OPTIMAL CONSUMPTION WITH STOCHASTIC INCOME: DEVIATIONS FROM CERTAINTY EQUIVALENCE

by

Stephen P. Zeldes

20-86

RODNEY L. WHITE CENTER FOR FINANCIAL RESEARCH
The Wharton School
University of Pennsylvania
Philadelphia, PA 19104

OPTIMAL CONSUMPTION WITH STOCHASTIC INCOME: DEVIATIONS FROM CERTAINTY EQUIVALENCE

Stephen P. Zeldes

Finance Department
The Wharton School
University of Pennsylvania

latest revision: August 1986

I would like to thank Andy Abel, Larry Ball, Stan Fischer, Bob Inman, George Pennacchi, Mike Woodford, members of the Macro Lunch Group at the University of Pennsylvania, and members of the "Contracting and Financial Markets" group at the 1986 NBER Summer Institute for helpful comments; and Ananth Krishnamurthy for his work as Research Assistant. This paper is a revised version of chapter 1 of my M.I.T. Ph.D. dissertation.

Optimal Consumption with Stochastic Income:

Deviations from Certainty Equivalence

Abstract

Those who study consumption behavior routinely assume that labor income is stochastic and that the utility function exhibits constant relative risk aversion. No one has derived closed form solutions to this problem, however, and therefore we do not know what the resulting consumption function looks like. In this paper, a numerical technique is used to give accurate approximations to the consumption function in multiperiod models with income uncertainty. The resulting consumption function is often dramatically different than the certainty equivalence solution typically used, in which consumption is proportional to the sum of financial wealth and the present discounted value of expected future labor income. The results help explain three important empirical consumption puzzles: the excess sensitivity of consumption to transitory income, the high growth of consumption in periods of low real interest rates, and the under spending of the elderly.

In the last section of the paper, the numerical technique is applied to examples in which borrowing constraints are imposed. It is seen that future constraints, which bind only in certain states of the world, can have effects similar to those of current constraints that are binding.

Stephen P. Zeldes
Finance Department
The Wharton School
University of Pennsylvania
Philadelphia PA 19104-6367
(215) 898-3477

I. Introduction

A great deal of research has been done recently deriving and testing implications of the life cycle/permanent income hypothesis under uncertainty. One of the most important sources of uncertainty facing individuals is that of labor income. Yet closed form decision rules for optimal consumption in the presence of uncertain labor income have not, in general, been derived. It seems strange that so much theoretical and empirical work has been done studying consumption and yet we do not even know what the optimal level of consumption or sensitivity of consumption to income should be in most very simple settings. In this paper I use numerical methods to closely approximate the optimal consumption function and the corresponding value function for some simple multiperiod problems. I then examine how consumption behavior differs from that implied by the certainty or certainty equivalence models used by most authors who write down closed form decision rules for consumption with random labor income.

This technique enables me to address some important questions about the optimal consumption function that could previously not be addressed because of the lack of closed form solutions. First, I can calculate the optimal amount of precautionary saving--saving done to hedge against the uncertainty in future labor income. Second, I can calculate the optimal sensitivity of consumption to both permanent and transitory changes in income. Finally, I use this technique to examine the effects that current and future borrowing constraints have on the level and sensitivity of consumption.

A few puzzles have arisen in the consumption literature recently. First, Hall and Mishkin (1982) have found that consumption appears to exhibit "excess sensitivity" to transitory changes in income, relative to a certainty or certainty equivalent benchmark. Second, with constant relative risk aversion utility, a certainty model implies that the growth of consumption is equal to the difference between the interest rate and the rate of time preference. Yet both aggregate time series and panel data studies have found extended periods during which the growth of consumption was too high to be justified by the ex-ante real interest rate and plausible (positive) estimates of the rate of time preference. Third, studies have found that dissaving by the elderly given their wealth is too little relative to a benchmark that involves no uncertainty and no bequest function.

The results of this paper potentially shed light on each of these puzzles. I provide examples in which individuals have constant relative risk aversion utility functions; face uncertain (and non-tradable) labor income streams, a riskless technology for borrowing and lending, and no external borrowing constraints; and are fully optimizing. Relative to a certainty model, the individual optimally chooses to have the sensitivity of consumption to transitory income "too" high, the expected growth of consumption "too" high, and the level of consumption "too" low.

In the last part of the paper, I use the same technology to examine models with borrowing constraints. The results show the importance not just of current constraints, but of future constraints as well, on current consumption. In particular, I look at circumstances in which current constraints are not binding and future constraints are binding only with probabilities between zero and one, and show that certain aspects of the consumption function are quite similar to what would occur if current constraints were binding.

The paper is structured as follows. Section II defines the optimization problem faced by individuals in the economy and discusses the assumptions necessary to derive the certainty equivalence (or simple permanent income hypothesis) solution. In Section III, I briefly review the existing literature and discuss some of the theoretical issues involved. Section IV briefly discusses the empirical puzzles. Section V presents the numerical results in the absence of constraints, and Section VI presents the results on borrowing constraints. Section VII concludes the paper.

II. The Consumer's Problem

Consider the standard problem of a consumer who lives for T periods and chooses optimal current consumption and contingency plans for future consumption. The objective is to maximize the expected value of a lifetime time separable utility function with rate of time preference equal to δ, subject to a series of budget constraints which will be specified. The only source of uncertainty that I consider here is uncertainty about exogenous future labor income, and I assume that no markets exist in which individuals can hedge against this uncertainty by trading contingent claims. In other words, the focus here is on uncertainty in non-traded labor income, as opposed to uncertainty about the rate of return on traded assets.[1] We can think of labor income as the dividend on a share of human capital, but the individual is forced to hold a fixed number of these shares and cannot trade them. The apparent absence in modern economies of markets for human capital is most likely due to problems of moral hazard and adverse selection, but these are not incorporated directly in the model below.

The formal problem can be summarized by an objective function, a transformation equation, and an initial and terminal condition. In each period t ($t = 1, \ldots, T$, where $T < \infty$), the consumer chooses C_t in order to:

Maximize $$E_t \sum_{j=0}^{T-t} [1/(1+\delta)]^i U(C_{t+j}) \tag{1}$$

subject to:

$$W_{t+1} = (W_t - C_t)(1 + r_t) + Y_{t+1} \tag{2}$$

$$C_t \geq 0 \tag{3}$$

$$W_T - C_T \geq 0 \tag{4}$$

where:

W_t = financial wealth in period t (after receiving income and before consuming;

r_t = the real riskless interest rate between t and t+1;

U = the one period utility function;

Y_t = labor income in period t;

C_t = consumption in period t;

E_t = the expectation operator, conditional on information available at time t;

T = the non-stochastic date of death.[2]

Individuals are allowed to borrow and lend at the riskless rate r_t, i.e., all loans must eventually be paid back with probability one. Define J_{T-t+1} as the maximized value of (1). That is, J_{T-t+1} is the value function for a T-t+1 period problem in terms of period t utility. Income is received at the very beginning of the period and then consumption is chosen.[3,4] The remaining financial wealth earns a rate of return r_t between the current and subsequent periods.[5] Using recursive substitution on (2) and imposing (4) yields:

$$W_t + \sum_{j=1}^{T-t} Y_{t+j}\left[\prod_{i=1}^{j} [1/(1 + r_{t+i-1})]\right] \geq C_t + \sum_{j=1}^{T-t} C_{t+j}\left[\prod_{i=1}^{j} [1/(1 + r_{t+i-1})]\right] \tag{5}$$

Recall that Y_{t+j} is unknown at time t for $j > 0$. Equation (5) could be restated in terms of future value instead of present value, and would say that the future value at time T of wealth and actual realizations of income must be at least as great as the future value of consumption--the individual cannot die in debt. Note that (5) implies an expected value budget constraint, but not vice-versa.[6]

In this paper, I concentrate on the case where r is constant over time and equal to δ. This problem might appear at first glance to be an easy one to solve. For any concave utility function for the case of variable, but non-stochastic, future income, the following is the solution.[7]

$$C_{CEQ,t} = k_{T-t+1}[W_t + HW_t] \tag{6}$$

where:

$$k_{T-t+1} = \left(\frac{r}{1+r}\right)\left[\frac{1}{1-\left(\frac{1}{1+r}\right)^{T-t+1}}\right]$$

$$HW_t = E_t \sum_{j=1}^{T-t} (1+r)^j \cdot Y_{t+j} \ .$$

This smoothing solution is what is commonly referred to as the life cycle or permanent income hypothesis and is the consumption function that is routinely used in the literature (Flavin (1981), Hall and Mishkin (1982)). Consumption is proportional to the expected present value of lifetime resources, which consist of human plus non-human wealth. The constant of proportionality k_{T-t+1} is equal to the annual payment on a T-t+1 period $1 annuity, i.e., it is the equal amount that one could receive in each period left of life by putting up one dollar today. When $r = 0$, k is simply equal to 1/(T-t+1), i.e., the inverse of the number of periods left in life. HW is human wealth--the present discounted value of expected future labor income. This solution implies $C_t = E_t C_{t+j}$ for all $j \geq 0$, i.e., expected consumption is flat over one's lifetime. For any concave utility function, the above solution is correct in a certainty model.[8] However, this is not in general the correct solution to a maximizing model when income is stochastic. The necessary conditions for the above to be the general solution under stochastic income are:

(1) The period utility function $U(C_t)$ is quadratic;

(2) C_t is allowed to range from $-\infty$ to $+\infty$.

If these conditions hold, then consumption is identical to what it would be with no uncertainty. This is known as the certainty equivalence (CEQ) solution. I define C_{CEQ} as the optimal consumption (function) under certainty equivalence, the solution represented by (6).

Optimal time t consumption under certainty or certainty equivalence is plotted in figure 2.1 as a function of financial assets. The slope of the curve (dC_t/dW_t) is equal to k_{T-t+1}; if we give this individual one extra dollar of financial wealth, k_{T-t+1} of it would be spent today.

When riskless borrowing is freely allowed at rate r, any future income that will be received with certainty (i.e., a floor of future income) can, without loss of generality, be redefined as part of financial assets instead of income. Therefore, unless otherwise stated, W_t should be thought of as the "certain" component of lifetime resources, including the present value of the floor of future income.

Under certainty equivalence, the sensitivity of consumption to current income (dC_t/dY_t) is:

$$dC_{CEQ,t}/dY_t = k_{T-t+1}\Big[\sum_{j=0}^{T-t}(\partial E_t Y_{t+j}/\partial Y_t)(1/1+r)^j\Big] \tag{7}$$

As shown by Flavin (1981), under certainty equivalence the marginal propensity to consume (MPC) out of current income depends crucially on the time series properties of income and the extent to which current income signals changes in expected future income. For the case of i.i.d. income, $\partial E_t Y_{t+j}/\partial Y_t = 0$ for $j > 0$, and therefore the MPC out of current income is simply equal to k_{T-t+1}. In this case, consumers respond in the same way to an extra dollar of income as to an extra dollar of wealth.

As mentioned above, deriving this solution with income uncertainty requires some very restrictive assumptions: the utility function must be quadratic and consumption must be allowed to be negative. In the following sections, I examine what happens to optimal consumption when we relax each of these assumptions. We will see that in many circumstances the consumption functions look very different from those implied by certainty equivalence.

III. The Theoretical Literature

A. Non-quadratic Preferences

I will investigate two sources of deviations from CEQ: allowing for non-quadratic utility, and imposing the condition that consumption be non-negative. I begin with the former and examine optimal consumption when the third derivative of the utility function is positive.

The first papers on this subject were by Leland (1968), Sandmo (1970), and Dreze and Modigliani (1972). They used a two period model, and their results imply that with time separable preferences a necessary and sufficient condition for optimal first period consumption to be less with income uncertainty than without (and income set equal to its mean) is that the third derivative of the utility function be positive (i.e., that marginal utility be convex). The uncertainty about future income generates an additional reason for saving, and Leland called this the "precautionary demand for saving."[9]

Sibley (1975) and Miller (1976) independently extended these results in two ways. They examined a multi-period model (and assumed an i.i.d. income process) and showed that if U''' > 0, increased income uncertainty (in a well-defined sense which includes a Rothschild/Stiglitz (1970) mean preserving spread) leads to lower current consumption. That is, they showed that in a multiperiod model with a utility function with a positive third derivative, the precautionary demand for saving increases with income uncertainty, beginning at any initial level of uncertainty.[10] This means that the consumption function when U''' > 0 lies everywhere below the consumption function in Figure 2.1.

How reasonable is the assumption that U''' is positive? The mere existence of risk aversion (a negative second derivative) obviously need not imply a positive third derivative. Decreasing absolute risk aversion, a commonly accepted assumption, does however imply a positive third derivative.[11]

What can be said about the sensitivity of consumption to wealth or transitory income? Elsewhere (Zeldes, 1984), I use a two period model and a second order Taylor expansion of marginal utility to show that with constant relative risk aversion and starting at zero uncertainty, adding uncertainty makes consumption more sensitive to wealth. In other words, when a small amount of uncertainty is added to a certainty model, the consumption function in Figure 2.1 becomes more steeply sloped, implying a greater sensitivity of consumption to transitory income than under certainty equivalence. It is *not* the case, however, that this is due solely to a positive third derivative. The effects on the level and the slope of the consumption function are generated by two different properties of the utility function. This result of "excess sensitivity" depends on higher derivatives of the utility function than the third. For example, for constant absolute risk aversion utility, the consumption function in Figure 2.1 would be shifted down in a parallel way when uncertainty is added.

A closed form solution for consumption with random, non-traded, labor income has not been derived except for some specific examples. Solutions have been derived for certain income processes for constant absolute risk aversion (exponential) utility (Merton (1971), Cantor (1985), Roëll (1984)).[12] None of these, however, impose the non-negativity condition on consumption. No closed form solutions have been derived for constant relative risk aversion utility. Therefore, in order to examine and quantify the effects of a labor income uncertainty in a multiperiod model, numerical techniques will be used in Section V.

B. Imposing the Non-negativity Condition on Consumption

The second source of departure from CEQ mentioned above is the imposition of a non-negativity condition on consumption, and is intimately related to the treatment of borrowing against future labor income. The CEQ model assumes that people borrow against the expected value of future labor income, and if income turns out to be low ex-post individuals pay by having negative consumption. Since negative consumption has no counterpart in reality, I will investigate here what happens if we do not let it occur in our models. There are two possible approaches to imposing the non-negativity condition on consumption. First, one could impose the constraint that all borrowing must be paid back with certainty, allowing people to borrow against the certain part of future income, but effectively disallowing risky borrowing. Second, one could allow for risky borrowing with carefully specified default clauses. Because of the possibility of not being able to pay back a loan, borrowing against the uncertain part of future income is equivalent to a short position not in a riskless asset, but in an asset whose payoff is contingent on the outcome of income in subsequent periods. Risky borrowing in this case is a negative position in an asset with a stochastic rate of return. As long as the non-negativity condition on consumption is imposed, it is inconsistent to assume that income is stochastic and borrowing against the expected value of future income is allowed, but that returns on assets are non-stochastic. In this paper, I will take the former approach of disallowing risky borrowing. It should be pointed out, however, that while loan contracts with default clauses will satisfy an expected value budget constraint, they will not lead to results equivalent to CEQ.[13]

One of the first results on this topic was due to Yaari (1976). In his model, $r = \delta = 0$, income is i.i.d., the utility function must be such that utility is defined over all consumption values between $-\infty$ and $+\infty$, and consumption is not restricted to be non-negative.[14] In his model, the fact that income is independently distributed over time enables people to diversify away, across time, the risk inherent in income. By using the law of large numbers, he shows that as the horizon is extended, optimal consumption approaches the mean of income--the same solution as in the certainty case with income equal to the mean (CEQ). Yaari's result seems to depend crucially on all of the above assumptions.

Schectman (1976) adds the restriction that consumption must be non-negative, and shows that Yaari's result disappears.[15] When this restriction is added (continuing to allow no loans with default clauses), there can be no borrowing at all against the risky part of income. This gives the result that the limit of optimal consumption as the time horizon is extended is strictly less than the mean of income. He also shows that as both the time horizon and wealth go to infinity, optimal consumption approaches mean income. What this says is that when the non-negativity condition is imposed on consumption, (1) uncertainty about future income, even in the infinite horizon i.i.d. case, reduces consumption below what it would be under certainty, and (2) as wealth goes to infinity this effect disappears, i.e., as wealth rises, consumption approaches what it would be under certainty. This result holds up (Bewley (1977)) even when income is allowed to be a general stochastic process.

IV. The Empirical Puzzles

Three empirical consumption puzzles have arisen in the literature recently.[16] The first of these relates to the sensitivity of consumption to innovations in income. Hall and Mishkin (1982) found that food consumption responded too much to transitory changes in income relative to permanent changes in income. Their estimate of the ratio of the MPC out of transitory changes in income to the MPC out of permanent changes in income was .29. They compared this to a theoretical value (given their income processes) based on certainty equivalence of about .1.

The second puzzle concerns the expected growth of consumption over time. Under any certainty model of consumption, if the interest rate is less than the rate of time preference the growth in consumption must be negative.[17] Studies have pointed out that there have been long periods of time in which consumption growth has been positive, despite real interest rates that were very low (close to zero) and rates of time preference that were assumed to be positive.[18]

The third puzzle relates to the savings behavior of the elderly. Mirer (1979), Danziger, et al (1983), and others have found that the elderly dissave too little relative to a benchmark with certainty and no bequest motive.

Rather than assuming certainty or certainty equivalence as has been done for documenting each of these puzzles, I will look at models with random labor income and constant relative risk aversion utility. The results of Section V may help explain each of these puzzles.

V. Numerical Results on the Deviations from Certainty Equivalence

In this section, I calculate the optimal consumption for an individual who lives for a large number of periods and faces a stochastic income stream. The goal is to determine the empirical importance of possible deviations from certainty equivalence. In order to do so, it is necessary to calibrate the model to roughly match the income processes that we observe in the world.

A. Calibrating the Model

A number of models for individual earnings have appeared in the literature. Hall and Mishkin (1982), assume that income can be decomposed into the sum of two separate components, one of which follows a random walk (the "permanent" component) and the other of which follows an MA(2) (the "transitory" component), and that the shocks to these components are separately observable. MaCurdy (1982) assumes that the log of earnings follows an IMA(1,2), and thus that there is only one type of shock to income. In both cases, the conditional variance of the j step forecast of income grows with j. Hall and Mishkin's formulation is convenient for the purposes of this paper because it allows us to distinguish between the effects on consumption of transitory and permanent disturbances.

I use two sets of examples in this section. The first is a multiplicative version of Hall and Mishkin's process. YL_t is the "lifetime" component of income. It follows a geometric random walk and is hit each period by the iid shock EL_t. EL_t is meant to capture the effects of raises, job changes, health changes, and other persistent factors. Total labor income

Y_t is equal to this component times an i.i.d. shock ES_t. ES_t is meant to capture the effects of one time bonuses, unemployment spells, and other transitory factors. This gives:

$$Y_t = YL_t \cdot ES_t$$
$$YL_t = YL_{t-1} \cdot EL_t \quad (8)$$

ES and EL are i.i.d., have mean equal to one, and are assumed to be <u>separately observable</u>. Expected income in period t+j conditional on time t information is just equal to YL_t. I calibrate the income process so that the j step ahead coefficient of variation is roughly in line with that implied by both Hall and Mishkin and MaCurdy. To keep things simple, I use the following three point distributions for ES and EL (I call this distribution #2A):

EL	outcome	prob	ES	outcome	prob
	.9	.25		0	.05
	1.0	.50		1.0	.45
	1.1	.25		1.1	.50

In any period, the permanent component of income can rise or fall by ten percent (this corresponds to raises or permanent pay cuts). Also there is a 5 percent chance that income will equal zero in a given period (temporary unemployment), a 45 percent chance that income will equal the permanent component, and a 50 percent chance that income will be 10 percent higher than the permanent component (a transitory bonus).[19] Figure 5.1 plots the coefficient of variations for the j year ahead forecasts of the level of earnings for this process against those implied by the estimates in Hall and Mishkin and MaCurdy. The coefficients of variation lie everywhere below MaCurdy's and generally below Hall and Mishkin's and thus provide a conservative estimate of the income uncertainty facing households.

In the second set of examples, I assume income is i.i.d. (EL ≡ 1) and choose the income process so that the coefficient of variation (= .55) is roughly that of an average horizon with either the Hall and Mishkin or MaCurdy process.[20]

Distribution #1A:

ES	prob
0	.15
1.0	.70
2.0	.15

There are two important differences between the two income processes that I use. First, for the non i.i.d. process, a change in current income changes expectations about the level of future income, and this is clearly not the case with i.i.d. income. Second, the variance of the j step ahead forecast of income grows with j in the non i.i.d. case I examine, whereas it is constant when income is i.i.d. In both cases, however, future income is uncertain, and in both cases we can examine the sensitivity of consumption to one time changes in wealth or transitory income. The results for the two income processes in fact turn out to be quite similar.

The method used to calculate the optimal consumption is stochastic dynamic programming.[21] The technique for the case of i.i.d. income is described in detail in Barsky, Mankiw, and Zeldes (1986). The problem is formulated as a one state (wealth), one control (consumption), one disturbance (income) stochastic control problem. The state space is discretized into an S element grid using a technique suggested in Bertsekas (1976).[22] Beginning in the last period, backwards induction is used to solve for the value function and corresponding optimal consumption. The computer was programmed to search, for each state and time, over all feasible levels of consumption and choose the one that maximizes the sum of current utility and the discounted expected value of next period's value function. The results are two SxT matrices which give optimal consumption and the value function at all possible levels of wealth and times left to live.

The procedure is similar for the more complicated income process except that another state variable--the lagged value of the random walk component of income--has to be added, yielding two three-dimensional matrices. This is conceptually straightforward, but unfortunately, due to the "curse of dimensionality," it requires dramatically more computer memory and CPU time, such that only relatively short horizons could be tested.

While the numbers calculated with this technique are an approximation to the actual solution, the approximation error can be made arbitrarily small by narrowing the width of the grid used for discretization (Bertsekas, 1976).[23] In part because of the high cost of this computer search,[24] I restrict my attention to relatively simple problems.

B. Results with Constant Relative Risk Aversion Utility and No Constraints

In what follows, I compare the consumption function with constant relative risk aversion (CRRA) utility to the consumption function with quadratic utility (CEQ). Thus $U(C) = \frac{C^{1-A}}{1 - A}$ were A is the coefficient of relative risk aversion. For simplicity, I assume $r = \delta = 0$. Since $r = \delta$, the consumption function with quadratic utility (C_{CEQ}) is identical to the consumption function with CRRA utility <u>under certainty</u>. Thus these results can also be thought of as comparing consumption under uncertainty to consumption under certainty (both with CRRA utility). I examine the level of consumption, the sensitivity of consumption to permanent and to transitory changes in income, and the expected path of consumption. I then investigate how these are affected by financial wealth, the time left to live, the variance of the income process, and the coefficient of relative risk aversion.

The first example uses distribution #2A, the combination random walk/i.i.d. process. The time horizon is 15 periods, and the coefficient of relative risk aversion is set equal to three.[25] Consumption as a function of initial financial assets is plotted in figure 5.2. Both initial income and expected income in all future periods are equal to 100. The results, especially at low levels of assets, are quite striking. First, notice that the slope of the curve, which is equal to the marginal propensity to consume out of transitory changes in income, is considerably larger than that predicted by certainty equivalence. For example, a household with two years worth of expected income in the form of assets would have an MPC of over twice what would be predicted by a certainty or certainty equivalence model. A family with one years worth of expected income as assets would have an MPC 7 times as great as under CEQ. This fully optimizing unconstrained household exhibits dramatic "excess sensitivity" relative to the certainty equivalence benchmark, suggesting that this is not a terribly good benchmark.

For a household with very high assets relative to expected future income (e.g., four to five years worth) the MPC is almost the same as that implied by the simple smoothing model.[26] This means that we should expect MPCs out of transitory changes in income to be much larger for households with low current (certain) assets relative to expected future (uncertain) income, than for the rest of the population. Note that this does not necessarily correspond to "poor" vs. "wealthy" households, because what is relevant is not the absolute amount of current wealth, but current assets <u>relative</u> to expected future (uncertain) labor income.

Next, consider the sensitivity of consumption to changes in expected future income. Figure 5.3 is a plot of optimal time t consumption as a function of YL_t, the random walk component of income, holding current assets constant at 200. A one unit increase in YL_t increases expected income in every future period by one unit. The slope of the consumption line under CEQ is (T-1)/T.[27] In fact, however, the sensitivity of consumption to expected future income is only about 60% as large as the CEQ benchmark. If instead initial wealth is equal to 400 (not shown), the sensitivity of consumption to expected future income is much closer to the CEQ benchmark.

Overall, what we are seeing is quite interesting. Relative to the certainty equivalence benchmark, individuals optimally "overrespond" to changes in current income or wealth and "underrespond" to changes in expected future income. This is especially true when current assets are relatively low. In these cases, current assets play a much more important role in the consumption function relative to risky future labor income than would be predicted under CEQ.

Finally, I examine the level of precautionary saving defined as the difference between optimal consumption and what it would be in the absence of uncertainty (C_{CEQ}). At a wealth of 200, precautionary saving is 20% of optimal consumption. That is, if there were no income uncertainty consumption would be 20% higher than it is with the uncertainty. At low levels of assets, the certainty equivalent benchmark dramatically overstates the optimal level of consumption. At a wealth of 500, precautionary saving is still about 7% of the optimal consumption level. This suggests that significant fraction of the capital accumulation that occurs in the U.S. may be due to precautionary savings.

Figure 5.4 is a plot of the expected growth of consumption between this period and the next as a function of initial assets.[28] Recall that since $r = \delta$, under CEQ expected future consumption would be equal to current consumption, i.e., the growth in expected consumption would equal zero. What we see here, however, is that the expected growth of consumption is considerably greater than zero.[29] At a wealth of 100, consumption is expected to grow by 25% between the first and second period. At a wealth of 200, this number drops down considerably to 2.5%.

Figure 5.5 is a plot of the annualized expected growth of consumption between the current period and the last period of life. At an initial wealth of 200, expected consumption grows at an average rate of 4% per year, between the current period and the end of life.[30] Even at an initial wealth of 500, this number only drops to 2½% per year.

In order to be able to experiment with longer horizons, the i.i.d. income model is used next. Figure 5.6 shows the optimal consumption for a fifteen period horizon compared to the case with a random walk/i.i.d. income process. The forecast variance is roughly that of a 30 year horizon for the non-i.i.d. examples. The optimal consumption profile is similar in the two cases.

To examine whether the time horizon is important to the results, Figure 5.7 plots the consumption function when the time horizon is extended to 30 and 60 periods, and Figure 5.8 plots the percentage deviation from CEQ. It is interesting to note that for a given asset level, the percentage difference of optimal consumption from the corresponding CEQ amount is roughly the same, independent of the time horizon. Increasing the time horizon does not diminish the effects of income uncertainty on the level of consumption.

Next, I experiment with different coefficients of relative risk aversions (A). Figure 5.9 compares the consumption function for individuals with log utility (A = 1), to those of more risk averse individuals (A = 3, A = 6). More risk averse individuals have a higher amount of precautionary saving. (At a wealth of 300, precautionary saving as a fraction of C_{CEQ} is 10%, 17%, and 29% respectively.) They do not necessarily have higher MPCs out of transitory income, however. The rankings for the MPCs depends on the level of initial wealth.

Finally, a mean preserving spread is performed on the income distribution. (See figure 5.10.) The higher variance income process (#1B) is:

Y =	0	100	200
prob =	.25	.50	.25

This raises the coefficient of variation on income to .71 from .55. A more risky income stream leads to lower consumption levels (Miller (1976), Sibley (1975)). At a wealth 300, precautionary saving is 25 percent of C_{CEQ}, as opposed to 17 percent for the less risky income stream.

C. Results from Imposing the Non-Negativity Condition on Consumption

In this section, I examine the optimal consumption for an individual with quadratic utility whose consumption is restricted to be non-negative in all periods. This effectively rules out borrowing against the risky part of future labor income. I restrict myself to the case of i.i.d. income for simplicity.[31] The consumption function for a thirty period example is plotted in Figure 5.11.

Since it turns out that this is equivalent to imposing a weak form of a borrowing constraint, I will leave most of the discussion for section VI. What we observe, however, is that even if the consumer has a quadratic utility function, when the reasonable assumption that consumption be non-negative is imposed the consumption function for individuals at low wealth levels differs considerably from the certainty equivalence benchmark.[32] The level of consumption is lower and the sensitivity of consumption to wealth is higher, than the certainty equivalence benchmark. At relatively high levels of wealth (in this example, more than four years' worth of expected income), however, the optimal consumption function is very well approximated by the certainty equivalence benchmark.

D. Implications for Empirical Consumption Functions

The standard consumption function posits a linear relationship between consumption and "permanent income," defined as the annuity value of the sum of non-human wealth and the present discounted value of expected future labor income. The results here indicate that such a consumption function is likely to be severely misspecified, especially at low levels of wealth.

One possible remedy to this problem would be to put a weight of less than one on human wealth before adding it to non-human wealth, or to discount expected future income at a higher discount rate. This would be an improvement, because it would shift the certainty equivalence line downwards in a parallel way. However, it would still miss important aspects of the consumption function, because the weight that should be placed on human wealth is not constant, but depends on the ratio of human wealth to non-human wealth. This suggests the following as a better approximation:

$$C = k \cdot \left[NHW + x(NHW/HW) \cdot HW\right]$$

where x is the weight ($0 \leq x \leq 1$) placed on human wealth, and is an increasing function of the fraction of lifetime resources in non-human wealth. This will help capture the curvature in the consumption function seen in the preceding figures. While the resulting function is still not correct, it gives a significantly better approximation to the true optimal consumption function.[33]

E. Relationship to Consumption Puzzles

Overall, when we impose the non-negativity condition on consumption, and especially when we examine constant relative risk aversion utility, we see fairly dramatic deviations of optimal consumption from what is implied by a simple smoothing or certainty model. This is especially noticeable when current assets are low relative to expected future income.

Consider now how this relates back to the consumption puzzles described in section IV. Using a similar income process to that assumed by Hall and Mishkin, an optimizing model with constant relative risk aversion utility yields a consumption function that exhibits "excess" sensitivity as defined by Hall and Mishkin. They find that the ratio of the MPCs out of transitory and permanent income was .29 instead of a CEQ benchmark of .10 (about 2.9 times as great). In the 15 period model examined here, the CEQ benchmark is 1/15 = .067, and the optimizing consumer with a wealth of two years expected income in fact had a ratio of .16, about 2.4 as great as under CEQ. While this is by no means conclusive evidence, it at least suggests that their results may have been consistent with optimizing behavior.

The second puzzle related to the expected growth in consumption. In the examples presented here, the rate of interest and the rate of time preference were both equal to zero. In a certainty model this would imply that the growth of consumption must be equal to zero in all periods. When uncertainty is explicitly modeled, however, the expected growth in consumption for an individual with two years worth of income in the form of wealth and a fifteen year horizon was found to be $2\frac{1}{2}$ percent over the first year and 4 percent per year over the entire horizon. Again, this suggests that the high growth in consumption observed simultaneously with low real interest rates may still be consistent with optimizing behavior.

Finally, consider the puzzle relating to the saving behavior of the elderly. While I have not tried to calibrate the model for this case explicitly, we can model uninsured medical expenses or other emergency expenses that the elderly must sometimes occur as equivalent to negative draws of income. The possibility of having to incur severe costs for catastrophic illness is likely to be an important source of uncertainty for the elderly. The results of this section show that the level of precautionary saving in the face of plausible amounts of income uncertainty is quite high. This suggests that the same may be true for the elderly facing uncertainty about expenses. More work clearly needs to be done calibrating such a model, but again the results here suggest that the high savings by the elderly may be consistent with optimizing behavior under uncertainty.

VI. Borrowing Constraints

One of the sources of deviation from CEQ investigated in previous sections was the imposition of the non-negativity condition on consumption. If individuals cannot have negative consumption and cannot default on loans, then this implies that they cannot borrow against the risky part of future income. This is nothing other than a weak form of a borrowing constraint. The same machinery that has already been developed can be used to analyze the effects of stronger borrowing constraints on consumption in a multiperiod model with random labor income. I begin by discussing some theoretical issues concerning borrowing constraints, and then proceed to the numerical results.

Most of the theoretical work on borrowing constraints has been done either in the context of a two period model, or a multiperiod model under certainty (e.g., Mariger (1986), Heller and Starr (1979)). The advantage of looking at multiperiod models is that we can look at the separate effects of present and future constraints. The advantage of examining an uncertainty model is that we can look at what happens when future constraints are binding only in certain states of the world. To date, a closed form solution has not been derived for a multiperiod model with income uncertainty. I therefore use the numerical techniques described earlier to learn more about the optimal consumption function. In what follows, I assume that there is a floor on the distribution of future income, that all borrowing and lending is riskless, and that there is a limit on the amount of net indebtedness. Define A_t as end of period wealth = $(W_t - C_t)$. Thus, the individual faces the optimization problem (1) - (4), subject to the additional set of constraints:

$$A_{t+j} \geq -B_{t+j} \qquad \forall\ j \geq 0 \tag{9}$$

where B_t are the limits on net indebtedness and are assumed to be exogenous.[34] Constraints (9) can be rewritten as $C_{t+j} \leq W_{t+j} + B_{t+j} \equiv X_{t+j}$. X_t is the maximum resources available for consumption in period t, and therefore I sometimes refer to X_t as wealth in what follows. In any period t, the current constraint is defined to be binding if removing the time t constraint (and leaving future constraints unchanged) alters time t consumption. We will see that the existence of borrowing constraints can alter current consumption even if the current constraint is not actually binding, because of the possibility of future constraints being binding.

The weakest form of borrowing constraint considered here is

$$B_t = \sum_{i=1}^{T-t} \left(\frac{1}{1+r}\right)^i \cdot YF_{t,t+i} \qquad \forall t . \tag{10}$$

where $YF_{t,t+i}$ equals the minimum possible level of income in t+i given information available in time t.[35] In this case, the consumer can borrow against any future income that will be received with certainty, but no more.

The strongest form of borrowing constraint considered is $B_t = 0 \; \forall \; t$. In this case, the individual is prohibited from borrowing, on net, against any future income. In other words, net financial assets can never be negative.

In the following analysis, I assume that utility is quadratic and compare the CEQ solution (no restriction on B and consumption allowed to be negative), to the solution when a borrowing constraint is imposed. Let C_{CEQ} represent the optimal consumption function of the unconstrained problem, and C represents the optimal consumption function of the constrained problem.[36] I again assume for simplicity that $r = \delta = 0$ and that income is i.i.d.

Before proceeding to the numerical results, I briefly describe some qualitative characteristics of the solution to the multiperiod liquidity constraint problem. Consider the function C(X) plotted against X, holding B constant. I divide spendable resources X into four regions as in figure 6.1. In the lowest wealth region (1), the current constraint is binding and $C(X) = X$. The MPC out of additional unit of financial wealth equals one.

In both the second and third regions, the current constraint is not binding ($C(X) < X$), but consumption is nevertheless below the unconstrained optimum ($C(X) < C_{CEQ}(X)$) and the MPC out of wealth is greater than in a world with no constraints ($MPC > MPC_{CEQ}$).[37] This is due to consumers' fears of hitting future constraints. In region 2, $C(X) < X < C_{CEQ}(X)$, so that if there were no future constraints, the current constraint would be binding and C(X) would equal X. In region 3, $C(X) < C_{CEQ}(X) < X$, so the future constraints lower consumption even though C_{CEQ} is a feasible consumption choice. The current constraint does not bind, and would not bind even if all future constraints were removed.

The highest wealth region is defined such that there are no possible chains of future events that will lead to consumption being constrained. Wealth is high enough such that even if the worst realization occurs for income in each period, consumption can equal what it would in a totally unconstrained world in all periods. In this region, clearly $C(X) = C_{CEQ}(X)$.

In a two period model, regions (2) and (3) do not exist, for in such a model either the current constraint is binding or the consumer is totally unconstrained. In a multiperiod world, however, regions (2) and (3) do exist, and their existence depends only on risk aversion. Fear of hitting future constraints forces consumption to be lower and the sensitivity of consumption to wealth to be higher than in a world without constraints, even though the current constraint is not binding. The inability to borrow in the future can alter consumption today even if there is no desire to borrow today and even if there is only a small probability that the borrowing constraint will be binding sometime in the future.[38] People, knowing that they will not be able to borrow in the future and that income might come in low before then, choose to consume less than they otherwise would, holding more in reserve in the form of wealth. It should be pointed out that here, as opposed to in the results on nonquadratic preferences, low consumption and high sensitivity are linked together and caused by a common factor.

The Euler equation between t and t + 1 is not satisfied in region (1), but is satisfied in regions (2), (3), and (4), despite the fact that consumption is still globally constrained in regions (2) and (3). In other words, if the Euler equation is satisfied between t and t + 1 (i.e. we are not in region (1)), this does not imply that consumption is as it would be in a world with no constraints (i.e. it does not imply that we are in region (4)). This is discussed in more detail in Zeldes (1985).

Numerical Results with Borrowing Constraints

In order to learn more about how large each of the above regions are and exactly how much future constraints influence current consumption, I use the numerical methods described above to calculate the optimal consumption function. To keep the analysis tractable, I restrict myself to the case of i.i.d. income. Income follows the following simple three point distribution (#1C):

Y =	10	110	210
prob =	.15	.70	.15

This process has the same variance as that used in the previous section (#1A), but a certain income of 10 per period has been added, in order to be able to examine strong borrowing constraints. The minimum value that income can take on in any period (the floor on income) is 10. The utility function is quadratic, so that the CEQ solution holds in the absence of constraints.[39] Again, $r = \delta = 0$. Both a strong and weak form of borrowing constraint are examined. The strong form allows no borrowing, and the weak form allows borrowing against the sum of the floors of future income. The time horizon is 30 periods and the results are presented in figure 6.2.

The four regions described above are indicated in figure 6.2 for the strong borrowing constraint. Region one runs from an asset level of zero to an asset level of 74 (2/3 of a year of expected income). In this region, all of the financial assets are consumed, implying an MPC of one. Between wealth levels of 74 and 110 (region 2), consumption is less than wealth, i.e., the borrowing constraint is not binding even though optimal consumption under CEQ is greater than wealth. For example, at a wealth of 100, CEQ would yield consumption equal to 109 2/3, but optimal consumption (equal to 82) is in fact only 75% of that. The fear of not being able to borrow in the future against future risky income reduces consumption by 25%. The MPC in this region equals .27, about 8 times as high as under CEQ (.033).

In the third region, consumption is less than CEQ even though the CEQ level is feasible. For example, at a wealth of 165 (1.5 periods of expected income), optimal consumption is only 96, about 16 units or 14% less than C_{CEQ} and the MPC is about 4 times as high as CEQ. Again, the lower consumption and higher MPC is due to the probability of hitting a no borrowing constraint some time in the future, even though the constraint is not currently binding. As we increase the level of wealth, consumption approaches CEQ. At a wealth of 400, optimal consumption is only 6 units or 5% below CEQ and consumption is 1.5 times as sensitive to income. At a wealth of 770, the level and MPC of consumption are indistinguishable from what they would be under CEQ. Thus, at a wealth of 7 periods' worth of expected income, consumption is approximately what it would be in region 4, even though region 4 does not actually begin until a wealth level of about 9000. In other words, after a wealth of 9000, the CEQ and C lines are the same, but they are very close to each other for a large region before that.

These results show us the importance of future constraints on current consumption. In regions 2 and 3, the MPC is less than one, but still quite high. Thus, even though the Euler equation is satisfied between this period and the next in this region, consumption is still dramatically less than in a world with constraints at all, and the MPC is dramatically higher. The mere possibility that future constraints may bind can produce effects very similar to those produced by binding current constraints.

Comparing consumption with strong constraint to that with a weak constraint, we see that strengthening the borrowing constraint decreases the level of current consumption and increases the sensitivity of consumption to current assets. This is true not just at asset levels where the new constraint is binding today, but also at asset levels much higher than this. Even when the constraint is not binding today, a stronger borrowing constraint increases the probability that an individual will hit the constraint in the future, and this leads to lower current consumption and a greater sensitivity of consumption to current income. For example, for an individual with two periods' worth of expected income stored in current assets, strengthening the liquidity constraint from the weak to the strongest form reduces consumption by about 8%, and increases the sensitivity of consumption by about 100%. The result would be stronger for income processes with a larger certain component or for utility functions with more risk aversion.

Figure 6.3 shows the effects of a higher variance of income (assuming a strong constraint).[40] A more risky income stream raises the probability of hitting future constraints. The inability to borrow in the future has a greater dampening effect on current consumption when the variance of income is higher. Since future constraints are relatively more important in this case, region 1 (current constraint binding) is smaller and region 2 is larger.

Consider next the expected one period growth in consumption (not shown). For asset levels where the current constraint is binding (region 1) the expected growth in consumption is very high. For all other regions, since the constraint is not binding between today and tomorrow, the expected growth in consumption between today and tomorrow is equal to zero.

Figure 6.4 shows the annualized growth in expected consumption between the current and final period of life. Here, expected consumption growth is higher in regions 2 and 3 as well. Looking at the case of the strong constraint, if initial assets are 110, the current constraint is not binding, but the possibility of hitting future constraints causes annualized expected consumption growth to be about 2% per year, versus 0% under CEQ.

VII. Conclusions

The simple life cycle/permanent income model described in most macroeconomic texts says that individuals base their current consumption on the sum of their financial assets and the expected discounted value of their future labor income, in such a way that consumption is expected to be constant over their lifetimes. This result can easily be derived in a certainty model, but when future labor income is random this becomes much more difficult. In order to strictly get this result, it is necessary to assume that the utility function is quadratic and that consumption can be negative. When the restriction is added that consumption be non-negative, or when the utility function is instead assumed to be of the standard constant relative risk aversion form, the optimal consumption function looks quite different from the common benchmark described above. This is especially true when the "certain" component of lifetime resources is small relative to the risky components of lifetime resources, i.e., when financial assets are small relative to human capital. In such circumstances, the level of precautionary saving calculated for individuals is large, suggesting that precautionary saving may represent a significant fraction of the total savings of U.S. households. In addition, the results show that current assets (which include the income just received) and non-stochastic future receipts are optimally given much more weight than future random labor income in making the current consumption decision. This cannot be compensated for merely by discounting future labor income at a higher rate, because the appropriate rate would vary with the level of initial assets.

The numerical results indicate that rational, unconstrained individuals with constant relative risk aversion utility will optimally exhibit "excess" sensitivity to transitory income, save "too" much, and have expected growths of consumption that are "too" high, relative to the simple permanent income hypothesis benchmark. This suggests that we should rethink our presumption that the certainty equivalent model is the appropriate benchmark, especially at low levels of financial wealth. These results have the potential to at least partially explain three important empirical puzzles in the consumption literature: the excess sensitivity of consumption to transitory changes in income, the high growth in consumption with low real interest rates, and the high savings rate of the elderly.

The results on liquidity constraints show that the fear of hitting future constraints can significantly influence current consumption behavior. The level of consumption can be significantly lower and the sensitivity to income higher than in an unconstrained world, even if the current constraint is not actually binding and the future constraints are binding only in certain states of the world.

FOOTNOTES

[1]When all assets are traded, closed form solutions for consumption can be derived. See, for example, Merton (1971).

[2]In this paper I investigate arbitrarily long, but finite and known, horizons. The techniques used could be modified, however, to incorporate infinite or random horizons.

[3]Labor supply is taken as exogenous in this paper. If instead workers faced a given wage and were allowed to choose their labor supply, this would provide another channel for hedging wage uncertainty. The effects of wage uncertainty on consumption would be very similar to, but somewhat less dramatic than, those presented here.

[4]It will be helpful for later discussions if, without loss of generality, I divide income into certain and uncertain components. The certain component is equal to the minimum that it is possible to receive in a given period, i.e., any positive floor on income.

[5]In a model with many assets with uncertain rates of return, the problem would look quite similar except the consumer would also choose portfolio shares. In that case, r_t would be defined as the arithmetic weighted average of the realized return on the financial (non-human capital) portfolio.

[6]In a model with uncertain rates of return, the future value calculations in the budget constraint would be based on ex-post (realized) rates of return on the individual's portfolio.

[7]The proportionality hypothesis which underlies the simple life cycle or permanent income model says that in any period consumption is proportional to the present value of lifetime resources, with the constant of proportionality independent of lifetime resources. Yaari (1964) showed for a certainty model, that if the proportionality hypothesis holds (for all r and δ), then the utility function must be a constant relative risk aversion utility function. When $r = \delta$, the proportionality hypothesis holds under certainty for any concave utility function. For this reason, I focus on the $r = \delta$ case throughout this paper.

[8]In a certainty model, the expectations operator is obviously inoperative.

[9]To see the intuition behind this, imagine beginning in a certainty situation with endowment pattern (Y1e, Y2e), such that the optimal decision involves equalizing marginal utility across periods. Now add uncertainty into second period income, keeping the mean at Y2e. If first period consumption were to remain unchanged, then the convexity of marginal utility would imply $E[U'(C2)] > U'[E(C2)] = U'(C1)[(1 + \delta)/(1 + r)]$. Therefore, this cannot be an optimum. At the new optimum, first period marginal utility must be higher, and therefore first period consumption must be lower, than under certainty.

[10]Grossman, Levhari, and Mirman (1979) showed that the value function inherits the positive third derivative property of the period utility function.

[11]Recall that decreasing absolute risk aversion means that the dollar amount that an individual would be willing to pay to avoid a fair gamble of a given size decreases as the level of consumption rises.

[12]Merton (1971) uses a poisson income process that can only jump up.

[13]In fact, I have shown for certain two period models with quadratic utility that imposing the non-negativity condition by replacing riskless with risky borrowing (with defaults) leads to current consumption greater than under CEQ.

[14]Note that utility is not defined over negative consumption for constant relative risk aversion utility and therefore Yaari's proof does not apply to that case.

[15]Schectman assumes $U'(0) <$ infinity. See also Schectman and Escudero (1977).

[16]See, for example, Dornbusch and Fischer (1984).

[17]This assumes time separable utility.

[18]See, for example, Singleton (1985).

[19]As discussed previously, the present value of any positive floor of future income should be included as part of financial assets, leaving future income with a floor of zero. As discussed in Barsky, Mankiw and Zeldes (1986), however, it may also be reasonable to assume a minimum survival level of consumption (C_S) and to define utility as a function of consumption in excess of that level (e.g., $U = (C - C_S)^{1-A}/(1 - A)$. In this case, the results are the same as in the text, with W interpreted as including only tangible non-human wealth and not including the present value of the floor of future income.

[20]The reason for doing the set with i.i.d. income is that I can experiment with longer horizons without bumping into memory constraints on the computer.

[21]I thank Andy Abel for helpful suggestions about the technique used here.

[22]Barsky, Mankiw, and Zeldes (1986) use the same basic technique, but assume i.i.d. income. They focus on the relevance of Ricardian equivalence, whereas the focus here is on the optimal consumption function. A more complicated technique has to be used here, in order to examine the effects of both permanent and transitory income.

[23]The grid solutions are very close to the analytical solutions for some simple three period problems that I checked.

[24]The problems with i.i.d. income were calculated at relatively low cost on a VAX 8600. However, the fifteen period example with non i.i.d. income involves considerable memory requirements and CPU time. This example requires creating 2 matrices with about 625,000 elements each. The optimal consumption (and value function) then had to be determined for each of the 625,000 possible nodes, for each of the fifteen periods.

[25]Estimates of A cover a wide range. They include Friend and Blume (1975) (2), Mankiw (1981, 1985) (4, 3 respectively), Hansen and Singleton (1983) (1), Shapiro (1983) (.5), Mankiw, Rotemberg and Summers (1982) (.5), and Hall (1981) (15).

[26]Recall that what is referred to as assets here is equal to the sum of financial wealth and the present discounted value of the floor of future income. See footnote 19.

[27]The change in expected lifetime income equals T-1 (one for each period left to receive income), and this is divided by T, the number of periods left to consume.

[28]The jagged nature of the curve is due to approximation errors of the numerical method. The surrounding area was shaded in to highlight this.

[29]One can see that that must be positive by looking at a Taylor expansion of the standard Euler equation (see Zeldes, 1985).

[30]That is, $\left(\frac{E_1 C_T}{C_1}\right)^{1/(T-1)} - 1 = 4\%$.

[31]The quadratic utility function is written as $U = a - bC^2$ or alternatively: $U = 1 - (b/a)C^2$. This is only an increasing function for the range $C < a/2b$. Thus a/b must be chosen large enough such that $a/2b > C_{max}$, where C_{max} is the largest possible value that consumption might take. Unfortunately, the coefficient of relative risk aversion $(= [(a/2bC) - 1]^{-1})$ is an exact declining function of a/2b, so that the choice of C_{max} puts a ceiling on the coefficient of relative risk aversion. It would have been nice for the sake of comparison to choose a/2b such that the coefficient of relative risk aversion at C = 100, for example, was equal to 3, but the utility function would then only be well behaved for C < 133.

The quadratic function used for all of the examples in the paper is $U = 1 - C^2/3000$, implying a coefficient of risk aversion = C/(1500 - C), considerably smaller than 3 for most of the relevant range of C. (For example, it equals .25 for C = 300.) Recall that quadratic utility exhibits both increasing relative and absolute risk aversion.

[32]As discussed above, another way to impose the non-negativity condition would be to allow (risky) borrowing with a payback contingent on the individual's income in a given period. This is beyond the scope of this paper. This approach would give different results than those presented here, but would not bring us back to certainty equivalence.

[33]One might also want to model x as a function of the amount of individual income uncertainty and utility function parameters.

[34]Note that this rules out debt limits that depend on last period's outstanding debt or borrowing. This type of model implies that borrowing and lending are never done simultaneously. Rotemberg (1984) uses a model where debt limits depend on past levels of borrowing and shows that borrowing and lending can occur simultaneously as a hedge against future borrowing constraints.

[35]Note that $YF_{t,t+i}$ can increase with t (given t+i) but not decrease, i.e., the income floor for a certain period in the future can be greater than but never less than what it was thought to be in the past.

[36]Note that for an individual with a utility function with $U'(0) = \infty$ the weak form of liquidity constraint considered above would never be binding, i.e. the consumer would always be at an interior solution and the solutions to the two cases above would be equivalent. However, stronger liquidity constraints would still be binding on such an individual for some initial quantities of wealth.

[37]Some of the intuition behind this result comes from looking at a certainty model. Imagine that a liquidity constraint is going to be binding in some period M between period 1 and the terminal date (T). An extra unit of current (period 1) wealth will therefore be split up over the first M periods, so $MPC = 1/M > 1/T$. (Heller and Starr (1979) and Mariger (1986) show this in a certainty model. Levhari, Mirman, and Zilcha (1980) provide a similar explanation in the context of a different certainty model.) When uncertainty is added, consumption is less today because of a positive probability of a future constraint. For a given probability of future liquidity constraints, $MPC > MPC_{CEQ}$ as in the certainty case. As wealth tends to infinity, the probability of hitting future constraints goes to zero, so this effect disappears and the MPC tends toward MPC_{CEQ}.

[38]I believe that the previous analysis also applies when the future imposition of the borrowing constraint is itself stochastic.

[39]For a description of the utility function, see footnote 31.

[40]The higher variance income process used (#1D) is:

Y =	10	110	210
prob =	.25	.50	.25

REFERENCES

Barsky, Robert, N. Gregory Mankiw, and Stephen P. Zeldes, "Ricardian Consumers with Keynesian Propensities," American Economic Review, forthcoming, September 1986.

Bertsekas, Dimitri, Dynamic Programming and Stochastic Control, Academic Press, New York, 1976.

Bewley, Truman, "The Permanent Income Hypothesis: A Theoretical Formulation," Journal of Economic Theory, 16, 1977, 252-292.

Cantor, Richard, "A Consumption Function Example of a Closed-Form Solution to Stochastic Control Problems with an Exponential Criterion," Economics Letters, 1985.

Danziger, Sheldon, Jacques Van Der Gaag, Eugene Smolensky, and Michael Taussig, "The Life-cycle Hypothesis and the Consumption Behavior of the Elderly," Journal of Post Keynesian Economics, Winter 1982-83, 208-227.

Dornbusch, Rudiger and Stanley Fischer, Macroeconomics, Third Edition, 1984.

Dreze, Jacques H. and Franco Modigliani, "Consumption Decisions under Uncertainty," Journal of Economic Theory, 5, 1972, 308-335.

Flavin, Marjorie A., "The Adjustment of Consumption to Changing Expectations about Future Income," Journal of Political Economy, 89, October 1981, 974-1009.

Friend, Irwin and Marshall Blume, "The Demand for Risky Assets," American Economic Review, December 1975, 900-923.

Grossman, Sanford, David Levhari and Leonard J. Mirman, "Consumption under Uncertainty," in General Equilibrium, Growth, and Trade, Academic Press, 1979.

Hall, Robert E., "Intertemporal Substitution in Consumption," NBER Working Paper No. 720, July 1981.

Hall, Robert E. and Frederic S. Mishkin, "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households," Econometrica, 50, March 1982, 461-481.

Hansen, Lars Peter and Kenneth J. Singleton, "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns," Journal of Political Economy, April 1983, 249-266.

Heller, Walter P., and Ross M. Starr, "Capital Market Imperfection, the Consumption Function, and the Effectiveness of Fiscal Policy," Quarterly Journal of Economics, 82, 1963, 465-473.

Leland, Hayne E., "Saving and Uncertainty: The Precautionary Demand for Saving," Quarterly Journal of Economics, 82, 1968, 465-473.

Levhari, David, Leonard J. Mirman and Itzhak Zilcha, "Capital Accumulation under Uncertainty," International Economic Review, 21, October 1980, 661-671.

MaCurdy, Thomas E. "The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis," Journal of Econometrics, 18, 1982, 83-114.

Mankiw, N. Gregory, "The Permanent Income Hypothesis and the Real Interest Rate," Economic Letters, November 1981, 307-311.

Mankiw, N. Gregory, "Consumer Durables and the Real Interest Rate," Review of Economics and Statistics, August 1985, 353-362.

Mankiw, N. Gregory, Julio J. Rotemberg and Lawrence Summers, "Intertemporal Substitution in Macroeconomics" Quarterly Journal of Economics, February 1985.

Mariger, Randall, Consumption Behavior and the Effects of Government Fiscal Policies, Harvard Press, 1986.

Merton, Robert L., "Optimum Consumption and Portfolio Rules in a Continuous Time Model," Journal of Economic Theory, 3, December 1971, 373-413.

Miller, Bruce L., "The Effect on Optimal Consumption of Increased Uncertainty in Labor Income in the Multiperiod Case," Journal of Economic Theory, 13, 1976, 154-167.

Mirer, Thad, "The Wealth-Age Relation Among the Aged," American Economic Review, June 1979, 435-443.

Roëll, Ailsa, "Precautionary Saving and the Evaluation of a Risky Income Stream: A Simple Example," manuscript, May 1984.

Rothschild, Michael and Joseph E. Stiglitz, "Increasing Risk I: A Definition," Journal of Economic Theory, 1970, 225-243.

Rotemberg, Julio, "Consumption and Liquidity Constraints," manuscript, MIT, February 1984.

Sandmo, Agnar, "The Effect of Uncertainty on Saving Decisions," Review of Economic Studies, 37, 1970, 353-360.

Schectman, Jack, "An Income Fluctuation Problem," Journal of Economic Theory, 12, 1976, 218-241.

Schectman, Jack, and Vera L. S. Escudero, "Some Results on 'An Income Fluctuation Problem,'" Journal of Economic Theory, 16, 1977, 151-166.

Shapiro, Matthew D., "The Permanent Income Hypothesis and the Real Interest Rate: Some Evidence from Panel Data," Economic Letters, 14, 1984, 93-100.

Sibley, David S., "Permanent and Transitory Income Effects in a Model of Optimal Consumption with Wage Income Uncertainty," Journal of Economic Theory, 11, 1975, 68-82.

Singleton, Kenneth J., "Specification and Estimation of Intertemporal Asset Pricing Models," manuscript, Carnegie-Mellon University, August 1985.

Yaari, Menahem E., "On the Consumer's Lifetime Allocation Process," International Economic Review, 5, September 1964, 304-317.

__________, "A Law of Large Numbers in the Theory of Consumer's Choice under Uncertainty," Journal of Economic Theory, 12, 1976, 202-217.

Zeldes, Stephen P., "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," unpublished MIT Ph.D. Dissertation, 1984.

__________, "Consumption and Liquidity Constraints: An Empirical Investigation," Rodney L. White Center Working Paper #24-85, The Wharton School, University of Pennsylvania, November 1985.

FIGURE 2-1

Optimal First period Consumption
Under Certainty Equivalence

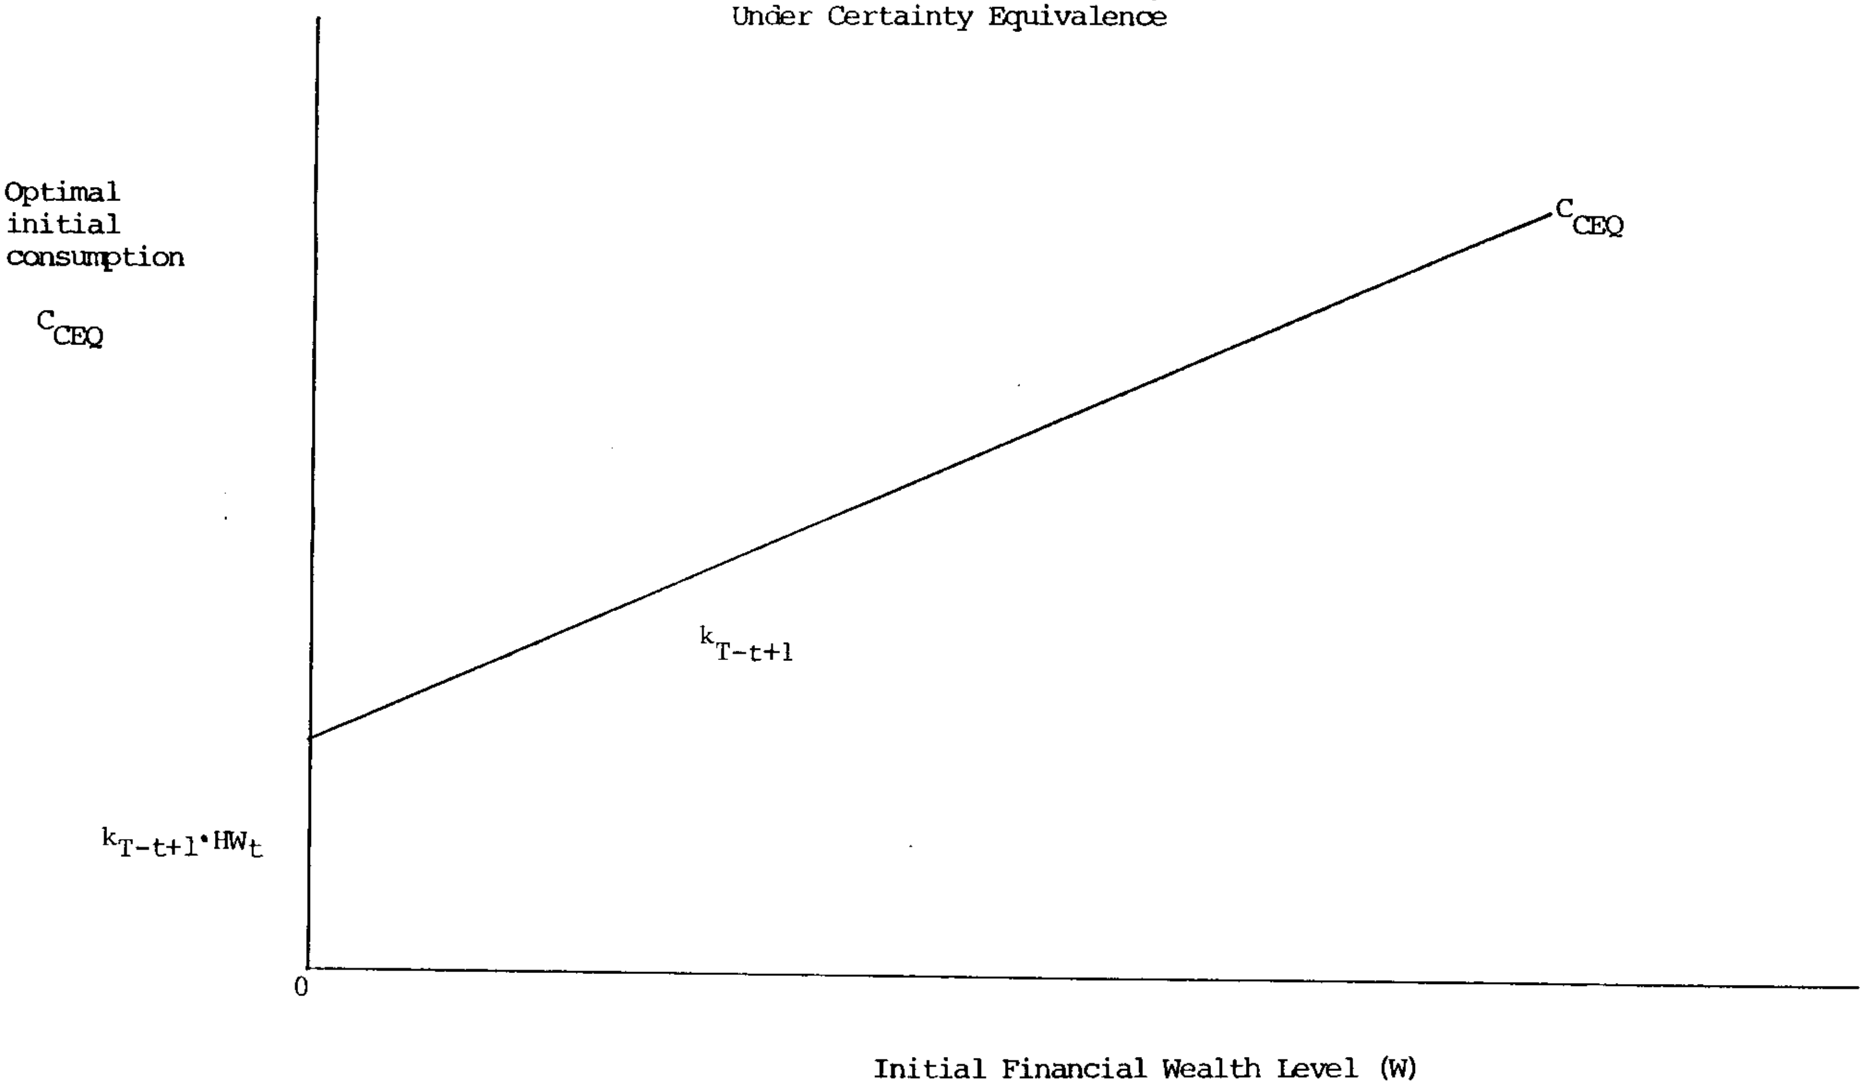

FIGURE 5-1

COEFFICIENT OF VARIATION: j STEP AHEAD FORECAST

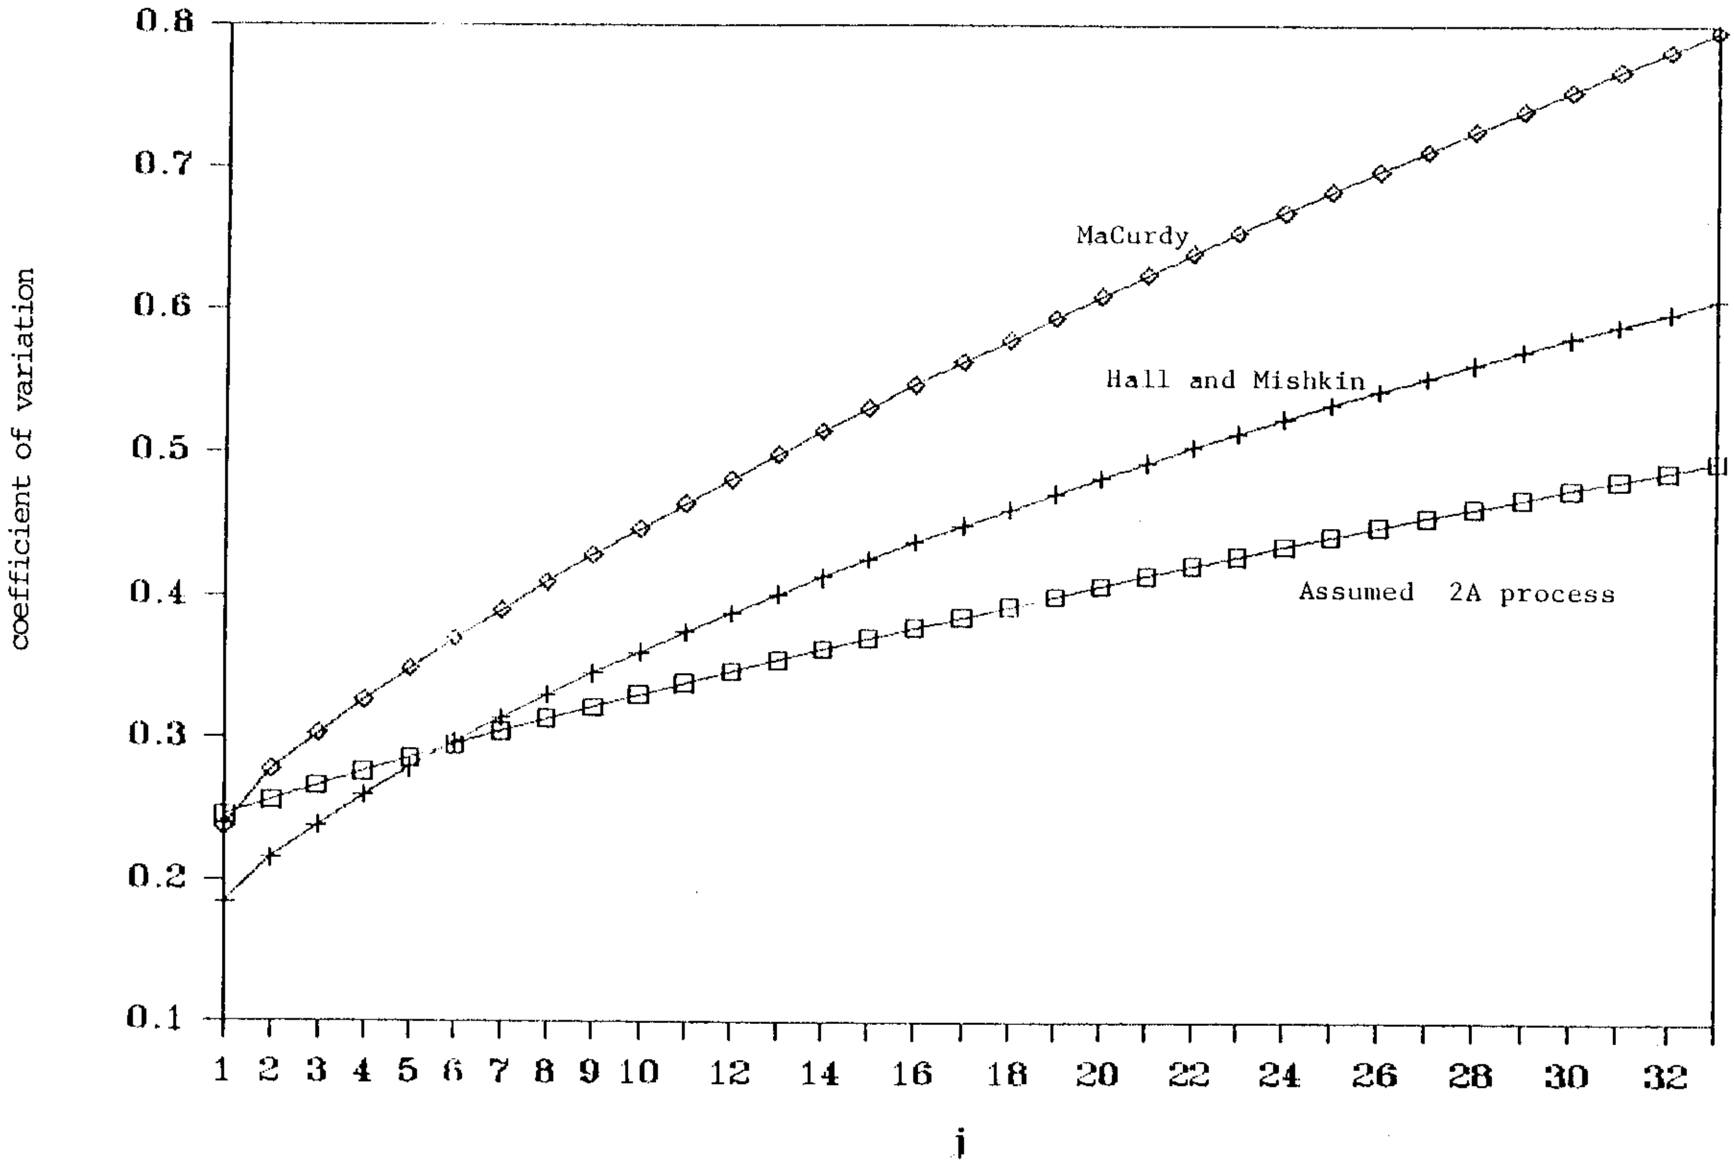

FIGURE 5-2

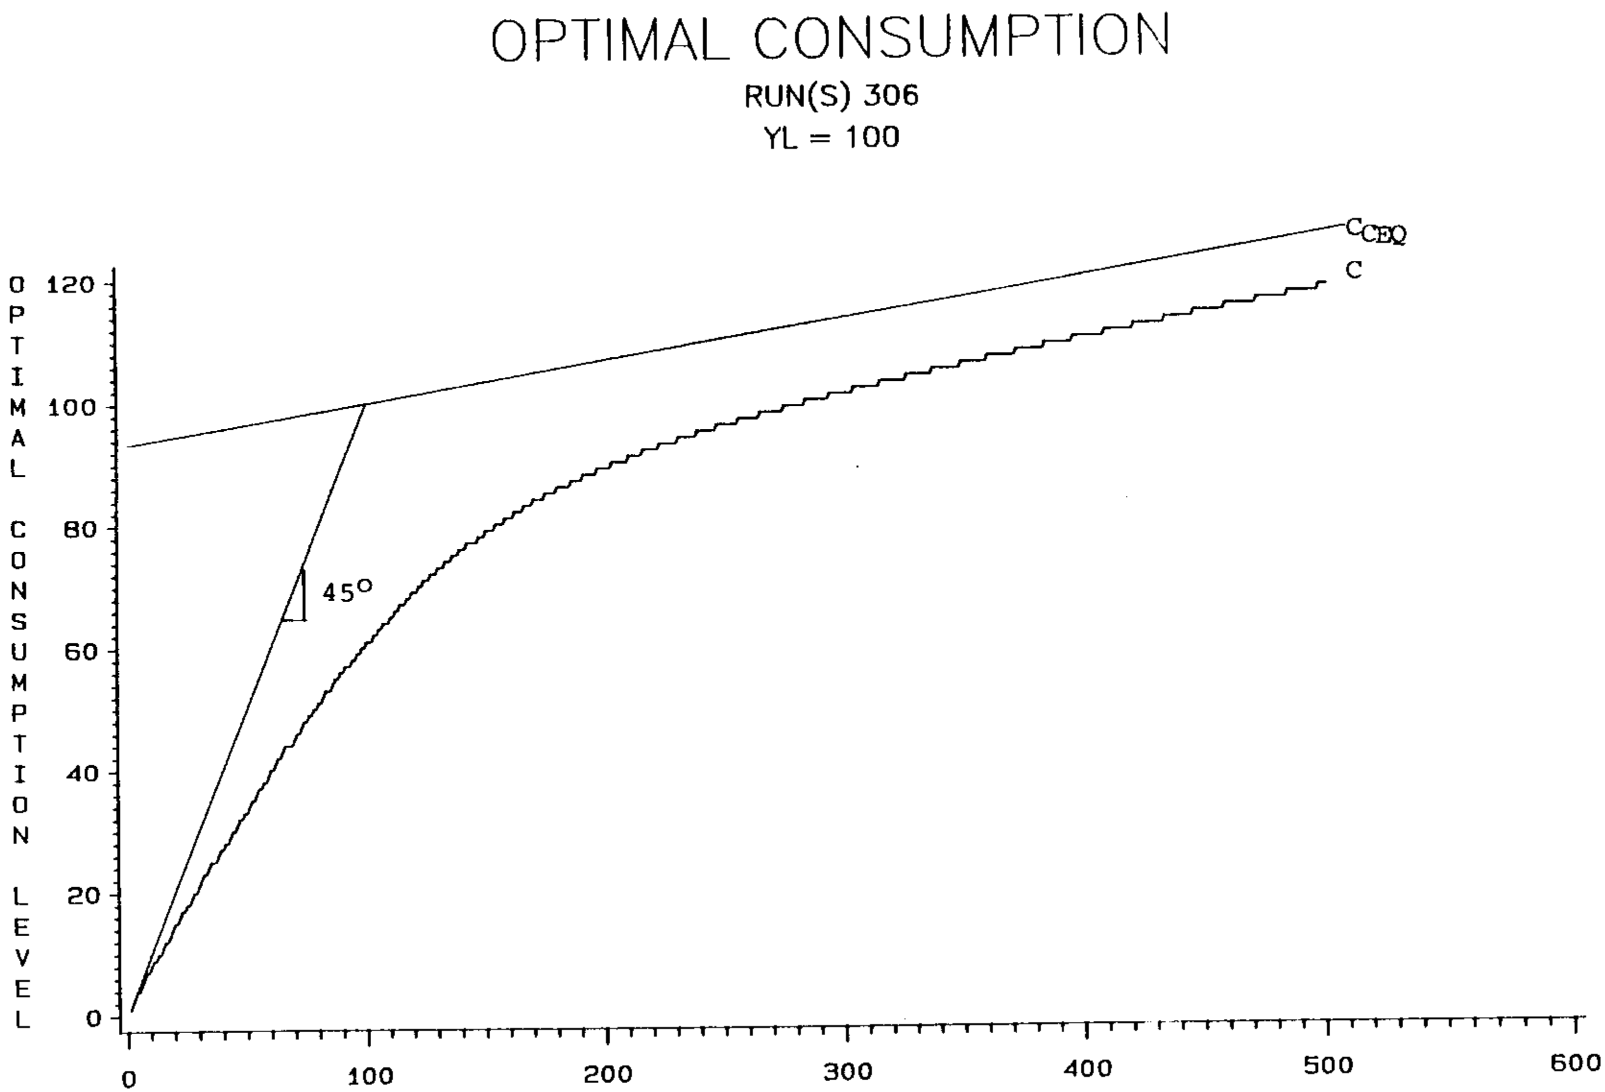

T = 15
Income Process: #2A

Expected future income equals 100 per period

$U = \frac{C^{1-A}}{1-A}$ A=3

FIGURE 5-3

OPTIMAL CONSUMPTION

as a function of expected future income (Run 306)

WEALTH = 200

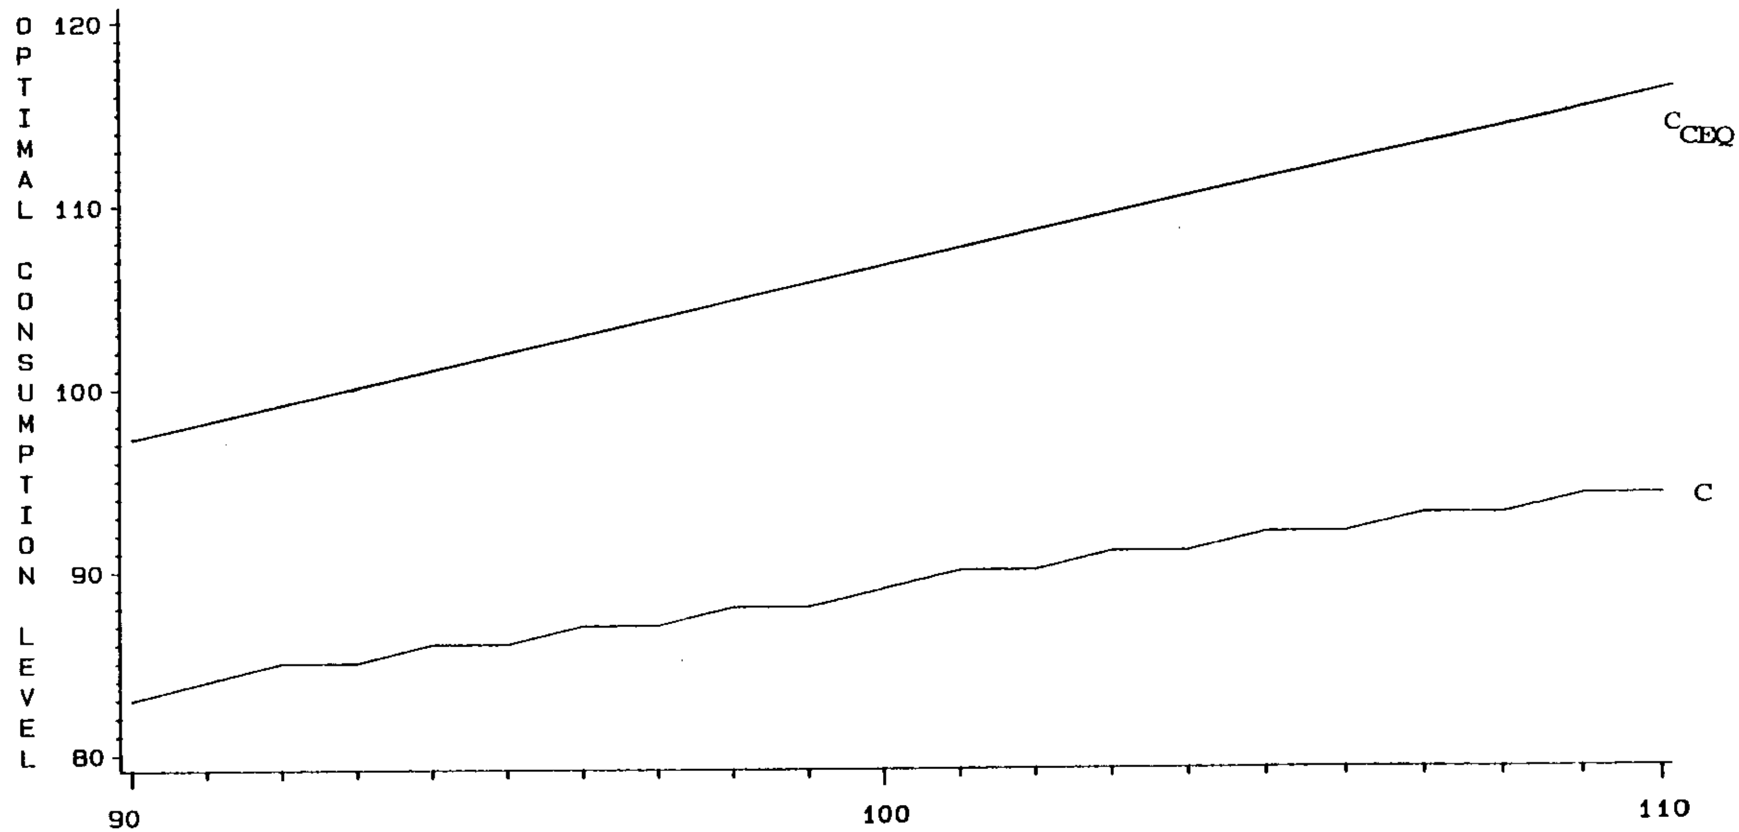

T = 15
Income Process: #2A

$$U = \frac{C^{1-A}}{1-A} \quad A = 3$$

FIGURE 5-4

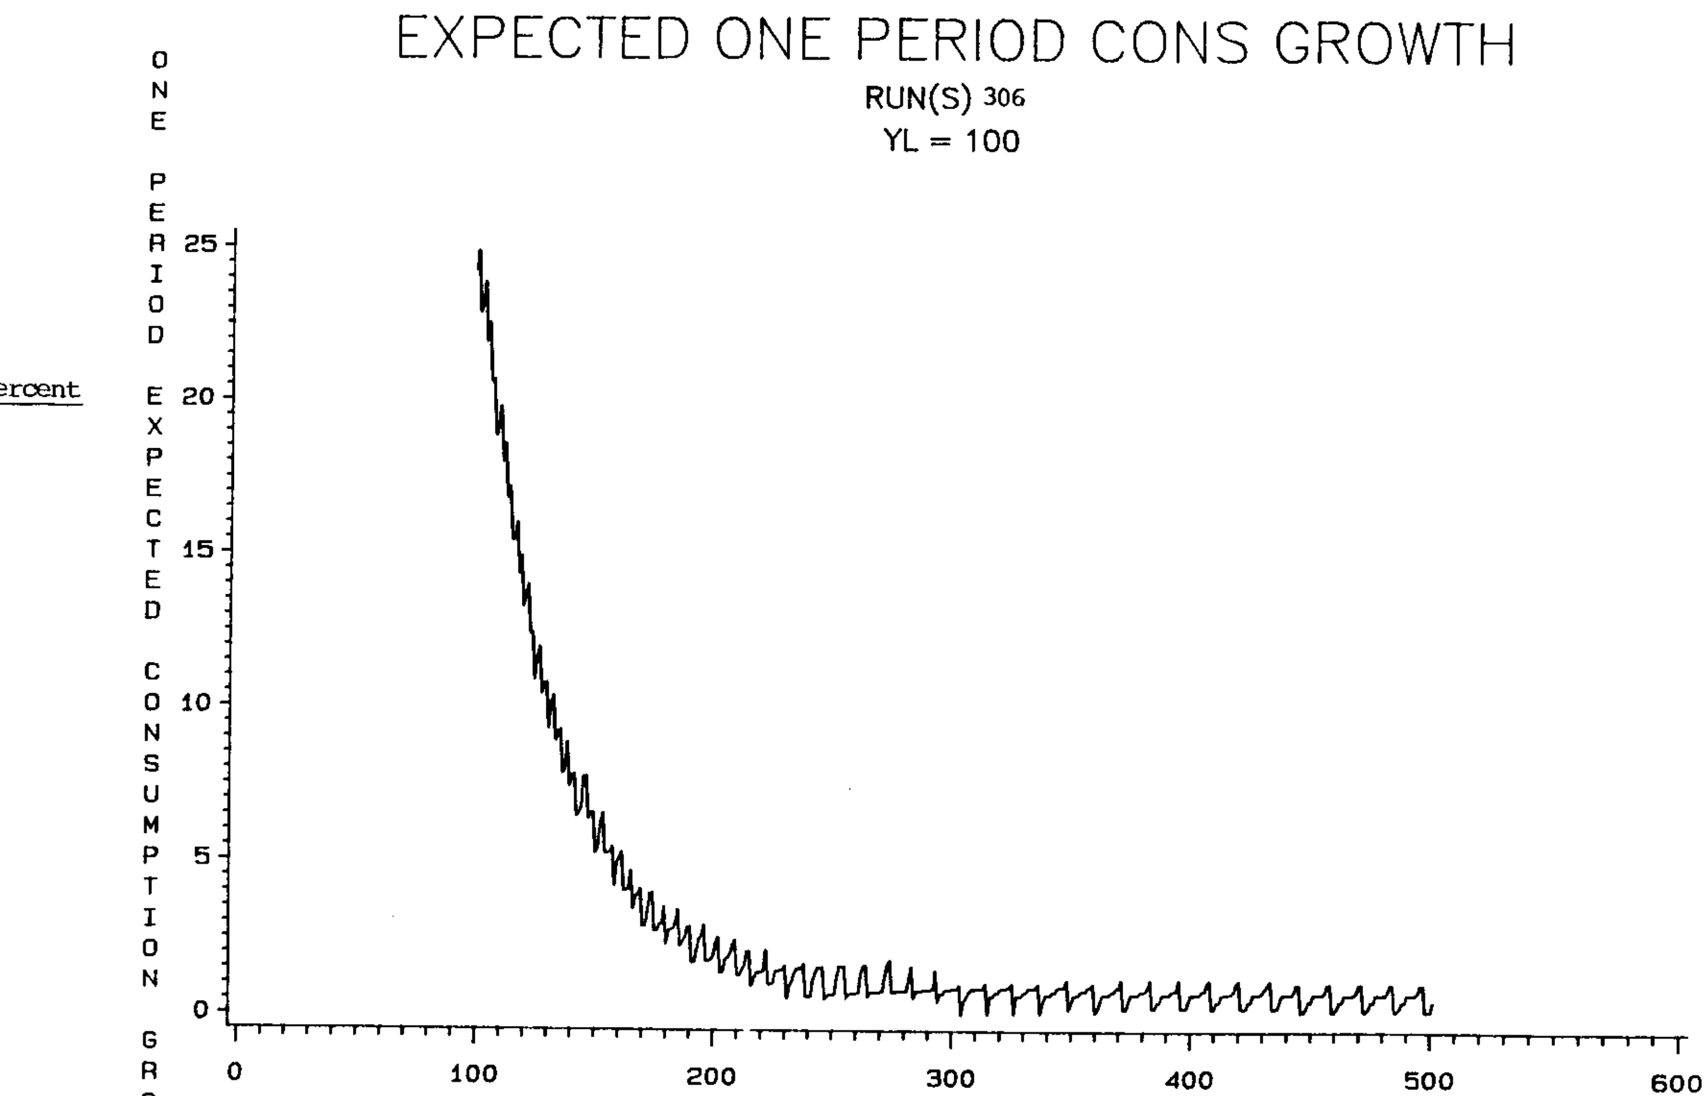

T = 15
Income process: 2A
$U = \frac{C^{1-A}}{1-A}$ A = 3 Expected futire income equals 100 per period

FIGURE 5-5

EXPECTED LIFETIME CONS GROWTH

RUN(S) 306

YL = 100

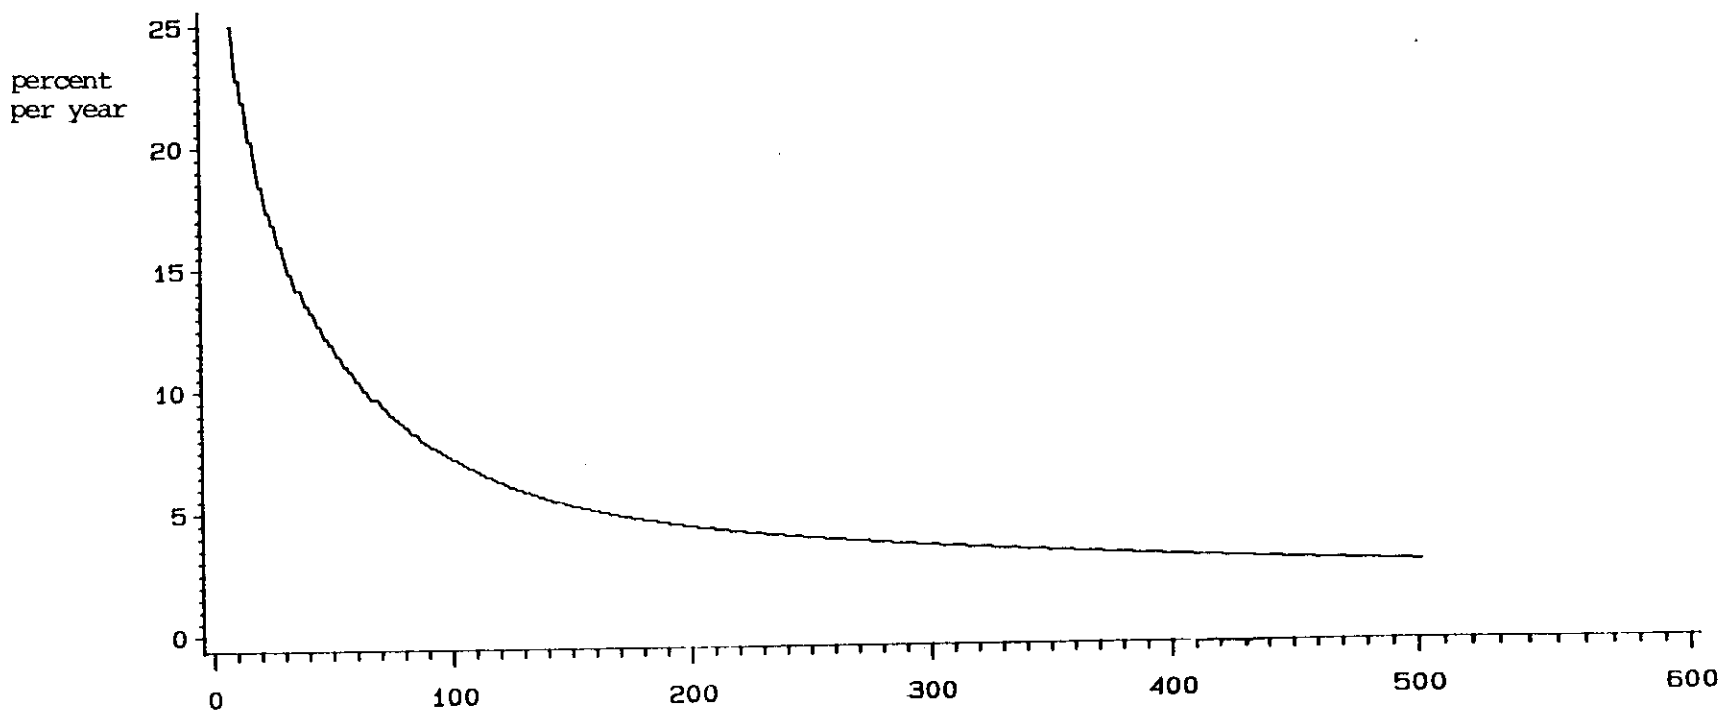

T = 15

Income Process: #2A

$U = \frac{C^{1-A}}{1-A}$ A = 3   Expected future income equals 100 per period

FIGURE 5-6

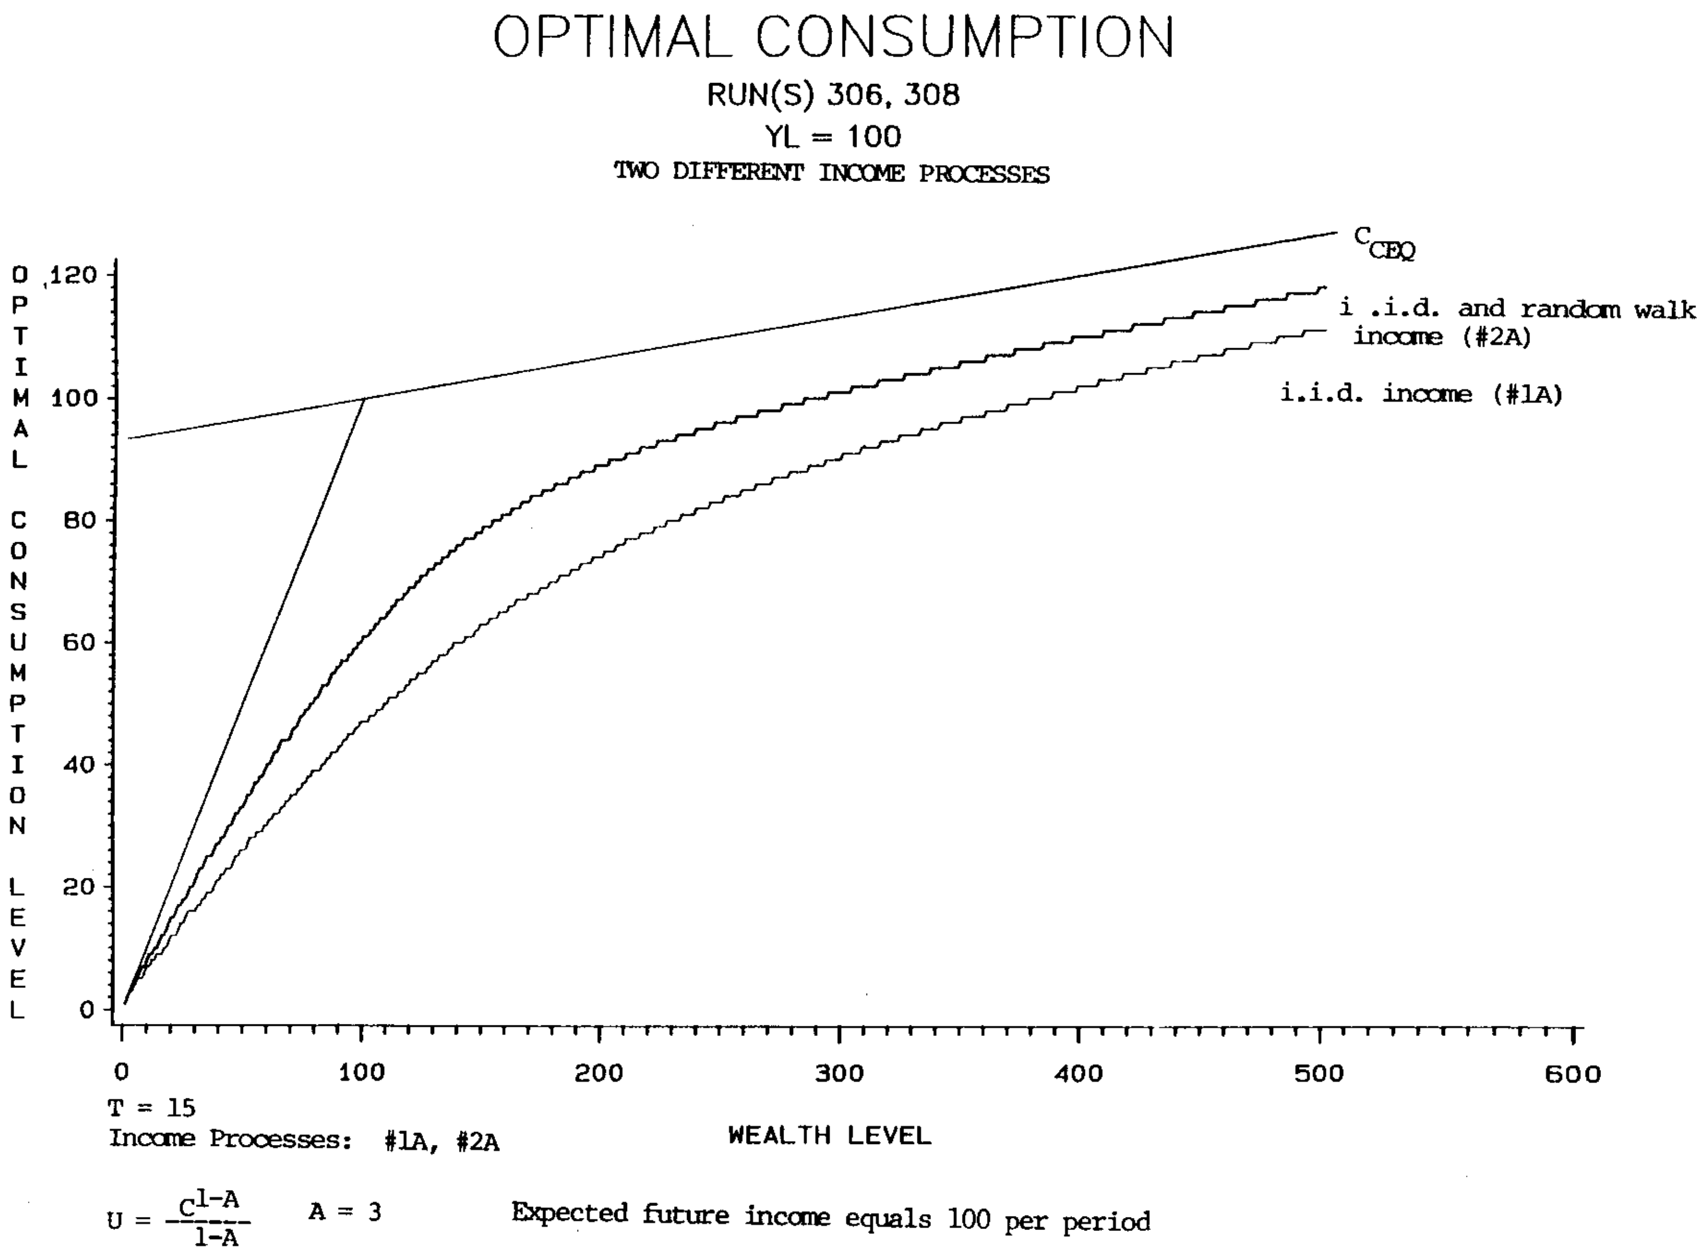
OPTIMAL CONSUMPTION
RUN(S) 306, 308
YL = 100
TWO DIFFERENT INCOME PROCESSES
OPTIMAL CONSUMPTION LEVEL
120
100
80
60
40
20
0
100
200
300
400
500
600
C_{CEQ}
i .i.d. and random walk income (#2A)
i.i.d. income (#1A)
WEALTH LEVEL
T = 15
Income Processes: #1A, #2A
$U = \frac{C^{1-A}}{1-A}$ A = 3
Expected future income equals 100 per period

FIGURE 5-7

OPTIMAL CONSUMPTION

RUN(S) 308,309,310

DIFFERENT HORIZONS

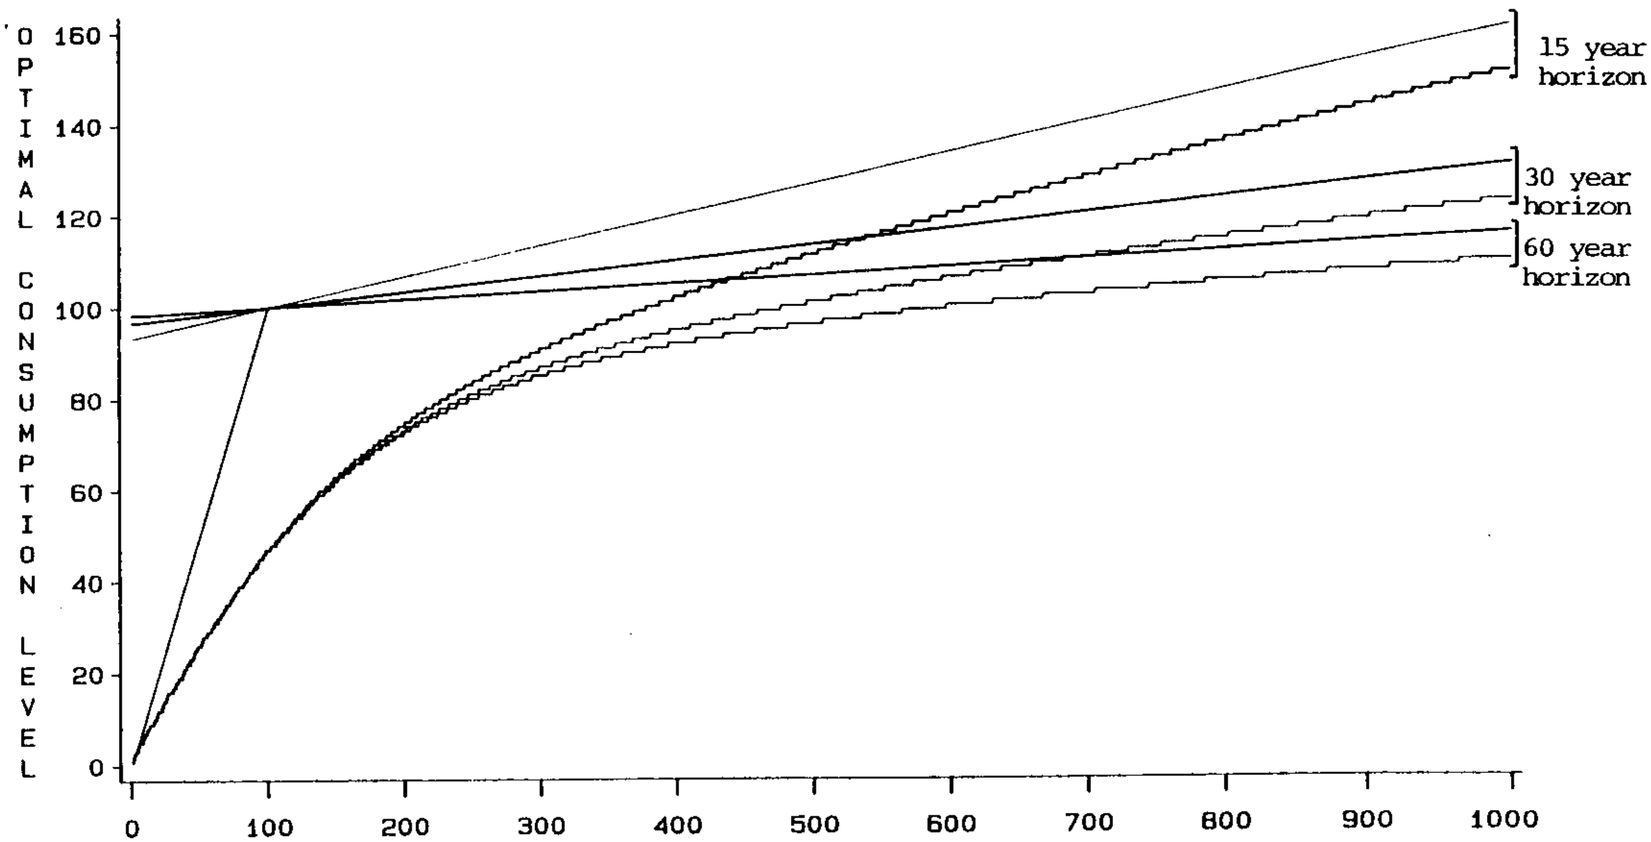

T = 15, 30, 60
Income Process: #1A

$U = \frac{C^{1-A}}{1-A}$ A = 3

Expected future income equals 100 per period

FIGURE 5-8

OPTIMAL CONSUMPTION

PERCENT DIF FROM CEQ

RUN(S) 308,309,310

DIFFERENT HORIZONS

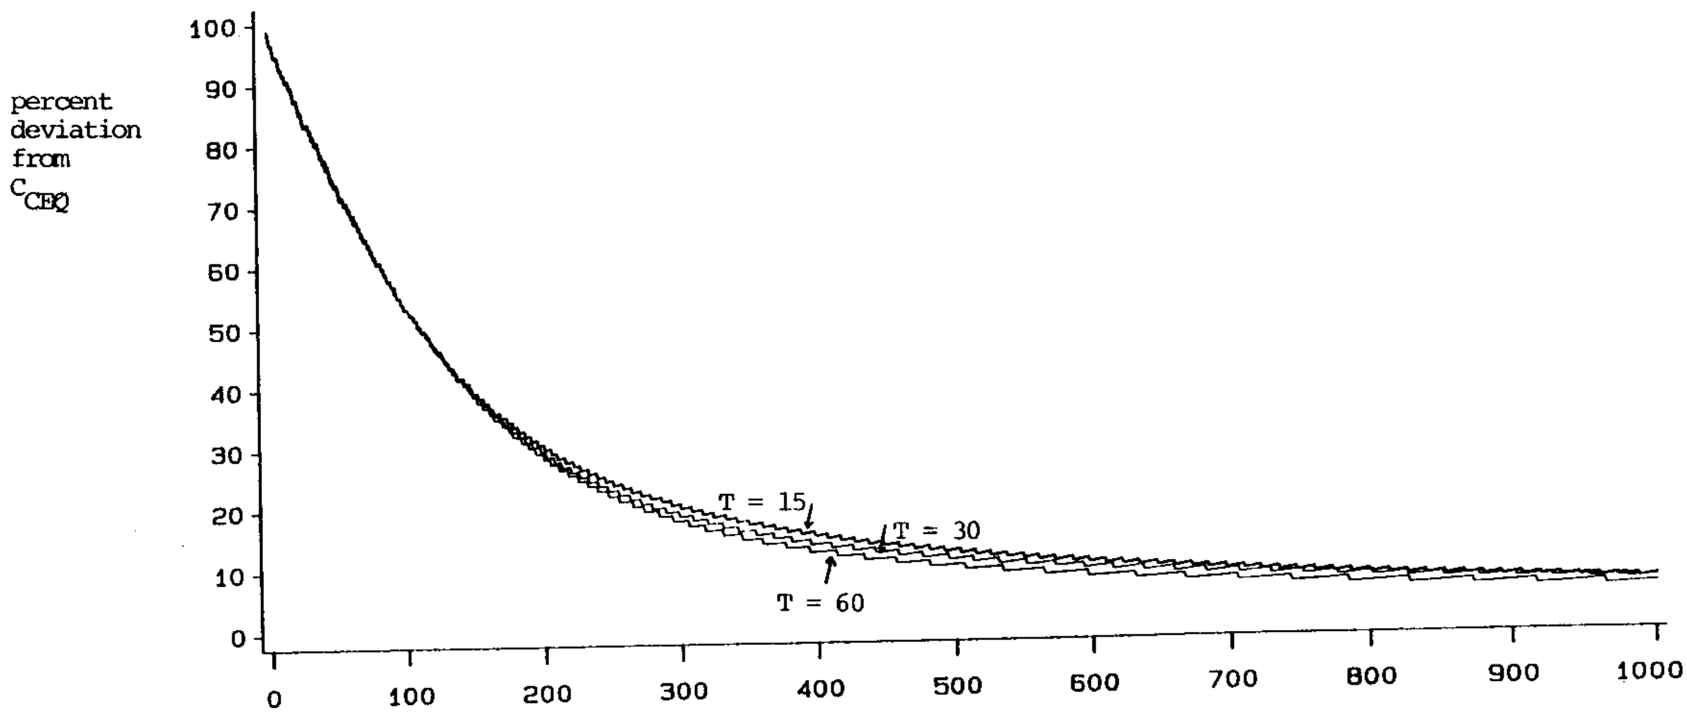

T = 15, 30, 60

Income Process: #1A

$U = \frac{C^{1-A}}{1-A}$ A = 3 Expected future income equals 100 per period

FIGURE 5-9

OPTIMAL CONSUMPTION

RUN(S) 309,316,317

DIFFERENT COEFFICIENTS OF RISK AVERSION

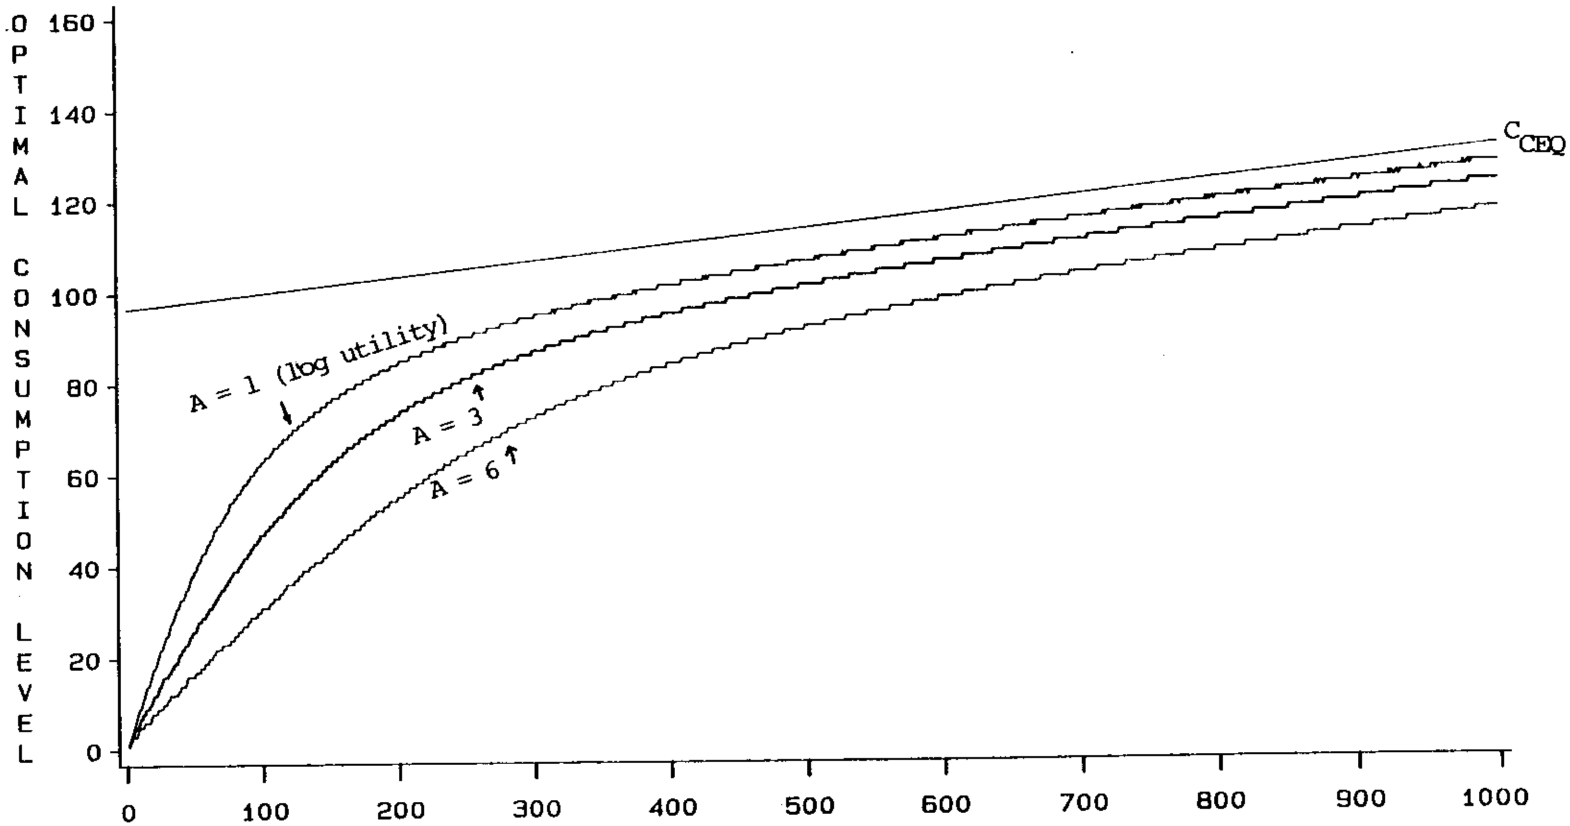

T = 30
Income Process: #1A

$U = \frac{C^{1-A}}{1-A}$ A = 1,3,6 Expected future income equals 100 per period

FIGURE 5-10

OPTIMAL CONSUMPTION

RUN(S) 309,315

DIFFERENT AMOUNTS OF INCOME RISK

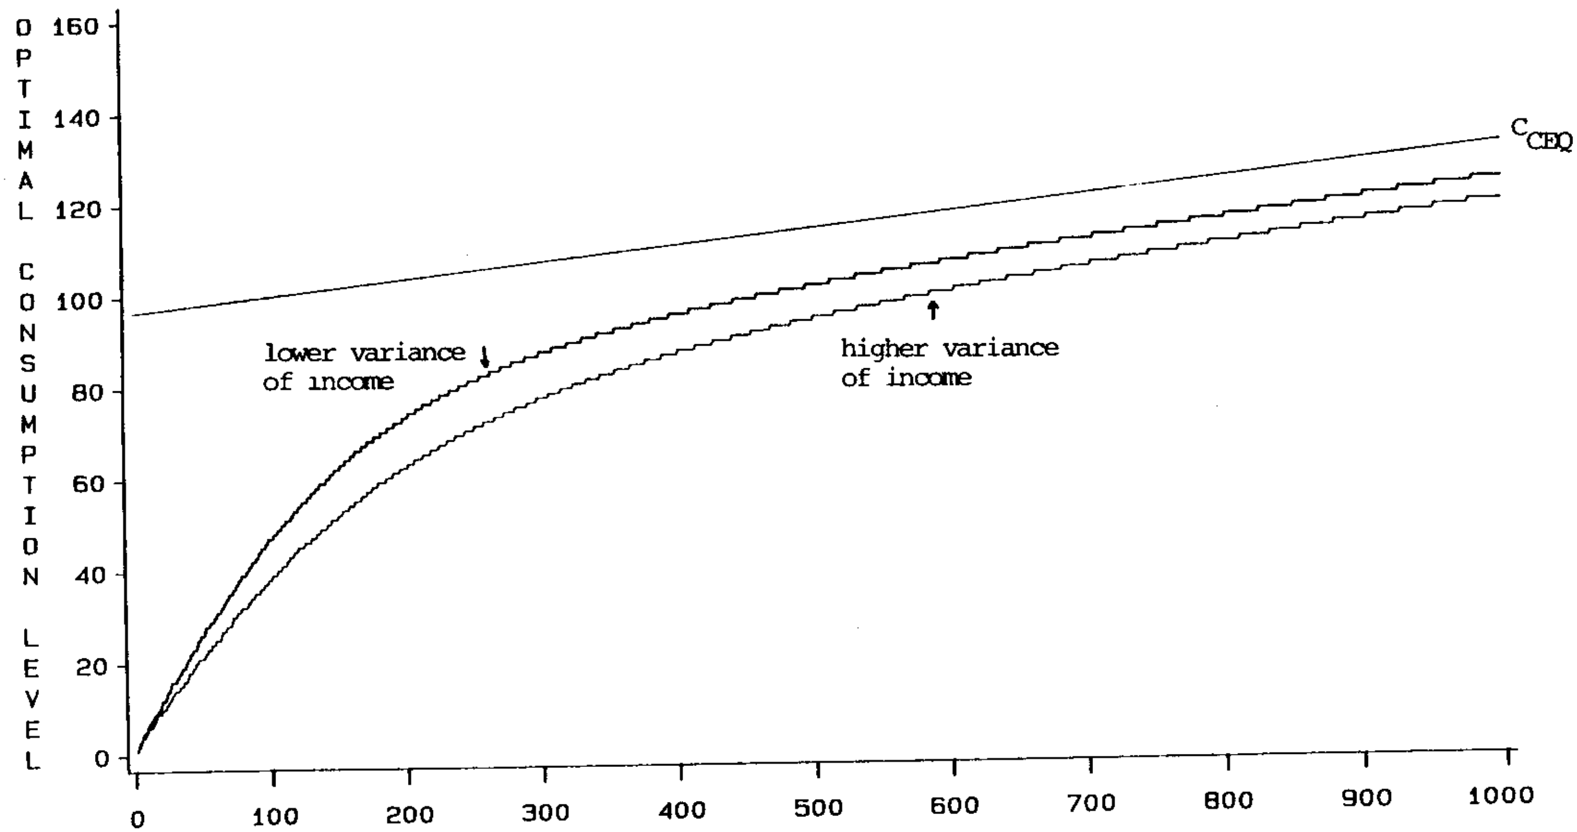

T = 30
Income Processes: #1A, #1B

$U = \frac{C^{1-A}}{1-A}$ A = 3

Expected future income equals 100 per period

FIGURE 5-11

OPTIMAL CONSUMPTION

RUN(S) 312 (0 floor)

Non-negativity constraint imposed on consumption

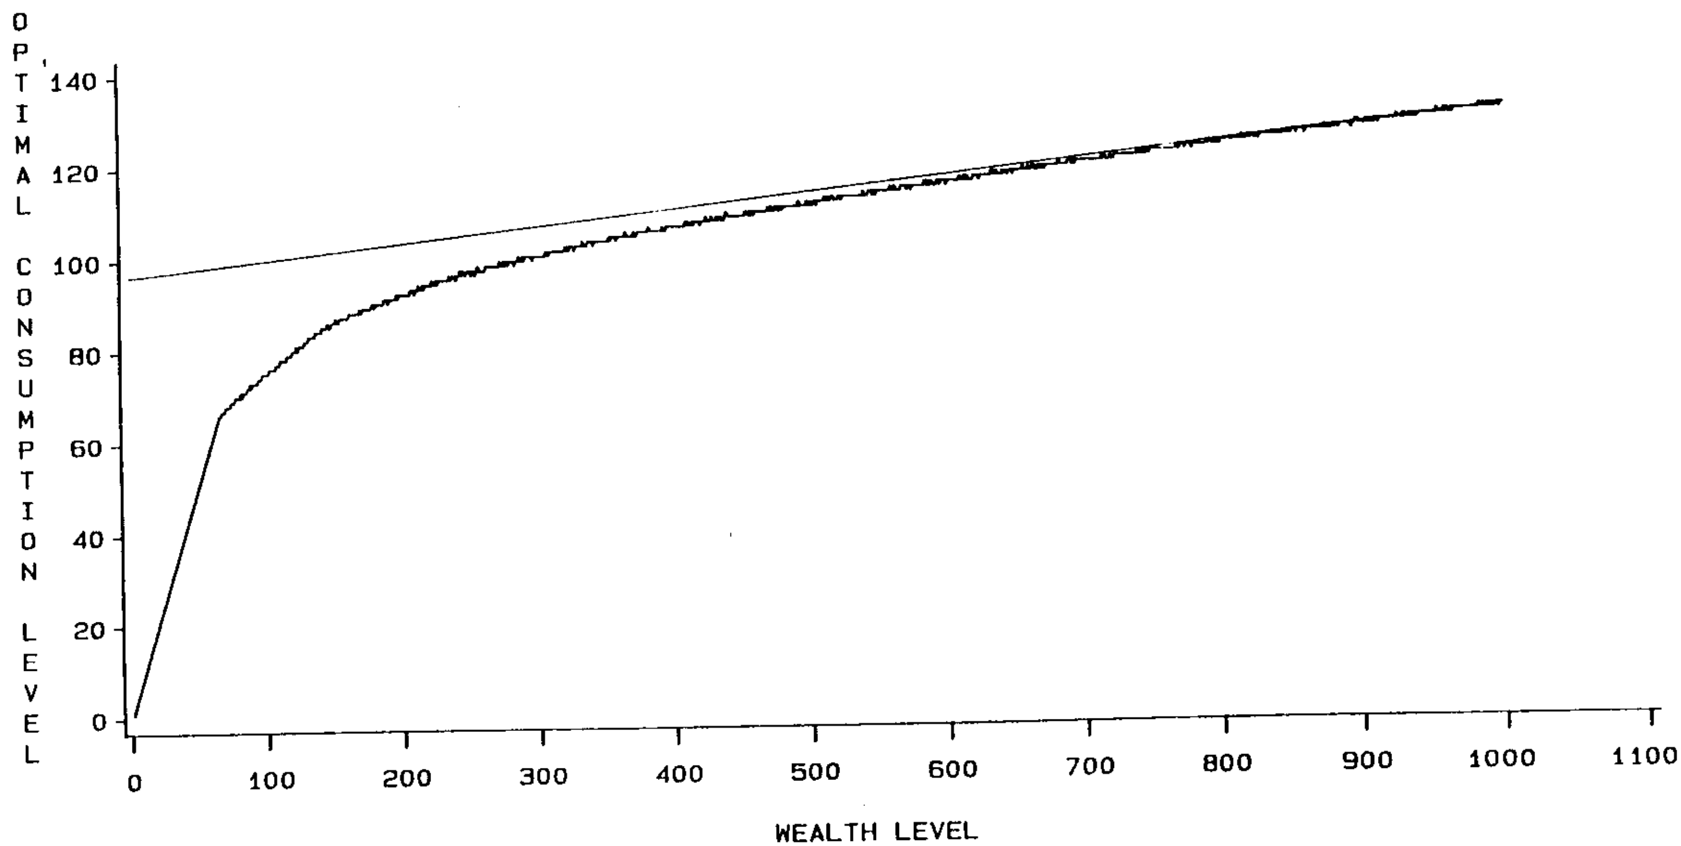

T = 30
Income process: #1A

Expected future income equals 100 per period

$U = C - C^2 / 3000$

FIGURE 6-1

CONSUMPTION WITH A BORROWING CONSTRAINT: FOUR REGIONS

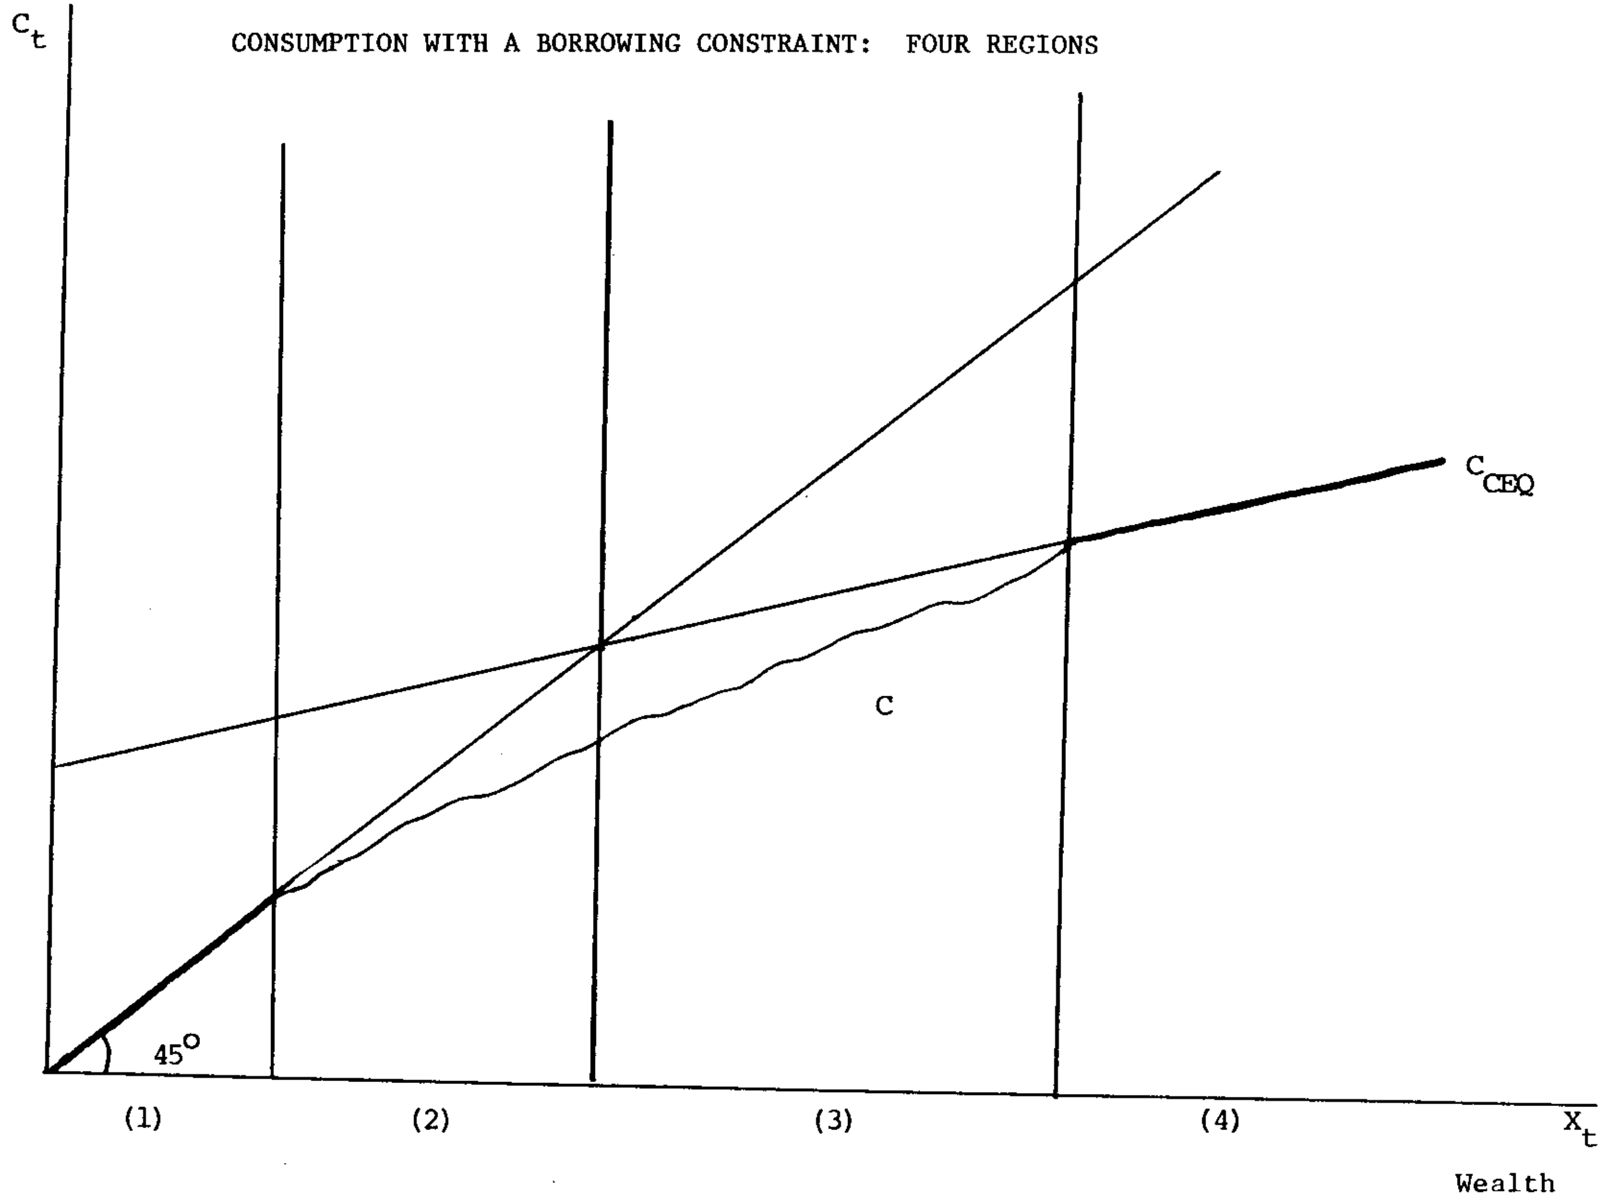

FIGURE 6-2

OPTIMAL CONSUMPTION

RUN(S) 312,314

With weak and strong borrowing constraints

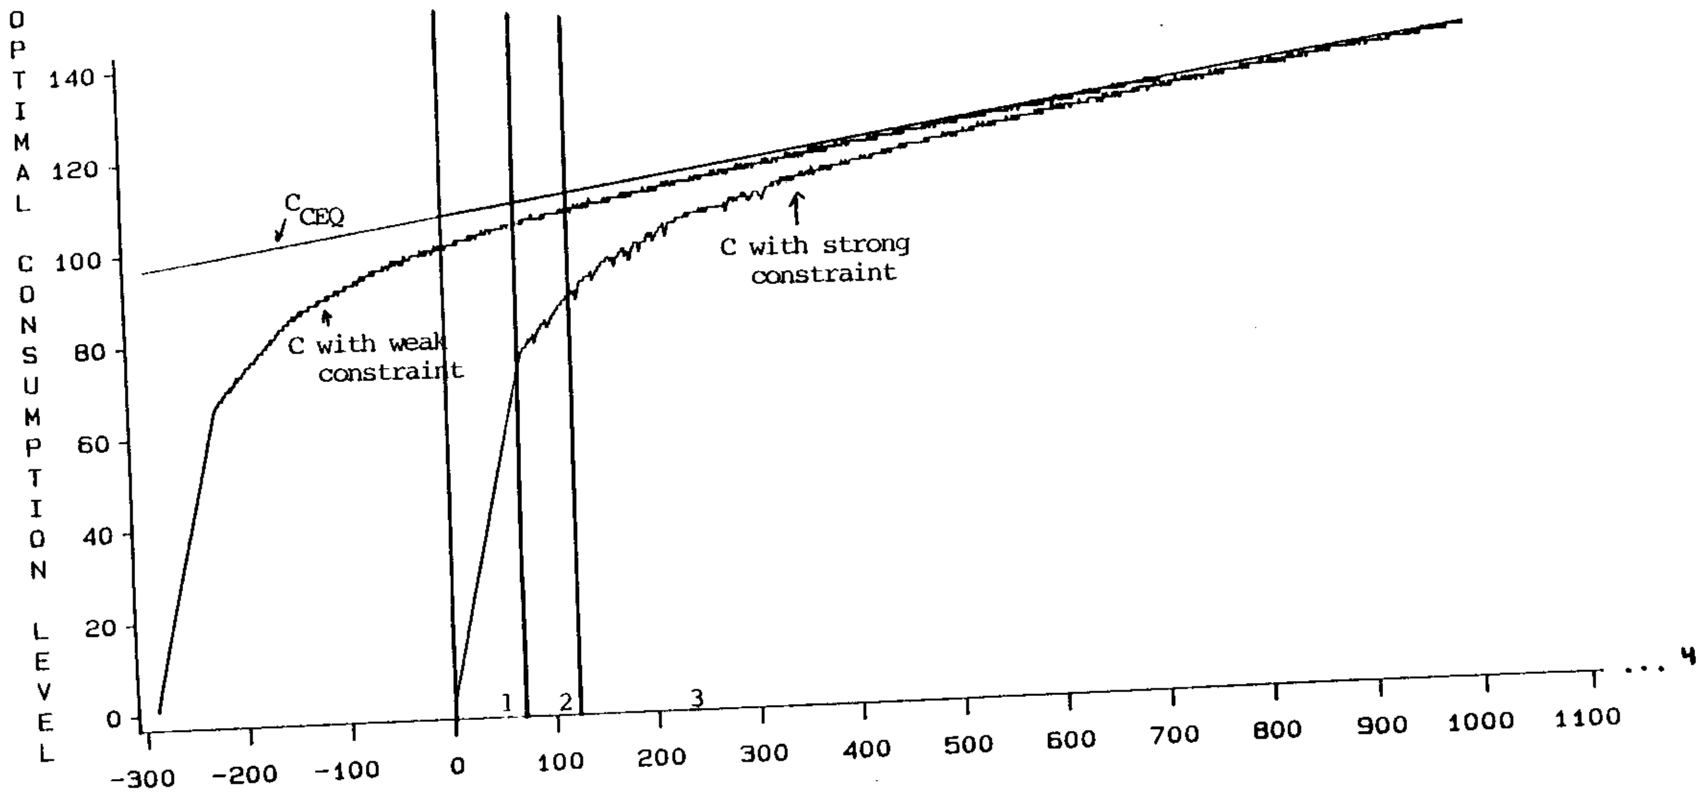

T = 30
Income process: #1C

Expected future income equals 110 per period
Income floor equals 10 per period

$U = C - C^2/3000$

Regions marked are for strong constraint

FIGURE 6-3

OPTIMAL CONSUMPTION

RUN(S) 314,318

DIFFERENT AMOUNTS OF INCOME RISK.
STRONG BORROWING CONSTRAINT.

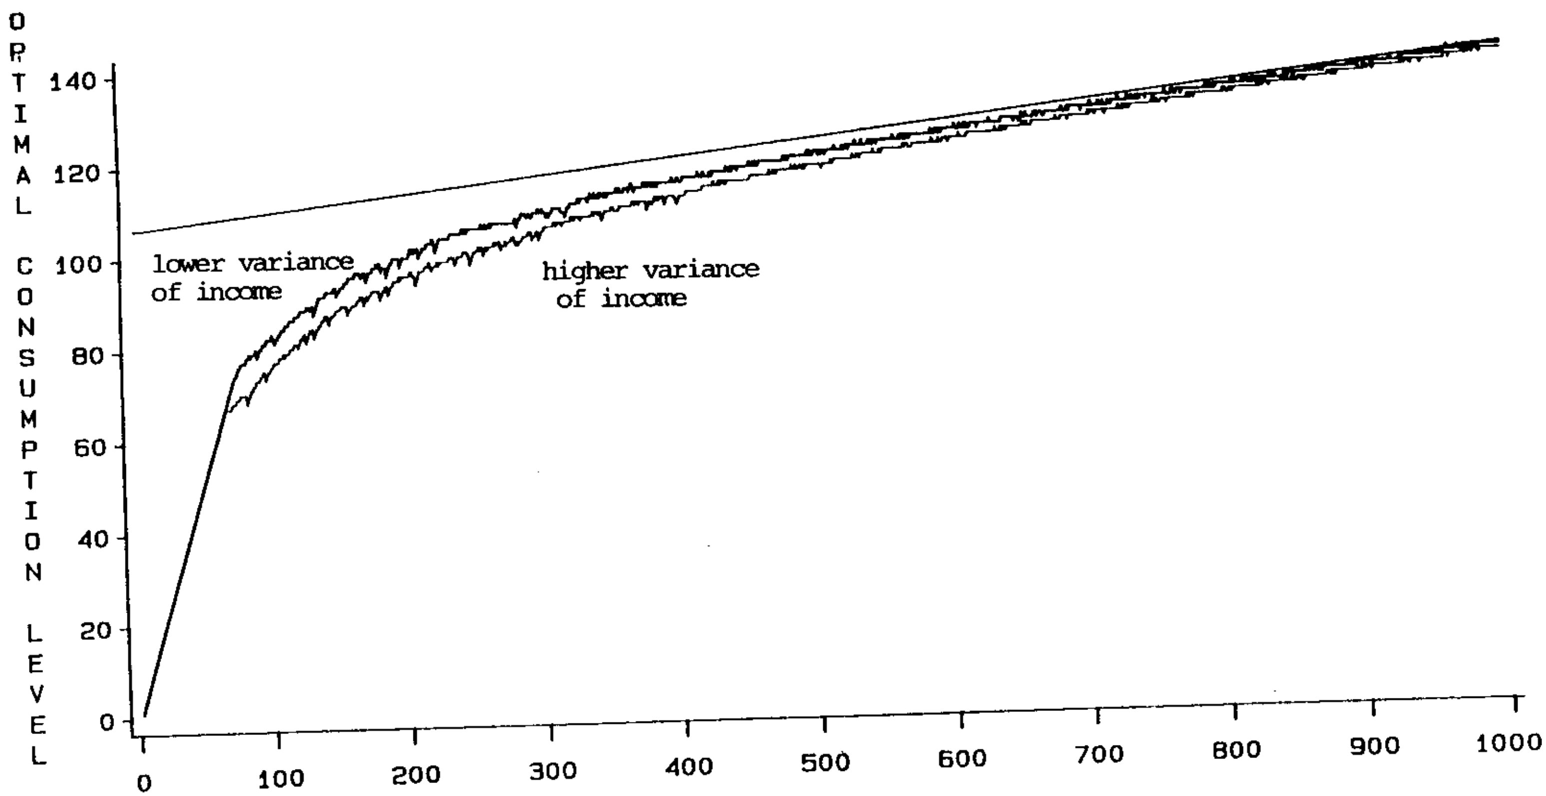

T = 30
Income Processes: #1C, #1D

$U = C - C^2/3000$

Expected future income equals 110 per period
Income floor equals 10 per period

FIGURE 6-4

EXPECTED LIFETIME CONS GROWTH

RUN(S) 312,314

WITH WEAK AND STRONG BORROWING CONSTRAINTS

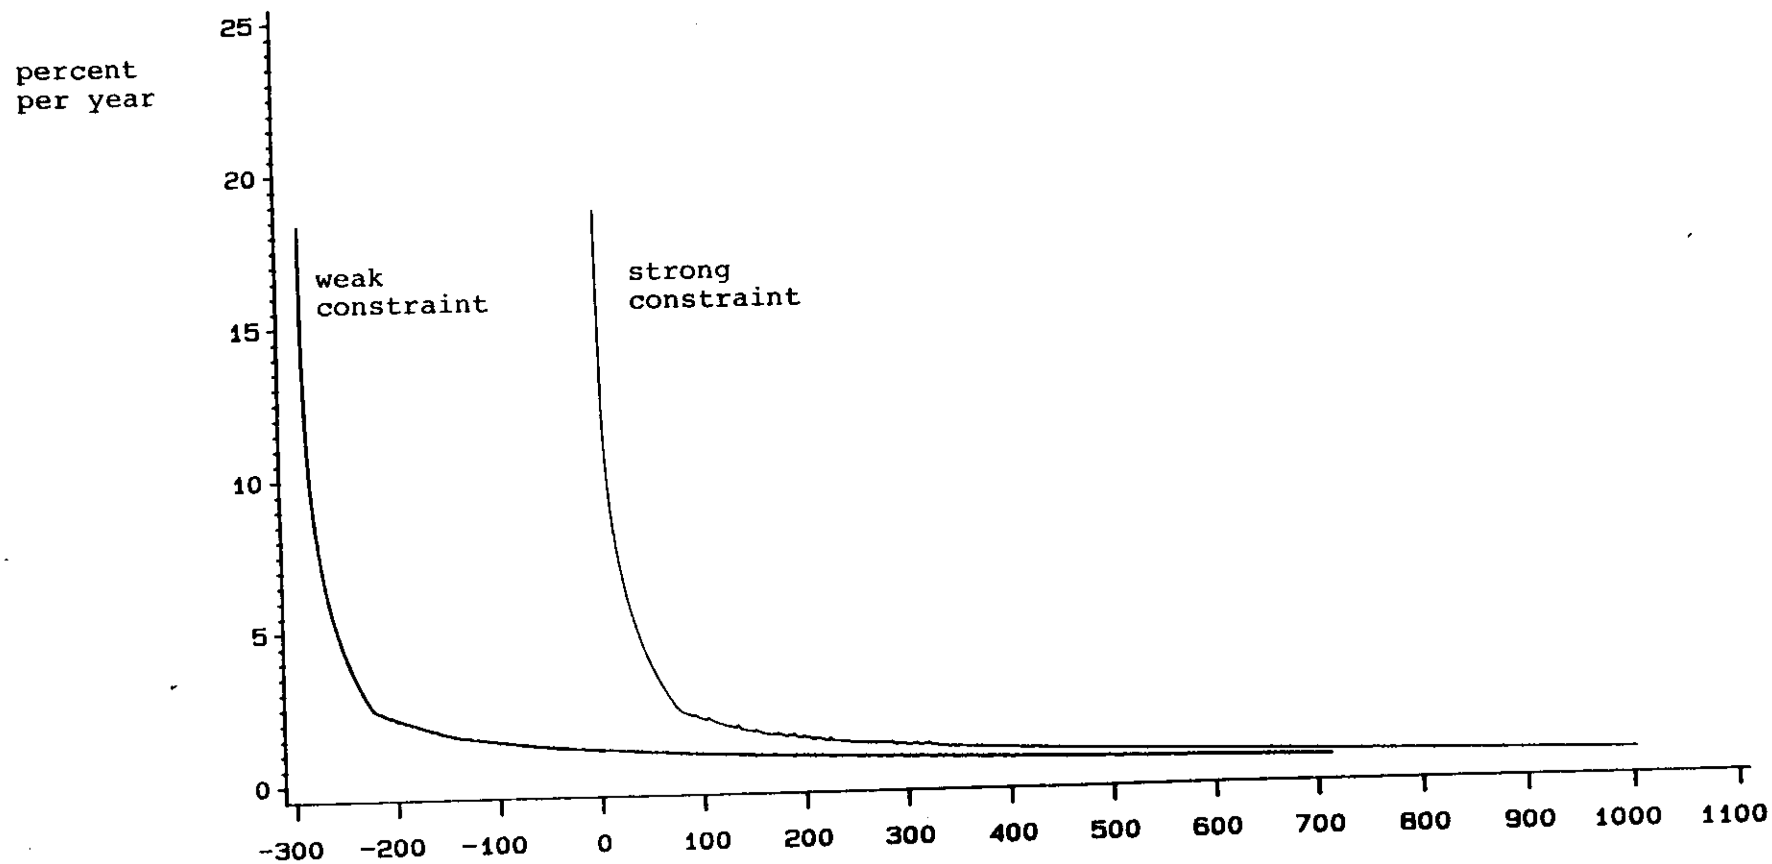

T = 30
Income Process: #1C

U = C $-C^2/3000$

Expected future income equals 110 per period

Income floor equals 10 per period