

Dynamic, or Life-Cycle Labor Supply

1. The setup

Dymo lives for T periods. Let his utility in period t be given by the function:

$$U(C_t, L_t) \tag{1}$$

where C is consumption and L is leisure. We'll assume U has the usual properties: it is increasing in both C and L , and strictly concave. Strict concavity implies negative (own) second derivatives, but (as noted in the static case) does not restrict the sign of the cross-partial derivative, U_{CL} . Plausible stories can be told for both positive and negative U_{CL} 's and both are consistent with well-behaved solutions to the problem.¹ If utility is intertemporally separable and ρ is the intertemporal (subjective) discount factor, total lifetime utility is:

$$W = \sum_{t=1}^T \frac{U(C_t, L_t)}{(1 + \rho)^t} \tag{2}$$

Note that the *function* U is not indexed by t ; in this baseline case we are therefore not allowing tastes for consumption or leisure to vary systematically with age (this is easy to modify).

Dymo can work as many hours as he wants in each period of his life at fixed rate of pay per hour. But the wage, w , Dymo can get isn't the same in each period: instead it is indexed by t . Suppose that Dymo knows, even at the start of his working life, what w will be in each period (this assumption is also fairly easily relaxed). If Dymo also faces a known stream of nonlabor income, G_t , $t = 1, \dots, T$, we can write his income in period t as:

$$\text{period } t \text{ income} = w_t(1 - L_t) + G_t \tag{3}$$

(Note that we are normalizing the total amount of time available in each period to equal one.) In the above definition of income, G_t does **not** include interest income on money that was saved earlier in life (the amounts of interest income Dymo actually ends up earning in each period of his life will be *endogenous* outcomes of Dymo's utility-maximizing choice of a lifetime income and consumption plan). G_t as defined in (3) includes only "exogenous" items of income such as inheritances, lottery winnings, demogrants, etc. (It is quite OK to think of all or most of the G 's as zero, but keeping track of them helps in the interpretation of some results below).

¹ If you can enjoy your spending more effectively when you have more leisure (e.g. a dollar yields more marginal utils when spent on vacation in a beautiful place) then $U_{CL} > 0$. If your need to spend at the margin is greater when you are working (e.g. you need to have better clothes, pay for commuting, eat out more, and don't have time to do your own home repairs), then $U_{CL} < 0$.

2. Static Labor Supply in this context.

Suppose there were no capital markets, i.e. no way for Dymo to use income earned in one period of life for consumption in another. Then Dymo's lifetime consumption and work plan would have to satisfy T separate constraints, given by:

$$C_t \leq w_t(1 - L_t) + G_t, \quad t = 1, \dots, T \quad (4)$$

Now, the consumption and labor supply plan, $C_1, \dots, C_T; (1 - L_1), \dots, (1 - L_T)$, that maximizes lifetime utility (2), consists of the C_t and L_t that maximize each *period's* utility (1) separately given each *period's* own budget constraint (4). This just means that Dymo solves the static labor supply problem we have already analyzed, over and over again, each period of his life. As we have already shown, this theory does *not* predict that Dymo will work harder in periods when his wage is higher: in fact if leisure is normal there is a good chance Dymo will take it easier in periods when wages are high.

3. Saving, Borrowing and the Intertemporal Budget Constraint

Now we'll consider the polar opposite case to a total absence of capital markets: perfect capital markets (clearly reality is somewhere between these extremes). If we allow Dymo to save *or* borrow as much as he wants at the same fixed real interest rate r (allowing r to vary over time in these models is a favorite hobby of macroeconomists; we won't pursue it here), so long as he pays off all his debts before he dies, the T separate budget constraints in (4) can be replaced by one, single lifetime constraint of the form:

$$\sum_{t=1}^T \frac{C_t}{(1+r)^t} \leq \sum_{t=1}^T \frac{w_t(1-L_t) + G_t}{(1+r)^t} \quad (5)$$

which says that the present value of lifetime consumption cannot exceed the present value of lifetime income. Borrowing and lending allow consumption to exceed, or fall short of, income in each period of life so long as the budget balances over Dymo's whole life.

To see where (5) comes from, here's an **exercise** (which will also help you understand what is going on). Define $S_t \equiv Y_t - C_t$ as net savings in period t (it can be negative, denoting borrowing). Next, convince yourself that Dymo's net assets (or net worth, exclusive of future earnings and future G 's) at any date t must be given

by $A_t \equiv \sum_{\tau=1}^t S_\tau (1+r)^{t-\tau}$. Finally show that (5) follows from the requirement that Dymo is not allowed to die in debt, $A_T \geq 0$.

An intermediate case between perfect and nonexistent capital markets would be characterized by the T constraints: $A_t \geq 0, \forall t = 1, \dots, T$. These force the consumer to have a non-negative *asset position* at all points in life; i.e. he/she can be a net saver but not a net borrower. (Another way to express the same thing is to say that consumers can only

borrow if they have physical or financial collateral, e.g. a house or car. Borrowing against *human* capital, as in a student loan, *would* allow A as defined here to go negative). The $A_t \geq 0, \forall t = 1, \dots, T$ case is examined by Deaton (*Econometrica* Sept 1991); fyi the solution divides the lifetime into two types of periods—those where one of the above “liquidity” constraints binds, where behavior is characterized by a static labor supply model—, and those where only the overall lifetime constraint is binding (where behavior is characterized by the perfect capital markets examined below).

4. Lifetime Utility Maximization under Perfect Capital Markets

Under perfect capital markets, the consumption and work plan that maximizes lifetime utility (W) can be expressed as:

Choose C_1, \dots, C_T and L_1, \dots, L_T to maximize:

$$W = \sum_{t=1}^T U(C_t, L_t)(1 + \rho)^{-t} + \lambda \sum_{t=1}^T (w_t(1 - L_t) + G_t - C_t)(1 + r)^{-t} \quad (6)$$

Note we have incorporated the *single* lifetime budget constraint via a Lagrangian with a *single* (not indexed by t or anything else) multiplier λ . There are a total of T first-order conditions for the C 's, each of which takes the form:

$$U_C(C_t, L_t)(1 + \rho)^{-t} - \lambda(1 + r)^{-t}, \quad t = 1, \dots, T \quad (7)$$

Similarly there are T FOC's for the L 's:

$$U_L(C_t, L_t)(1 + \rho)^{-t} - \lambda w_t(1 + r)^{-t}, \quad t = 1, \dots, T \quad (8)$$

Equations (7) and (8), plus the constraint (5), together constitute a system of $2T+1$ equations in $2T+1$ unknowns (the C 's, the L 's, and λ). Every one of these equations must be satisfied if lifetime utility, W , is at a maximum. There are a total of $2T+2$ exogenous *parameters* in this equation system, namely the interest rate r , the subjective discount factor ρ , the entire lifetime wage stream w_1, \dots, w_T ; and the entire lifetime nonlabor income stream G_1, \dots, G_T .

In principle, (7), (8) and (5) could amount to a real mess. After all, in most optimization problems like this, the optimal solution for any one of the endogenous variables (say, labor supply at age 47) depends on the values of all the exogenous parameters (wages at *every* age, G at *every* age, and the interest rate). For example, if you know you will inherit 20 million dollars at age 52, you might choose not to work at age 47 even if your wage is high then). Not only would this be cumbersome theoretically, it would make empirical work next to impossible (how can we possibly control for expected future wages and nonlabor income in a labor supply regression?).

Luckily, two assumptions (intertemporal separability of utility and perfect capital markets) greatly simplify the solution. To see this, start by re-arranging (7) and (8):

$$U_C(C_t, L_t) = \lambda \frac{(1+\rho)^t}{(1+r)^t}, \quad t = 1, \dots, T \quad (9)$$

$$U_L(C_t, L_t) = \lambda w_t \frac{(1+\rho)^t}{(1+r)^t}, \quad t = 1, \dots, T \quad (10)$$

Recall that the *function* U is not indexed by time and note that the derivatives of U on the LHS of (9) and (10) depend only on each period's own C and L . Note also that $(1+\rho)^t / (1+r)^t \equiv \theta(t)$ is an exogenous, deterministic function of time, t . (It will be increasing in t if the consumer is less "patient" than the market ($\rho > r$); decreasing otherwise). Solving (9) and (10) for C_t and L_t in terms of the remaining items yields "demand" functions for consumption and leisure of the form:

$$C_t = C(w_t, \lambda, \theta(t)) \quad (11)$$

$$L_t = L(w_t, \lambda, \theta(t)). \quad (12)$$

Each period's optimal C and L depends on wages (w) and endowed incomes *outside* the current period *only* via the Lagrange multiplier, λ . The functions $C(\cdot)$ and $L(\cdot)$ are known as Frisch demand functions, after Ragnar Frisch. More clumsily, they are also called "constant-marginal-utility-of-income" demand functions and happen to be mathematically identical to the input demand functions of a profit-maximizing firm facing a fixed price of output (see Browning, Deaton and Irish, *Econometrica* May 1985). To see the origins of the "constant MU" terminology, apply the envelope theorem to (6). This reveals that, at an optimum, λ is the marginal utility of lifetime income,

$$\text{i.e. } \lambda = \frac{dW}{d \sum_{t=1}^T \frac{w_t(1-L_t) + G_t}{(1+r)^t}} = \frac{dW}{dG_1} > 0. \quad \text{In words, } \lambda \text{ gives the effect on maximized lifetime}$$

income of an exogenous one-dollar increase in the present value of endowed lifetime income (an easy way to conceptualize such an increase is as a one-dollar increase in first-period G). Note that λ is a scalar (therefore not indexed by t).

Treating λ and θ as constants and totally differentiating the system (11) and (12) yields:

$$U_{CC}dC + U_{CL}dL = 0 \quad (13)$$

$$U_{CL}dC + U_{LL}dL = \lambda \theta dw \quad (14)$$

In matrix form this becomes:

$$\begin{bmatrix} U_{CC} & U_{CL} \\ U_{CL} & U_{LL} \end{bmatrix} \begin{bmatrix} dC/dw \\ dL/dw \end{bmatrix} = \begin{bmatrix} 0 \\ \lambda\theta \end{bmatrix} \quad (15)$$

Using Cramer's rule to solve the system gives

$$dL/dw = \frac{\begin{vmatrix} U_{CC} & 0 \\ U_{CL} & \lambda\theta \end{vmatrix}}{\begin{vmatrix} U_{CC} & U_{CL} \\ U_{CL} & U_{LL} \end{vmatrix}} = \frac{\lambda\theta U_{CC}}{pos} < 0 \quad (16)$$

In contrast to the static labor supply case, leisure *unambiguously* falls (and therefore labor supply unambiguously rises) with the wage, when λ and θ are held constant. But what does this mean? There are at least two distinct implications, both explored and explained in more detail in the practice problem. They are:

1. Responses to expected (sometimes called “evolutionary”) wage changes over the life cycle are unambiguously related to wage changes in the sense that $\Delta H = g(\Delta w)$, where g is a monotonically-increasing function. This is easiest to see in the case where $\rho = r$, in which case θ does not vary with t . In this case, labor supply rises between any two periods t and $t+1$ iff the wage rises between those two periods. When $\theta \neq 1$, this remains true after adjusting for an underlying time trend in H .

2. Responses to unexpected wage changes that are temporary enough not to materially change λ are also unambiguously positive. The idea of this is to suppose Dymo planned for one wage path but at some point in the middle of his life was surprised to encounter a wage path that was different in only one of very many periods. If there are enough periods, then this wage change is approximately the same as an evolutionary wage change in the sense that the shadow price of lifetime income, λ , isn't materially affected.

The intuition for both these results has to do with the fact that a one-period wage increase generates little or no income effects when income is pooled across all periods of the lifetime. In contrast to the static, or liquidity-constrained case where Dymo might decide to take it easy when wages rise, Dymo should always “make hay when the sun shines” if he has access to capital markets. That responses to temporary wage changes should be more likely to raise labor supply than permanent wage changes also conforms with simple introspection for most people: Suppose your wage was raised to \$50,000 per hour for just one day. My guess is that you would work 24 hours that day. But suppose your wage was raised to \$50,000 per hour, effective for the rest of your life. Comparing your remaining total lifetime labor supply under that scenario to what you are currently planning, I bet it will be less, not more (due to income effects).

Note also that we have, in a sense, generalized a model (the static labor supply model), making it much more complex. Usually this reduces the number of clear, testable predictions we can derive from a model. In this case, the opposite occurred: Ambiguous predictions for the effects of wages on labor supply become unambiguous.

Exercise: derive dC/dw . Show that it depends on the sign of the cross-partial derivative U_{CL} and *interpret*. In particular, show that if $U_{CL} < 0$, (i.e. if the marginal utility of consumption is greater when you have less leisure) people will spend more in periods when they have higher incomes (i.e. when they work more). Thus to take positive correlation between incomes and spending across periods as evidence of liquidity constraints (or of a Keynesian “consumption function”) would be incorrect.

5. Empirical Implementation and Testing

Suppose we had panel data on the wages and labor supply of a group of individuals over time. Consider a linear version (perhaps in logs) of the Frisch labor supply equation, (12), where we have now included a vector of time-varying characteristics X that affect tastes for work (such as the presence of small children or effects of age other than those related to discounting and impatience), a time-invariant tastes-for-work fixed effect for each person, A , and an optimization-error term for each person/year, ε :

$$H_{it} = 1 - L_t = \alpha_0 + \alpha_1 w_{it} + \alpha_2 \lambda_i + \alpha_3 \theta(t) + \alpha_4 X_{it} + A_i + \varepsilon_{it} \quad (17)$$

Collecting terms:

$$H_{it} = \alpha_0 + [\alpha_2 \lambda_i + A_i] + \alpha_3 \theta(t) + \alpha_1 w_{it} + \alpha_4 X_{it} + \varepsilon_{it} \quad (18)$$

This can be estimated by OLS, provided we incorporate a fixed effect for each individual in the sample. This fixed effect picks up variation across individuals in the lifetime marginal utility of income, λ , as well as in unobserved tastes for work, A , (so this kind of omitted variables bias isn't a problem any more). The function θ should be a smooth function of time. Theory predicts that $\alpha_1 > 0$; this parameter is generally called the “intertemporal substitution elasticity” and plays a key role in many real-business-cycle macro models.

The beauty of (18) is that, even though labor supply at any t depends on wages and (expected) endowed incomes in all periods of one's life (which would normally lead to an impossible estimation problem with hopeless data requirements), the effects of “non-current” w 's and G 's operate exclusively through the λ term. Provided the labor supply equation is linear, this will be absorbed completely in the person fixed effect (i.e. it is “differenced out” when comparing the same person's labor supply at two different times).

Much to the disappointment of real business cycle modellers, empirical estimates generally reveal very small, though positive estimates of α_1 , which turns out to make it very hard to explain business cycles with labor supply models. All kinds of devices have been employed to try to get around this problem (see Hotz, Kydland and Sedlacek, *Econometrica* 1988 for an particularly strained example, in my opinion...).

A final, and key issue that arises in testing the life-cycle labor supply model on panel data is distinguishing between evolutionary or unexpected temporary wage changes on the one hand (which are predicted to generate a positive labor supply response) and unexpected, permanent wage changes (such as for example a promotion) which do not generate an unambiguous predicted response. While some interesting ideas have been tried here (including case studies of taxi drivers and baseball stadium vendors), nothing truly convincing (in my opinion) has yet been done.