Housing, Adjustment Costs, and Macro Dynamics

Marjorie Flavin

This version: June 12, 2011

Economics Department, 0508
UCSD
9500 Gilman Dr.
La Jolla, CA 92093-0508

Housing, Adjustment Costs, and Macro Dynamics

ABSTRACT

When utility depends on a single, frictionlessly adjustable consumption good, the household's willingness to substitute consumption intertemporally is solely determined by the curvature of the utility function. When the utility specification is generalized from one good to two, however, the curvature parameter then specifies the curvature of the utility function with respect to a composite good. If *both* goods are frictionlessly adjustable, then each of the individual goods will all have the same dynamics, and the intuition from the one-good case -- that the elasticity of intertemporal substitution is determined solely by the curvature parameter -- remains valid. In this paper, however, the two goods are interpreted as housing services and non-housing goods. Nonhousing consumption can be adjusted frictionlessly, but housing services are subject to a nonconvex adjustment cost. The paper explores the dynamic behavior of nondurable consumption by solving numerically for the optimal level of nonhousing consumption, conditional on an assumed path for housing consumption and for the marginal utility of wealth. The results indicate that the intertemporal behavior of nonhousing consumption depends crucially on the intratemporal substitutability of the two goods as well as the curvature of the utility function with respect to the composite good.

JEL: E21, E32, D91

Keywords: housing, adjustment costs, intratemporal substitution, consumption dynamics

Section 1: Introduction

When a change in the household's stock of housing (and thus its consumption of housing services) is subject to a nonconvex adjustment cost, the effect of the adjustment cost on the consumption of housing services is conceptually straightforward and easy to identify in household level data. Given that adjusting the consumption of housing services by moving from one residence to another requires substantial costs in terms of time and effort, in addition to any direct pecuniary costs for moving services or real estate commissions, changes in the consumption of housing services will occur infrequently.

Because modeling the dynamics of consumption of housing services is greatly complicated by the adjustment costs on housing, the vast majority of macro models abstract from housing altogether and instead focus on the determination of nondurable consumption, or nonhousing consumption. In the case of nondurable consumption, the assumption of costless adjustment is much more plausible. In narrowing the consumption concept to nondurable consumption, however, these models are implicitly assuming that the household's utility function is separable between nondurable consumption and consumption of housing services.

It seems unlikely, *a priori*, that the utility function is exactly separable in housing and nonhousing consumption. Nevertheless, unless models based on a more general nonseparable utility function generate some important insights or dramatic improvements in the model, one could argue that the assumption of separability should be maintained, given the degree to which it simplifies the analysis.

In previous papers, Flavin and Nakagawa (2008) and Flavin and Yamashita (2010), I have argued that when an otherwise completely standard model of the lifetime consumption/saving and portfolio allocation problem is altered by introducing two elements – 1) a utility function that depends nonseparably on housing and on nondurable consumption, and 2) a nonconvex adjustment cost on housing – the more general model does indeed offer some new insights and provide a better fit to the data.

In the previous papers, I have proposed a generalized version of the standard macro/finance model of the household's lifetime optimization problem that incorporates the housing in its dual role as both an argument of the utility function and as an asset. The model is as follows: The household chooses the optimal level of housing, holdings of financial assets, and the level of nonhousing consumption in a continuous time setting. Adjustment of the quantity of housing requires the payment of a nonconvex adjustment cost, while nondurable consumption and financial assets can be adjusted frictionlessly. Because of the adjustment cost on housing, the solution to the general problem has a recursive structure: at each moment, the household considers whether or not to sell the house, pay the adjustment cost, and choose a new quantity of housing. Most of the time, it is not optimal to incur the adjustment cost. Having decided not to sell the house at that instant, the household then chooses the optimal level of nonhousing consumption and the optimal holdings of financial assets conditional on the current level of housing. When, very infrequently, it is optimal to sell the house, the household optimally chooses the size of the new house.

Flavin and Yamashita (2010) explore the implications of the housing model for portfolio allocation. In order to model the household's portfolio constraints in a realistic manner, these papers assume that the family holds nonnegative amounts of the stocks, bonds, and the riskless asset (Treasury bills), and can borrow only in the form of a mortgage. Further, the maximum amount borrowed in the form of a mortgage is limited by the value of the house. The paper shows that the portfolio allocation decision can be treated as a well-defined subproblem within the household's overall intertemporal optimization problem. Under the assumptions of the model[1], the optimal portfolio is the outcome of a constrained mean-variance optimization problem, with the constraints reflecting the housing collateral constraint on the amount of borrowing, and the assumption of nonnegative holdings of financial assets other than the mortgage. While the model is designed to assess the portfolio decision of a homeowner, it nevertheless applies to renters as a special case. That is, renters and homeowners both derive utility from their consumption of housing services; likewise, renters and homeowners both choose a portfolio allocation that maximize a function of the mean and variance of the portfolio return subject to nonnegativity and collateral constraints. For a renter, the collateral constraint simply says that no borrowing in the form of a mortgage is allowed.

The implications of the housing model for the elasticity of intertemporal substitution (EIS) of nonhousing consumption are examined in Flavin and Nakagawa (2008). A basic feature of the standard, frictionless, one-good model (that is, the standard model without housing)) is the tight link between risk aversion and intertemporal substitution. In contrast to the standard frictionless model, Flavin and Nakagawa (2008) shows that the EIS of nonhousing consumption is not necessarily equal to the reciprocal of the degree of relative risk aversion in the housing model. The current paper further pursues the implications of the housing model for nonhousing consumption by looking explicitly at the implied dynamic behavior of nonhousing consumption for a range of assumptions on the utility function. As with the implications for portfolio allocation, the model applies equally well to homeowners and to renters. In both cases, the argument of the utility function is the flow of housing services consumed by the household in a given period; whether the household owns or rent simply determines the manner in which the housing services are financed. In terms of the implications of the model for the dynamics of nonhousing consumption, it is not the physical durability of the house that matters. Instead, the

[1] Normality of asset returns is one of the assumptions of the model.

crucial aspect of housing that determines the effect on the dynamics of nonhousing consumption is the fact that adjusting the consumption of housing services entails a nonconvex, or lumpy, adjustment cost. Since renters, like homeowners, incur a nonconvex adjustment cost in order to change their consumption of housing services, the implications of the model for the dynamics of nonhousing consumption applies to both renters and homeowners.

In particular, the previous paper shows that the standard model, augmented only by assuming that the utility function is nonseparable and that housing is subject to a nonconvex adjustment cost, can explain both the smoothness (that is, lack of volatility) of nonhousing consumption and a low elasticity of intertemporal substitution (EIS) without invoking either habit persistence or an implausibly high degree of curvature of the utility function.

My objective in this paper is to explore the implications of the more general model in order to characterize the effect of the adjustment cost on housing on 1) the dynamics of nondurable consumption, and 2) the joint time series behavior of nondurable consumption and the consumption of housing services. To anticipate the eventual result, it turns out that the dynamic, or *intertemporal*, stochastic behavior of nondurable consumption depends crucially on the degree of *intratemporal* substitutability of nondurables and housing in a given period.

More generally, I argue that there is an essential interaction between the household's preferences regarding the substitutability of goods within a period, and the household's preferences regarding the substitutability of goods across time. Because of this interaction, I argue that, for the same reason that our understanding of the dynamic (intertemporal) behavior of consumption is not complete without a specification of the household's preferences on intratemporal substitution, neither is our understanding of intratemporal substitution complete without specification of preferences regarding intertemporal substitution.

In order to build my case that there is an essential interaction between the intratemporal and intertemporal dimension of the household's allocation problem, I start by investigating the determinants of intratemporal substitution in an intertemporal setting in Section 2. This discussion will provide support and intuition for the results in Section 3, which considers the effect of intratemporal preferences on the intertemporal, or dynamic, behavior of consumption.

Section 2: Intratemporal Substitution in an Intertemporal Setting

Once we relax the assumption of separability, a natural specification of the within-period utility function is:

$$(1) \quad U = \frac{(C^{\alpha} + \gamma H^{\alpha})^{\frac{1-\rho}{\alpha}}}{1-\rho} \qquad \alpha \leq 1 \ , \ \rho \geq 0 \ , \rho \neq 1$$

Here C denotes nondurable (or nonhousing) consumption, which is assumed costlessly adjustable. Aside from nondurable consumption, current period utility depends on the flow of housing services. The flow of housing services is assumed to be proportional to the stock of housing, H, and current period utility can therefore be written as a function of C and H.

An alternative, equivalent, parameterization of the utility function in equation (1) is

$$(1') \quad U = \frac{\left(C^{1-\frac{1}{\sigma}} + \gamma H^{1-\frac{1}{\sigma}}\right)^{\frac{1-\rho}{1-\frac{1}{\sigma}}}}{1-\rho} \qquad \sigma > 0 \ , \ \rho \geq 0 \ , \rho \neq 1$$

While I will work with the parameterization in equation (1), one can, of course, always map the results into the parameterization in (1') by using $\sigma = \frac{1}{1-\alpha}$.

Assume that the household maximizes a utility function specified over four goods:

$$(2) \quad U = (1-\rho)^{-1}(C_1^{\alpha} + \gamma H_1^{\alpha})^{\frac{1-\rho}{\alpha}} + \beta(1-\rho)^{-1}(C_2^{\alpha} + \gamma H_2^{\alpha})^{\frac{1-\rho}{\alpha}}$$

The choice of notation and the asymmetry between the treatment of the (C_1, H_1) pair of goods and the (C_2, H_2) pair of goods immediately suggests a two period problem with two distinct physical goods. Nevertheless, the preference specification in equation (2) could equally well apply to a one-period problem with four distinct physical goods. While our primary interest is in characterizing the intratemporal substitution of goods in an intertemporal setting, the atemporal, or one-period, interpretation of the preference specification will occasionally be invoked in order to provide intuition.

In a previous paper, I suggested that a pair of goods be classified as either complements or substitutes on the basis of the sign of the cross derivative of the utility function. A criterion based on the cross derivative seems intuitive: if an increase in the quantity of good 1 increases the marginal utility of good 2, it seems natural to think of the two goods as complements. Likewise, if an increase in the quantity of good 1 reduces the marginal utility of good 2, the two goods can logically be considered substitutes. For utility functions of the form of equation (2), the sign of the cross derivative depends on all of the parameters of the utility function. In particular,

(3) $$\text{sign}\left[\frac{\partial^2 U}{\partial C_1 \partial H_1}\right] = \text{sign}\,[\gamma(1-\rho-\alpha)] = \text{sign}\,(1-\rho-\alpha)$$

where the last inequality follows from the assumption that $\gamma > 0$. Under this criterion, the two goods would be complements if $1-\alpha > \rho$ and substitutes if $1-\alpha < \rho$.

Other authors have taken issue with the classification of goods as complements or substitutes based on the cross derivative of the utility function. For example, Piazzesi, Schneider and Tuzel (2007) in their influential paper on the role of housing in asset pricing, use the utility function parameterized as in equation (1'), but argue that the intratemporal substitutability of goods depends only on the parameter σ (or, equivalently, α): [2]

> "We use standard Hicksian language here: two goods are substitutes if and only if $\sigma > 1$. This property can be inferred from data on relative prices and quantities, and has nothing to do with the agent's intertemporal concern for smoothing consumption. Some papers refer to $u_{12} > 0$ [a positive cross derivative] as the case in which numeraire and shelter are 'substitutes' while $u_{12} < 0$ is the case in which these goods are "complements." We refrain from this language here, since the second derivative of the utility function captures both intertemporal and intratemporal tradeoffs."

There is, of course, no great difficulty that arises when different authors use different terminology as long as the terms are clearly defined. However, in explaining their criterion for distinguishing substitutes from complements, Piazzesi et al. seem to argue that the household's behavior regarding intratemporal substitution of goods and its behavior in terms of intertemporal substitution are distinct phenomena, governed by different parameters. This view – that preferences regarding intratemporal substitution are governed by a single parameter (σ or α, depending on the parameterization) and preferences regarding intertemporal substitution are governed by a different parameter (ρ) – is common in the literature.

In this paper, I argue that both the intratemporal substitutability of the two goods and the intertemporal substitutability of either of the goods depend on an *interaction* of both parameters; one cannot disentangle the two aspects of behavior in the way that Piazzesi et al. suggest. In this section I start by showing that the intratemporal behavior of two goods (C_1 and H_1) in equation (2)) is not determined by α [or σ] alone, but instead depends of the value of α *relative* to ρ. The interaction is a two-way street; in the next section, I argue that the dynamic, or intertemporal, behavior of nondurable consumption also depends on *both* α and ρ.

Consider a setting with no uncertainty. Households maximize utility (equation (2)) subject to a budget constraint:

---

[2] Piazzesi, et al (2007), p 537.

(4) $P_1^C C_1 + P_1^H H_1 + P_2^C C_2 + P_2^H H_2 = W$

Under the intertemporal interpretation of the problem, all four quantities are purchased simultaneously in period 1; the period two prices are the prices paid in period 1 for delivery of C_2 and H_2 in period 2.

Using λ to denote the Lagrange multiplier on the budget constraint, the four first order conditions are:

(5) $(C_1^\alpha + \gamma H_1^\alpha)^{\frac{1-\rho-\alpha}{\alpha}} C_1^{\alpha-1} \quad = P_1^C \lambda$

(6) $\gamma(C_1^\alpha + \gamma H_1^\alpha)^{\frac{1-\rho-\alpha}{\alpha}} H_1^{\alpha-1} \ = P_1^H \lambda$

(7) $\beta(C_2^\alpha + \gamma H_2^\alpha)^{\frac{1-\rho-\alpha}{\alpha}} C_2^{\alpha-1} \ = P_2^C \lambda$

(8) $\beta\gamma(C_2^\alpha + \gamma H_2^\alpha)^{\frac{1-\rho-\alpha}{\alpha}} H_2^{\alpha-1} = P_2^H \lambda$

Consider an increase in P_1^C , with the prices of the three other goods held constant, which will cause the household's demand for C_1 to fall. I am interested in investigating whether, in response to a change in the price of C_1, the consumption of H_1 moves in the same direction as C_1 (that is, falls), or moves in the opposite direction as C_1.

Before turning to the four-good system, consider the nature of intratemporal substitution in the one-period, two-good case, which can be represented by the utility function in equation (2) when $\beta = 0$. Eliminating the Lagrange multiplier by taking the ratio of the first order conditions in equations (5) and (6), and rearranging implies:

(9) $$\frac{C_1}{H_1} = \left(\frac{\gamma P_1^C}{P_1^H}\right)^{\frac{1}{\alpha-1}}$$

or, in logs,

(10) $$\ln\left(\frac{C_1}{H_1}\right) = \left(\frac{1}{\alpha-1}\right)\ln\gamma + \left(\frac{1}{\alpha-1}\right)\ln\left(\frac{P_1^C}{P_1^H}\right)$$

Differentiating with respect to P_1^C, and rearranging, gives:

$$(11) \quad \frac{\frac{\partial\left(\frac{C_1}{H_1}\right)}{\partial P_1^C}}{\frac{\left(\frac{C_1}{H_1}\right)}{P_1^C}} = \frac{1}{\alpha - 1} \qquad \text{(which equals} - \sigma \text{ in the parameterization } (1'))$$

Equation (11) tells us that the percent change in the *ratio*, or relative quantities, of C_1 and H_1 in response to a 1% change in P_1^C is $(\alpha - 1)^{-1}$. However, knowing that the ratio of C_1 to H_1 declines when P_1^C rises does not, by itself, establish that the quantity of H_1 consumed rises as the quantity C_1 falls. A decline in the ratio of C_1 to H_1 is also consistent with the quantities of the two goods moving in the same direction, as long as the magnitude of the decline in H_1 is smaller than the decline in C_1 .

To determine the effect of the change in price on the level of H_1, as opposed to the ratio of C_1 to H_1, it is necessary to introduce a restriction on the levels of the optimal quantities of the two goods – either a budget constraint or a level curve of the utility function. The algebra is simplified by using the level curve of the utility function and thus considering the income compensated demand system. Assuming that the household achieves the same level of utility before and after the price change, the effect of the change in P_1^C on the levels of C_1 to H_1 is constrained by:

$$(12) \quad 0 = \frac{\partial U}{\partial C_1}\frac{\partial C_1}{\partial P_1^C} + \frac{\partial U}{\partial H_1}\frac{\partial H_1}{\partial P_1^C}$$

which of course implies:

$$(13) \quad \frac{\partial H_1}{\partial P_1^C} = -\left[\frac{\frac{\partial U}{\partial C_1}}{\frac{\partial U}{\partial H_1}}\right]\frac{\partial C_1}{\partial P_1^C}$$

Since the marginal utility of each good is strictly positive, the term in square brackets is positive for all permissible values of the parameters of the utility function. Based on equation (13), we can conclude that the quantities of the two goods move in opposite directions in response to a change in P_1^C. Further, note that in deriving both the expression for the elasticity, equation (11), and the expression showing that the quantities of the two goods move in opposite directions, equation (13), the marginal utilities of the two goods appear only in the form of the ratio of marginal utilities. Since only the *ratio* of marginal utilities of the two goods matters, the parameter governing the curvature of the utility function with respect to the composite good (ρ) cancels out. That is, the effect of an increase in P_1^C on the optimal quantities of C_1 and H_1 depends only on the parameters that determine the aggregation of C_1 and H_1 into the composite good (α and γ) and is independent of the parameter governing the curvature of the utility function with respect to the composite good (ρ). In the one-period problem, the curvature parameter, ρ, would determine the household's degree of risk aversion in a setting with uncertainty, but plays no role in determining the intratemporal behavior of the two goods in response to a change in the relative price.

When we return to the 4-good case (i.e., the utility function in equation (2) with $0 < \beta < 1$), some but not all of the equations for the 2-good case continue to hold. The first order conditions for C_1 and H_1, equations (5) and (6), still generate equations (9), (10), and (11); thus we can conclude that the ratio of C_1 and H_1 falls in response to an increase in P_1^C, as before. However, equations (12) and (13) do not hold in the 4-good case, so it is not immediately obvious that the quantities of C_1 and H_1 move in opposite directions in response to a change in P_1^C, as before.

In the 4-good case, the assumption that the household achieves the same level of utility before and after the price change implies that

$$(14)\quad 0=\frac{\partial U}{\partial C_1}\frac{\partial C_1}{\partial P_1^C}+\frac{\partial U}{\partial H_1}\frac{\partial H_1}{\partial P_1^C}+\frac{\partial U}{\partial C_2}\frac{\partial C_2}{\partial P_1^C}+\frac{\partial U}{\partial H_2}\frac{\partial H_2}{\partial P_1^C}$$

which, compared to equation (12), includes two additional unknowns, $\frac{\partial C_2}{\partial P_1^C}$, and $\frac{\partial H_2}{\partial P_1^C}$. However, using the marginal conditions for C_2 and H_2 (equations (6) and (7)), we know that C_2 and H_2 are linked by:

$$(15)\quad \frac{H_2}{C_2}=\left(\frac{\gamma P_2^C}{P_2^H}\right)^{\frac{1}{1-\alpha}}\equiv\chi$$

Equation (15), in conjunction with the fact that P_2^C and P_2^H are held constant in this thought experiment, implies that the ratio of H_2 to C_2 will remain constant when P_1^C varies. Thus equation (14) can be written as a restriction on three unknowns:

$$(16)\quad 0=\frac{\partial U}{\partial C_1}\frac{\partial C_1}{\partial P_1^C}+\frac{\partial U}{\partial H_1}\frac{\partial H_1}{\partial P_1^C}+\left[\frac{\partial U}{\partial C_2}+\chi\frac{\partial U}{\partial H_2}\right]\frac{\partial C_2}{\partial P_1^C}$$

In order to determine whether the quantities of C_1 and H_1 move in the same direction, or the opposite direction, in response to a change in P_1^C, we need an additional (intertemporal) restriction that allows us to write $\frac{\partial C_2}{\partial P_1^C}$ in terms of $\frac{\partial C_1}{\partial P_1^C}$ and $\frac{\partial H_1}{\partial P_1^C}$ and obtain the analog of equation (13) for the 4-good case. Consider the ratio of the first order conditions for H_1 and H_2 (equations (6) and (8)), which implies:

$$(17)\quad P_2^H\frac{\partial u}{\partial H_1}=P_1^H\beta\frac{\partial u}{\partial H_2}$$

In equation (17), lower case "u" denotes the within-period utility function; that is, $U=u(C_1,H_1)+\beta u(C_2,H_2)$.

Differentiating equation (17) with respect to P_1^C , and using equation (15) gives:

$$(18)\quad P_2^H\left[\frac{\partial^2 u}{\partial H_1\partial C_1}\frac{\partial C_1}{\partial P_1^C}+\frac{\partial^2 u}{\partial H_1^2}\frac{\partial H_1}{\partial P_1^C}\right]=P_1^H\beta\left[\frac{\partial^2 u}{\partial H_2\partial C_2}+\chi\frac{\partial^2 u}{\partial H_2^2}\right]\frac{\partial C_2}{\partial P_1^C}$$

Dividing equation (18) by equation (17) implies:

(19) $$\left(\frac{\partial u}{\partial H_1}\right)^{-1}\left[\frac{\partial^2 u}{\partial H_1 \partial C_1}\frac{\partial C_1}{\partial P_1^C}+\frac{\partial^2 u}{\partial H_1^2}\frac{\partial H_1}{\partial P_1^C}\right]=\left(\frac{\partial u}{\partial H_2}\right)^{-1}\left[\frac{\partial^2 u}{\partial H_2 \partial C_2}+\chi\frac{\partial^2 u}{\partial H_2^2}\right]\frac{\partial C_2}{\partial P_1^C}$$

Assume that at the initial vector of prices was such that $C_1 = C_2 = C_0$, and likewise, $H_1 = H_2 = H_0$. Since all the derivatives of the utility function are evaluated at the common point (C_0, H_0), we have:

(20a) $$\frac{\partial^2 u}{\partial H_1 \partial C_1}=\frac{\partial^2 u}{\partial H_2 \partial C_2}$$

(20b) $$\frac{\partial^2 u}{\partial H_1^2}=\frac{\partial^2 u}{\partial H_2^2}$$

(20c) $$\frac{\partial u}{\partial H_1}=\frac{\partial u}{\partial H_2}$$

From equation (15), the ratio of H_2 to C_2 is equal to a constant denoted by χ. Thus for $C_2 = C_0$, $H_2 = H_0$,

(21) $$\frac{H_0}{C_0}=\left(\frac{\gamma P_2^C}{P_2^H}\right)^{\frac{1}{1-\alpha}}\equiv\chi$$

In addition, multiply both sides of equation (19) by $\frac{\partial u}{\partial H_1}C_0$, and solve for $\frac{\partial C_2}{\partial P_1^C}$ to get:

(22) $$\frac{\partial C_2}{\partial P_1^C}=\left[\frac{\frac{\partial^2 u}{\partial H \partial C}C_0}{\frac{\partial^2 u}{\partial H \partial C}C_0+\frac{\partial^2 u}{\partial H^2}H_0}\right]\frac{\partial C_1}{\partial P_1^C}+\left[\frac{\frac{\partial^2 u}{\partial H^2}C_0}{\frac{\partial^2 u}{\partial H \partial C}C_0+\frac{\partial^2 u}{\partial H^2}H_0}\right]\frac{\partial H_1}{\partial P_1^C}$$

In equation (22), $\frac{\partial^2 u}{\partial H \partial C}$ is notation used to reflect the common value of $\frac{\partial^2 u}{\partial H_1 \partial C_1}$ and $\frac{\partial^2 u}{\partial H_2 \partial C_2}$ when evaluated at $C_1 = C_2 = C_0$ and $H_1 = H_2 = H_0$. Similarly, $\frac{\partial^2 u}{\partial H^2}$ is notation used to reflect the common value of $\frac{\partial^2 u}{\partial H_1^2}$ and $\frac{\partial^2 u}{\partial H_2^2}$.

Consider equation (16), which describes the joint behavior of C_1, H_1, and C_2 on a level curve of the utility function. After substituting out the constant χ, we get:

$$(23) \quad 0 = \frac{\partial u}{\partial C_1}\frac{\partial C_1}{\partial P_1^C} + \frac{\partial u}{\partial H_1}\frac{\partial H_1}{\partial P_1^C} + \beta\left[\frac{\partial u}{\partial C_2} + \frac{H_o}{C_0}\frac{\partial u}{\partial H_2}\right]\frac{\partial C_2}{\partial P_1^C}$$

Using equation (22) to eliminate $\frac{\partial C_2}{\partial P_1^C}$, and simplifying, we finally obtain an answer to the question of whether the quantity H_1 moves in the same direction, or in the opposite direction, as the quantity C_1 in response to a change in the price of C_1:

$$(24) \qquad \frac{\partial H_1}{\partial P_1^C} = -\left[\frac{\frac{\partial u}{\partial C_1}}{\frac{\partial u}{\partial H_1}}\right]\frac{\left[(1+\beta) - \beta\left(\frac{1-\alpha}{\rho}\right)\right]}{\left[(1+\beta) + \beta\left[\frac{\frac{\partial u}{\partial C_1}}{\frac{\partial u}{\partial H_1}}\frac{C_0}{H_0}\right]\left(\frac{1-\alpha}{\rho}\right)\right]}\frac{\partial C_1}{\partial P_1^C}$$

By setting the discount factor, β , to zero, equation (24) can be specialized to reflect the atemporal, two-good case. With $\beta = 0$, equation (24) coincides with equation (13) and indicates that a change in the price of C_1 causes the quantities of H_1 and C_1 to move in opposite directions, for all allowable values of the parameters of the utility function; if the set of goods is limited to just two goods, the two goods must be substitutes.

However, in the more general, four-good case (i.e., for $0 < \beta \leq 1$), equation (24) indicates that the quantities of H_1 and C_1 can either move in opposite directions or in the same direction, depending on the values of the parameters. Both the ratio of marginal utilities and the denominator must be positive for all allowable parameter values. Therefore, it is the sign of the expression in square brackets in the numerator that determines whether H_1 rises, falls, or remains unchanged when the price of C_1 rises. The quantities demanded of the two goods will move in opposite directions (that is, act like substitutes) if $\frac{1+\beta}{\beta} > \frac{1-\alpha}{\rho}$, but move in the same direction (that is, act like complements) if the inequality goes the other way. Since the permissible range of α is $-\infty \leq \alpha \leq 1$, this tells us that the two goods will act as substitutes when α is in the upper part of the range, $1-\left(\frac{1+\beta}{\beta}\right)\rho < \alpha \le 1$, and like complements when α is in the lower part of the range, $-\infty \le \alpha < 1-\left(\frac{1+\beta}{\beta}\right)\rho$.

To summarize, let's return to the two-period, two good utility function in equation (2), and assume that there is no uncertainty, either in terms of the budget constraint, given by W, or in the prices of the goods. The price of C_1 rises, while the other prices remain constant. We know immediately that the quantity demanded of C_1 falls, as C_1 is now more expensive relative to the three other goods. We also know that the *ratio* of C_1 to H_1 falls, and that the magnitude of the effect on the ratio is determined solely by the parameter α. However, the magnitude of the fall in quantity demanded of C_1, and both the direction and magnitude of the effect on the quantity demanded of H_1 depends on the value of α *relative* to the value of ρ .

To understand the intuition behind the result, consider the schematic in Figure 1.

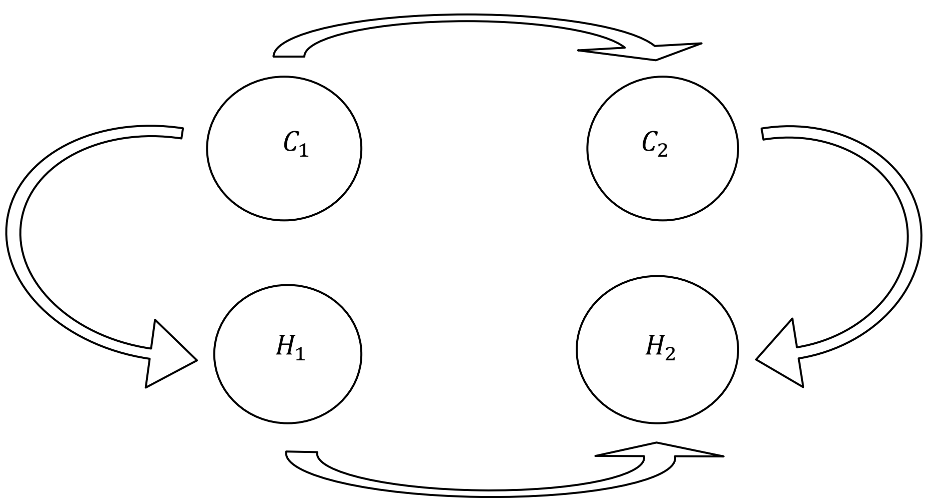

Figure 1: Low degree of curvature with respect to intertemporal reallocation;
high degree of curvature with respect to intratemporal reallocation

The horizontal arrows representing reallocation of consumption across time are drawn fairly flat (that is, with a low degree of curvature) to represent the assumption that the curvature of the utility function with respect to the composite good (as determined by a value of ρ) is small. In contrast, the “highly curved” vertical arrows represent the assumption that the curvature of the utility function with respect to a change in the composition of the composite good (as determined by the value of α) is large. When intratemporal reallocation is difficult *relative* to intertemporal reallocation, as depicted in Figure 1, the values of α and ρ would satisfy

$$-\infty \leq \alpha < 1 - \left(\frac{1+\beta}{\beta}\right)\rho \; .$$

As P_1^C rises (holding constant the prices of the other goods), the household will reduce consumption of C_1 and increase consumption of some or all of the other goods. Within period 2, the ratio $\frac{C_2}{H_2}$ is unchanged. Within period 1, the ratio $\frac{C_1}{H_1}$ falls as a result of the increase in $\frac{P_1^C}{P_1^H}$. However, if the curvature of the utility function with respect to the composition of the composite good is high, as assumed in Figure 1, the magnitude of the increase in the ratio $\frac{C_1}{H_1}$ is small. In an intertemporal context, the household responds to the increase in P_1^C not only by changing the composition of the composite good, but also by increasing consumption of both goods in period 2. If, as depicted in Figure 1, the curvature of the utility function with respect to the composite good is low, the increase in P_1^C will induce a large increase in the consumption of C_2 . Since the second period price ratio, $\frac{P_2^C}{P_2^H}$, has not changed, the second period ratio of quantities $\frac{C_2}{H_2}$ is unchanged. Therefore the increase in P_1^C induces a large increase in H_2 as well as C_2 .

When intratemporal substitution is “difficult” (shorthand for a high degree of curvature of the utility function with respect to the composition of the composite good), the ratio of quantities within period 1, $\frac{C_1}{H_1}$, falls only slightly when P_1^C rises. If at the same time, intertemporal substitution is “easy” (shorthand for a low degree of curvature of the utility function with respect to the composite good), the optimal response of the household is to substantially increase consumption of the second period composite good (that is, *both* H_2 and C_2) while reducing consumption of the first period composite good. Thus while the intratemporal substitution effect causes the consumption of H_1 to increase in response to an increase in P_1^C, the intertemporal substitution effect works in the opposite direction, inducing a reduction in H_1 (and C_1) as consumption of H_2 and C_2 increase. When intratemporal substitution is “difficult” and intertemporal substitution is “easy”, intertemporal effect dominates; both C_1 and H_1 will decrease in response to an increase in P_1^C.

The diagram in Figure 2 depicts a different configuration of parameters; intertemporal reallocation is “difficult”, while intratemporal reallocation is “easy”.

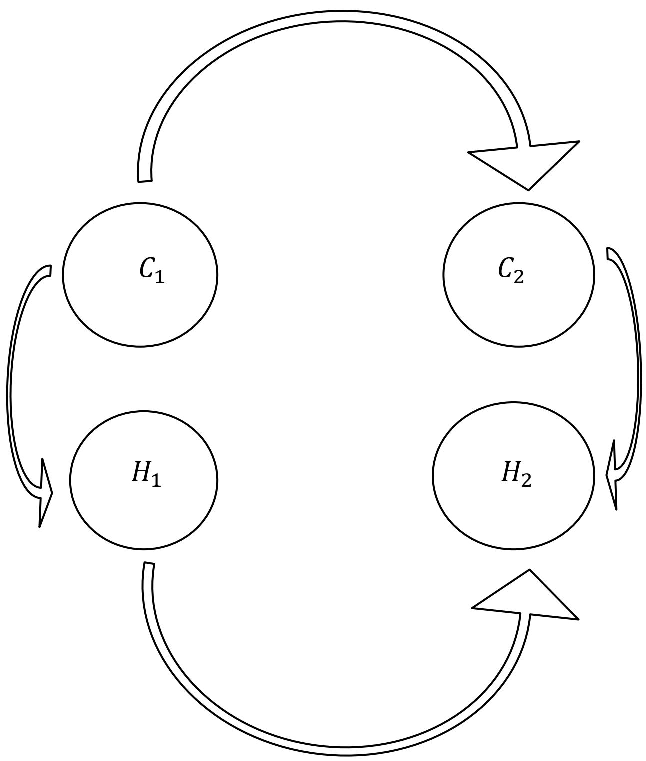

Figure 2: High degree of curvature with respect to intertemporal reallocation; low degree of curvature with respect to intratemporal reallocation

As before, the ratio $\frac{C_1}{H_1}$ falls as a result of the increase in $\frac{P_1^C}{P_1^H}$. However, because the curvature of the utility function with respect to the composition of the composite good is low (intratemporal substitution is “easy”, the magnitude of the decrease in the ratio $\frac{C_1}{H_1}$ is large. At the same time, the magnitude of the intertemporal reallocation from H_1 and C_1 to C_2 and H_2 is sharply limited by the “difficulty” in intertemporal substitution. As before, the intratemporal substitution effect and the intertemporal substitution effect have opposing effects on the consumption of H_1. For a configuration of parameters as depicted in Figure 2 (“easy” intratemporal substitution and “difficult” intertemporal substitution), however, the intratemporal effect dominates and an increase in P_1^C will cause the consumption of H_1 to increase.

Intuitively, one could think of two goods as complements if, in response to an increase in the price of one good, the quantity demanded of both goods declines (that is, the consumption of both goods moves in the same direction). According to this approach, two goods would be substitutes is an increase in the price of one good causes the quantity demanded of the other good to increase (that is, the consumption of the two goods moves in opposite directions). In an atemporal, or one-period, setting with two goods (utility function (2) with $\beta = 0$), the two goods must be substitutes with the magnitude of the increase in H_1 in response to an increase in P_1^C determined solely by α. In an intertemporal setting, in contrast, whether two goods in a given period behave as complements, with the quantities of both goods moving in the same direction in response to an increase in P_1^C, or as substitutes, with the quantities of the two goods moving in opposite directions, does not depend on the value of α alone, but instead depends on the value of α *relative* to the value of ρ.

Section 3: Dynamics of nondurable consumption

Assuming that nondurable consumption is costlessly adjustable, an optimizing household will continuously equate the marginal utility of nondurable consumption to the marginal utility of wealth. In a completely specified model, the marginal utility of wealth is, of course, determined endogenously as a result of the household's optimization decisions in response to exogenous driving variables such as asset returns, wage rates, and labor market shocks. In this section, I simply assume a stochastic process for the marginal utility of wealth, and calculate the optimal level of nondurable consumption required to continuously equate the marginal utility of nondurable consumption with the marginal utility of wealth. While this approach obviously does not provide a complete solution to the household's problem, it provides a useful characterization of the influence of intratemporal substitutability on consumption dynamics.

Using μ to denote the marginal utility of nondurable consumption, a first order Taylor series approximation can be used to relate marginal utility in t+1 to marginal utility in t:

$$(25) \qquad \mu_{t+1} = \mu_t + \frac{\partial^2 u}{\partial C_t^2}\Delta C_{t+1} + \frac{\partial^2 u}{\partial C_t \partial H_t}\Delta H_{t+1}$$

where u is the notation for the one-period utility function, as before. For any time period (t, t+1) during which the level of housing consumption does *not* change, the last term is equal to zero, and we can solve for the change in nondurable consumption as a function of the change in the marginal utility:

$$(26) \qquad \Delta C_{t+1} = \left[\frac{\partial^2 u}{\partial C_t^2}\right]^{-1} \Delta\mu_{t+1}$$

Dividing both sides by C_t gives

(27) $$\frac{\Delta C_{t+1}}{C_t} = \left[\frac{\frac{\partial^2 u}{\partial C_t^2}\, C_t}{\frac{\partial u}{\partial C_t}}\right]^{-1} \frac{\Delta \mu_{t+1}}{\mu_t}$$

or

(28) $$\Delta \ln C_{t+1} = \left[\frac{\frac{\partial^2 u}{\partial C_t^2}\, C_t}{\frac{\partial u}{\partial C_t}}\right]^{-1} \Delta \ln \mu_{t+1}$$

As in the standard model, the growth rate of consumption is related to the growth rate of marginal utility by the inverse of the curvature of the utility function. That is, if we abstract from housing for a moment, and think of utility as depending on just nondurable consumption, for example,

(29) $$\tilde{u} = \frac{C_t^{1-\rho}}{1-\rho}$$

equation (26) becomes the familiar

(30) $$\Delta \ln C_{t+1} = \frac{-1}{\rho}\, \Delta \ln \mu_{t+1}$$

since the curvature of the power utility function in equation (29) is ρ . The greater the curvature of the utility function with respect to consumption, the smaller is the growth rate of consumption for a given percentage change in the marginal utility of wealth.

When adding housing as an argument of the utility function, the natural generalization of the standard one-good power utility function is:

(31) $$u(C_t, H_t) = \frac{\left(C_t^{\alpha} + \gamma H_t^{\alpha}\right)^{\frac{1-\rho}{\alpha}}}{1-\rho}$$

For the utility function in equation (31), the curvature of the utility function with respect to the composite good is still given by the parameter ρ, while the aggregation of the two physical goods into the composite good is controlled by the parameter α.

Suppose that *both* goods are frictionlessly adjustable, and, further, that the relative price of housing in terms of numeraire (nondurable consumption) is constant and denoted by P. Using the intratemporal first order condition, the utility maximizing quantity of housing will be proportion to the quantity of nondurable consumption:

$$(32) \qquad H_t = \left(\frac{P}{\gamma}\right)^{\frac{1}{\alpha-1}} C_t$$

which allows us to substitute out housing and write utility in terms of nondurable consumption and the relative price:

$$(33) \qquad u(C_t, H_t) = \frac{\left\{\left[1 + \gamma\left(\frac{P}{\gamma}\right)^{\frac{\alpha}{\alpha-1}}\right] C_t^{\alpha}\right\}^{\frac{1-\rho}{\alpha}}}{1-\rho} = \left[1 + \gamma\left(\frac{P}{\gamma}\right)^{\frac{\alpha}{\alpha-1}}\right]^{\frac{1-\rho}{\alpha}} \frac{C_t^{1-\rho}}{1-\rho}$$

From equation (33), note that if both goods are frictionlessly adjustable with a constant relative price, the curvature of the utility function with respect to nondurable consumption is equal to ρ, the curvature of the utility function with respect to the composite good. Thus equation (30) describes the growth rate of nondurable consumption in both the standard one-good case (that is, with utility given by equation (29) and in the two good case (utility given by equation (31) when both goods are frictionlessly adjustable. Further, the dynamics of nondurable consumption depend only on the parameter governing the curvature of the utility function with respect to the composite good, and does not depend on α, the parameter that determines the aggregation of the two goods into the composite good.

However, if we assume that housing is subject to a nonconvex adjustment cost (while nondurable consumption remains costlessly adjustable), the solution to the household's lifetime utility maximization problem will result in infrequent, large adjustments in the quantity of housing, and continuous, small changes in nondurable consumption. The intratemporal marginal condition relating the marginal utilities of the two goods to the relative price will no longer hold. As long as nondurable consumption is costlessly adjustable, however, the optimizing household will continue to equate the marginal utility of nondurable consumption to the marginal utility of wealth. Thus equation (28) provides a valid characterization of the dynamics of nondurable consumption. However, for any interval within which the quantity of housing remains constant (due to the adjustment cost), the curvature of the utility function with respect to nondurable consumption is calculating by taking partial derivatives, holding H constant.

For the two-good utility function in equation (31), calculating the inverse of the curvature of the utility function with respect to nondurable consumption, and simplifying, yields

$$(34) \qquad \Delta \ln C_{t+1} = \frac{-1}{(1-\kappa_t)(1-\alpha)+\kappa_t \rho} \Delta \ln \mu_{t+1} \quad \text{for} \quad \kappa_t = \frac{C_t^{\alpha}}{C_t^{\alpha} + \gamma H_t^{\alpha}}$$

Note that if we eliminate housing (by setting $\gamma = 0$), then $\kappa_t = 1$ and equation (34) coincides with equation (30). In the general case, though, κ_t obeys the restriction $0 \leq \kappa_t \leq 1$.

Using equation (34), we can now illustrate the dynamics of nondurable consumption under the plausible assumption that housing is subject to a nonconvex adjustment cost. The stochastic process for the log of the marginal utility of wealth is assumed to be

$$(35) \qquad \Delta \ln \mu_{t+1} = -\delta + \sigma_{\epsilon} \epsilon_t$$

where ϵ_t is independently and identically distributed standard normal. The drift in the marginal utility of wealth is assumed to be negative (reflecting a positive drift in wealth itself; σ_ϵ is a scale parameter).

In the figures that follow, housing is assumed constant at the level H=5 for the first 20 periods. In period 21, the level of housing changes to H=8 and remains constant at that level for the next 20 periods. Nondurable consumption is assumed to be equal to 5 units in period 1. Interpreting the time periods as quarters, the value of the drift parameter, δ, is set to .005; the scale parameter is set at $\sigma_\epsilon = .02$. A sample path for the marginal utility of wealth is generated by using 40 draws from the standard normal distribution as inputs in equation (35). Given the assumed path of housing, the initial value of nondurable consumption in period 1, and the sample path for the marginal utility of wealth, equation (34) is then used to generate the path of nondurable consumption over the 40 periods. In this "comparative dynamics" exercise, the value of ρ is held constant in order to isolate the effect of the parameter α on the stochastic behavior of nondurable consumption. In particular, ρ is held constant at the moderate value of 2.

Having fixed the value of ρ at $\rho = 2$, it seems natural to start with the assumption that $\alpha = -1$, as the utility function is separable in this case. That is, when $\alpha = 1 - \rho$, the utility function becomes

$$(36) \qquad u(C_t, H_t) = \frac{C_t^{1-\rho} + \gamma H_t^{1-\rho}}{1-\rho}$$

and the marginal utility of nondurable consumption no longer depends on the level of housing consumption. Given the separability of the utility function when $\alpha = -1$, the path of nondurable consumption coincides with the path predicted by the standard, one-good model.

Figure 3: Dynamics of nondurable consumption, separable case: α = 1 - ρ = -1

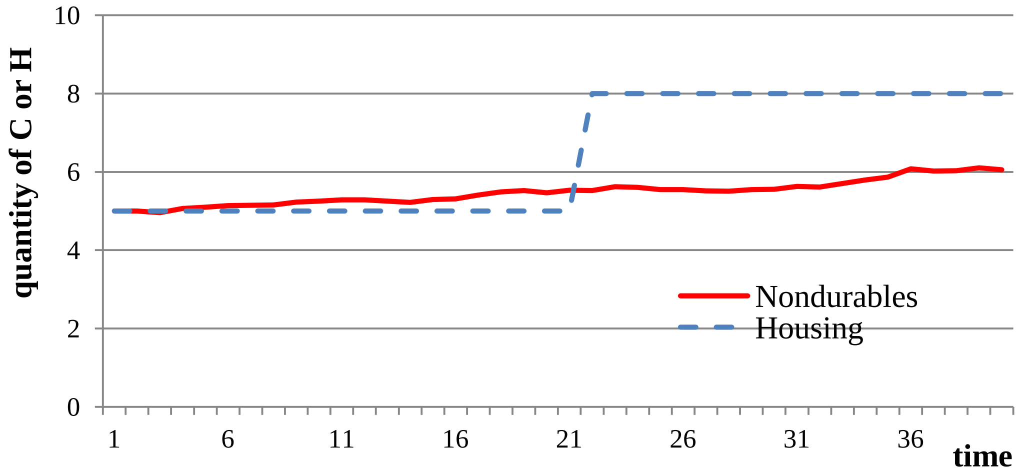

For simplicity, assume that the units of measurement of the two goods are such that the relative price of housing is unity. In this case, we can consider total consumption by simply adding the quantities of the two goods. While total consumption is not plotted in Figure 3, it is clear that the discontinuity in housing consumption would be translated directly into total consumption.

Next, consider the utility function when $\alpha = 1$. In this case the utility function is simply

$$(34) \qquad u(C_t, H_t) = \frac{\left(C_t + \gamma H_t\right)^{1-\rho}}{1-\rho}$$

and utility depends solely on the weighted sum of the two goods. While the path of nondurable consumption prior to, and after, the discrete change in H at period 21 is qualitatively similar to the path in the separable case, note that it is twice as volatile. When the two goods are perfect substitutes, the standard deviation of the growth rate of *total* (weighted) consumption will be equal to the standard deviation of the growth rate of total consumption when both goods are frictionless adjustable. However, since the quantity of housing is constant over the first 20 periods due to the adjustment cost, the volatility of the growth rate of nondurable consumption must be larger since all of the variation in total consumption comes from nondurable consumption alone.

Figure 4: Dynamics of nondurable consumption,
C and H are perfect substitutes:
α = 1 > 1 - ρ = -1

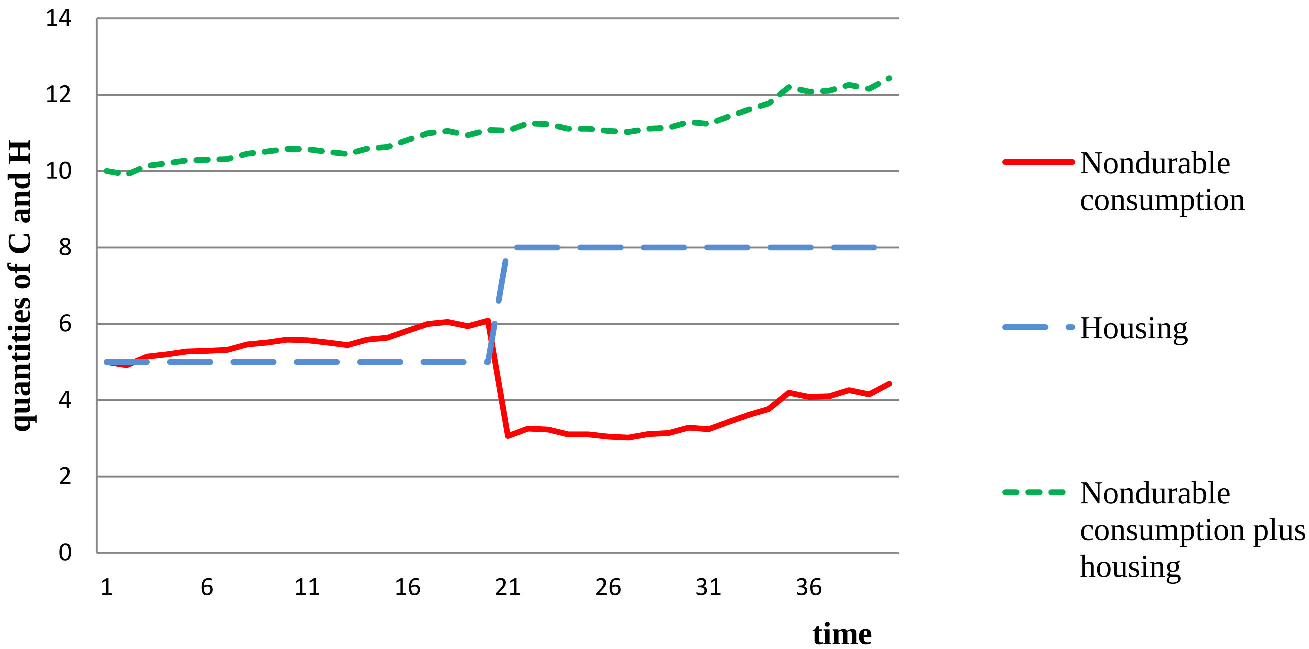

More dramatic is the discontinuity in the path of nondurable consumption that arises when the two goods are perfect substitutes. When $\alpha = 1$, for as long as the level of housing consumption is constant at the original level, the household responds to the stochastic process on the marginal utility of wealth by varying nondurable consumption, but when the level of housing consumption jumps, nondurable consumption exhibits a jump equal in magnitude but in the opposite direction. Thus when the two goods are perfect substitutes, total consumption (C+H) follows the path implied by the standard frictionless model, even though neither the path of nondurable consumption nor that of housing consumption is consistent with the standard frictionless model.

Next, consider a value of $\alpha = .5$. While the two goods are not perfect substitutes, for this value of α the curvature of the utility function with respect to changes in the composition of the composite good is small. Compared to the previous case (of perfect substitutes), the volatility of nondurable consumption is slightly smaller within the intervals during which the quantity of housing is fixed. During the first 20 periods, nondurable consumption varies in order to keep the marginal utility of nondurable consumption equal to the marginal utility of wealth. However, in contrast to the case in which the two goods are perfect substitutes, marginal utility declines as nondurable consumption becomes a larger fraction of total consumption, and this response of the marginal utility of consumption to the composition of the composite good implies that nondurable consumption varies less, over the first 20 periods, than in the previous case of perfect substitutability.

Figure 5: Dynamics of nondurable consumption,
$\alpha = .5 > 1 - \rho = -1$

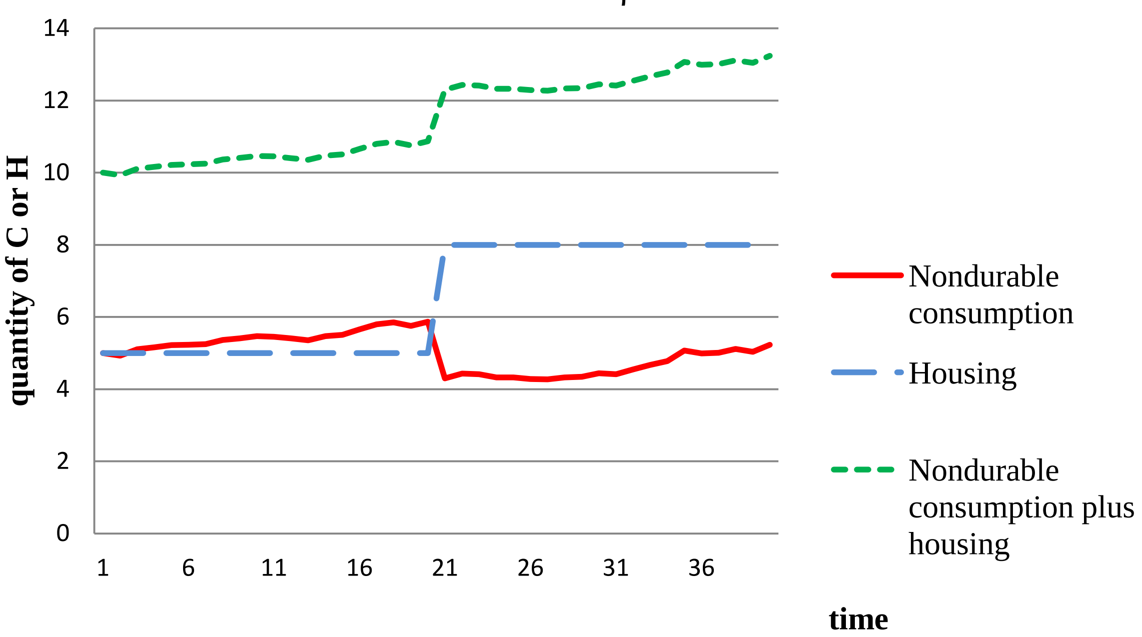

When the quantity of housing moves discontinuously in period 21, note that while nondurable consumption moves in the opposite direction, as in the previous case, the magnitude of the discontinuity (or “jump”) in nondurable consumption is now smaller than the contemporaneous jump in the quantity of housing. Since the jump in housing consumption is only partially offset by an opposite direction jump in nondurable consumption, the discontinuity appears in total consumption as well.

In Figure 4 (for α = 1) and Figure 5 (for α = .5), intratemporal substitution is “easy” *relative* to intertemporal substitution in the sense that, given the assumed value of ρ = 2, the two parameters satisfy the inequality: $\alpha > 1 - \rho$. When intratemporal substitution is “easy” relative to intertemporal substitution, the discontinuous jump in nondurable consumption is in the opposite direction of the jump in housing consumption. In Figures 6 and 7, intratemporal substitution is “difficult” relative to intertemporal substitution in the sense that α and ρ are now related by $\alpha < 1 - \rho$. Note that the consequence of intratemporal substitution being difficult relative to intertemporal substitution is that the jump in nondurable consumption is now in the same direction as the jump in housing consumption. While the figures themselves are simply illustrative, it is easy to show analytically that the sign of the covariance of the jumps in the two goods depends on the value of α relative to ρ by considering equation (25). Rearranging slightly, equation (25) implies:

$$(35) \qquad C_{t+1} = \frac{\frac{-\partial^2 u}{\partial C_t \partial H_t}}{\frac{\partial^2 u}{\partial c^2}} \Delta H_{t+1} + \frac{1}{\frac{\partial^2 u}{\partial c^2}} \Delta \mu_{t+1}$$

For the assumed utility function, the coefficient on ΔH_{t+1} is:

$$(36) \qquad \frac{\frac{-\partial^2 u}{\partial C_t \partial H_t}}{\frac{\partial^2 u}{\partial c^2}} = \left[\frac{1 - \rho - \alpha}{\frac{\rho}{\gamma}\left(\frac{C}{H}\right)^{\alpha} + (1 - \alpha)} \right] \frac{C}{H}$$

For all allowable values of the parameters, the coefficient of ΔH_{t+1} is: zero if $\alpha = 1 - \rho$, negative if $\alpha > 1 - \rho$, and positive if $\alpha < 1 - \rho$.

Figure 6: Dynamics of nondurable consumption,
α = -2 < 1 - ρ = -1

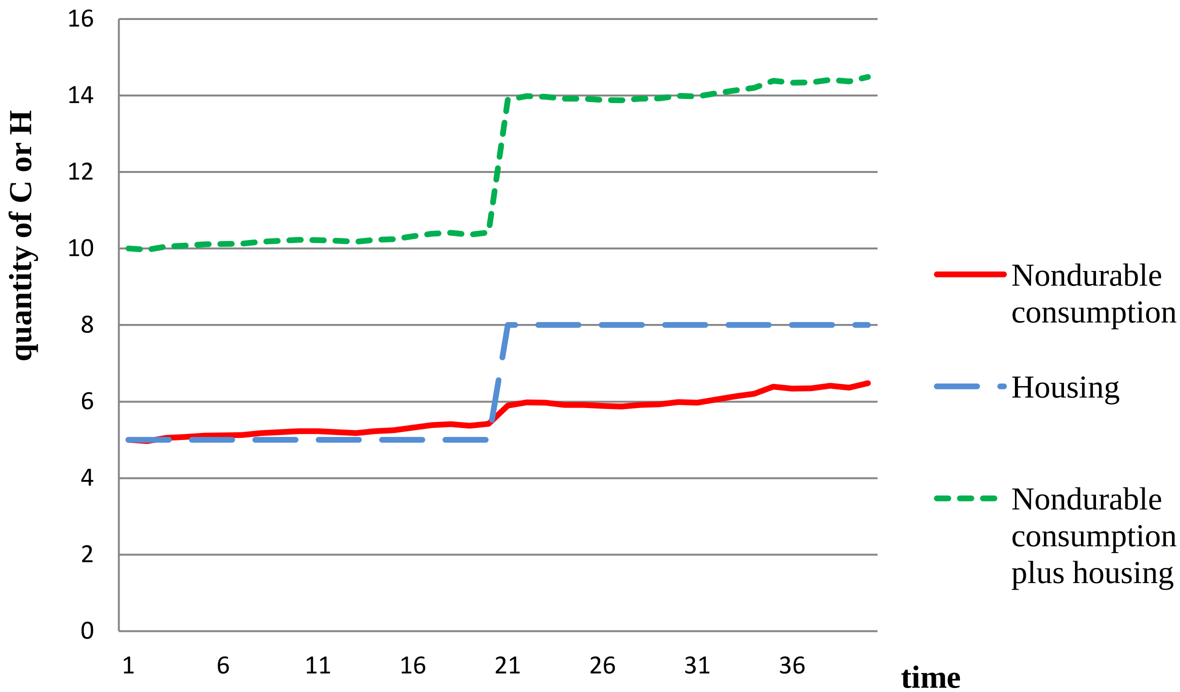

For α = -2, the nondurable consumption exhibits a small jump in the same direction as the discontinuity in housing consumption in period 21. In this case intratemporal substitution is “difficult” relative to intertemporal substitutution, in the sense that -2 = $\alpha < 1 - \rho <$ -1 ,but because it is only slightly more difficult, the behavior of nondurable consumption in Figure 6 differs only slightly from the path of nondurable consumption in Figure 3, which illustrates the separable case.

Figure 7 illusrates the case in which α = -6; intratemporal substitution is “very difficult” relative to intertemporal substitution in the sense that α is substantially smaller than $1 - \rho$. Note that the magnitude of the jump in nondurable consumption between periods 20 and 21 is large, and the response of nondurable consumption to variation in the marginal utility of wealth

Figure 7: Dynamics of nondurable consumption,
α = -6 < 1 - ρ = -1

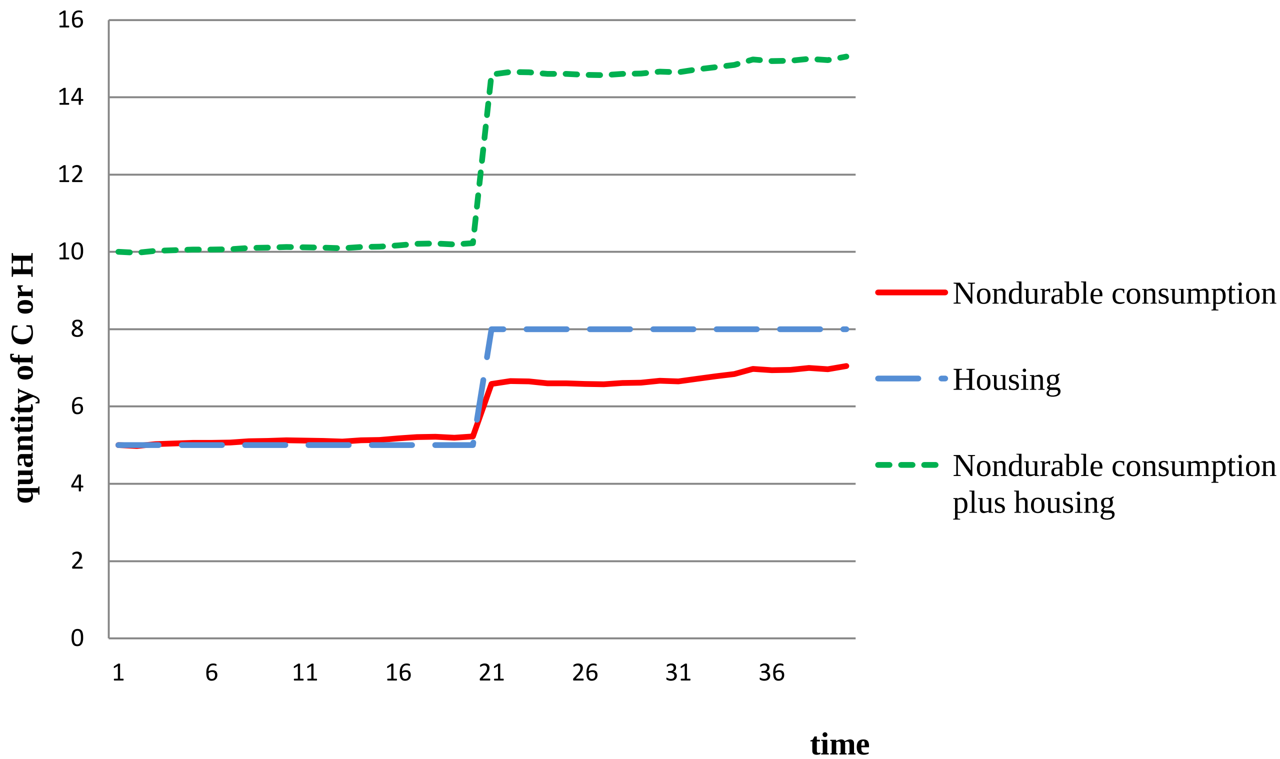

within periods 1 to 20 and 21 to 40 is correspondingly small. While housing is assumed to be held constant at H = 5 units for the first 20 periods, the household's optimal choice of nondurable consumption varies only slightly within the first 20 periods because the vale of α of -6 imples that the marginal utility of the composite good falls off very dramatically as the composition of the composite good changes. As wealth drifts upward within the first 20 periods, even though the household is free to expand total consumption by altering its consumption of nondurables, the optimal decision is to increase consumption of nondurables only very slightly. That is, within the first 20 periods, almost all of the stochastic increments to wealth are saved. When the jump in housing consumption occurs between period 20 and 21, nondurable consumption follows with a jump of the same direction and similar magnitude. In effect, the household is intertemporally reallocating nonduruable consumption from the first 20 periods to the second 20 periods, compared to the allocations that result when the utility function is separable ($\alpha = 1 - \rho$ as in Figure 3) or, more dramatically, when the two goods are perfect substitutes ($\alpha = 1$ as in Figure 4).

Section 4: Conclusions

In the simplest intertemporal setting -- that is, when utility depends on a single, frictionlessly adjustable consumption good -- the household's willingness to substitute consumption intertemporally is solely determined by the parameter governing the curvature of the utility function. When the basic utility specification is then generalized from one good to two, as in the common utility function given by equation (1), however, the curvature parameter, ρ , now specifies the curvature of the utility function with respect to a composite good. If we assume that *both* goods are frictionlessly adjustable, then the composite good and each of the individual goods will all have the same dynamics, and the intuition developed from the one-good case -- that the household's willingness to substitute intertemporally is solely determined by the curvature parameter, ρ , remains valid. In this paper, however, the two goods are interpreted as housing services, H, and non-housing goods (or nondurable goods), C. Any frictions or costs involved in changing the level of nondurable consumption may be sufficiently small that abstracting from these frictions is a valid modeling choice. While retaining the conventional assumption of frictionless adjustment of nondurable consumption, a basic premise of the paper is that a substantial adjustment cost – in particular a nonconvex, or lumpy, adjustment cost – is an essential feature of housing services. In a fully specified model of the lifetime optimization problem, the household will endogenously choose 1) the timing of the infrequent changes in the quantity of housing services consumed, 2) the level of housing services consumed, and 3) the level of nondurable consumption. During an interval within which the household chooses to leave the level of housing services unchanged, the level of nondurable consumption will be determined by the (optimally) predetermined level of housing consumption and the marginal utility of wealth. In this paper, I numerically solve the second part of the problem – finding the optimal level of nondurable consumption, conditional on an assumed exogenous path for housing consumption and for the marginal utility of wealth. While the assumed path of housing is not generated endogenously from a complete specification of the household's lifetime optimization problem, it captures the essential feature that emerges from a complete specification of the problem: long periods of during which housing consumption is constant, and large, infrequent changes in the level of housing services.

The presence of a nonconvex adjustment cost on housing has a dramatic effect on the dynamics of nondurable consumption. Further, the numerical results indicate that dynamic, or intertemporal behavior of nondurable consumption depends crucially on the parameter governing the intratemporal substitutability of the two goods (α in my parameterization) as well as the parameter governing the curvature of the utility function with respect to the composite good (ρ). Holding constant the variance of the growth rate of the marginal utility of wealth, the paper demonstrates that the variance of the growth rate of nondurable consumption depends on not only the curvature parameter, ρ , but also on the parameter governing intratemporal substitutability. Further, the influence of the relative ease of intratemporal substitution on the dynamic, or intertemporal behavior of nondurable consumption is not a one way street. The paper also argues that the intratemporal behavior of the two goods is determined by the ease of intratemporal substitution *relative* to the ease of intertemporal substitution, rather than on the intratemporal substitution parameter alone.

References

Flavin, Marjorie, and Shinobu Nakagawa, 2008, "A Model of Housing in the Presence of Adjustment Costs: A Structural Interpretation of Habit Persistence", *American Economic Review*, 98(1): 474-495.

Flavin, Marjorie, and Takashi Yamashita, 2010, "Owner-Occupied Housing: Life-Cycle Implications of the Household Portfolio", working paper, UCSD.

Grossman, Sanford, and Guy Laroque, 1990, "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods", *Econometrica*, 58(1): 25-51.

Hall, Robert E., 1988, "Intertemporal Substitution in Consumption", *Journal of Political Economy*, 96(2): 339-57.

Piazzesi, Monika, Martin Schneider, and Selale Tuzel, 2007, "Housing, Consumption, and Asset Pricing", *Journal of Financial Economics*, 83(3): 531-69.