

The Basic Static Labor Supply Model

Consider a single individual with a utility function $U(y, \ell)$ where y is income and ℓ is leisure. Both y and ℓ are “goods”, i.e. the consumer prefers more of each: $U_1 > 0$; $U_2 > 0$.

Suppose this person has non-labor income of G , and can work as many hours, h , as she wishes at a wage of w per hour. Total time available for the only two possible activities, work (h) and leisure (ℓ) is T .

If she allocates her time between work and leisure to maximize her utility, what can we say about her decisions, and about how these decisions will respond to changes in the exogenous parameters, w and G ?

Using the constraints $h + \ell = T$ and $y = wh + G$, this problem can be written as a single-variable maximization problem without constraints:

$$\text{Choose } h \text{ to maximize: } U(wh + G, T - h). \quad (1)$$



First-order conditions for a maximum are:

$$wU_1(wh + G, T - h) - U_2(wh + G, T - h) = 0 \quad (2)$$

or simply:

$$w = \frac{U_2}{U_1} = \frac{MU(\text{leisure})}{MU(\text{income})} \quad (3)$$

Diagrammatically, (3) represents the familiar tangency condition between the slope of the budget constraint (w) and the slope of an indifference curve (U_2/U_1), shown for individual #1 below:



Of course, it is also possible (e.g. when G is high and w low) for the solution to be at a “corner” where it is optimal not to work at all, as shown for individual #2. We could show this formally using Kuhn-Tucker conditions in the maximization but the intuition is perfectly clear from the diagram.

To understand the comparative-static predictions of the model, assume an interior solution (as for individual #1), and totally differentiate equation (2), yielding:

$$(U_1 + wU_{11}h - U_{21}h)dw + (w(U_{11}w - U_{12}) - (U_{21}w - U_{22}))dh + (wU_{11} - U_{21})dG = 0, \text{ or simply:}$$

$$Adw + Bdh + CdG = 0 \quad (4)$$

To help interpret this, note that, by the second-order condition for a maximum (U as defined in equation 1 must be concave in h), $B < 0$. (btw the SOC also require $U_{11} < 0$ and $U_{22} < 0$). Note that w and G are exogenous and h endogenous, so we can define ceteris-paribus thought experiments where we introduce small changes in w and G one at a time, and consider their consequences for the utility-maximizing level of h :

Effects of nonlabor income on labor supply:

$$\frac{dh}{dG}(\text{holding } dw = 0) = \frac{C}{-B} = \frac{wU_{11} - U_{21}}{\text{pos}}, \quad (5)$$

which can be greater, less than, or equal to zero. It will be negative for sure if $U_{21} > 0$, i.e. if the utility function is such that additional income always raises the marginal utility of leisure. It will also be negative as long as U_{21} is not “too” negative, i.e. as long as additional income does not reduce the marginal utility of leisure “too much”. From here on we will *define* leisure as a *normal good* if the partial derivative defined by (5) is negative. (This is the standard definition: a good is normal if its optimal consumption increases when the budget constraint shifts outward without changing slope). Of course, if (5) is positive, we will say that leisure is *inferior*.

Effects of the hourly wage rate on labor supply:

$$\begin{aligned} \frac{dh}{dw}(\text{holding } dG = 0) &= \frac{A}{-B} = \frac{U_1 + h(wU_{11} - U_{21})}{\text{pos}} \\ &= \frac{U_1}{\text{pos}} + h \frac{(wU_{11} - U_{21})}{\text{pos}} \\ &= \frac{U_1}{\text{pos}} + h \left(\frac{dh}{dG}(\text{holding } dw = 0) \right) \end{aligned} \quad (6)$$

The first term in this expression is always positive. It is called the *substitution effect* (for reasons we shall see diagrammatically later). The second, called the *income effect*, is of indeterminate sign, but by definition it is negative if leisure is a normal good. All together, equation (6) constitutes the *Slutsky equation* for labor supply.

In sum, the Slutsky equation for labor supply tells us the following:

First, utility-maximizing behavior does NOT predict positive labor supply responses to a higher wage. Despite what tax-cut proponents in Washington might say, theory says it can go either way, even if (in fact only when) leisure is a normal good.

Second, despite this, the theory is not devoid of testable predictions, because Slutsky implies a very specific relation between two observable parameters: the response of labor supply to a (*ceteris paribus*) increase in nonlabor income (G), and its response to a (*ceteris paribus*) increase in the hourly wage (w).


For example, (easily but not very interestingly), suppose we determined that a particular person (or organism) responded to increases in G by working harder (i.e. that leisure was an inferior good for that person/organism). Then (by the Slutsky equation), it must be the case that that person/organism responds to higher hourly wages by working harder. If he/she/it did not, the theory would be proved wrong.

The Slutsky equation for labor supply differs in an important way from the Slutsky equation you are probably more familiar with – the one for consumption of purchased goods. In the labor supply case, ambiguous predictions for the effects of prices on quantities arise when the good in question (leisure) is normal. In the case of purchased goods, this kind of ambiguity only arises when the purchased good (say apples) is *inferior* (a negative effect of price on consumption is guaranteed if the good is normal). Thus there is a fundamental sense in which the predictions of basic economic theory regarding labor supply are “more ambiguous” than its predictions regarding other consumer goods. The reason, as we shall illustrate diagrammatically below, because people “sell” leisure, while they buy goods like apples (unless they of course own an orchard). This distinction means that an increase in the price of leisure (the wage rate) has opposite effects on consumer wellbeing (making people better off) than an increase in the price of apples (which makes people worse off).

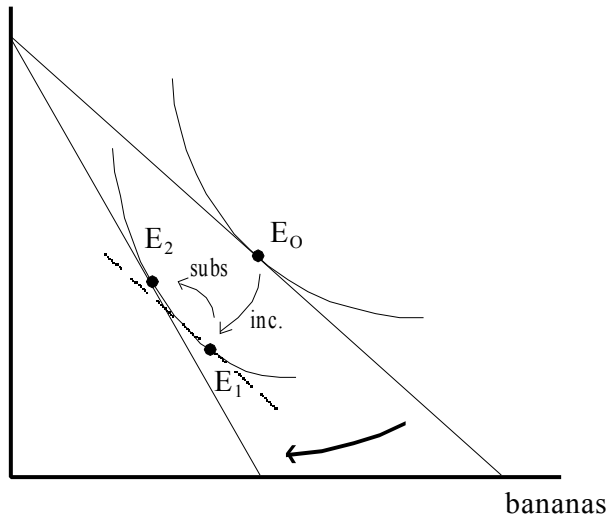
The income and substitution effects in the Slutsky equation are illustrated in the following five diagrams, which show why they interact in different ways for goods such as leisure which are “sold” to in factor markets than for goods such as apples which are purchased in product markets.

Income and substitution effects for “purchased” goods:

The figure below illustrates the effect on a consumer’s budget constraint of an increase in the price of bananas. An increase in the price of bananas moves the budget constraint towards the origin, making the consumer worse off because it restricts the set of bundles that can be afforded:




Decomposing the effects of this price increase into a substitution and an income effect shows that utility-maximizing consumption of bananas *must* fall with their price if bananas are a normal good.:



The movement from E_0 to E_1 is the income effect. It must involve a lower consumption of bananas if bananas are normal. The movement from E_1 to E_2 is the substitution effect. It must be negative because (by construction) utility is the same at both points. Because both the income and substitution effects work in the same direction (towards reduced consumption) the total effects of a price increase (the movement from E_0 to E_{12}) on consumption must be negative.


Income and substitution effects for labor supply (or the consumption of leisure):


The figure below illustrates the effect on a consumer's budget constraint of an increase in the hourly wage rate (i.e. the price of leisure). (To reduce clutter we have assumed zero nonlabor income, G):



Decomposing the effects of this price increase into a substitution and an income effect shows that utility-maximizing consumption of leisure may either rise *or* fall with its price if leisure is a normal good. This is because income and subs effects work in opposite directions. We show two cases: (a) labor supply falls with the wage (income effect outweighs subs effect), (b) labor supply rises with the wage (subs effect outweighs income effect).

CASE A:



CASE B:

Note that, in both cases A and B, leisure is a normal good.

Other predictions of the Static Labor Supply Model:

So far, our examination of the static labor supply model has yielded only one testable prediction, which is distinctly unuseful. It is:

1. Labor supply responses to wage increases must be positive if leisure is an inferior good. In other words, suppose that we found an economic agent whose labor supply responds to increases in G are positive. Then the theory predicts that that agent's labor supply response to a wage increase must be positive as well. Since there is strong evidence that leisure is a normal good in the vast majority of cases, this is not a useful prediction for testing the theory.


There are, however, at least three more, much more useful predictions, which are:

2. Substitution effects must be positive. Consider the four separate elements of the Slutsky equation in (6):

$$\frac{dh}{dw} (\text{holding } dG = 0), \quad \frac{U_1}{pos}, \quad h, \quad \text{and} \quad \frac{dh}{dG} (\text{holding } dw = 0).$$

In principle, three of these can be determined empirically for a particular economic agent. The first and last are responses of hours worked to small changes in w and G respectively. The third is the (initial) level of hours before the small changes in w and G are implemented.¹ Beginning with these three quantities, the fourth, U_1/pos , or the substitution effect, can be calculated using the Slutsky equation. The theory says it must be negative. If it isn't, the theory is wrong. This is the logic behind the exercise in, for example, Ashenfelter and Heckman, 1974).

3. Wage effects on labor force participation are always positive. Unlike the decision on how many hours a worker should work, offered hourly wages have an *unambiguous* predicted effect on the decision to work or not. There are two ways to see this. Mathematically, set initial hours of work, h to zero in (6). Then income effects disappear and the response of hours to wages must be positive. Diagrammatically, the following figure shows the budget constraints for two different wages. At the lower wage, the individual is indifferent between working and not (this is called the *reservation wage*); by construction this budget constraint is tangent to the indifference curve when $h=0$. At all wages below the reservation wage, the consumer also prefers not to work. At all wages above the reservation wage, he/she will want to work positive hours.



¹ Actually, our theory is framed in terms of derivatives, i.e. the effects of infinitesimal changes in w and G , so initial and final levels of h do not need to be distinguished. In reality, of course, all observable changes in w and G are finite in magnitude; all the main results of the theory go through if we reformulate it in terms of finite changes but it's messier to work with so we don't bother here.


Note that, as wages continue to rise above the reservation wage, this individual might eventually curtail his/her hours of work relative to that shown for the higher wage in this figure due to income effects. But a higher wage –no matter how high-- will never induce a utility-maximizing individual person to stop working completely. (This would be killing the goose that lays the golden egg!). Thus, an increase in offered wages will sometimes raise, but never reduce the labor force participation rate.

4. Labor Supply Responses to a *Compensated Wage Increase* must be Positive.

Suppose that I control both a subject's hourly wage, w , and his/her nonlabor income, G . I want to induce my subject to work harder, but I know from the preceding discussion that simply raising w might not work, because this makes the subject better off, and subject might decide to consume some of this increased utility in the form of greater leisure. (This is the intuition of income effects).

How can I get around this problem? By combining the increase in w with a *cut* in G (thereby preventing the subject from becoming much better off, and eliminating the income effect). In fact, it is easy to show the following combined change in w and G (henceforth called a “compensated wage increase”) is guaranteed to raise labor supply: increase w , but at the same time cut G by an amount such that, if the subject did not change his/her labor supply behavior, his income would remain exactly the same as before.

The following figure shows two budget constraints: the “initial” one has a higher level of nonlabor income (G_1) but a lower wage. The indifference curve shows a utility-maximizing point, **a**, along this budget constraint. The second budget constraint is steeper (therefore has a higher hourly wage) but a lower nonlabor income (G_2). The level of G_2 is chosen such that this budget constraint passes through the original utility-maximizing point. Note that, by construction, the utility-maximizing point along the new, steeper budget constraint *must* lie to the left of the original equilibrium, i.e. it must involve more hours of work than point **a**.




In experimental settings, it is possible to implement exactly-compensated wage increases and observe subjects' responses to them, thus testing the theory. Applying this same logic to nonexperimental settings (where we don't have such precise control over w and G), note that (as long as leisure is normal), an “over-compensated” wage increase (i.e. if we accidentally cut G below G_2) also always raises labor supply, because income and substitution effects both work towards greater work hours in this case. Situations like this sometimes occur in the design and reform of income-support programs, thus providing another way to test the theory.

Relatedly, in a recent series of papers starting with Prescott (2004), some macroeconomists (including Rogerson and Sargent) have argued that higher levels of income taxes explain why Europeans spend a smaller share of their lifetimes working than Americans. A key (and perhaps underappreciated) element of these papers is the assumption that tax revenues are refunded to households in the form of a lump-sum subsidy; thus higher taxes effectively amount to a *compensated wage decrease* in their models. If, instead of refunding tax revenues to households, European governments wasted the tax money they collected, we would no longer necessarily expect higher taxes to reduce labor supply.

Extensions to the Static Labor Supply Model

1. Enjoyable Work. Suppose that, at $\ell = T$, leisure is a “bad”, not a good. Does anything important change?

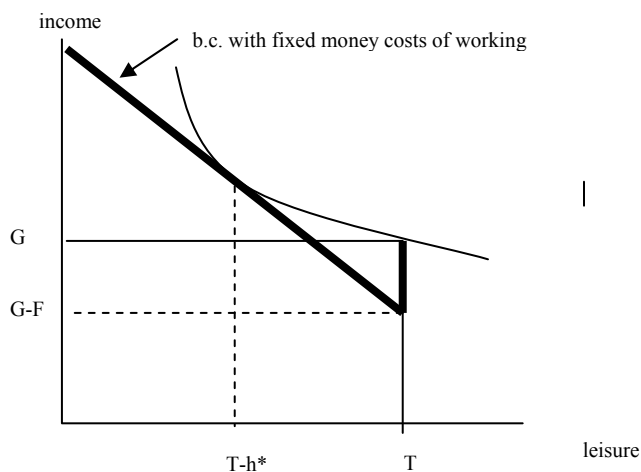


If no paid work is available, people will work positive hours.

If paid work is available, nothing really changes because if the household is optimizing, leisure will always be a good *at the margin*.

2. Fixed Costs.


One unrealistic feature of the “standard” model above is that, when the offered wage rises marginally above the reservation wage, the household enters the labor market by working a level of hours that is marginally above zero. This is easily avoided by allowing for fixed money or time costs of work. Below, F is a fixed money cost of going to work; the dark line is the budget constraint *at the reservation wage*. At the reservation wage, the individual is indifferent between working 0 hours and $h^* > 0$ hours.



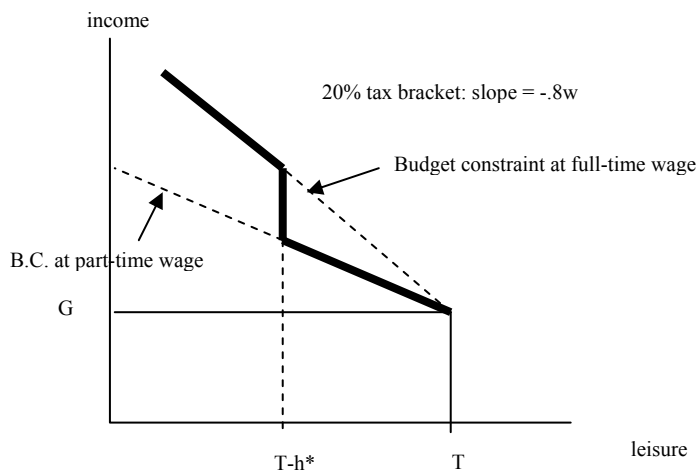
3. Proportional Taxes and Unconditional Transfers. Introducing a proportional *earnings* tax into the model is of course straightforward; it leaves G unchanged but lowers the slope of the budget constraint. Income taxes lower G as well as the slope, and unconditional transfers of income just raise G . Two testable features of the basic model in this simple context are that households should react to increases in taxes on earnings just like they react to reductions in wages, and they should react to unconditional transfers as they do to any other increase in income.

4. Piecewise-linear budget constraints. Many government tax and transfer programs, and some features of firms' salary/incentive policies, generate budget constraints that are piecewise-linear. Here are a few examples:

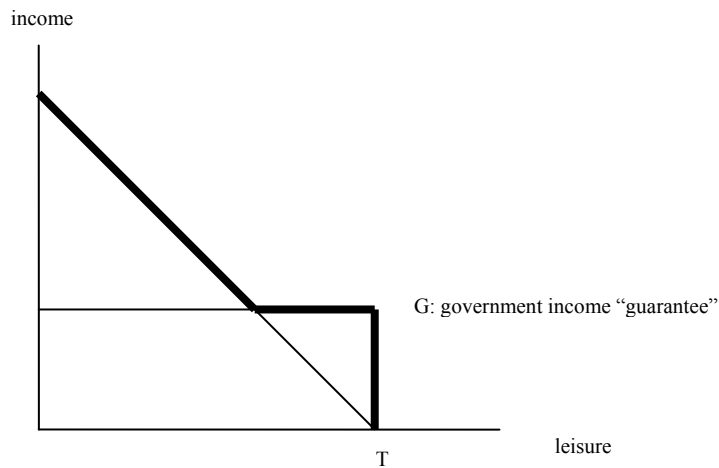
a) Progressive Income Tax with Brackets:



b) Full-Time Wage Exceeds Part-Time Wage: (h^* is min hours needed to qualify for FT wage)




c) Classic “Welfare” (100% “tax-back” of income up to the ‘guarantee’ level)




Prediction: Introducing welfare will either reduce labor supply or leave it unchanged. Introducing welfare will raise income (for the very poor), reduce income (for the ‘borderline poor’), or leave it unchanged (for everyone else).

d) Unemployment Insurance with 60% replacement rate and W^* work weeks needed to qualify:



Prediction: Introducing UI can either raise or lower weeks worked per year. For persons initially working fewer than W^* weeks, UI either raises labor supply or leaves it unchanged. For persons initially working more than W^* weeks, UI either reduces labor supply or leaves it unchanged. Thus, in general, UI “pulls” people out of the tails of the annual weeks worked distribution, into the ‘middle’, specifically, those weeks-worked categories most highly subsidized by the UI system (See Kuhn-Riddell, ILRR forthcoming 2010).

e) Earned Income Tax Credit (EITC):




Predictions: quite similar to UI. See Eissa and Liebman (1996); Meyer and Rosenbaum (2001), Eissa and Hoynes (2006).

Two very general predictions with ‘kinked’ budget constraints:

a) “Holes” People should avoid “inward” kinks in budget constraints, e.g. working very small number of hours in the presence of fixed costs.

a) “Bunching”: Persons facing the same budget constraint but with quite different tastes have a strong incentive to choose exactly the same allocation, at ‘outward’ kinks.

Consider the example of ‘classic welfare’. In the absence of a welfare program, person **a** (solid indifference curves) locates at point **a**. Introducing welfare moves him to point **c**, reducing his income. In the absence of a welfare program, person **b** (dashed indifference curves) locates at point **b**. Introducing welfare moves her to point **c** as well, raising her income. So, two people who would have behaved very differently in the absence of welfare choose exactly the same allocation in its presence. See Friedberg ReStat 2000; Chetty et al., “The Effect of Adjustment Costs and Institutional Constraints on Labor Supply Elasticities: Evidence from Denmark” 2009 (available on course website); and Green and Riddell, “Qualifying for Unemployment Insurance: An Empirical Analysis” *Economic Journal*, 1997. Volume 107, Number 440, pp. 67-84.



5. Multiple Types of Labor. Ashenfelter and Heckman (1974) consider the case of a two-person household, maximizing the joint utility function:

$$U = U(L_m, L_f, Y) \quad (7)$$

where L_m is the husband's "leisure" (or non-market-time), L_f is the wife's, and Y is total household income, subject to the budget constraint:

$$Y = W_m(T - L_m) + W_f(T - L_f) + G \quad (8)$$

where W_m and W_f are the husband's and wife's hourly wages respectively, and G is household nonlabor income.

Maximizing (7) subject to (8) yields a pair of labor supply equations:

$$H_m = H_m(W_m, W_f, G) \quad (9)$$

$$H_f = H_f(W_m, W_f, G) \quad (10)$$

Notice that, in general, the household's optimal policy will exhibit some cross-elasticity, i.e. the husband's labor supply will depend on how much the wife is capable of earning, and *vice versa*. In general, these can be either sign, just like the responses to W and G by a single-person household.

As before, though, this does not mean the theory is without empirical content. Generalizing the Slutsky equation in (6) to a two-person household, AH show that (once again for $dG = 0$):

$$\frac{\partial H_i}{\partial W_j} = S_{ij} + H_j \frac{\partial H_i}{\partial G} \quad (11)$$

where S_{ij} is the pure substitution effect of spouse j 's wage on spouse i 's labor supply, H_j is spouse j 's initial level of labor supply, and $\partial H_i / \partial G$ is the effect of household nonlabor income on spouse i 's labor supply. Note that the income effect of W_j on H_i is now weighted by the initial labor supply of spouse j —i.e. the spouse whose wage rate is being changed. Thus, if my wife is not working at the initial optimum, and her wage rises, that wage increase has no income effects on either her *or* my labor supply decision. $\partial H_i / \partial W_j$ is a 2×2 matrix of ("uncompensated" or Marshallian) wage effects on labor supply; H_j are just the two initial levels of labor supply; and $\partial H_i / \partial G$ are the labor supply responses of the two spouses to changes in household nonlabor income. From these empirically-observed quantities we can calculate the 2×2 matrix of pure substitution effects:

$$\mathbf{S} = \begin{bmatrix} S_{hh} & S_{hw} \\ S_{wh} & S_{ww} \end{bmatrix},$$

which must have the following properties if the household maximizes equation (7) (we'll demonstrate this explicitly for the case of labor demand in a couple of weeks):

\mathbf{S} is negative semidefinite. (thus, among other things, $S_{hh} \leq 0$ and $S_{ww} \leq 0$), and:

\mathbf{S} is symmetric (thus $S_{hw} = S_{wh}$).

AH test these restrictions using estimates of $\partial H_i / \partial W_j$, H_j and $\partial H_i / \partial G$ derived from a cross-section of U.S. households in the 1960 U.S. Census. We will talk about how credible these estimates are in class shortly.

6. But “Leisure” isn’t Leisure!: Household Production

In his famous 1965 article, “An Economic Theory of the Allocation of Time”, Gary Becker introduced another way to think about labor supply decisions.

Rather than deriving utility directly from income and leisure, Becker thought of households as deriving utility from a vector of m ‘final consumption goods’, \mathbf{z} , whose elements, Z_i included things like meals, child care, entertainment, etc. Each of these Z_i ’s, in turn, is produced from two kinds of inputs, \mathbf{x}_i (a vector of market goods such as groceries, kitchen appliances, or payments to restaurants), and \mathbf{T}_i , a vector of time inputs (elements of \mathbf{T}_i can differentiate among persons supplying the time, and/or among ‘types’ of time, e.g. days, evenings, weekends), via a production function, so $Z_i = f_i(\mathbf{x}_i, \mathbf{T}_i)$.

So the household maximizes

$$U = U(\mathbf{z}) = U(Z_1, Z_2, \dots, Z_m) = U[f_1(\mathbf{x}_1, \mathbf{T}_1), f_2(\mathbf{x}_2, \mathbf{T}_2), \dots, f_m(\mathbf{x}_m, \mathbf{T}_m)], \quad (12)$$

subject to constraints (13), (14) and (15) below. The first of these is just that the household can afford all the goods inputs it uses:

$$\sum_{i=1}^m \mathbf{p}_i \cdot \mathbf{x}_i \leq Y \quad (13)$$

(note that total money/goods cost for each final consumption good, Z_i , is the vector cross-product of all the possible goods used in its production and their prices). Second is a vector of time constraints, one for each type of time. For the special case where the household has just two types of time (husband’s and wife’s) available (each has a total of T hours available), these are:

$$T^h \equiv \sum_{i=1}^m T_i^h \leq T, \text{ and } T^w \equiv \sum_{i=1}^m T_i^w \leq T. \quad (14)$$

Finally, the household gets money with which to purchase the goods inputs by working in the labor market, so:

$$Y \equiv w^w(T - T^w) + w^h(T - T^h) + G, \quad (15)$$

where w^w and w^h are the wife’s and husband’s market wage rates, respectively. $G = G^w + G^h$ is nonlabor income. Note that a household’s optimal decisions in this model will only depend on the total, G , not on G^w and G^h individually.

Among other applications, household production theory helps us think about how economic factors, such as market wages, affect the household’s internal division of labor, and how changes in the home production technology (i.e. the f_i functions) affect market work decisions (for an example of the latter, see Greenwood, J. and G. Vandenbroucke NBER working paper 11629, 2005). It also encourages us to distinguish ‘true’ leisure from work done in the home (see for example Aguiar and Hurst QJE August 2007) using time-use data. An important application of household production theory is in developing countries, where households often divide their time between, for example, tending their own land and providing agricultural labor to the market. Importantly, some predictions of home production theory are *independent of the form of the utility function*, thus for example it does not matter if leisure is normal or inferior: goods should be produced in an efficient way, regardless of the tastes of the consumer.

7. But most workers can’t choose their hours! For treatments of **hours constraints** and search frictions, see Lundberg (ReStat); Kahn and Lang; Chetty et al 2009, and the section of the reading list on search models.

8. But Husbands and Wives don't always agree!: Nonunitary Models of Household Behavior.

Here, we'll sketch a simple version of the nonunitary model in Bobonis (2009). It models consumption, *not* labor supply. (Tests of the unitary model are cleaner in this case).

Consider a two-person household, with members A and B , who have exogenous "private" incomes y^A and y^B , and jointly-held income y^O . The utilities of persons A and B are given respectively by:

$$u^A = u^A(q^A, q^B, C^A, C^B, C^H, K) \quad (16)$$

$$u^B = u^B(q^A, q^B, C^A, C^B, C^H, K) \quad (17)$$

where q^A and q^B are private, "assignable" goods (e.g. clothing, food), K is a household public good (e.g. children), and $C^A + C^B + C^H = C$ are all other goods, which can be consumed either privately or publicly.

The household's budget constraint is given by:

$$p^A q^A + p^B q^B + \mathbf{p}^C \mathbf{C} + K \leq y^A + y^B + y^O \quad (18)$$

where the p 's denote prices and the price of K has been normalized to one.

In the unitary model of household consumption, there is only one utility function (pick either 16 or 17), maximized wrt q^A, q^B, C^A, C^B, C^H , and K , subject to (18). The exogenous variables are $p^A, p^B, \mathbf{p}^C, y^A, y^B$, and y^O . It is straightforward to show that optimal consumption decisions depend only on $Y = y^A + y^B + y^O$, not on the individual components of total household income, Y . In other words, the consumption decisions of a household in which the husband 'owns' all the income are identical to the consumption decisions of a household in which the wife owns it all—the source of the income does not matter for consumption decisions.

In contrast, the collective model of the household assumes that the two household members bargain over the how the household's income is spent, and that such bargains are Pareto-efficient, i.e. they are on the spouses' utility-possibility frontier. (For models in which efficiency is not necessarily attained, see Lundberg and Startz's work on 'separate spheres' intrahousehold bargaining). By the definition of Pareto-optimality, this means that households' consumption decisions satisfy:

$$\text{Max } u^A + \lambda u^B, \quad (19)$$

subject to (18), and to $\lambda = \lambda(\mathbf{z})$, where \mathbf{z} is a vector of 'distribution factors' that affect the relative bargaining power of the spouses. These might include the jurisdiction's laws about divorce and the division of property in case of divorce, enforcement of domestic violence laws, the sex ratio in the marriage or remarriage market, and perhaps religion and custom. Further, if any of y^A, y^B , or y^O enter \mathbf{z} individually (rather than as a sum), then the source of the household's income *will* affect its consumption choices.

The above property suggests a simple test of the unitary model: if a household's consumption decisions depend on the source of its income (and if they do so in a plausible direction, e.g. men get more male-specific goods when their share of the household's income goes up), this rejects the unitary model in favor of a more general, nonunitary model. Most credible tests of this nature consider the effects of exogenous changes in the *nonlabor* income of one spouse, since wage rates are likely to affect labor supply, which in turn is likely to affect consumption decisions directly, confounding the test.

Testing among different types of nonunitary models (e.g. Pareto-efficient versus not) is more challenging, but Bobonis (2009) has recently implemented the following test of efficiency, which requires the existence of two distinct 'distribution factors', \mathbf{z} :

$$\frac{\partial q^A / \partial z_1}{\partial q^A / \partial z_2} = \frac{\partial q^B / \partial z_1}{\partial q^B / \partial z_2} = \frac{\partial c_i / \partial z_1}{\partial c_i / \partial z_2} = \frac{\partial K / \partial z_1}{\partial K / \partial z_2} = \frac{\partial \lambda / \partial z_1}{\partial \lambda / \partial z_2} \quad (20)$$

In words, suppose that z_1 is the share of the household's assets that a wife automatically gets to keep in the event of divorce in the jurisdiction where the household lives. Suppose that z_2 is the wife's share of this household's income (holding total household income constant). Say that both of these raise the wife's bargaining power, so we expect both of them to raise the wife's consumption of private goods that are specific to her (q^A if she is person A), and to reduce the husband's consumption of private goods, q^B .

What (20) says is that, if a one-unit increase in z_1 raises q^A twice as much as a one-unit increase in z_2 , then –if intrahousehold bargains are efficient– a one-unit increase in z_1 must *reduce* q^B by twice as much as a one-unit increase in z_2 . Further, whatever z_2 does to c_i and to K , z_1 must do the same, but *twice as strongly*. These are unexpected and strong restrictions (akin in that way to symmetry of the substitution matrix) that can be tested empirically. The intuition is that the *only* reason that any of these consumption decisions are affected by the z 's is because the z 's change the distribution of intrahousehold bargaining power, λ . (They cause a movement along the household's utility-possibility frontier). Thus the ratio of their effects on (observable) consumption decisions for all items must equal the ratio of their effect on the ratio of bargaining powers.

For treatments of household *labor supply* in a nonunitary context, see Chiappori 1992; Blundell, Chiappori and Meghir 2005; and Blundell, Chiappori, Magnae and Meghir 2007.