Long-Run Labor Supply and the Elasticity of Intertemporal Substitution for Consumption

Susanto Basu

Miles S. Kimball

University of Michigan

December 2000

Abstract: The fact that permanent increases in the real wage have very little effect on labor supply implies a parameter restriction in the consumption Euler equation augmented by predictable movements in the quantity of labor. This parameter restriction is not rejected by aggregate U.S. data. The implied estimate of the elasticity of intertemporal substitution is around .35, and is significantly different from zero. This estimate is robust to different instrument sets and normalizations. After accounting for the effects of predictable movements in labor implied by the restriction, there is no remaining evidence in aggregate U.S. data of excess sensitivity of consumption to current income.

We would like to thank Tom Bishop for able research assistance. Susanto Basu's work on this paper was supported by the National Science Foundation. Miles Kimball's work on this paper was supported by National Institute of Aging Program Project grant 2PO1-AG10179.

KEYWORDS: Consumption, Euler equation, Labor Supply, King-Plosser-Rebelo utility function.

I. Introduction

Traditional theology wrestles with the inconsistency of three propositions that are mutually inconsistent: (A) God is all powerful; (B) God is all good; (C) evil is real. A logically consistent theology must qualify at least one of these propositions. Consumption theory poses a significantly less important conundrum with the same logical structure. Currently, we suspect there are many economists thinking about consumption who would like to maintain simultaneously that (A) consumption and labor are additively separable in an additively time-separable utility function, (B) the elasticity of intertemporal substitution for consumption is relatively low—well below 1, and (C) labor supply is not totally inelastic (that is, income and substitution effects are not both zero), but income and substitution effects on labor supply cancel, so that a permanent increase in the real wage will have little effect on the labor supply of a household that relies entirely on labor income.

To see the logical inconsistency, assume additive separability of consumption and labor from each other and additive separability across time. For much of consumption theory this has long been the default assumption. Usually additive separability between consumption and labor is an implicit assumption made by omitting labor from the analysis as anything other than a source of income. Given (A) the assumption of additive separability between consumption and labor and across time, (B) empirical estimates of the elasticity of intertemporal substitution have repeatedly found quite low values. To be specific, as recommended by Hall (1988), consider the IV estimation of the equation

$$\Delta \ln(C_t) = s(r_t - \rho) + \epsilon_t + \theta\epsilon_{t-1} \qquad (1)$$

using as instruments appropriately lagged variables that should be uncorrelated with the time-averaged rational expectations error $\epsilon_t + \theta\epsilon_{t-1}$. C is consumption, r is the real interest rate, ρ is the utility discount rate. The parameter s is the elasticity of intertemporal substitution for consumption. Hall (1988) gets point estimates of s equal to .1 or .2 that are not significantly different from zero statistically. With the same maintained assumptions of additive separability, but identifying s by asking respondents to choose among hypothetical consumption paths rather than by observed responses to actual interest rate fluctuations, Barsky, Kimball, Juster and Shapiro (1995) find a value for the elasticity of intertemporal substitution s in the same range as Hall (1988). Let us take these results at face value and choose $s = .2$, which on the high end of what can be justified by either Hall (1988) or Barsky, Kimball, Juster and Shapiro (1995). With $s = .2$, and the maintained assumption of additive separability between consumption and labor, the implied period utility function would be of the form $u(C,N) = -\frac{1}{4C^4} - v(N)$, where $v(N)$ is a convex function of N. The implied real consumption wage is $\frac{W}{P_C} = -\frac{u_N(C,N)}{u_C(C,N)} = C^5 v'(N)$. Per capita consumption C has roughly doubled in the 35 years since 1960 (a growth rate of approximately 2% per year). The average number of work hours per person N has stayed fairly constant—or if anything has slightly increased over that period (which in interaction with the convexity of $v(N)$ would slightly increase $v'(N)$). Thus, this functional form implies, counterfactually, that the real consumption wage should have increased by a factor of $2^5 = 32$ over that time period! Even if s were as high as .333, this kind of exercise would imply an eight-fold increase in the real wage.

An alternative way to state the same problem is that with consumption and the real wage both roughly doubling over this time period, the swift decline in the marginal utility of consumption implied by $s = .2$ should have led to a marked reduction in work hours as households satisfied the most pressing consumption needs and then turned to additional leisure when it became difficult to find additional attractive consumption opportunities. Such an outcome, with the income effect of the higher wage exceeding the substitution effect was quite conceivable, but it didn't happen. Indeed, Keynes predicted a large increase in leisure in the century following his 1930 essay "Economic Possibilities for our Grandchildren." He still has another 30 years to go on his prediction, but there are few signs of the great leisure boom he predicted. (His other prediction that "in the long-run we are all dead" has done better.) No *a priori* principle prevents the income effect from exceeding the substitution effect as Keynes guessed it would, but the lack of a strong trend in labor hours in the face of an enormous joint trend in wages and consumption indicates something close to cancellation between income and substitution effects on labor supply.

The macroeconomic literature on home production[1] questions the standard interpretation of income and substitution effects cancelling. The other alternative for explaining trend labor supply facts is that the rate of technological progress in home production is the same as the rate of technological progress in market production. We will discuss the issue of home production more below. But the most important point is that the hypothesis of technological progress in home production at just the right rate can only explain the trend facts. In addition to long-run growth facts, a great deal of both cross-sectional and panel evidence analyzed by labor economists indicates that the elasticity of labor supply with respect to a permanent increase in the real wage is very small.

---

[1] See for example Benhabib, Rogerson and Wright (1991), Greenwood and Hercowitz (1991), Greenwood, Rogerson and Wright (1995), Campbell and Ludvigson (1997), McGratten, Rogerson and Wright (1997), Canova and Ubide (1998) and Baxter and Jermann (1999).

To summarize, the typical approach has been to maintain additive separability between consumption and labor. That maintained assumption leads to an estimate of the elasticity of intertemporal substitution in consumption which would make the income effect of a permanent wage increase much stronger than the substitution effect of a permanent wage increase. This implication is at variance with at least three types of evidence about long-run labor supply: (1) the lack of a strong trend in weekly hours in the face of a dramatic trend in the real wage, (2) the fact that households, say, at the 75th percentile of wages work on average about as much as those with wages in the 25th percentile, and (3) the fact that permanent wage shocks to an individual on average do not appear to have much effect on work hours.

In the face of the impressive evidence for approximate equality of the income and substitution effects on labor supply of a permanent increase in the real wage, our approach is to make this equality of income and substitution effects on labor supply a maintained assumption when estimating the elasticity of intertemporal substitution in consumption. With this maintained assumption, we find a different value for the elasticity of intertemporal substitution than one estimates under the maintained assumption of additive separability between consumption and labor, but the elasticity of intertemporal substitution is still significantly different from 1. Thus, we reject additive separability between consumption and labor.

While an income effect of permanent wage increases that is much larger than the substitution effect is at serious violence with evidence on long-run labor supply, we see no serious problem with abandoning additive separability between consumption and labor. Indeed, as we will discuss in the Conclusion, additive nonseparability between consumption and labor helps to make sense of a wide variety of economic phenomena beyond those that motivate us to consider this nonseparability in the first place.

Historically, we suspect that one of the greatest recommendations of the assumption of additive separability between consumption and labor has been simplicity. In this paper, we hope to show among other things that the price in added complexity of our approach is quite reasonable. In order to make the issues introduced by additive nonseparability with income and substitution effects on labor supply cancelling as clear as possible, we illustrate our approach in the context of what is an otherwise simple log-linearized consumption-Euler-equation estimation with additive time separability in the style of Hall (1988). At this point, more complicated consumption empirics would obscure the intuition for the issues we most want to clarify. As for theory, the key idea here of imposing the parameter restrictions implied by cancellation of income and substitution effects on labor supply can be applied to much more general models[2] (and we hope it will be), but the case of additive time separability is the obvious baseline case.

Because we hope to have readers come away with a new perspective on the consumption Euler equation, we will present our theory (Section II) and evidence (Section III) before discussing the extensive previous literature on the consumption Euler equation (Section IV) and on home production in macroeconomics (Section V).

II. Theory

As alluded to above, in order to focus on the main issue of interactions between consumption and labor, we maintain the assumption of a representative consumer who has an additively time-separable von Neumann-Morgenstern utility function with a constant utility discount rate:

$$V_t = \operatorname*{E}_t \sum_{j=0}^{\infty} e^{-\rho j} u(C_{t+j}, N_{t+j}).$$

For convenience in the estimation, the time interval will be one quarter. Time aggregation up from continuous time will be handled in the usual way in the estimation by lagging the instruments an extra quarter and allowing for an MA(1) structure to the error term of the Euler equation. However, for clarity of exposition, this theory section will use discrete time and ignore time aggregation.

Operationally, "imposing cancellation between (nonzero) income and substitution effects" means choosing from the set of utility functions that yield a real wage proportional to consumption times some function of the quantity of labor[3]:

$$\frac{W}{P_C} = -\frac{u_N(C,N)}{u_C(C,N)} = Cv'(N) \tag{2}$$

Integrating this partial differential equation indicates that felicity (the period utility function) must be of the form

$$u(C,N) = \Phi(\ln(C) - v(N))$$

for some monotonically increasing function Φ. At the risk of belaboring the obvious, note that with this form of the felicity function,

---

[2] More general models of interest include models with consumer durables, models with habit formation in consumption and habit formation or durability of leisure.

[3] It is a key part of the argument for Equation (2) that since labor supply is not totally inelastic, the supply-side real wage $-u_N(C,N)/u_C(C,N)$ is a continuous function of the quantity of labor N.

$$-\frac{u_N}{u_C} = -\left(\frac{-v'(N)\Phi'(\ln(C) - v(N))}{\frac{1}{C}\Phi'(\ln(C) - v(N))} \right) = Cv'(N).$$

Thus, the monotonically increasing function Φ does not affect the within-period first-order condition. Φ only affects intertemporal substitution between now and the future, not atemporal substitution between consumption and labor.[4]

The reasonable additional assumption of a constant elasticity of substitution in consumption when the quantity of labor is held constant narrows the utility function down to the King-Plosser-Rebelo form, which we write conveniently as

$$u(C, N) = \frac{C^{1-\gamma}}{1-\gamma} e^{(\gamma-1)v(N)}. \tag{3}$$

We also write

$$s = 1/\gamma,$$

where s now represents the *labor-held-constant* elasticity of intertertemporal substitution in consumption, as will be apparent below.

The marginal utility of consumption is

$$u_C(C_t, N_t) = C_t^{-\gamma} e^{(\gamma-1)v(N)}.$$

For the consumption Euler equation, we will need the natural logarithm of u_C:

$$\ln(u_C(C_t, N_t)) = -\gamma \ln(C_t) + (\gamma - 1)v(N) \tag{4}$$

While it is helpful to assume constancy of the elasticity of substitution s as consumption C trends upward, the lack of a strong trend in N allows us to deal with the function $v(N)$ non-parametrically by using a Taylor expansion around the average value of N, which we label N^*. Given good data on the fluctuations in W, we could be slightly more exact ("slightly" in the sense of second-order),

[4] Within a period, the indifference curves between labor and consumption are unaffected by the outer function Φ. With N on the horizontal axis and C on the vertical axis, the restriction to the form $u(C, N) = \Phi(\ln(C) - v(N))$ means that each indifference curve is a vertical (C-direction) multiple of every other indifference curve. The restriction we use empirically can be seen on this graph as follows. On the graph, the marginal utility of consumption shows up as the reciprocal of the vertical distance between indifference curves with a fixed (small) difference in utility. When moving along a given indifference curve, this gap is proportional to the level of consumption, meaning that the marginal utility of consumption is inversely proportional to consumption *when changing C and N in a way that holds total utility constant.* Locally, knowing the wage as seen by the household is enough to know how to change C and N in a way that holds total utility constant.

but we consider the short-run fluctuations in the observed real wage to be unreliable as indicators of the short-run fluctuations in the marginal disutility of labor (primarily, in our view, because of the existence of long-term relationships between workers and firms in which observed wage payments often include implicit borrowing and lending—as well as implicit insurance premia and payouts—between workers and firms). Our approach has the advantage of relying only on the long-run average value of (after-tax) $\frac{WN}{P_C C}$.

Optimal choice of consumption implies the Euler equation

$$u_C(C_{t-1}, N_{t-1}) = E_{t-1} e^{(r_t - \rho)} u_C(C_t, N_t).$$

Programmatically, we want to focus on just the first-order terms of the Taylor expansion. Moreover, assuming homoscedasticity of the stochastic process for $\ln(C)$ and N, even the second-order departures from certainty equivalence only contribute a constant to the right-hand side of this equation. Accordingly, we can act as if the natural logarithm can be interchanged with E_{t-1}, yielding

$$\ln[u_C(C_{t-1}, N_{t-1})] = E_{t-1}\{r_t - \rho + \ln[u_C(C_t, N_t)]\} + \text{higher-order terms.}$$

Substituting in the King-Plosser-Rebelo form of the utility function yields

$$-\gamma \ln(C_{t-1}) + (\gamma - 1)v(N_{t-1}) = E_{t-1}[r_t - \rho - \gamma \ln(C_t) + (\gamma - 1)v(N_t)] + \text{higher order terms}$$

Dividing through by γ, rearranging, writing $\frac{1}{\gamma} = s$, and making rational expectations error term ϵ_t explicit,

$$\Delta \ln(C_t) = s(r_t - \rho) + (1 - s)\Delta v(N_t) + \epsilon_t + \text{higher order terms.} \tag{5}$$

(Note that the rational expectations error term involves surprises in all three terms: consumption growth, the *ex post* real interest rate and a function of labor.)

The final step is to use a first-order Taylor expansion of $v(N)$ in $\ln(N)$ around the trend level of labor N^*. Since $v(N) = v(e^{\ln(N)})$, the chain rule implies

$$v(N) \approx v(N^*) + N^* v'(N^*)[\ln(N) - \ln(N^*)].$$

Assuming that in the long run, the household can optimize labor supply, the intratemporal first-order condition $W/P_C = Cv'(N)$ applied to the long run implies

$$N^*v'(N^*) = \left(\frac{WN}{P_C C}\right)^* = \tau, \tag{6}$$

where it is important that W be the after-tax wage seen by the household. The right-hand side, which we denote as τ, is calibrated as a long-run average value of $\frac{WN}{P_C C}$. Thus, τ is treated as a constant in the estimation. τ is a number known from long-run labor supply facts. For example, below, our preferred value is $\tau = .8$. The rest of the estimation takes place conditional on a particular value of τ.

Substituting the constant τ known from long-run labor supply facts into the log-linearized Euler equation, and using small c and n to denote the natural logarithms of consumption C and labor N, yields

$$\Delta c = s(r_t - \rho) + \tau(1-s)\Delta n + \epsilon_t + \text{higher order terms.} \tag{7}$$

One more rearrangement shows that this is a very simple IV estimation:

$$\Delta c - \tau\Delta n = \text{constant} + s[r_t - \tau\Delta n] + \epsilon_t. \tag{8}$$

In Equation (8), we have finally omitted the higher-order terms, except those that can be absorbed into the constant term.

It is clear now that our estimation does more than simply add the growth in labor to the consumption Euler equation. There is a non-trivial linear restriction between the coefficient on the real interest rate and the coefficient on the growth in the quantity of labor. This restriction comes from facts about long-run labor supply. What is surprising about this restriction is that (other than the need for nonzero income and substitution effects for our main argument to go through) the parameter restriction does not depend on the value of the labor supply elasticities indicating the size of the income and substitution effects. This lack of dependence on the values of labor supply elasticities is fortunate given the lack of consensus on the values of these labor supply elasticities.[5]

Before going on to estimation, it is worth pausing to ask if we can give a more intuitive explanation for this linear restriction on the consumption Euler equation with labor. In Equation (7), it is evident that the degree of nonseparability required—as indicated by the size of $(1-s)\tau$, the coefficient on Δn—becomes greater as the elasticity of intertemporal substitution drops further

[5] This lack of consensus about labor supply elasticities comes in large part from the concern alluded to above that the short-run fluctuations in the observed wage may not equal the short-run fluctuations in the shadow wage that governs the equilibrium quantity of labor.

below one. One way to understand this is as follows. As illustrated in the introduction, low values of the elasticity of intertemporal substitution in consumption mean that the marginal utility of consumption falls rapidly with growth in consumption. Without any interaction between consumption and labor in the utility function, this swift decline in the marginal utility of consumption would lead households to want more leisure unless the real wage increased markedly. What happens in the case of the King-Plosser-Rebelo utility function is that consumption and labor are complements, so that the increased level of consumption expenditures makes labor more pleasant. To tell a story, with the extra expenditures, things at home can be taken care of pretty well despite all of the hours spent at work; this makes households willing to continue working the same workweek even as they become richer.

Note that complementarity between consumption and labor goes both ways. To use introspection to check the plausibility of the complementarity that arises with the King-Plosser-Rebelo utility function when $s < 1$, one can consider equivalently (1) whether an increase in work hours would lead to an increase in the marginal utility of expenditures or (2) whether an increase in expenditures would reduce the marginal disutility of work. Both are reflections of the same cross-partial derivative inequality $u_{CN} > 0$.[6]

III. Evidence

Data

We use quarterly, seasonally-adjusted, aggregate U. S. data from 1949:1 to 1998:3. Our measure of consumption comprises non-durable consumption plus services, per capita. Our measure of per-capita hours is total hours worked by all persons (civilian and military), from unpublished BLS sources, divided by the non-institutional population over 16 plus members of the military. When we augment the regressions with disposable income, we use aggregate disposable personal income as defined by NIPA, per capita. The real interest rate is computed as the after-tax nominal rate on three-month U. S. Treasury bills minus inflation in the price index of non-durable consumption and services. We took our measure of the average marginal tax rate from Stephenson (1998); since Stephenson's calculations extend only to 1994, we assumed that the average marginal rate for all

[6] One way to see the intuition behind complementarity between consumption and labor is to think about work-related expenses. However, if one goes this route, one should think of "work-related expenses" in the broadest sense possible. For example, having less leisure may cause one to shift towards more expensive but less time consuming forms of entertainment, or to hire someone to fix a leaky faucet instead of doing it oneself. It is easy to miss these more subtle types of added expenditures if one tries to make a direct measure of "work-related expenses." The Euler equation approach can be seen as a way to capture the full range of "work-related expenses."

subsequent years equals the 1994 value.[7]

According to the theory, τ equals labor income divided by nominal consumption expenditure. Taking nominal wages and salaries from the National Income Accounts and dividing by nominal spending on non-durable consumption and services gives an average ratio of 0.90. But we should use prices as perceived by the consumer, so we should define τ using the after-tax wage. Multiplying the numerator by our average marginal tax series reduces the mean τ to 0.77. We thus use $\tau = 0.8$ as our preferred value, but check our results for several other values.

Results

Due to time aggregation of the data, the error term in our estimating equation has an MA(1) structure. Thus, instead of equation (7), we actually estimate:

$$\Delta c - \tau \Delta n = \mu + s(r - \tau \Delta n) + \epsilon_t + \theta \epsilon_{t-1} \tag{9}$$

(Consumption and labor are in logs, as indicated by the small letters c and n. The constant μ potentially includes higher order terms such as precautionary saving effects from a homoscedastic time series process for c.) Given the potential ambiguities about the precise value of τ, we estimate the equation for four values of τ ranging from 0.6 to 1.2. We use twice-lagged values of Δc, Δn and r as instruments. The results are in Table 1.

We consistently estimate values of the EIS significantly greater than zero, unlike Hall (1988) and most subsequent work in this area. Depending on the value of τ the estimate of the elasticity of intertertemporal substitution s ranges from 0.3 to 0.5. The value s=.5 corresponds to the utility function

$$-\frac{e^{v(N)}}{C},$$

while $s = .33\bar{3}$ corresponds to the utility function

$$-\frac{e^{2v(N)}}{2C^2}.$$

We also report the first-stage F-statistics in Table 1. According to Staiger and Stock (1997), these indicate that we are unlikely to have weak-instrument problems. The Hansen J-statistic indicates that we cannot reject the overidentifying restrictions, except, perhaps, for $\tau = 1.2$, indicating that this value of τ may be significantly too high.

[7] Since average marginal rates are very stable, this procedure should not create significant problems.

We now check the robustness of our results. First, we show that our results are robust to alternative instruments sets. (We use some of these instruments later in the paper, when we investigate the robustness of our result to adding disposable income to the consumption Euler equation.) In addition to the variables we use as instruments in Table 1, we use the twice-lagged change in disposable income, $\Delta y(-2)$, and the twice-lagged ratio of consumption to disposable income $c(-2) - y(-2)$. The results are in Table 2.

We find that the estimated s is not sensitive to the instrument set used. All the results say that s is about one-third; three of the four point estimates are within 0.02 of one another, and the standard errors are all roughly 0.10. The first-stage fit is reasonably good in all cases, and in no case can we reject the overidentifying restrictions at conventional significance levels.

Second, we investigate the restrictions imposed by the King-Plosser-Rebelo functional form that we have assumed so far. In particular, it is possible that the significant EIS that we have estimated so far is due to the correlation between Δn and Δc, and does not reflect the effect of the real interest rate on consumption growth. We check this hypothesis by estimating several less-constrained versions of our equation, for our preferred value of $\tau = 0.8$. We report the results in Table 3.

We fail to reject the null hypothesis that the coefficient on r makes no additional contribution beyond the combination $r - \tau\Delta n$ at the 5% level. The p-values for the test of the null hypothesis hover around 10% for our four instrument sets (.090, .105, .143 and .081). To the extent that the point estimates do not obey the restriction, they disobey in the direction of unpredictability of Δc.[8]

We then run our regressions "in reverse"—that is, we regress $r - \tau\Delta n$ on $\Delta c - \tau\Delta n$, which normalizes the coefficient of $r - \tau\Delta n$ to 1, instead of normalizing the coefficient of $\Delta c - \tau\Delta n$ to 1 as in the previous regressions. Assuming additive separability, which in equation (8) is mechanically equivalent to setting τ equal to zero, Campbell and Mankiw (1989) find that both the forward regression and the reverse regression often yield coefficients close to zero. They interpret this finding as a specification test indicating that the standard rational representative-agent model of consumption is incorrect. By contrast, in Table 4B we find that our reverse regressions yield estimates of $1/s$ that are quite consistent with the estimates of s from the forward regressions in Table 4A. For example, for $\tau = 0.8$, we estimate s to be 0.36 from the forward regression and 0.47 from the reverse regression. Since 0.47 is within the 95 percent confidence interval of our forward

[8] In our regressions we can say that we obey both of Frank Fisher's laws. According to a personal communication with Greg Mankiw, Frank Fisher's *Iron Law of Econometrics* is: "Regression coefficients are always too small." Frank Fisher's *Iron Law of Non-Linear Econometrics*: is "Don't do it."

estimate, we do not think that the Campbell-Mankiw specification test rejects our model. On the other hand, the reverse estimate of the standard $\tau = 0$ consumption Euler equation, 0.75, is not in the confidence interval of its forward estimate, confirming Campell and Mankiw's original finding. However, we prefer estimates of s based on the forward regression, because the reverse regression has poor first stage fit (for all instrument sets we examine in Table 2, not just for the baseline result we report). The first-stage F-statistic for the reverse regression ranges from 2 to 3, which is well within the danger zone identified by Staiger and Stock (1997). The reason is fairly intuitive-in the reverse regression, the right-hand-side variable is approximately the change in consumption minus the change in labor hours. Both of these series are procyclical and both are less volatile than output, so it is not surprising that their difference is difficult to predict.

We now proceed to compare our model of consumption based on non-separable leisure to Campbell and Mankiw's (1989) hypothesis of rule-of-thumb consumption. Campbell and Mankiw augment the standard ($\tau = 0$) consumption Euler equation with disposable income, and find a positive, significant coefficient on disposable income. They interpret this coefficient as the fraction of consumption that is done by rule-of-thumb consumers who consume a fraction of their current disposable income (or, perhaps, the fraction of consumption due to liquidity-constrained consumers). They find estimates as large as 0.5, suggesting that up to half of all consumption is done by rule-of-thumb consumers. In Table 5 we augment our basic estimating equation (8) with disposable income, for $\tau = 0.8$. We use the various instrument sets we explored in Table 2, since some of the additional instruments may be better at predicting income growth than our basic set of variables. The results generally support our model, and do not support the hypothesis of significant rule-of-thumb consumption. In all cases the disposable income variable is insignificant. In three of the four cases the coefficient is negative, which has no meaningful interpretation in the rule-of-thumb context. In three of the four cases, our estimate of s is significant and quite close to the values we estimate for our basic specification. The one exception is for our standard instrument set, which is not surprising since that instrument set is not chosen for its ability to predict future income. The largest instrument set, which adds the twice-lagged consumption-income ratio to our basic instruments, yields an estimate of s that is exactly the same as our original estimate in Table 1 and an insignificant, negative coefficient on disposable income.

IV. Relationship to the Consumption Euler Equation Literature

Campbell and Ludvigson (2000), in summarizing previous empirical literature on the consumption Euler equation, write that "... aggregate data offer no evidence of any important nonseparability between market consumption and labor hours .... For example, Campbell and Mankiw find that although there is substantial predictable variation in hours, it is not significantly related to predictable consumption growth as it should be if utility over leisure and consumption were additively nonseparable. This evidence suggests that consumption and nonmarket hours can be well characterized by an additively separable utility function over consumption and nonmarket time, or, more generally, over consumption and some function of nonmarket time, as would be the case in models with home production."[9]

We disagree. First, even when we do not impose the parameter restriction implied by King-Plosser-Rebelo preferences, we find a significant relationship between consumption growth and both the real interest rate and predictable movements in labor. This shows up in Table 3 as a coefficient on $\tau\Delta n$ significantly different from -1, since the left-hand-side variable is $\Delta c - \tau\Delta n$.

To stack the deck against ourselves even more, we performed unrestricted "horse-race" IV regressions of Δc on both Δn and Δy with a wide variety of instrument sets and both with and without the real interest rate r in the regression. To summarize a large number of results, without the real interest rate r in the regression, Δn and Δy do equally well by a t-statistic metric, but neither Δn or Δy ever have a coefficient significantly different from zero for any of our four instrument sets. (The minimum p-value is greater than .11.) With r in the regression, Δn does better than Δy.[10] Both are insignificant for the first two instrument sets, but in the last two instrument sets, which include the lagged consumption/income ratio, Δn has a two-tailed p-value of .056 and .050, while Δy has a p-value of .362 and .688.[11] The relative performance of Δn and Δy for the last instrument set[12] can also be seen by comparing the last two scatter plots. Based on these results, we maintain that an analyst who was indifferent between Δn and Δy *a priori* would have no reason to prefer Δy to Δn based on these horse-race regressions.

But second, we do not think one should be indifferent *a priori* between augmenting the consumption Euler equation with Δn or with Δy. The burden of the first two sections of this paper is

---

[9] In thinking about the early consumption Euler equation research, it is important to remember that certain aspects of current empirical practice were not yet solidly in place during this period. First, some of the early work on consumption Euler equations addresses time aggregation and some does not. Second, almost all of the early consumption Euler equation research using instrumental variables used what would now be considered too many instruments for reliable results.

[10] Strangely enough, there is some tendency for Δy to do better when using instrument sets that emphasize lagged Δn and for Δy to do better when using instrument sets that emphasized lagged Δy and the lagged consumption to income ratio.

[11] The instrument set $\Delta c(-2)$, $r(-2)$, $\Delta y(-2)$, $c(-2) - y(-2)$ makes Δn significant with or without r (p-values of .011 with r and .031 without r) while Δy is insignificant (p-values of .345 with r and .096 without r).

[12] We consider the last instrument set most appropriate for this exercise since it includes both $\Delta n(-2)$ and $c(-2) - y(-2)$.

to show why theory makes it almost mandatory on theoretical grounds to have Δn appear in the consumption Euler equation.

By contrast, the closest reasonable empirical substitute—adding Δy to the consumption Euler equation—involves taking the grave step of abandoning the Permanent Income Hypothesis. Campbell and Mankiw (1989), believing that one could not reasonably avoid including Δy in the consumption Euler equation, write "The failures of the representative consumer model documented here are in some ways unfortunate. This model held out the promise of an integrated framework for analyzing household behavior in financial markets and in goods markets." We believe that appropriate inclusion of labor in the formula for the marginal utility of consumption holds great promise for improving the empirical performance of the Permanent Income Hypothesis in a wide variety of economic contexts.

In the consumption Euler equation context, theory not only mandates the inclusion of Δn, it mandates the size of the coefficient on Δn given the coefficient on the real interest rate r. We are not adding any free parameters. Therefore, we believe that the appropriate test for "excess sensitivity" of consumption to current income in aggregate data is the one reported in Table 5.

V. Relationship to the Macroeconomic Literature on Home Production

Although there is no doubt that home production is an important area of research for *microeconomics*, we will argue in this section that *macroeconomic* analysis can ordinarily ignore the details of home production without serious loss.

To see this, consider a household with an underlying utility function given by $U(X)$, where X is a vector of goods produced in home production. Some examples of possible elements of X are being well rested, being well fed, being entertained, etc. The household production function for the vector X is

$$X = F(Q, L, Z),$$

where Q is a matrix with rows showing all the different ways of using up the vector of marketed consumption goods, L is a row vector of different possible ways to spend time away from work, and Z is a vector of the technology for home production. Conditional on the total quantity vector C of marketed consumption good used and time N spent at work and the technology for home production Z, the household solves

$$\max_{Q,L} \; U(F(Q, L, Z))$$

s.t.

$$\sum_j Q_{ij} = C_i$$

and

$$\sum_j L_j + N = T,$$

where T is the total time endowment. Since the total time endowment per unit time is fixed, the maximum value is only a function of C and N. Thus, for this household engaged in home production, we can define $u(C, N, Z)$ by

$$u(C, N, Z) = \max_{Q,L} \; U(F(Q, L, Z))$$

s.t.

$$\sum_i Q_i = C$$

and

$$\sum_i L_i + N = T.$$

We maintain that the reduced form utility function $u(C, N, Z)$ contains all of the information needed for macroeconomics. The only importance for considering home production is the theoretical one of establishing the *a priori* plausibility for various forms of the utility function $u(C, N, Z)$. For example, thinking about home production may make one more eager to allow for more than one consumption good and may legitimate the exogenous effects of Z, which in the reduced form utility function look like preference shocks. But considerations about the household production function are on a par with any other reasoning about the form of the reduced form utility function $u(C, N, Z)$.

For example, the main costs and benefits of going from one type of marketed consumption good to two (say adding consumer durables to the nondurables and services we have been concentrating on) are there in roughly the same degree whether or not one worries about home production, and can be considered by looking directly at the behavior implied by different forms of the reduced form utility function $u(C, N, Z)$. And from the macroeconomic point of view, the key costs and benefits of adding preference shocks to a reduced form model of household behavior remain little changed by household production considerations.

If, for the sake of parsimony, the reader grants us the simplification of having only a single marketed consumption good, the argument we give in the first two sections for the King-Plosser-Rebelo utility function remain. Since the evidence from labor economics about the effects of a permanent increase in the real wage on N indicates that households at a given point in time, with a given value of Z, show cancellation of income and substitution effects, which we think of operationally as the equation $W = Cv_N(N, Z)$, the utility function must have the form

$$u(C, N, Z) = \frac{C^{1-\gamma}}{1-\gamma} e^{(\gamma-1)v(N,Z)}. \tag{10}$$

Following through with the same exercise as in section II, we arrive at the same estimating equation, except that linear terms involving the household technologies Z are added to the error term:[13]

$$\Delta c = \mu + sr + (1-s)\tau\Delta n + (1-s)\sum_j \zeta_j \Delta z_j + \epsilon_t + \theta\epsilon_{t-1}. \tag{11}$$

The difficulties preference shocks can cause for Euler equation estimation are well known. If one believes in important household preference shocks, the best way to interpret our results is to think of our estimation as assuming that preference shocks (or equivalently, home production technology shocks) cannot be predicted by our instruments.

To round out this section on home production, we should mention that the paper closest in spirit to ours is Baxter and Jermann (1999): "Household Production and the Excess Sensitivity of Consumption." Indeed, our model is the special case of theirs when there is no home production. Given the large share of pure leisure in their utility function, which is modeled in the King-Plosser-Rebelo form, we suspect that many of their results are not so much due to anything special about home production but rather to the effects we discuss in this paper. We think readers of both our paper and theirs will find the two papers complementary, even though we would interpret their results differently than they do.

---

[13] The reason a first-order approximation should be OK for the effects of Z is that when $s \neq 1$ and $v_{N,Z} \neq 0$, any strong trend in Z should cause a trend in optimal N over time that is not observed. In the dimensions where $v_{N,Z} = 0$, so that v is additively separable in labor and the production technology, a trend in Z will not cause a trend in labor N, so we cannot rule out a trend in this dimension of Z, but the additive separability means that variation in ζ_j in that dimension would be hard to distinguish from heteroscedasticity in that dimension of Z.

VI. What if the Cancellation Between Income and Substitution Effects is Not Exact?

So far, for the sake of a clarity, we have maintained an exact cancellation between the income and substitution effects of a permanent wage increase. What if that cancellation is not exact? Indeed, some of the stylized facts seem to suggest that the income effect is slightly stronger. For example, looking at the scatter plot Abel and Bernanke (1995) display of average weekly hours across countries suggests that richer countries have somewhat shorter work weeks: across countries, a tenfold higher real wage across countries seems to be associated with about a 12% reduction reduction in weekly hours, from 44 hours to 39 hours. Taking those numbers at face value, that represents an elasticity of labor hours with respect to permanent wage increases of about -.05. Could the gap between -.05 and 0 make a big difference to the empirical analysis?

To give some insight into this issue, let us take a totally nonparametric approach with a felicity function $u(C, N)$ of general form. (We will continue to assume additive time-separability with a constant utility discount rate ρ.) The way we have imposed the restriction of income and substitution effects cancelling is by assuming that $W = Cv'(N)$. In particular, this means that, holding labor constant, the elasticity of the real wage with respect to consumption is 1. Since the shadow real wage as seen by the household is

$$W(C, N) = \frac{-u_N(C, N)}{u_C(C, N)},$$

the elasticity of the wage with respect to consumption when labor is held constant is

$$\xi(C, N) = \frac{\partial \ln(W(C, N))}{d \ln(C)} = \frac{Cu_{NC}(C, N)}{u_N(C, N)} - \frac{Cu_{CC}(C, N)}{u_C(C, N)}.$$

Let's explore what happens when $\xi(C, N) \not\equiv 1$.

Define also the elasticity of intertemporal substitution for consumption $s(C, N)$ and the consumption-constant elasticity of labor supply $\eta(C, N)$ by

$$\frac{1}{s(C, N)} = -\frac{Cu_{CC}(C, N)}{u_C(C, N)},$$

and

$$\frac{1}{\eta(C, N)} = \frac{\partial \ln(W(C, N))}{\partial \ln(N)} = \frac{Nu_{NN}(C, N)}{u_N(C, N)} - \frac{Nu_{CN}(C, N)}{u_C(C, N)}$$

Finally, define

$$\tau(C, N) = \frac{W(C, N)N}{C} = \frac{-Nu_N(C, N)}{Cu_C(C, N)}.$$

Since the Hessian $\begin{bmatrix} u_{CC} & u_{CN} \\ u_{NC} & u_{NN} \end{bmatrix}$ is symmetric, the three elasticities $\xi(C,N)$, $s(C,N)$ and $\eta(C,N)$, plus $\tau(C,N)$ (which indicates the local ratio of marginal utilities) determine all local first and second derivatives of the felicity function $u(C,N)$ (except for overall scale of felicity, which has no economic meaning.)

A first-order Taylor expansion of $\ln(u_C)$ indicates that

$$\begin{aligned}\Delta \ln(u_C(C,N)) &\approx \frac{Cu_{CC}(C,N)}{u_C(C,N)}\Delta \ln(C) + \frac{Nu_{CN}(C,N)}{u_C}\Delta \ln(N) \\ &= -\frac{1}{s(C,N)}\Delta \ln(C) + \left(\frac{-Nu_N(C,N)}{Cu_C(C,N)}\right)\left(\frac{-Cu_{CN}}{u_N(C,N)}\right)\Delta \ln(N) \\ &= -\frac{1}{s(C,N)}\Delta \ln(C) + \tau(C,N)[\frac{1}{s(C,N)} - \xi(C,N)]\Delta \ln(N).\end{aligned} \tag{12}$$

Combining the results of Equation (12) with the consumption Euler equation, and using c and n to represent the logarithms of C and N, the estimation equation corresponding to Equation (9) is

$$\Delta c - \tau\Delta n = \mu + s(r - \xi\tau\Delta n) + \epsilon_t + \theta\epsilon_{t-1}. \tag{13}$$

Equation (13) reduces to Equation (9) when $\xi = 1$. Numerically, this estimate will be very similar to what we do in previous sections as long as ξ is reasonably close to 1 (say, in the range [.8, 1.2]), and the elasticities s and ξ—plus the ratio τ—do not vary too widely over the sample period.[14]

How far might ξ reasonably be away from 1 given the facts about long-run labor supply? For a numerical example, suppose that we take -.05 as the elasticity of labor with respect to permanent, equal increases in both the real wage W and consumption C. (See the discussion at the beginning of this section.) From the definitions of ξ and η,

$$d\ \ln(W) = \xi d\ \ln(C) + \frac{1}{\eta}d\ \ln(N). \tag{14}$$

If improvements in technology, coupled with the budget constraint causes consumption and the real wage to trend up together, then $d\ \ln(C) = d\ \ln(W)$ and Equation (14) reduces to

$$(1-\xi)d\ \ln(W) = \frac{1}{\eta}d\ \ln(W)$$

[14] The other way of arranging Equation (13)—

$$\Delta c = \mu + sr + (1 - s\xi)\tau\Delta n) + \epsilon_t + \theta\epsilon_{t-1}$$

—suggests that a value of ξ somewhat higher than 1 could help to explain why the freely estimated coefficient on Δn is somewhat smaller than it should be if $\xi = 1$.

or

$$\psi = \frac{d\ \ln(N)}{d\ \ln(W)} = \eta(1-\xi)$$

where ψ is the elasticity of labor supply with respect to permanent increases in the real wage (coupled with equal increases in consumption). Solving for ξ,

$$\xi = 1 - \frac{\psi}{\eta}.$$

If $\psi = -.05$, then $\xi = 1 + \frac{.05}{\eta}$. Alternatively, if one allows some leeway in the estimate of the long-run elasticity of labor with respect to permanent increases in the real wage, by setting $\frac{d\ \ln(N)}{d\ \ln(W)} = -.1$, then $\xi = 1 + \frac{.1}{\eta}$. Given an elasticity of labor supply with respect to permanent increases in the real wage, the only way ξ can be far away from 1 is if the *consumption-constant* elasticity of labor supply η is quite low (say .25 or below). Although some estimates of the consumption-constant elasticity of labor supply can indeed be quite low, these estimates are questionable, because they all rely on either (1) life-cycle variation in the real wage, which is confounded with age effects or (2) treating the observed short-run fluctuations in the real wage as if they were the true shadow wage determining labor allocations. If the true consumption-constant elasticity of labor supply were as low as, say .25, macroecomists of any stripe would have little hope of explaining the substantial variation in the quantity of labor over the business cycle. Departures of ξ from exact unity are the least of the problems macroeconomists would face with a low value of η.

In summary, we do not claim that ξ is literally equal to exactly 1. However, we believe it to be close enough to 1 that for practical purposes, it is best to treat it as equal to 1. It is worth addressing the issue once, as we have here, but we think that economic understanding will advance fastest if discussion of departures from exact cancellation of the income and substitution effects of permanent wage increases is, in the future, primarily relegated to footnotes. Indeed, following this dictum, we return in the next section to assuming a King-Plosser-Rebelo utility function.

VII. Aggregation of the Consumption Euler Equation for Individual Households

Many authors have discussed issues about consumption aggregate for the consumption Euler equation in the additively separable case. We will not attempt to repeat what they say here, but only to give a perspective on those issues and the new issues that arise from nonseparability.

One reason one might be concerned about the interaction of nonseparability with aggregation up from the household level is that changes between working and not working often involve a discontinuous jump in hours. These discontinuous jumps in hours suggest a utility function that is not always concave in hours. Close examination of the logical derivation behind the King-Plosser-Rebelo utility function shows that the same form should hold even when the utility function is not concave in hours. The inner function $v(N)$ can encode such nonconcavities:

$$u(C, N) = \frac{C^{1-\gamma}}{1-\gamma} e^{(\gamma-1)v(N)}.$$

At the same time, it is important to clarify the effects of more than one working member of the household. From here on, we will be careful to interpret N as a vector indicating the labor hours of each member of the household. With this reinterpretation, the form of the utility function again looks the same.

Nonconcavities in the felicity function give rise to three technical issues. First, in dealing with non-concave utility, the Pontryagin maximum principle that the Hamiltonian must be globally maximized at every point in time is very helpful. The maximum principle holds exactly only in continuous time. So there is one more reason to insist that the true underlying model is in continuous time. Second, to ensure that the Hamiltonian actually achieves a maximum, it is helpful to assume that $v(N)$ is lower hemi-continuous-that is, the argument N going to a limit can lead to a jump *down* in $v(N)$ (and up in felicity), but not to a jump *up* in $v(N)$. For example, there can be a fixed cost of going to work (away from an element of N being zero), but there is no sudden fixed felicity benefit of going to work. Third, given the possibility of discontinuous changes in labor hours, it is important to assume that wherever such a discontinuous change takes place, the marginal shadow wage is unaffected by that discontinuous change. In other words, if there is something like an overtime premium in the true shadow wage, that change in the marginal wage takes place away from the discontinuous jump in labor hours that might take place, for example, in going from not working at all to working positive hours.

Consider the household problem

$$\mathop{\mathrm{E}}_{t} \max_{C,N} \int_{t}^{\infty} e^{-\rho(t'-t)} \frac{C_{t'}^{1-\gamma}}{1-\gamma} e^{(\gamma-1)v(N_{t'})}\, dt'$$

s.t.

$$\frac{db_t}{dt} = r_t b_t + \Pi_t + W_t \cdot N_t - C_t,$$

and

$$\lim_{t \to \infty} e^{-\infty_{t_0}^{\infty} r_{t'} dt'} b_t = 0.$$

The variable b_t is household financial assets, while Π_t is non-labor income. The price of consumption goods has been normalized to 1, so that W is a vector of real wages. The Hamiltonian, which is maximized globally, is

$$H(C, N, \lambda, W) = \frac{C^{1-\gamma}}{1-\gamma} e^{(\gamma-1)v(N)} + \lambda[r_t b_t + \Pi_t + W_t \cdot N_t - C_t].$$

Even though the model is stochastic, there is no problem with speaking of the Hamiltonian. Even if we added risky asset choices to the model, maximization of the "Hamiltonian" written above is a necessary condition for the maximization implicit in the Bellman equation, where λ is just the derivative of the value function with respect to household assets b_t. (It is important, if one adds risky asset choices to assume that the value of risky assets follows a diffusion process, so that wealth and the marginal value of wealth and consumption λ follow a continuous, if jagged, path.)

Let us call the globally maximized Hamiltonian $H^\star$:

$$H^\star(\lambda, W) = \max_{C,N} H(C, N, \lambda, W)$$

The envelope theorem is robust to jumps in the control variables. Since $H(C, N, \lambda, W)$ is continuous in λ and W, the maximized Hamiltonian $H^\star(\lambda, W)$ is also continuous in λ and W. Moreover, whenever small changes in λ and W do not cause any large jumps in C or N,

$$dH^\star(\lambda, W) = H_\lambda(C, N, \lambda, W)d\lambda + H_W(C, N, \lambda, W) \cdot dW = [W \cdot N - C]d\lambda + \lambda N \cdot dW, \qquad (15)$$

evaluated at the optimal values of C and N. Around discontinuities, this result works in terms of right-hand and left-hand derivatives and their multidimensional analogs.

Regardless of the shape of $v(N)$, the Hamiltonian is strictly concave in consumption C. Therefore the first order condition for optimal consumption, as before, is

$$\lambda = C^{-\gamma} e^{(\gamma-1)v(N)} \qquad (16).$$

The result we want comes from using this first order condition in conjunction with the envelope theorem for the globally maximized Hamiltonian. Equation (15) allows us to rewrite the globally maximized Hamiltonian as

$$H^\star(\lambda, W) = \lambda \left[\frac{C}{1-\gamma} + W \cdot N - C\right] = \lambda \left[W \cdot N - \frac{C}{1-s}\right],$$

where $s = 1/\gamma$ as before.

At boundaries where either C or N is discontinuous, the continuity of $H^\star$ requires that

$$\Delta C = (1-s)W \cdot \Delta N, \tag{17}$$

in moving across the boundary. At points where both C and N are continuous, the envelope theorem (15) requires that

$$\left[W \cdot N - \frac{C}{1-s}\right] d\lambda + \lambda N \cdot dW + \lambda W \cdot dN - \frac{\lambda}{1-s} dC = [W \cdot N - C]d\lambda + \lambda N \cdot dW,$$

or equivalently,

$$dC = (1-s)W \cdot dN - sC\frac{d\lambda}{\lambda}. \tag{18}$$

Since the jumps in C and N of Equation (17) occur essentially with essentially no change in λ and what change in λ there is essentially occurs on either side of a jump boundary, Equations (17) and (18) are both consistent with the following single equation, appropriate for small changes in λ and W, whether or not C or N jumps:

$$\Delta C = (1-s)W \cdot \Delta N - sC\Delta \ln(\lambda) + \text{higher order terms} \tag{19}$$

where the value of C in $sC\Delta \ln(\lambda)$ is a time average over the path of λ from beginning to end.

The next step is to combine Equation (19) with the Euler equation. The Euler equation is

$$\Delta\lambda = [\rho - r]\lambda + \text{rational expectations error}.$$

Dividing by λ,

$$\Delta \ln(\lambda) = \text{constant} - r + \text{R.E. error} + \text{higher order terms}, \tag{20}$$

where the difference between $\Delta \ln(\lambda)$ and $\frac{\Delta\lambda}{\lambda}$ has been absorbed in the constant term and in the higher order terms. (Remember that the risky assets follow a diffusion process, so that the marginal value of wealth λ is continuous.) Combining Equations (19) and (20) yields

$$\Delta C = (\text{constant})C + srC + (1-s)W \cdot \Delta N + \text{R.E. error} + \text{h.o.t.} \tag{21}$$

Because it is linear in C, ΔC and ΔN, Equation (21) aggregates perfectly, aside from the higher-order terms. To be explicit, if i indexes households, when s, r and W are constant across households,

$$\sum_i \Delta C_i = (\text{constant}) \sum_i C_i + sr \sum_i C_i + (1-s) W \cdot \sum_i \Delta N_i + \text{R.E. error} + \sum_i \text{h.o.t.}_i \tag{22}$$

Dividing through by aggregate consumption, and rewriting the third term with the steady state value of the aggregate after-tax labor income to consumption ratio $\tau = \frac{W^* \cdot \bar{N}^*}{\bar{C}^*}$ made explicit yields an equation very close to what we estimate empirically:

$$\frac{\sum_i \Delta C_i}{\sum_i C_i} = \text{constant} + sr + (1-s) \left(\frac{W \cdot \sum_i N_i}{\sum_i C_i} \right) \left(\frac{W \cdot \sum_i \Delta N_i}{W \cdot \sum_i N_i} \right) + \text{R.E. error} + \sum_i \text{h.o.t.}_i$$

m

What if the intertemporal elasticity of substitution is not constant across households? Then s in Equation (22) is replaced by the consumption-weighted average of household s_i

$$\frac{\sum_i C_i s_i}{\sum_i C_i}.$$

What if wages are not constant across households? Because it allows for a vector of different types of labor, the term

$$(1-s) W \cdot \sum_i \Delta N_i$$

in Equation (22) can handle this case with appropriate reinterpretation. Let the vector of different types of labor N include as components all distinctions of labor types that would ever command a different wage.[15] This approach to differing wages accords well with empirical approaches that distinguish different types of labor to the extent possible and weights changes in each type of labor hours by the corresponding wage.

Finally, what about the higher order terms? The higher order terms consist of effects such as precautionary saving effects, which are just as much a problem when consumption and labor are additively separable. They are not new here. We are not trying to claim that that our consumption

[15] Because the household function $v_i(N)$ is unconstrained, it can represent the absence of a certain type of labor in the household by a large enough drop in utility if the quantity of that type of labor departs from zero.

Euler equation aggregates perfectly, but that it aggregates as well as the consumption Euler equation with felicity additively separable between consumption and labor, where the aggregation issues caused by the higher order terms are relatively well understood.

The one big gap in our understanding of aggregation for the consumption Euler equation with nonseparable labor is the dearth of research on precautionary saving with nonseparable labor. This is not the place to try to fill that gap in any detail, but it is worth pointing out that with King-Plosser-Rebelo utility with $\gamma > 1$, the negative sign of u_{CCN}, coupled with the positive empirical correlation between consumption and labor makes it possible for nonseparability between consumption and labor to generate *smaller* precautionary saving effects than when felicity with the same value of γ is additively separable. Let us present a simple numerical example of the effect of second-order terms on the marginal utility of consumption. Consider a one-earner household operating at a point where $u(C, N)$ is thrice differentiable in all variables, the second-order conditions hold strictly, $\gamma = 3$ (corresponding to $s = .333$), $\tau = .8$ and $\eta = \frac{v'(N)}{Nv''(N)} = .8$. When felicity is additively separable between consumption and labor, and $\gamma = 3$, the proportionate effect of consumption variance on the marginal utility of consumption is

$$\mathrm{E}\,\frac{u_{CCC}(C,N)}{u_C(C,N)}\frac{\Delta C^2}{2} \approx 6\frac{\sigma_C^2}{C^2}.$$

For comparison, with King-Plosser-Rebelo felicity, $\gamma = 3$, $\tau = .8$ and $\eta = .8$, the proportionate effect of consumption variance on the marginal utility of consumption is

$$E\frac{u_{CCC}(C,N)\Delta C^2 + 2u_{CCN}\Delta C\Delta N + u_{CNN}\Delta N^2}{2u_C(C,N)} \approx 6\frac{\sigma_C^2}{C^2} - 4.8\frac{\varrho\sigma_C\sigma_N}{CN} + 2.28\frac{\sigma_N^2}{N^2},$$

where ϱ is the correlation between ΔC and ΔN. The overall precautionary saving effects for King-Plosser-Rebelo felicity, as gauged by the proportionate rise in the expected future marginal utility of consumption are smaller than for additively separable felicity whenever[16]

$$\varrho\frac{\sigma_C}{C} > .458\frac{\sigma_N}{N}.$$

What is much more important (both because of the size of the consumption variance involved and cross-sectional differences in this consumption variance), when possible movements in and

[16] When this inequality is violated, but only by a little, the precautionary saving effects will be slightly larger in terms of the effect on the expected future marginal utility of consumption u_C, but the effect in reduced current consumption will still be less, since precautionary saving comes from increased current work hours as well as reduced current consumption, *both* of which raise *current* u_C, since $u_{CC} < 0$ *and* $u_{CN} > 0$ for King-Plosser-Rebelo utility with $\gamma > 1$. Thus, the overall drop in current *consumption* induced by a given increase in expected future marginal utility is less in the nonseparable case.

out of work generate the possibility of jumps in C and N according to the logic of a nonconcave King-Plosser-Rebelo utility function, the discussion above makes clear that the marginal utility of consumption, equal to λ, is unaffected *ex post* by the realized jumps and is therefore unaffected *ex ante* by the possibility of these optimal jumps. Yet the associated variance of consumption would lead an economist thinking in terms of additively separable utility to believe that the expected marginal utility of future consumption was being driven up, adding significantly to precautionary saving.

VIII. Relationship to the Literature on Microeconomic Consumption Empirics

A cursory examination of recent work on the consumption Euler equation in micro data provides substantial evidence for nonseparability between consumption and labor and confirms the point that allowing for nonseparability between consumption and labor leaves very little evidence for liquidity constraints from the consumption Euler equation.[17] Attanasio and Weber (1995) even use a utility function that is multiplicatively separable between consumption and labor, as we recommend, but use the long-run equality of the real wage with the marginal rate of substitution to impose our restriction. Imposing our restriction in micro-data would add discipline that would help in making the case that nonseparability between consumption and labor is at work rather than the ability of labor participation and hours to proxy for badly measured components of disposable income.

The estimates of the elasticity of intertemporal substitution for consumption vary in different studies (from close to zero to close to one), but our estimate is well within the range of estimates obtained.

Overall, we read the existing microeconomic evidence as wholly consistent with our claim for nonseparability between consumption and labor of a substantial and well-defined magnitude.

IX. Conclusion

Departing from the assumption of additive separability, which is typically made more for convenience than from conviction, we estimate the elasticity of intertemporal substitution while imposing the King-Plosser-Rebelo functional form needed for balanced trend growth of consumption and the real consumption wage. This mode of estimation is *not* the same as simply including the quantity of labor in a consumption Euler equation in an unrestricted way. Additive separability is easily rejected, but our restriction to King-Plosser-Rebelo preferences is not rejected. We find that

[17] See for example Browning and Meghir (1991), Attanasio and Weber (1993), Attanasio and Weber (1995), Attanasio and Browning (1995), and papers cited in Attanasio (1998).

such an estimate, which effectively combines information about the responsiveness of consumption to fluctuations in the real interest rate and the quantity of labor and information about the low-frequency behavior of labor and the real wage, yields an estimate of the elasticity of intertemporal substitution of about .6. Moreover, we find that the interactions between consumption and labor induced by the lack of additive separability allow one to solve a problem one finds in a Hall (1988)-type regression: estimates of the elasticity of intertemporal substitution that differ dramatically depending on which variable in the instrumental variable regression has its coefficient normalized to 1. This problem can be viewed as a particularly serious failure of the overidentifying restrictions of the IV estimation. In view of our results, omitted variable bias from leaving the quantity of labor out of the equation can account for this failure. Omitted variable bias can also account for Campbell and Mankiw's (1989,1991) finding that predictable movements in disposable income are related to predictable movements in consumption. We find no evidence for such a relationship once labor is properly included in the regression.

In our thinking, we come to reject additive separability—and arrive at the view that the utility function exhibits complementarity between consumption and labor (or equivalently, substitutability between consumption and leisure)—by considering facts about long-run labor supply. But this view also points to predictions in other areas that seem tantalizing and beg for further investigation.

First, as some economists have already noticed, complementarity between consumption and labor means that households should plan to have their consumption drop at retirement. The observed drop in consumption at retirement has been considered a minor mystery, but if one takes the King-Plosser-Rebelo utility function as a touchstone, it is easy to get optimal planned drops in consumption at retirement that are quite large. From our perspective, the mystery may be why consumption doesn't drop even more at retirement than it does. But it is likely to be much easier to modify a nonseparable model to moderate the drop in consumption at retirement than to get a significant drop in consumption at retirement out of a model with additive separability.

Second, complementarity between consumption and labor provides a straightforward channel for a monetary expansion to cause an increase in consumption. Business cycle theorists have puzzled over how to get interest-rate effects alone to cause an increase in consumption in response to monetary shocks. No solution of this type has been entirely successful. But when labor and consumption are complentary, the increase in labor when aggregate demand increases is enough to cause consumption to increase (as long as interest and wealth effects are not too large–a condition easy to satisfy).

In conclusion, let us reemphasize that we are arguing not for the particular utility function that we use in this paper, but that any utility function used (though it might be more complex than the one here) should pass the test of implying that a permanent increase in the real wage will have very little effect on labor hours.

Bibliography

Abel, Andrew B, and Ben S. Bernanke, *Macroeconomics*, 2d. Edition (1995), Addison Wesley, U.S.A., Figure 8.11, p. 281.

Attanasio, Orazio P., "Consumption Demand," NBER Working Paper # 6466 (March 1998).

Attanasio, Orazio P., and Martin Browning, "Consumption Over the Life Cycle and over the Business Cycle," *American Economic Review* 85 (December 1995), 1118–1137.

Attanasio, Orazio P., and Guglielmo Weber, "Consumption Growth, the Interest Rate and Aggregation," *Review of Economic Studies* 60 (1993), 631–649.

Attanasio, Orazio P., and Guglielmo Weber, "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey," Journal of Political Economy 103 (1995), 1121–1157.

Barsky, Robert, Miles Kimball, Thomas Juster and Matthew Shapiro, "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study," *Quarterly Journal of Economics* 112 (May 1997), 537–579.

Baxter, Marianne and Urban J. Jermann, "Household Production and the Excess Sensitivity of Consumption to Current Income," *American Economic Review* 89 (September 1999), 902–920.

Benhabib, Jess, Richard Rogerson and Randall Wright, "Homework in Macroeconomics: Household Production and Aggregate Fluctuations," *Journal of Political Economy* 99 (December 1991), 1166–1187.

Browning, Martin, and Costas Meghir, "The Effects of Male and Female Labor Supply on Commodity Demands," *Econometrica* 59 (July 1991), 925–951.

Campbell, John Y., and Sydney Ludvigson, "Elasticities of Substitution in Real Business Cycle Models with Home Production," (June, 2000), electronic file available at www.ny.frb.org-/rmaghome/economist/ludvigson/ludvigson.html.

Campbell, John Y., and N. Gregory Mankiw, "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence," in O. Blanchard and S. Fischer eds., *NBER Macroeconomics Annual* (1989), Cambridge, MA: MIT Press, 185–216.

Campbell, John Y., and N. Gregory Mankiw, "The Response of Consumption to Income," *European Economic Review* 35 (1991): 723–767.

Campbell, John Y., and N. Gregory Mankiw, "Permanent Income, Current Income and Consumption," *Journal of Business & Economic Statistics* 8 (July 1990), 265–279.

Canova, Fabio and Angel Ubide, "International Business Cycles, Financial Markets and Household

Production," *Journal of Economic Dynamics and Control* 22 (April 1998), 545-572.

Eichenbaum, Martin S., Lars Peter Hansen and Kenneth J. Singleton "A time Series Analysis of Representative Agent Models of Consumption and Leisure Choice Under Uncertainty," *Quarterly Journal of Economics* 103 (February 1988), 51-77.

Greenwood, Jeremy and Zvi Hercowitz, "The Allocation of Capital and Time over the Business Cycle," *Journal of Political Economy* 99 (December 1991), 1188-1214.

Greenwood, Jeremy, Richard Rogerson and Randall Wright, "Household Production in Real Business Cycle Theory," in T. Cooley, ed., *Frontiers of Business Cycle Research.* Princeton, NJ: Princeton University Press (1995), 157-174.

Hall, Robert E., "Stochastic Implications of the Life-Cycle Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy* 86 (December 1978), 971–989.

Hall, Robert E., "Intertemporal Substitution in Consumption," *Journal of Political Economy* 96 (April 1988): 339–357.

Keynes, John Maynard, "Economic Possibilities for Our Grandchildren," *Nation and Athenaeum*, October 11 and 18, 1930. (Reprinted in *The Collected Writings of John Maynard Keynes*, Volume IX, *Essays in Persuasion*, London, 1972.)

King, Robert G., Charles I. Plosser, and Sergio T. Rebelo, "Production, Growth and Business Cycles: I. The Basic Neoclassical Model," *Journal of Monetary Economics* 21 (March 1988): 195–232.

Mankiw, N. Gregory, "The Permanent Income Hypothesis and the Real Interest Rate," *Economics Letters* 7 (1981), 307-311.

Mankiw, N. Gregory, Julio J. Rotemberg and Lawrence J. Summers, "Intertemporal Substitution in Macroeconomics," *Quarterly Journal of Economics* 100 (February 1985), 225–251.

McGratten, Ellen, Richard Rogerson and Randy Wright, "An Equilibrium Model of the Business Cycle with Household Production and Fiscal Policy," *International Economic Review* 38 (May 1997), 267-290.

Table 1: Estimates of s

$$\Delta c - \tau \cdot \Delta n = \mu + s(r - \tau \cdot \Delta n) + \varepsilon_t + \theta \varepsilon_{t-1}$$

τ	Estimated s	p-value of restrictions	First-stage F
0.6	0.30 (0.11)	0.40	13.6
0.8	0.36 (0.11)	0.24	11.6
1.0	0.42 (0.10)	0.12	10.4
1.2	0.49 (0.10)	0.06	9.7

Notes:
Instruments are $\Delta c(-2)$, $\Delta n(-2)$, and $r(-2)$.
p-value is for the test of over-identifying restrictions.

Table 2: Different Instrument Sets

$$\Delta c - \tau \cdot \Delta n = \mu + s(r - \tau \cdot \Delta n) + \varepsilon_t + \theta \varepsilon_{t-1}$$

Instrument Set	Estimated s	p-value of restrictions	First-stage F
$\Delta c(-2)$, $\Delta n(-2)$, $r(-2)$	0.36 (0.11)	0.24	11.6
$\Delta c(-2)$, $\Delta y(-2)$, $r(-2)$	0.34 (0.10)	0.39	10.9
$\Delta c(-2)$, $r(-2)$, $c(-2) - y(-2)$	0.30 (0.12)	0.31	10.2
$\Delta c(-2)$, $r(-2)$, $\Delta n(-2)$, $c(-2) - y(-2)$	0.35 (0.10)	0.34	8.8

Notes:
p-value is for the test of over-identifying restrictions.

Table 3: Robustness Checks

Included Variables

$r - \tau\Delta n$	0.52 (0.16)	
r	-0.35 (0.20)	0.30 (0.12)
$\tau\Delta n$		-0.61 (0.18)

Notes:

Dependent variable is $\Delta c - \tau \cdot \Delta n$.

τ is set to 0.8.

Instruments are $\Delta c(-2)$, $\Delta n(-2)$, and $r(-2)$.

All regressions include a constant.

Table 4A: Forward Regressions

$$\Delta c - \tau \cdot \Delta n = \mu + s(r - \tau \cdot \Delta n) + \varepsilon_t + \theta \varepsilon_{t-1}$$

τ	Estimated *s*
0	0.30 (0.11)
0.8	0.36 (0.11)

Table 4B: Reverse Regressions

$$r - \tau \cdot \Delta n = \mu + (1/s)(\Delta c - \tau \cdot \Delta n) + \varepsilon_t + \theta \varepsilon_{t-1}$$

τ	Estimated 1/*s*	Implied *s*
0	1.32 (0.44)	0.75
0.8	2.13 (0.54)	0.47

Notes:
Instruments are $\Delta c(-2)$, $\Delta n(-2)$, and $r(-2)$.
All regressions include a constant.

Table 5: Adding Disposable Income

$$\Delta c - \tau \cdot \Delta n = \mu + s(r - \tau \cdot \Delta n) + \beta \Delta y + \varepsilon_t + \theta \varepsilon_{t-1}$$

Instrument Set	Estimated s	Estimated β
$\Delta c(-2)$, $\Delta n(-2)$, $r(-2)$	0.53 (0.31)	-0.67 (0.62)
$\Delta c(-2)$, $\Delta y(-2)$, $r(-2)$	0.41 (0.18)	-0.47 (0.46)
$\Delta c(-2)$, $r(-2)$, $c(-2) - y(-2)$	0.28 (0.14)	0.02 (0.27)
$\Delta c(-2)$, $r(-2)$, $\Delta n(-2)$, $c(-2) - y(-2)$	0.36 (0.14)	-0.13 (0.27)

Note:

τ is set to 0.8.

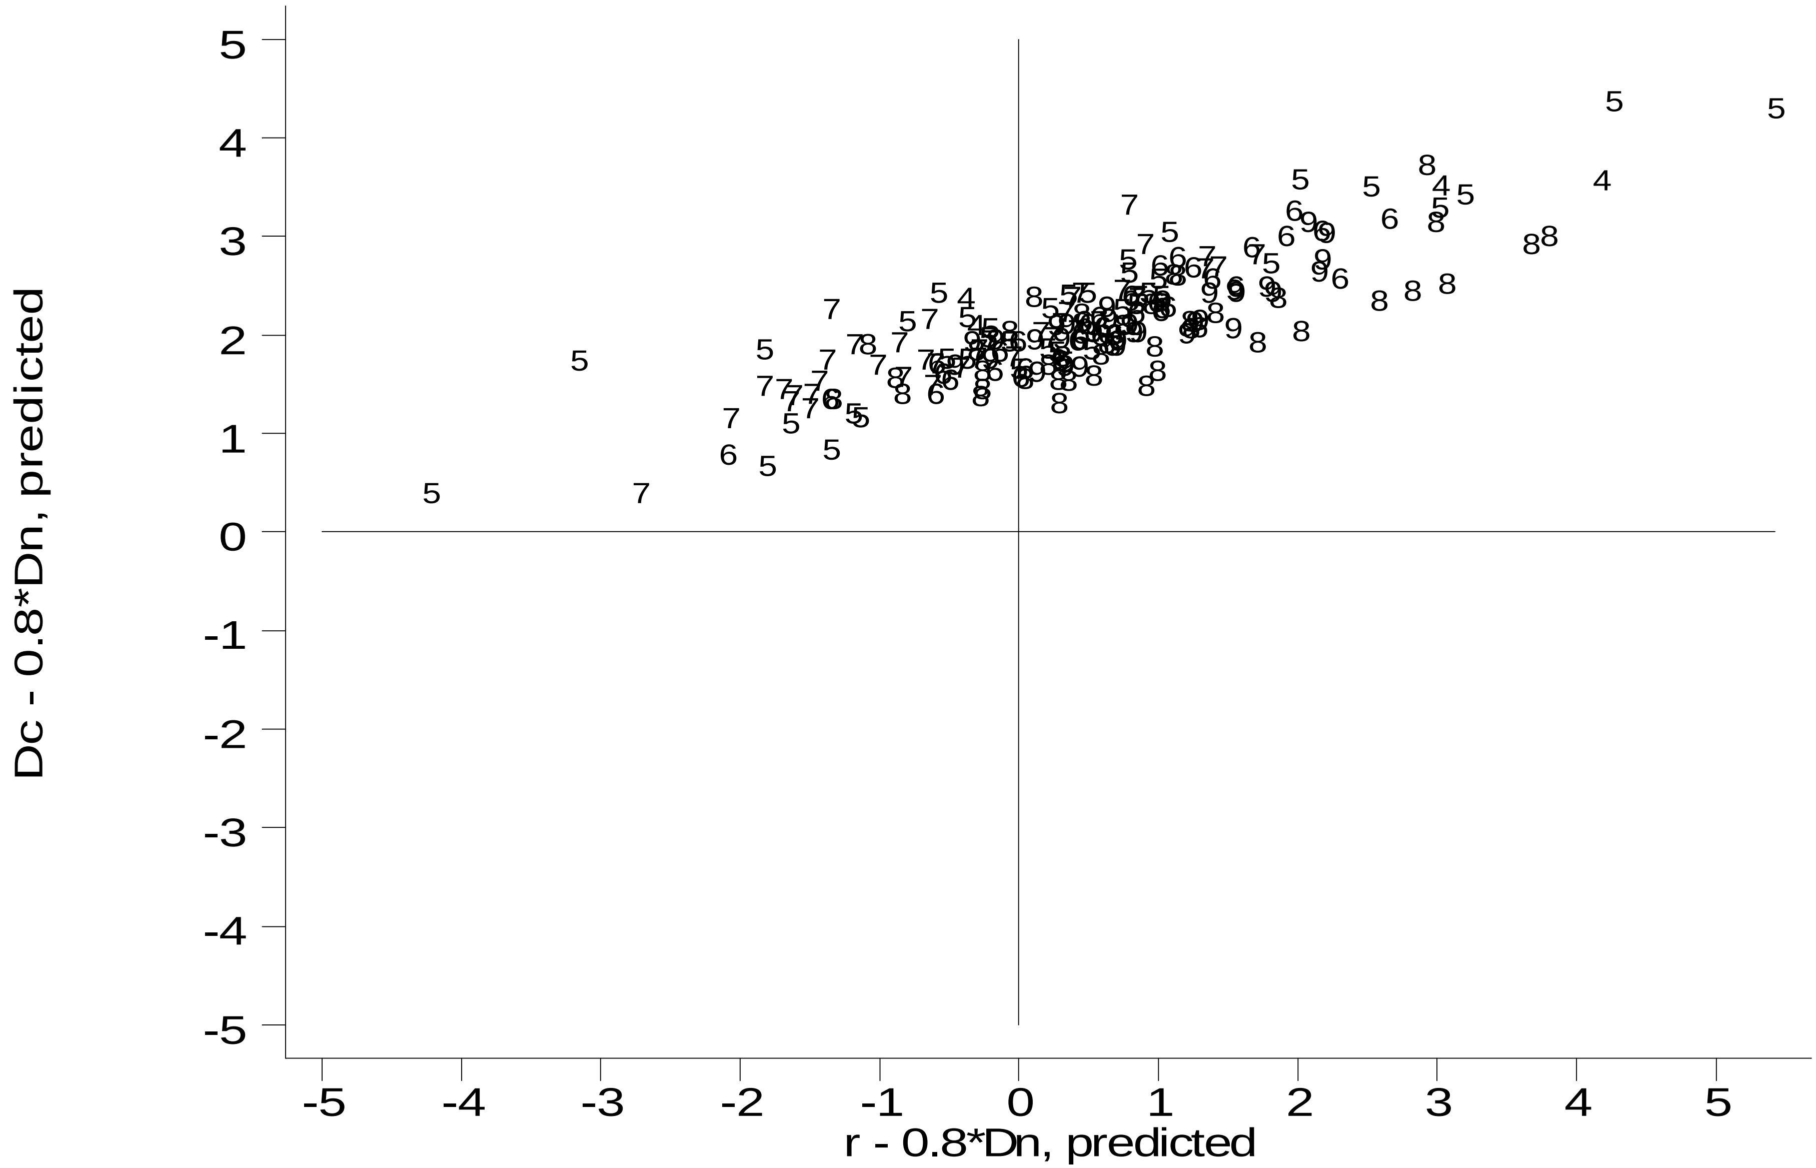

Dc - 0.8*Dn, predicted
r - 0.8*Dn, predicted

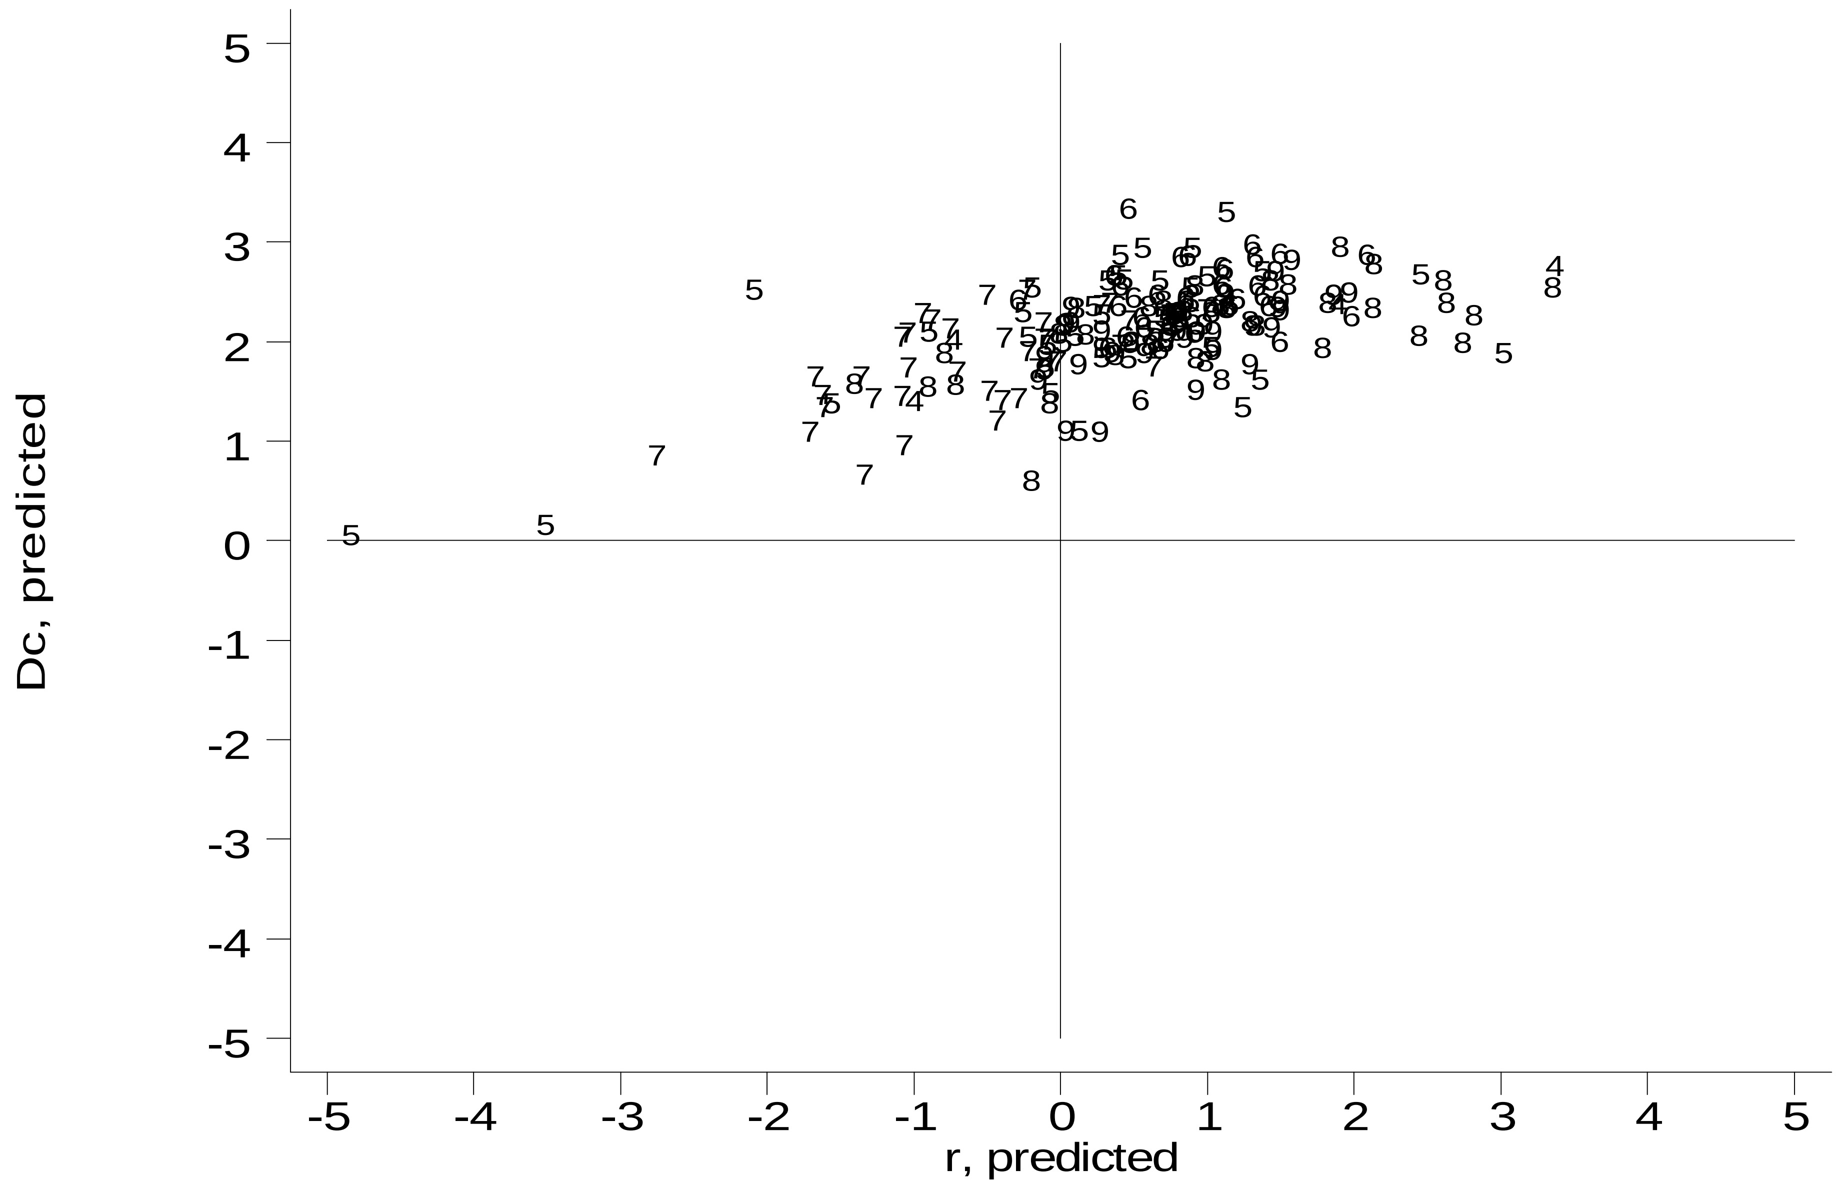
Dc, predicted
r, predicted

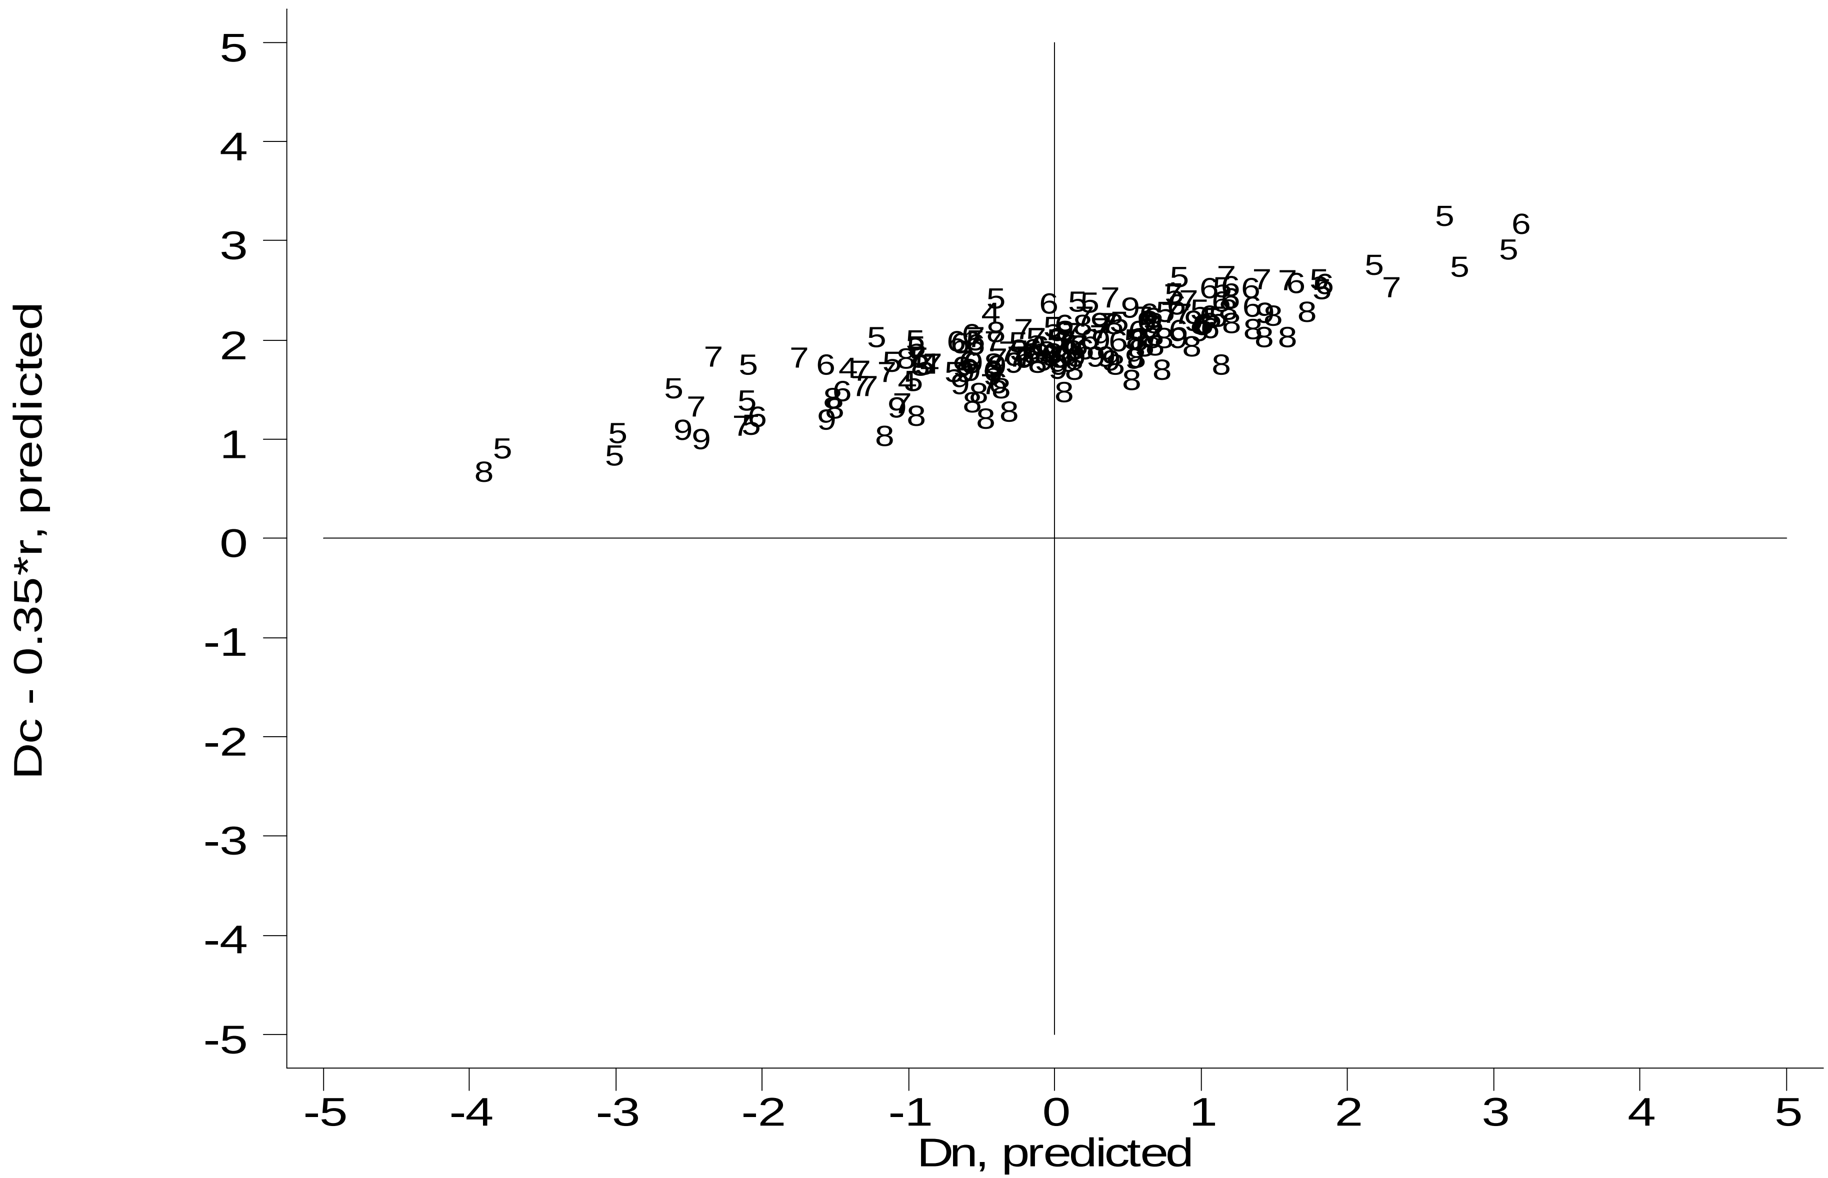
Dc - 0.35*r, predicted
Dn, predicted
5
4
3
2
1
0
-1
-2
-3
-4
-5
-5
-4
-3
-2
-1
0
1
2
3
4
5

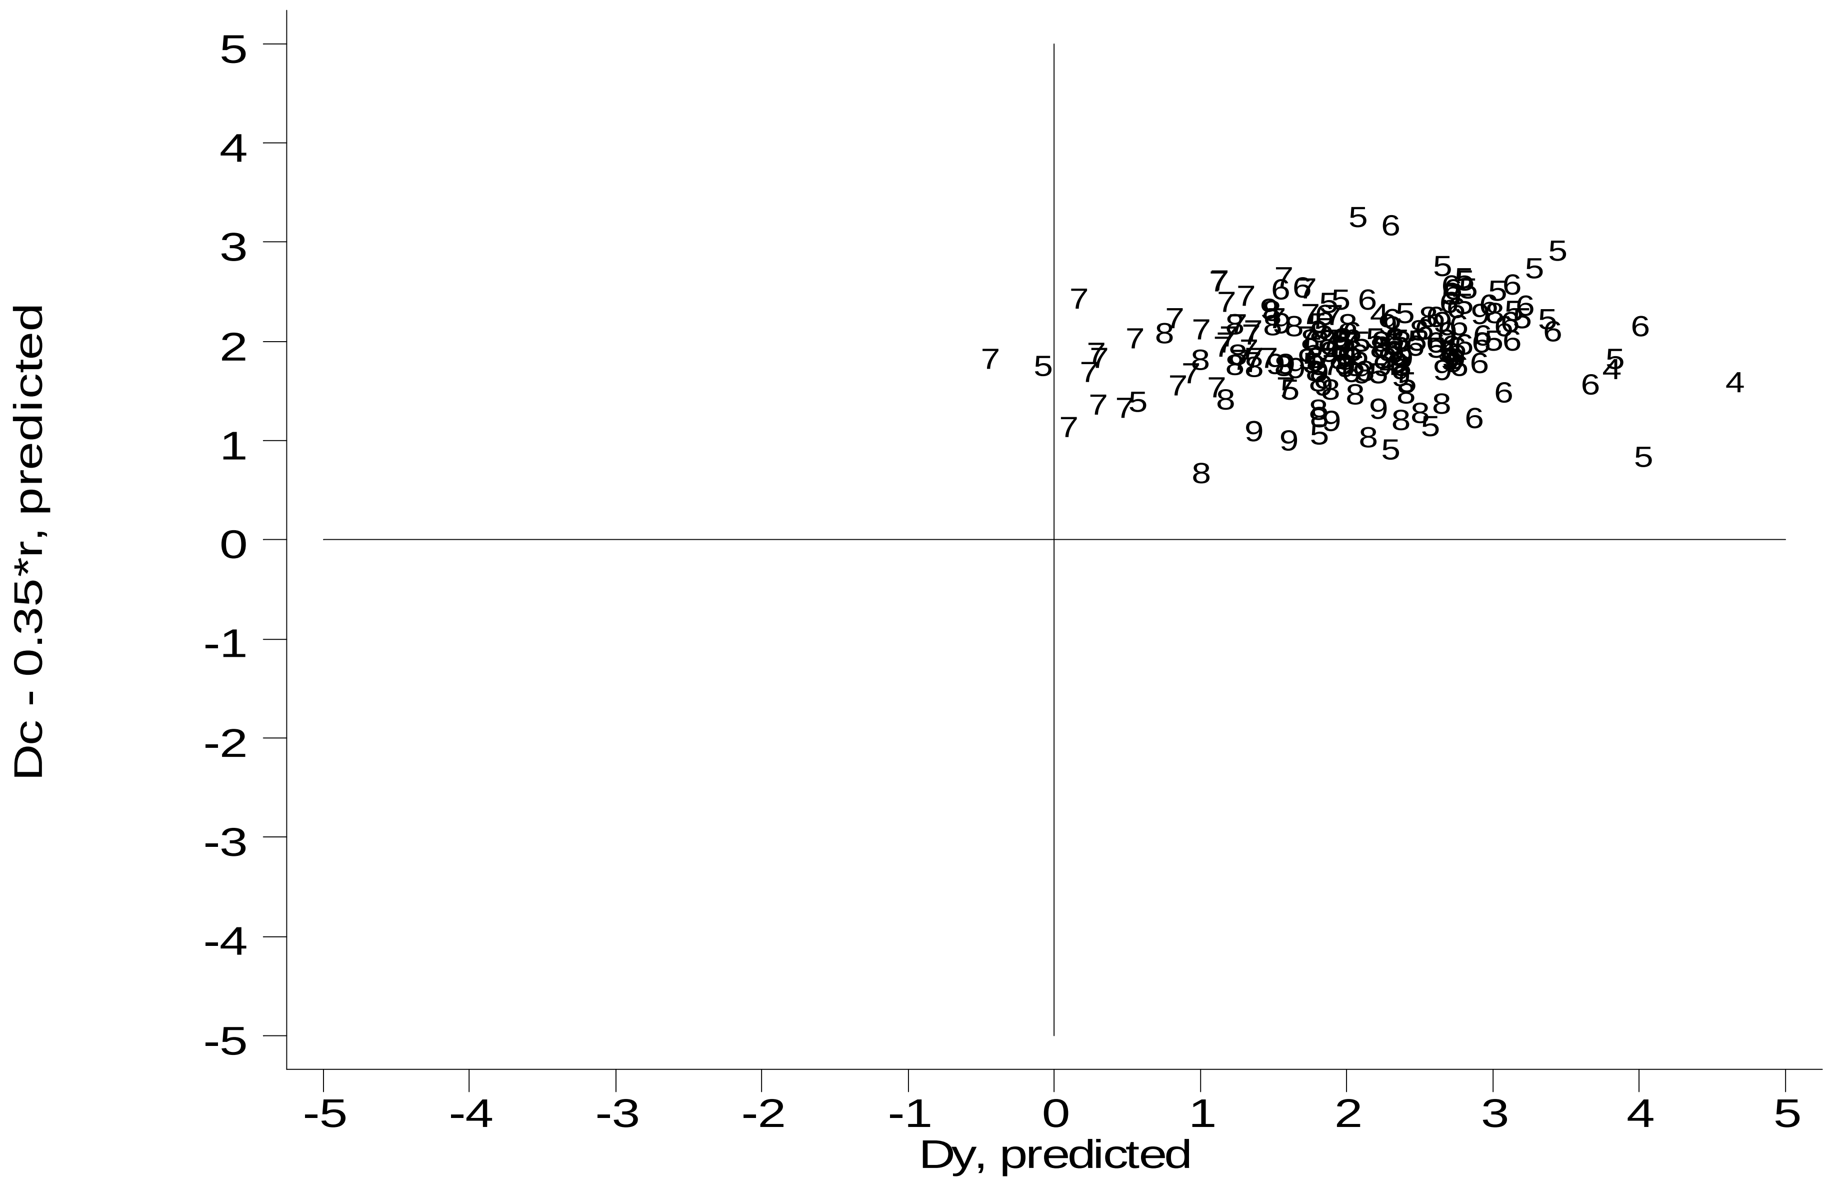

Dc - 0.35*r, predicted
Dy, predicted
5
4
3
2
1
0
-1
-2
-3
-4
-5