-order autoregressive* [AR(1)] model for y_{ht} (*autoregressive* because y_{ht} depends on its own previous values, or *lags; first-order* because here there is just one lag). Here, ϵ_t is an unpredictable random variable with zero mean, and μ and ϕ are constants. Also, $E_{t-1} y_{ht} = \mu + \phi y_{ht-1}$, so that ϵ_t is the innovation in y_{ht}. It is easy to confirm that here $E_t y_{ht+j} - E_{t-1} y_{ht+j} = \phi^j \epsilon_t$ for $j = 0, 1, 2, \ldots$ (that is, $\psi_i = \phi^i$, $i \geq 0$).[14]

For concreteness, suppose that $\phi = 0.5$. Then this model of labor income says that when there is a \$100 innovation in y_{ht} (when $\epsilon_t = 100$), households are induced to raise their forecast of next period's income by \$50, the following period's by \$25, and so on. The revision to the outlook of income far into the future is close to zero since $\phi^j \epsilon_t$ gets smaller as j increases, as long as ϕ is less than one in absolute value. Evidently, for $0 < \phi < 1$, the situation is intermediate to those in Examples 1 and 2; when $\phi = 1$, it is exactly Example 1;

[14]To see this, simply note that y_{ht+j} can be expressed as $y_{ht+j} = \mu(1+\phi+\phi^2+\ldots+\phi^j) + \phi^{j+1} y_{ht-1} + \epsilon_{t+j} + \phi\epsilon_{t+j-1} + \phi^2\epsilon_{t+j-2} + \ldots + \phi^j\epsilon_t$. Then $E_t y_{ht+j} = \mu(1+\phi+\phi^2+\ldots+\phi^j) + \phi^{j+1} y_{ht-1} + \phi^j\epsilon_t$ since $E_t\epsilon_{t+s} = 0$ for $s > 0$. Also, $E_{t-1} y_{ht+j} = \mu(1+\phi+\phi^2+\ldots+\phi^j) + \phi^{j+1} y_{ht-1}$.

and when $\phi = 0$, it is Example 2.

For this model, the annuity value of an innovation to income is, from equation (18),

$$\Psi = r(1+r)^{-1}\Sigma_{i=0}^{\infty}[\phi(1+r)^{-1}]^i = r(1+r-\phi)^{-1}.$$

The dependence of Ψ on r and ϕ shows that the annuity value of an innovation (or the relative volatility of consumption and labor income) depends on the interest rate and the persistence properties of labor income which, in this example, are governed by ϕ. The more persistent the innovation—that is, the greater ϕ—the higher is Ψ. When $\phi = 1$, the change in consumption is, as in Example 1, the whole of the innovation in income $[\Psi(1,r) = 1]$. When $\phi = 0$, the change in consumption, as in Example 2, is a tiny fraction, approximately r, of the innovation in income $[\Psi(0,r) = r/(1+r)]$.

A notable feature of Ψ here, which will play a role later, is that it is very sensitive to changes in ϕ when ϕ is near 1. If, for example, $r = 1$ percent and ϕ moves only slightly, from 0.900 to 0.980 to 0.990 to 0.999, Ψ leaps at the same time from 0.090 to 0.330 to 0.500 to 0.909. This is not surprising. The fact that ψ_i is a function of the ith power of ϕ implies that a very small change in the value of ϕ has a growing effect on the ψ_i's as the time horizon increases (for larger values of i). Moreover, for values of ϕ near 1, the effect is already appreciable after only two years (for $i \geq 7$). For example, $(0.9)^7 = 0.48$, but $(0.999)^7 = 0.99$. The relative weight of ψ_7 in Ψ is substantial, being equal to $(1+r)^{-7} = 0.93$ for $r = 1$ percent. [See equation (18).]

A possibility that will arise with the actual U.S. data is that the annuity value of the innovation in income could exceed the innovation itself ($\Psi > 1$). This would happen when the innovation induced households to raise their forecast of future income by an amount sufficiently larger than the innovation. The following example is designed to illustrate this possibility in the simplest setting. The example also illustrates a way to analyze the effects on consumption of an income innovation by decomposing the innovation into parts. This method is useful to study the effects of the type of innovations in y_{ht} that are in the actual U.S. data.

EXAMPLE 4. *A More-Than-Permanent Innovation*

Suppose that the current period's innovation in labor income is \$100 and that this induces households to raise their expectations about future income by more than \$100—say, by \$150 ($\psi_i = 1.5, i > 0$). If the interest rate is 1 percent, the annuity value of this innovation is

$$0.01[(1.01)^{-1}(\$100) + (1.01)^{-2}(\$150) + (1.01)^{-3}(\$150) + \ldots] = \$149.50.$$

This example can be thought of as a combination of Examples 1 and 2. Specifically, it is equivalent to the household receiving a permanent positive innovation in income of \$150 (which causes consumption to jump \$150) coupled with a temporary negative income innovation of \$50 (which causes consumption to drop \$0.50).

□ The Real World

The AR(1) model for y_{ht} in Example 3 is too simple to represent the actual U.S. labor income data. That data suggest a more complicated model:

$$(19) \quad y_{ht} = \mu + \alpha t + \rho_1 y_{ht-1} + \rho_2 y_{ht-2} + \epsilon_t$$

where μ, α, ρ_1, and ρ_2 are constants and, again, ϵ_t is a random variable with zero mean. This is a second-order autoregressive [AR(2)] model with trend for y_{ht}. It will be convenient to write (19) in terms of ϕ_1 and ϕ_2, defined as the values of ϕ which solve $\phi^2 - \rho_1\phi - \rho_2 = 0$. These are called the *AR roots* of y_{ht}. Since $\phi_1 + \phi_2 = \rho_1$ and $\phi_1\phi_2 = -\rho_2$, (19) can be rewritten as

$$(20) \quad y_{ht} = \mu + \alpha t + (\phi_1+\phi_2)y_{ht-1} - \phi_1\phi_2 y_{ht-2} + \epsilon_t.$$

Note that (20) can be rewritten as

$$(21) \quad y_{ht} - \phi_1 y_{ht-1} = \mu + \alpha t + \phi_2(y_{ht-1} - \phi_1 y_{ht-2}) + \epsilon_t.$$

Now we can see that the representation of labor income can take at least two very different forms. When $\phi_1 = 1$, $|\phi_2| < 1$, and $\alpha = 0$, equation (19) is an AR(1) model in first differences of y_{ht}: $y_{ht} - y_{ht-1}$. This type of model is usually written in the form of equation (21). [The AR(1) model in first differences is Deaton's preferred model.] It is known as a *difference stationary* representation because in this form y_{ht} is not covariance stationary (its variance is not defined), but its first difference is. This means, roughly, that the changes in y_{ht} fluctuate with constant amplitude about a constant mean over time, while the levels of y_{ht} follow no particular trend. When ϕ_1 and ϕ_2 are both less than one in absolute value and α is possibly nonzero, equation (19) is a *trend stationary* representation for y_{ht}. In this form, the levels of y_{ht} fluctuate with constant amplitude about a linear trend in time. Evidently, these two forms are special cases of the AR(2) model with trend.

It is obvious from (19) [or (20),(21)] that the random variable $\epsilon_t = y_{ht} - E_{t-1}y_{ht}$, the surprise change in labor income. In addition, it can easily be verified that

$$(22) \quad E_t y_{ht+j} - E_{t-1}y_{ht+j} = [\phi_1/(\phi_1-\phi_2)]\phi_1^j\epsilon_t - [\phi_2/(\phi_1-\phi_2)]\phi_2^j\epsilon_t$$

so that

$$(23) \quad \psi_j = [\phi_1/(\phi_1-\phi_2)]\phi_1^j - [\phi_2/(\phi_1-\phi_2)]\phi_2^j$$

for $j = 0, 1, 2, \ldots$. Note that when either ϕ_1 or ϕ_2 are zero, this model reduces to the one in Example 3. Note also that when ϕ_1 and ϕ_2 are both nonzero, the effect on the forecast revision of ϵ_t looks like the sum of the effects of two innovations like those in Example 3. The first is $[\phi_1/(\phi_1-\phi_2)]\epsilon_t$ with persistence parameter ϕ_1, and the second is $[-\phi_2/(\phi_1-\phi_2)]\epsilon_t$ with persistence parameter ϕ_2. (This sort of possibility was suggested by Example 4.) Not surprisingly, the annuity value of ϵ_t is just the sum of the annuity values of these two innovations:

$$(24) \quad \Psi(\phi_1,\phi_2,r) = [r/(1+r-\phi_1)][\phi_1/(\phi_1-\phi_2)] + [r/(1+r-\phi_2)][-\phi_2/(\phi_1-\phi_2)].$$

Two things are worth noting about Ψ. First, when one of the ϕ's is 1 and the other is positive, then (20) is an AR(1) in first differences form and Ψ exceeds 1. Thus, consumption is more volatile than income. This may be seen by noting that when, say, $\phi_1 = 1$, (24) becomes $(1+r)/[1+r-\phi_2]$, which always exceeds 1 if $\phi_2 > 0$. For example, $\Psi(1.0,0.5,0.01) = 1.98$. This case is similar to the one in Example 4 in that an innovation in income is more than permanent. To see this, substitute $\phi_1 = 1, \phi_2 = 0.5, r = 0.01$, and $\epsilon_t = \$100$ into (23) to find that given a \$100 innovation in labor income the household revises its forecast of this income j periods in the future by $(2 - 0.5^j) \times \$100$. For $j = 0$, 1, 2, and 3, this is \$100, \$150, \$175, and \$187.50. The decomposition method in Example 4 can also help understand this more-than-permanent feature. When one of the roots (say, ϕ_1) of y_{ht} equals 1 and the other is between 0 and 1, then an innovation has a permanent part, $[1/(1-\phi_2)]\epsilon_t$, and a temporary negative part, $-[\phi_2/(1-\phi_2)]\epsilon_t$. Because the temporary part is negative, the revision in the outlook for income several periods in the future exceeds the innovation in income. (See Example 4.)

A second notable feature of Ψ is that it is very sensitive to changes in one of the roots of y_{ht} (the ϕ's) when that root is near 1. The relative volatility of consumption and income, that is, depends a lot on how long the income innovation is expected to last. As in Example 3, this is not surprising once one notes from (23) that ψ_i is a function of the ith power of each of the roots of y_{ht}. A consequence of this is that very small changes in a root when its value is close to 1 produce large changes in ψ_i's starting as soon as two years out ($i \geq 7$). Since those ψ's have substantial weight in Ψ, it follows that their sensitivity to the roots of y_{ht} translates into a sensitivity for Ψ. For example, if $r = 1$ percent, $\phi_2 = 0.5$, and ϕ_1 starts at 0.9, then small changes in ϕ_1—from 0.900 to 0.980 to 0.990 to 0.999—produce large increases in Ψ—from 0.18 to 0.66 to 0.99 to 1.80. To put this more concretely, if $r = 1$ percent, $\phi_1 = 0.9$, and $\phi_2 = 0.5$, then this permanent income model predicts that households spend just \$18 of a \$100 innovation in income and invest the rest. But if $\phi_1 = 0.999$ and nothing else is different, then the model has a dramatically different prediction: households respond to the \$100 income innovation by increasing consumption \$180—\$80 more than the income increase, which they get by reducing investment.

The Deaton Paradox . . . ?

It is well known that changes in income from one quarter to the next are positively correlated. So when Deaton estimated the AR(1) first difference representation using labor income data, the coefficient on the change in y_{ht-1} came out positive, implying that $\Psi > 1$. More precisely, his estimate of Ψ is close to 1.6. Thus, the permanent income model together with the AR(1) model in first differences of labor income implies that consumption ought to be roughly one and a half times as volatile as labor income. This conflicts sharply with the empirical fact that consumption is only about half as volatile as labor income. This conflict is ironical—and paradoxical—because many economists have believed that the observed relative smoothness of consumption is the principal reason for taking the permanent income model seriously.

Despite Deaton's finding based on the AR(1) difference model, there continues to be debate on whether there really is a paradox. In part, this is due to the fact that other models of income, which appear to fit the labor income data equally well, do not yield the counterfactual implication Deaton got. The continued controversy also reflects the suspicion that any measure of the long-run effect of an innovation in income based on postwar U.S. data must be extremely imprecise and unreliable. According to the permanent income model, the 10–15-year effects of an income innovation have a significant influence on the consumption decision, yet the postwar data include only three or four non-overlapping intervals of this length. The continuing controversy over the existence of a paradox indicates the possibility that, in the end, the traditional assumption that income innovations are temporary may be vindicated.

Deaton's Results

Here I reproduce Deaton's empirical results using a slightly different measure of labor income than he used. My measure is 0.66 times gross output (y_t). The reason for the factor 0.66 is that it is roughly the fraction of gross output due to labor in the postwar U.S. data. I estimated a trend stationary model and an AR(1) difference model by the method of ordinary least squares, using the 111 quarterly observations on these data between the third quarter of 1956 and the first of 1984. The results:

Trend Model

$$y_{ht} = \underset{(2.2)}{172.29} + \underset{(1.8)}{0.78t} + \underset{(14.7)}{1.35y_{ht-1}} - \underset{(-4.3)}{0.39y_{ht-2}} + \hat{\epsilon}_t; \qquad \hat{\sigma}_\epsilon = 49.3 \tag{25}$$

AR(1) Difference Model

$$\Delta y_{ht} = \underset{(2.3)}{11.55} + \underset{(4.1)}{0.37\Delta y_{ht-1}} + \hat{\epsilon}_t; \qquad \hat{\sigma}_\epsilon = 49.8. \tag{26}$$

In (26), $\Delta y_{ht} \equiv y_{ht} - y_{ht-1}$. The numbers in parentheses beneath the equations' parameter estimates are t-statistics, and $\hat{\sigma}_\epsilon$ is the standard deviation of the fitted regression disturbances, $\hat{\epsilon}_t$. Note that the two models have approximately the same implication for the volatility of labor income, σ_ϵ: that it is roughly \$50 per quarter. This is consistent with Deaton's finding that both models fit the data equally well.

To determine the models' implications for the volatility of consumption, the permanent income model requires that Ψ be computed. This in turn requires the roots of y_{ht}. These are 0.92 and 0.43 in the trend model and 1.0 and 0.37 in the difference. The implied values of Ψ are, then,

$$\begin{aligned} &\text{Trend} \\ &\Psi(0.92,0.43,0.01) = 0.19 \\ &\text{Difference} \\ &\Psi(1.0,0.37,0.01) = 1.58. \end{aligned} \tag{27}$$

The implications of these two models for the response of consumption to an innovation in income are obviously very different. The trend model implies that households respond to a \$100 innovation in income by increasing spending only \$19 and investing the rest. The difference model implies instead that they increase spending \$158 and reduce investment \$58. The trend model's 0.19 is more realistic than the AR(1) difference model's 1.58 since 0.19 is considerably closer to the empirical amount of smoothness in consumption (roughly 0.5).

The different implications of the two models reflect their different maximal roots: 0.92 and 1.0. That is clear when the 0.92 root in the trend model is replaced by 1.0: its Ψ jumps from 0.19 to 1.74, just about the value of Ψ in the difference model. This sensitivity of Ψ to variations in the value of a root that is near 1 is striking. The reason for it is that such changes have a large effect on the trend or difference model's implication for the long-run effect of an innovation and the permanent income model assigns substantial weight to these long-run effects in the consumption decision. The effects of an income innovation on expectations of future income are measured by each model's ψ_i's. Chart 1 displays how the two estimated models measure these ψ_i's in the ten years after an innovation. Note how the two models' ψ_i's are similar in the first year (in the first few quarters), but then they diverge. Note also, in Chart 2, how slowly the consumption decision's weights on the ψ_i's die out. The weight on the ψ for the forecast revision ten years after the innovation is still about 70 percent of that for the initial quarter. [Recall from (18) that the relative weight on ψ_i is $(1+r)^{-i}$.]

Evidently, according to the permanent income model, the relationship between consumption and income depends on very long-run properties of the labor income process. But the trend and AR(1) difference models differ substantially on what those properties are. Yet, according to Deaton, it is impossible to determine which model fits the income data better. Essentially, this is because, with only 30 years of data, a root of 0.92 is virtually indistinguishable statistically from 1.0 (Christiano and Ljungqvist 1987).

Charts 1 and 2

Explaining the Trend and Difference Models' Different Implications for the Consumption/Income Relationship

Chart 1 Revisions to the Income Outlook Induced by a $1 Surprise Income Increase in the Two Labor Income Models

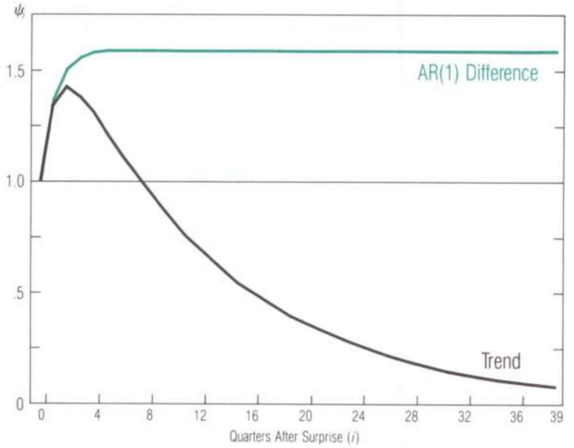

Chart 2 Weights on Revisions to the Income Outlook in the Consumption Decision of the Permanent Income Model

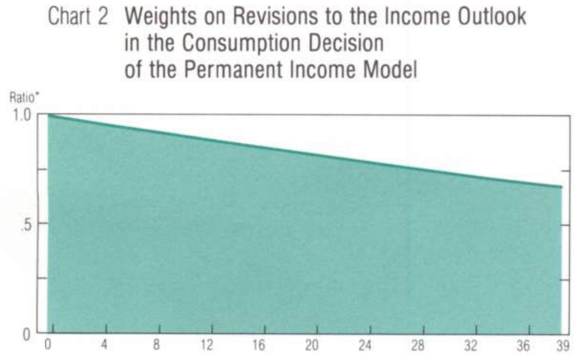

*Weight on ψ_i relative to weight on ψ_0

A Questionable Resolution

One potential resolution to the Deaton paradox is to replace Deaton's assumption that the long-run effect of an income innovation is strong with the traditional view that it is weak. The results above suggest that one way to do this is to replace Deaton's AR(1) difference representation of income by the trend model. Today we seem to have no statistical basis for ruling this out. Nevertheless, there are serious questions about the plausibility of the trend model which prevent it from supplying a convincing resolution to the paradox. The difference model is more plausible, and there are other difference models than Deaton's. Among them is at least one which implies the traditional assumption that an income innovation is relatively transient and, in particular, does not imply the paradox. Unfortunately, even this way to restore the traditional view of the income innovation is controversial.

□ *The Implausible Trend Model*

Despite its more realistic Ψ, the trend model cannot now be used to resolve Deaton's paradox. This is because there is empirical evidence that the trend model is implausible. This is clear in Chart 3, where you see quarterly levels of per capita labor income in the United States between the mid-1950s and the mid-1980s and a trend line computed from just the first 15 years of data and extended through the rest of the period. Evidently, a household that lived through the 1950s and 1960s, and believed the trend model of income, would have had serious doubts by the mid-1970s since its expectation for income to return to the old trend would have been repeatedly disappointed. The peak of the 1971–73 U.S. expansion just barely brought labor income back to the old trend line, and the 1975 recession seems to have driven it away permanently. Thus, the trend specification for income cannot be used without formally taking account of the possibility that households discarded their old trend specification and somehow settled on a new one.[15] That highly complex modeling problem is certainly beyond the scope of this paper.[16]

An alternative resolution to that problem is to modify the trend specification in (25) by adding a term in t^2, that is, by making the equation's trend quadratic,

[15]The trend in Chart 3 is $3{,}096.02 + 28.13t$ and was estimated by regressing y_{ht} on a constant and time trend over the period from the third quarter of 1956 to the third of 1970. Formal tests over the period from the third quarter of 1956 to the first of 1984 confirm the visual impression that there was a break in trend in the third quarter of 1970. This is based on the following regression:

$$y_{ht} = \underset{(3.3)}{398.90} + \underset{(2.1)}{131.41d_t} + \underset{(3.1)}{3.57t} - \underset{(-2.4)}{2.43d_t t} + \underset{(14.2)}{1.30y_{ht-1}} - \underset{(-4.7)}{0.42y_{ht-2}}$$

where d_t is a dummy variable that is zero before the fourth quarter of 1970 and one after the third quarter and numbers beneath the point estimates are t-statistics. The t-statistic on $d_t t$ is large enough to reject the null hypothesis that the associated coefficient is zero, indicating a significant change in trend. I repeated this test with log y_{ht} instead of y_{ht} in the above regression and got the same result.

[16]Perron (1987) argues that a trend model with an exogenous change in the coefficient on t fits postwar U.S. data better than a difference model.

so that it curves as the data appear to. Unfortunately, this has its own problems. I reestimated this quadratic trend model, and its implied trend appears in Chart 4.[17] Note the bow shape of the trend. The role of the bow is to capture the apparent deterioration in economic performance that occurred in the 1970s and early 1980s. Assuming households used this quadratic trend model to forecast income is equivalent to assuming that in the 1950s they confidently expected that 20 years later the economy would be weak. This seems implausible, to say the least. I think it is much more probable that in the 1950s households looking 20 years ahead found an upturn as likely as a downturn.

In sum, there are two problems with using the trend model to try to resolve Deaton's paradox. One is that the assumption underlying Deaton's own specification—that households used the same linear trend model to forecast income throughout the postwar era—seems implausible because it conflicts sharply with the properties of the income data themselves. The other problem is that a solution to the first problem—replacing the linear trend by a quadratic—is unattractive because it presumes an implausible amount of foresight by households.

□ *The Unreliable Difference Model*

The trend model is not the only alternative to Deaton's specification which rationalizes the traditional view about the effect of an income innovation. There are difference models other than Deaton's which also do this. Moreover, the problems encountered with the trend model do not occur if we instead assume that households think of labor income as a difference rather than a trend process. This can be seen in Chart 5, which plots the quarterly changes in the levels of y_{ht} since the mid-1950s. The two straight horizontal lines indicate the averages of the changes during two periods: from the third quarter of 1956 to the third of 1970 and from the fourth quarter of 1970 to the first of 1984. According to the difference model, when there is an unusually large or small change in income, all that can be expected is that the quarterly changes eventually return to some fixed underlying constant; the levels of y_{ht} do not necessarily return to any previous trend.

Under this view, the U.S. economy suffered several particularly bad shocks in 1975 and in the early 1980s, but after each one the quarterly changes returned to their underlying long-run average, which itself was fixed. This view is consistent with Chart 5's sawtooth pattern for the 1970s and 1980s. The contractions reflect the effects of bad shocks, and the expansions represent the resumption of the previous quarterly average change. Although the average change is slightly lower in the 1970s and 1980s than in the preceding two decades, it is not sufficiently smaller to shake the confidence of a believer in the difference specification.[18] Thus, we can plausibly assume that households used the same difference specification to forecast labor income throughout this period.

Therefore, the difference specification avoids the linear trend's problem in that the assumption that households used the same difference model to forecast income throughout the postwar period seems plausible. In addition, the difference model manages to capture the bow shape of postwar income data without assigning an implausible degree of foresight to households, as the quadratic trend model does. The difference model's account for the bow is that there were a couple of episodes of bad negative shocks which drove down the level of income, but did not alter its average change. According to the difference model, from the perspective of the 1950s, the economy 20 years later was just as likely to bow up as down, since a couple of episodes of good shocks are as likely as a couple of episodes of bad shocks.

Although a difference representation for y_{ht} appears more plausible than a trend representation, this does not mean that Deaton's AR(1) difference specification is necessarily the preferred one. An alternative is the unobserved components model studied by Clark (1987) and Watson (1986). The estimated version of that difference model implies that the long-run effect of

---

[17]The trend in Chart 4 is $2{,}936.51 + 37.93t - 0.16t^2$ and is the one implied by the following regression, estimated over the period from the third quarter of 1956 to the first of 1984:

$$y_{ht} = \underset{(3.1)}{336.81} + \underset{(2.6)}{3.94}t - \underset{(-2.1)}{0.02}t^2 + \underset{(14.2)}{1.31}y_{ht-1} - \underset{(-4.6)}{0.41}y_{ht-2}; \quad \hat{\sigma}_\epsilon = 48.5.$$

Numbers in parentheses are t-statistics. Note the statistical significance of the coefficient on t^2. I also tested the null hypothesis that the constant and coefficients on t and t^2 did not change in the third quarter of 1970. This joint null hypothesis fails to be rejected since the F-statistic has significance level 0.49. The quadratic trend model's Ψ is 0.09, about half the linear trend model's Ψ. This is not surprising because the deviations of the data from the quadratic trend are more temporary than the deviations from the linear trend. This is because the bow in the quadratic trend better mimics the apparent bow in the data.

[18]To test this formally, I estimated this AR(1) difference model for y_{ht} over the period from the third quarter of 1956 to the first of 1984:

$$y_{ht} - y_{ht-1} = \underset{(2.1)}{14.19} - \underset{(-0.6)}{5.26}d_t + \underset{(4.0)}{0.37}(y_{ht-1} - y_{ht-2})$$

where d_t is the dummy variable described in note 15 and numbers in parentheses are t-statistics. The small size of the t-statistic on the coefficient of d_t indicates there is no evidence against the null hypothesis that the mean quarterly change of labor income has remained unchanged. Again, I repeated this test for logs of y_{ht}, with the same result.

Charts 3–5

Three Views of Household Labor Income

3rd Quarter 1956–1st Quarter 1984 in the United States

Actual Data — Trend or Average

Chart 3 Levels With a Linear Trend

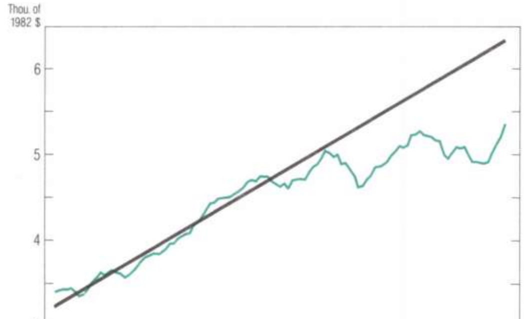

Chart 4 Levels With a Quadratic Trend

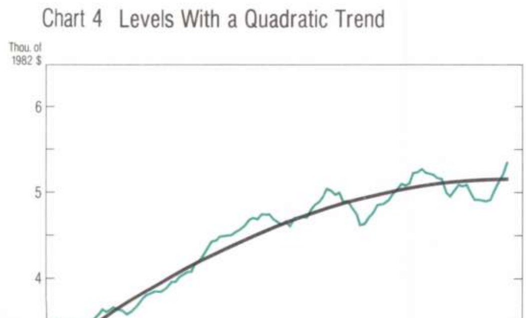

Chart 5 Changes With Averages

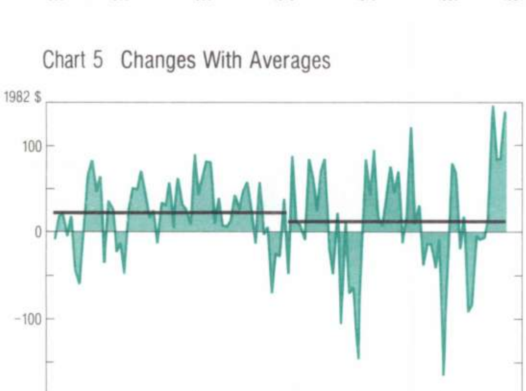

Sources of basic data: See Appendix.

a \$1 innovation to income is considerably less than \$1, so that with this representation of income the permanent income model implies an empirically plausible degree of consumption volatility. Obviously, either the AR(1) model or the unobserved components model (or both) must be giving a misleading estimate of the long-run impact of an income innovation. Further research is required to establish which one of these models yields the most plausible measure.

Cochrane (1987) argues that neither one is reliable for this measurement. Even though the reported standard errors for the long-run effect of an income innovation are small, he argues that the precision is spurious and reflects aspects of the models' structure, in which we can have little confidence. Cochrane proposes a procedure for measuring the long-run impact of an income innovation that does not depend on specifying a particular model of income. Application of this procedure to postwar U.S. data confirms that precisely estimating persistence with such a small data set is difficult (Campbell and Mankiw 1987a). This has led some researchers to expand the U.S. data set by considering data from other countries as well (Campbell and Mankiw 1987b, Kormendi and Meguire 1987).

In summary, we don't know yet which assumption is the most plausible: the traditional one, which is that income innovations are transient, or Deaton's, which is that they have a strong, permanent effect. Therefore, we cannot rule out the possibility that the paradox will ultimately be resolved by restoring the traditional view about the effect of an income innovation.

The Real Business Cycle Alternative

Again, one way to possibly resolve the Deaton paradox is to modify Deaton's AR(1) difference specification for labor income. Another way is to preserve that specification and consider the possibility that other parts of the permanent income model are misspecified. Here I do that by shifting to another equilibrium growth model, a close relative: a real business cycle model. I describe this model and show that, even though in it labor income is an AR(1) difference process, consumption is substantially less volatile than labor income.

The principal distinction between the permanent income and real business cycle models, which accounts for the latter's superior performance here, is that the permanent income model assumes the interest rate is fixed whereas the real business cycle model does not. Besides elaborating on this, I show here that there is a sense in which—aside from their family connection—these two are very similar models.

The Model

Like the permanent income model, the real business cycle model is a growth model with its own particular specification for preferences (the utility function u) and technology (the production function f):

$$(28)\quad u(c_t, l_t) = \ln c_t + \gamma l_t,\ \ f(k_t, z_t h_t) = (z_t h_t)^{(1-\theta)} k_t^{\theta}$$

where γ and θ are parameters and, recall, c_t is consumption, l_t leisure, k_t the capital stock, h_t hours worked, and z_t the factors affecting productivity. Two features distinguish this model from the permanent income model. One is that here l_t appears in the u function. This implies that the hours decision is modeled explicitly; labor income, y_{ht}, is determined by the model (is endogenous). The other distinguishing feature is that f is not linear in k_t in the real business cycle model. This implies that in this model the risk-free interest rate fluctuates rather than remains constant and the risk premium is not zero.

I adopt the following specification for the real business cycle model's z_t:

$$(29)\quad z_t = z_{t-1}\exp(x_t), \quad x_t = \mu + \rho x_{t-1} + \epsilon_t.$$

The random variable x_t is the growth rate of z_t with the indicated AR(1) representation. This is an AR(1) difference specification for $\log z_t$. To see this, note that (29) can be rewritten as $\log z_t = (1+\rho) \log z_{t-1} - \rho \log z_{t-2} + \mu + \epsilon_t$. Then

$$(30)\quad \log z_t - \log z_{t-1} = \mu + \rho(\log z_{t-1} - \log z_{t-2}) + \epsilon_t.$$

The real business cycle model has considerably more parameters than the permanent income model. I assigned these values to them: $\rho = -0.077$, $\mu = 0.0035$, $\gamma = 0.0026$, $\beta = 0.99$, $\delta = 0.018$, $\theta = 0.39$, and $\sigma_\epsilon = 0.019$. This value for δ is required if the gross investment series implied by $dk_t \equiv k_{t+1} - (1-\delta)k_t$ is to resemble the gross investment series published by the U.S. Department of Commerce. Values for θ, γ, and $\mu/(1-\rho)$ are chosen to roughly match the model's implications for the average values of h_t, c_t/y_t, and k_t/y_t with their empirical counterparts in U.S. data for the period from the second quarter of 1956 to the first of 1984. The implied averages (and empirical values) for these variables are 323.9 (320.4), 0.72 (0.72), and 11.32 (10.58), respectively. The values of ρ, μ, and σ_ϵ are based on regression analysis of the time series properties of z_t, which can be measured using data on y_t, k_t, and h_t given the value assigned to θ. (See Christiano, forthcoming, for a careful discussion of this method of selecting parameter values.)

The method used to solve the model is described in Christiano 1987a,b, and the algebraic formulas for the decision rules are reported in Christiano 1987c (n. 4). All the data used are discussed in the Appendix.

Partial Success

I compute the real business cycle model's implication for the relative volatility of consumption and income by forming the ratio of an innovation in consumption to an innovation in labor income. This ratio is Ψ, according to (17).

To facilitate the interpretation of the outcome of this calculation, I first discuss part of the real business cycle model's innovation response functions. Chart 6 shows the first 30 quarters of the responses of c_t, h_t, dk_t, and y_t to a one standard deviation innovation in the growth rate of the technology shock z_t in period 2, assuming that the system is on a steady-state growth path in periods 0 and 1 (that is, $\epsilon_2 = 0.019$ and $\epsilon_t = 0$ for $t = 0, 1, 3, 4, \ldots$). More precisely yet, the curves in Chart 6 are the quarterly percentage deviations in these variables from a baseline scenario in which $\epsilon_t = 0$ for $t = 0, 1, 2, 3, \ldots$.

With the assumed structure of z_t, the innovation in ϵ_t drives z_t 1.9 percent above its baseline growth path in period 2, after which it declines to a path 1.76 percent above the baseline. After the shock, all the model's variables except h_t and the risk-free rate of interest, r_t, end up 1.76 percent above their baselines. The exceptions at first are increased by the shock, but eventually return to the value they had before it.

As Chart 6 shows, consumption rises only gradually to its higher growth path. One way to explain this is to recall Robinson Crusoe. In particular, Crusoe chooses not to adjust consumption immediately after an unexpected productivity increase because this agent also wants to increase saving and investment in order to take advantage of the temporarily higher rate of return on investment. That higher rate is reflected in capital investment's strong shock response. The incentive the shock creates to delay consumption by increasing saving is known as the *substitution effect,* whereas the incentive to increase consumption because the long-run ability to consume has increased is called the *income effect.*

Though the interest rate, r_t, does rise in response to the innovation in x_t, the rise is quite small. It is so small, in fact, that it is not shown in Chart 6. In the steady state, $r_t = 0.01667$; after the shock, it jumps to 0.01728, then declines monotonically back to 0.01667. Thus, the effect on the interest rate is a negligible six one-hundredths of a basis point.

In Chart 6, note the early spikes in the responses of dk_t, h_t, and y_t. This reflects the fact that 7.15 percent of the initial 1.9 percent jump in z_t is only temporary. The lack of a spike in c_t reflects the small response of consumption to a temporary disturbance, which, in turn, explains the pronounced spike in capital investment. (Here the income effect on consumption is very small, whereas the substitution effect on investment is not.) Hours also respond fairly strongly to the temporary component in the productivity shock (as explained in G. Hansen 1985).

Chart 6

The Impact of a Technology Shock in the Real Business Cycle Model

Percent Deviation of Shocked Path From Steady-State, Unshocked Path*

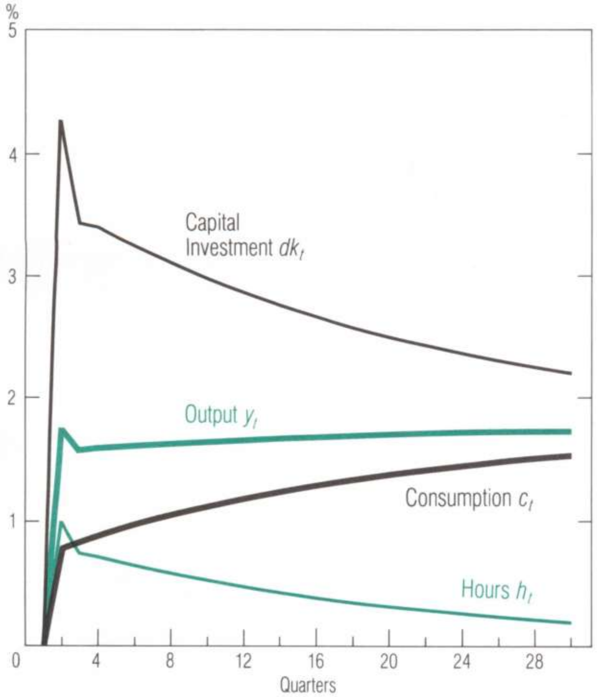

*The shock is an unexpected 1.9% increase in z_t in quarter 2; steady-state growth is assumed in quarters 0 and 1.

My goal here, remember, is to see how the model estimates the ratio of the jumps in consumption and labor income. The model's ratio of the jumps in consumption and total income in period 2 is 0.32; that is, consumption's innovation is about 32 percent of income's. In this model, $y_{ht} = (1-\theta)y_t = 0.61y_t$. Thus, the relevant relative volatility implied by the model is $0.32/0.61 = 0.52$. This real business cycle model seems to come close to matching the observed relative volatility of consumption and labor income, which is roughly one-half.

Unfortunately, though, I cannot claim to have entirely resolved the paradox raised by Deaton. That is because the model is clearly flawed: successive changes in income have correlation of about -0.119, reflecting the fact that the serial correlation properties of equilibrium output in the model closely mimic those of z_t. A naive application of the permanent income model's formula for Ψ thus yields $\Psi(1.0,-0.119,0.01) = 0.89$, considerably less than the value of 1.58 implied by U.S. data. To get a version of the real business cycle model with a more reasonable autocorrelation structure for changes in y_t, I changed ρ to 0.2, but left all other parameter values, including $\rho/(1-\mu)$, unchanged. With these parameter values, the model's first-order autocorrelation of labor income changes is about 0.35, and $\Psi(1.0,0.35,0.01) = 1.53$. Then the ratio of an innovation in consumption to an innovation in labor income is 0.84—higher than before, and than it should be, but about half as high as Deaton's paradoxical ratio. And unlike Deaton's ratio, the real business cycle model's is, like the actual ratio, less than 1; this model does, that is, correctly predict that consumption is less volatile than labor income.

The Essential Difference . . .

The reason for the real business cycle and permanent income models' sharply different implications seems to lie in the different nature of the shocks they assume. According to the real business cycle model, disturbances to output result from shocks that are expected to be permanent and which affect the rate of return on investment in capital. As we have seen, therefore, a rise in income signals not only that households' long-term ability to consume has risen, but also that the return on investment has. By itself, the increase in the long-run ability to consume motivates households to substantially increase consumption today. This income effect on consumption is partially offset, though, by the substitution effect as households take advantage of the increased return to saving.[19] In contrast, the permanent income model assumes that shocks affect the average product of capital, not the marginal product. Because of this and the linearity of f in k, the model implies a fixed interest rate, thus allowing only the income effect to operate on consumption. When shocks are partially permanent, the income effect is very strong, which accounts for the model's counterfactual implications when the difference specification is used.

. . . *And a Surprising Similarity*

Surprisingly, despite their very different implications for the volatility of consumption, the permanent income and real business cycle models have very similar implications for other dynamic characteristics of consumption. In particular, they both imply—at least roughly—that consumption follows a *random walk*. This means that consumption in the current period is the best predictor of consumption next period, so that any change in consumption is uncorrelated with information available in the current period or earlier.

The permanent income model implies that consumption is exactly a random walk. This can quickly be verified from equation (7) and the fact that $\beta(1+r) = 1$ [from (9) and (10)]. To see this, note that here (7) implies that $E_t u'(c_{t+1}) = u'(c_t)$, or $E_t c_{t+1} = c_t$ under the model's specification of u in (9). (See note 12.)

In the real business cycle model, r_t is not a constant; but recall how little it moved. This suggests that consumption is approximately a random walk in this model as well. To demonstrate this property, I used the model to simulate 1,000 sets of 112 observations for c_t, dk_t, h_t, and r_t. This was done using the approximate decision rules in Christiano 1987c (n. 4), starting the initial capital stock on a steady-state growth path, and using ϵ_t's drawn independently from a normal random number generator with mean zero and standard error 0.019. I needed 112 artificial observations because this is the length of my U.S. data set. For each artificial data set, I computed the correlation between the change in the log of consumption and the lag-one change in the log of consumption, income, and capital investment. In addition, I computed the correlation between the change in the log of consumption and lag-one hours worked and the lag-one real rate of interest. (I work in logs because the real business cycle model implies that is necessary to induce covariance stationarity.)

The accompanying table shows the means of the real business cycle model correlations across the 1,000 artificial data sets along with their associated standard deviations. The table also shows the correlations based on U.S. data. Note that all of the model correlations are small and have large standard deviations. In this sense, they are all close to zero, which is what they would be if consumption in the model were exactly a random walk. I computed (but don't show) the correlations up to lag four; they are also small, with large standard errors.[20]

As is obvious in the table, the model's random walk implication does not square well with the actual U.S. data. In Christiano 1987b, I argue that much of the discrepancy can be accounted for by the fact that the data are time averaged. A similar argument in the context of the permanent income model is made in Christiano, Eichenbaum, and Marshall 1987.

A Random Walk in the Real Business Cycle Model

First-Lag Variables Correlated With Consumption Growth*	Correlations With Consumption Growth*: Model Simulations**: Standard Deviation	Mean	U.S. Estimates†
Consumption Growth*	.108	.059	.271
Income Growth*	.098	.045	.204
Hours Worked	.129	.079	−.057
Investment Growth*	.093	.038	.161
Interest Rate	.130	.079	.104

*The growth of a variable x_t is defined as $\log x_t - \log x_{t-1}$.

**These are the results of 1,000 simulations, each 112 quarters long.

†The data period is from the 2nd quarter of 1956 to the 1st of 1984.

Sources of U.S. data: See Appendix.

[19] Another model that illustrates this possibility is in Christiano, Eichenbaum, and Marshall 1987. The model there which is most relevant is called the *discrete time stochastic labor requirement model*, which has two technology shocks: one affects average productivity, and one affects the intertemporal rate of return on investment. Here, as in Deaton's version of the permanent income model, the average productivity shock is specified to be an AR(1) in first differences. Econometric estimation of the model's parameters results in positive estimated correlation between the innovations in the average and marginal productivity shocks. This, in turn, results in the model having an empirically reasonable prediction for the relative volatility of consumption and income. The mechanism by which this is accomplished is identical to that in my real business cycle model.

[20] The real business cycle model has many other implications as well. Some of these are explored in Christiano 1987c, forthcoming.

Appendix
The U.S. Data Used in the Models

The U.S. data used in the models in the accompanying paper are primarily standard measures available from standard sources. (For a detailed discussion, see Christiano 1987d.)

The resource constraint, (1) in the paper, divides all output, y_t, into only two categories: consumption, c_t, and investment, dk_t. In view of this, I consolidate private and government spending. Thus, c_t and dk_t are private plus public consumption and investment, respectively. *Private consumption* here has three components: personal consumption expenditures on nondurables and on services and the service flow from the stock of durable goods held by households. My measure of the service flow is from the data base documented in Brayton and Mauskopf 1985. For the other two components of private consumption, I use the U.S. Department of Commerce National Income and Product Accounts (NIPA). *Public consumption* is measured as the NIPA government purchases of goods and services, but reduced by an update of the measure of government investment discussed in Musgrave 1980. *Investment* is the sum of that government investment, NIPA personal expenditures on durables, and NIPA private domestic investment. *Gross output* is the NIPA measure of gross national product plus the service flow from the stock of consumer durables minus net exports.

The *stock of capital*, k_t, is defined as the beginning-of-quarter stock of public and private equipment and structures plus the stock of consumer durables plus public and private residential capital. This definition conforms with the definition of investment. Given the U.S. data on dk_t and k_t, equation (3) in the accompanying paper can be used to measure the quarterly *rate of depreciation* in the capital stock, δ. I use G. Hansen's (1984) measure of *hours worked*, h_t.

Variables are converted to per capita terms by the working-age population, measured in a way that conforms with G. Hansen's (1984) measure of hours worked. All flow variables (c_t, dk_t, y_t, h_t) are measured at a quarterly rate.

A Feel for the Numbers

You may want an idea of the size of these measured variables. In the United States, between the second quarter of 1956 and the first quarter of 1984, the average values of h_t, c_t/y_t, and k_t/y_t are 320.4 hours, 0.72, and 10.58, respectively. In addition, the average value of the ratio of labor income to total income (y_{ht}/y_t) is roughly 0.66 and has been remarkably constant throughout the 20th century (Christiano, forthcoming, n. 3). Although no value of the quarterly depreciation rate results in an exact fit of the paper's equation (3), the value $\delta = 0.018$ seems to do best. This implies an annual depreciation rate of 7.4 percent. Since in the data k_t is composed of many different kinds of capital, $\delta = 0.018$ should be thought of as a kind of average depreciation rate across different kinds of capital, each with a different rate of depreciation. Included, for example, are both toasters and houses, which presumably have very different depreciation rates.

References

Aiyagari, S. Rao. 1987a. Intergenerational linkages and government budget policies. *Federal Reserve Bank of Minneapolis Quarterly Review* 11 (Spring): 14–23.

__________. 1987b. Overlapping generations and infinitely lived agents. Research Department Working Paper 328. Federal Reserve Bank of Minneapolis.

Brayton, Flint, and Mauskopf, Eileen. 1985. The MPS model of the United States economy. Manuscript. Board of Governors of the Federal Reserve System, Division of Research and Statistics, Econometrics and Computer Applications Section.

Campbell, John, and Deaton, Angus. 1987. Is consumption too smooth? Working Paper 2134. National Bureau of Economic Research.

Campbell, John Y., and Mankiw, N. Gregory. 1987a. Are output fluctuations transitory? *Quarterly Journal of Economics* 102 (November): 857–80.

__________. 1987b. International evidence on the persistence of economic fluctuations. Manuscript.

Chamberlain, Gary, and Wilson, Chuck. 1984. Optimal intertemporal consumption under uncertainty. Social Systems Research Institute Workshop Series 8422. University of Wisconsin, Madison.

Christiano, Lawrence J. 1987a. Dynamic properties of two approximate solutions to a particular growth model. Research Department Working Paper 338. Federal Reserve Bank of Minneapolis.

__________. 1987b. Intertemporal substitution and the smoothness of consumption. Manuscript. Federal Reserve Bank of Minneapolis.

__________. 1987c. Is consumption insufficiently sensitive to innovations in income? *American Economic Review* 77 (May): 337–41.

__________. 1987d. Technical appendix to "Why does inventory investment fluctuate so much?" Research Department Working Paper 380. Federal Reserve Bank of Minneapolis.

__________. Forthcoming. Why does inventory investment fluctuate so much? *Journal of Monetary Economics.*

Christiano, Lawrence J.; Eichenbaum, Martin; and Marshall, David. 1987. The permanent income hypothesis revisited. Research Department Working Paper 335. Federal Reserve Bank of Minneapolis.

Christiano, Lawrence J., and Ljungqvist, Lars. 1987. Money does Granger-cause output in the bivariate output-money relation. Research Department Staff Report 108. Federal Reserve Bank of Minneapolis.

Clark, Peter K. 1987. The cyclical component of U.S. economic activity. *Quarterly Journal of Economics* 102 (November): 797–814.

Cochrane, John H. 1987. How big is the random walk in GNP? Manuscript. University of Chicago.

Deaton, Angus. 1986. Life-cycle models of consumption: Is the evidence consistent with the theory? Working Paper 1910. National Bureau of Economic Research.

Hall, Robert E. 1978. Stochastic implications of the life cycle–permanent income hypothesis: Theory and evidence. *Journal of Political Economy* 86 (December): 971–87.

Hansen, Gary D. 1984. Fluctuations in total hours worked: A study using efficiency units. Working Paper. University of Minnesota.

__________. 1985. Indivisible labor and the business cycle. *Journal of Monetary Economics* 16 (November): 309–27.

Hansen, Lars Peter. 1985. Econometric modeling of asset pricing under rational expectations. Paper presented to the Fifth World Congress of the Econometric Society, in Cambridge, Mass.

Kormendi, Roger C., and Meguire, Philip G. 1987. The nonstationarity of aggregate output: A multicountry analysis. Manuscript. University of Michigan.

Kydland, Finn E., and Prescott, Edward C. 1982. Time to build and aggregate fluctuations. *Econometrica* 50 (November): 1345–70.

Long, John B., Jr., and Plosser, Charles I. 1983. Real business cycles. *Journal of Political Economy* 91 (February): 39–69.

Lucas, Robert E., Jr. 1985. On the mechanics of economic development. Paper prepared for the Marshall Lectures, Cambridge University, May.

Marshall, David A. 1987. Inflation and asset returns in a monetary economy with transaction costs. Manuscript. Graduate School of Industrial Administration, Carnegie-Mellon University.

Musgrave, John C. 1980. Government-owned fixed capital in the United States, 1925–79. *Survey of Current Business* 60 (March): 33–43.

Perron, Pierre. 1987. The Great Crash, the oil price shock, and the unit root hypothesis. Manuscript. University of Montreal.

Prescott, Edward C. 1986. Theory ahead of business cycle measurement. *Federal Reserve Bank of Minneapolis Quarterly Review* 10 (Fall): 9–22. Also in *Real business cycles, real exchange rates and actual policies,* ed. Karl Brunner and Allan H. Meltzer. Carnegie-Rochester Conference Series on Public Policy 25 (Autumn): 1–44. Amsterdam: North-Holland.

Romer, Paul M. 1986. Increasing returns and long-run growth. *Journal of Political Economy* 94 (October): 1002–37.

Sargent, Thomas J. 1986. Equilibrium investment under uncertainty, measurement errors, and the investment accelerator. Manuscript. Hoover Institution, Stanford University, and University of Minnesota.

Watson, Mark W. 1986. Univariate detrending methods with stochastic trends. *Journal of Monetary Economics* 18 (July): 49–75.

West, Kenneth D. 1986. The insensitivity of consumption to news about income. Discussion Paper 119. Woodrow Wilson School of Public and International Affairs, Princeton University.