

Chapter 2

Static Consumption-Leisure Model

In our study of consumer theory thus far, we have simply assumed that an individual has some given amount of labor income, which we denoted by Y , to spend on consumption goods. Doing so allowed us to focus attention on the tools and principles of consumer theory. Economics is at its core a set of theories about decision-making, and casual reflection reveals that individuals do have some control over how much labor income they earn. That is, at least to some degree, individuals “choose” how much income they earn just as they choose how much, say, good 1 and good 2 they consume. In this section, we extend our model of consumer theory to incorporate this feature of individual decision-making. As we will see, the tools of analysis and general principles of this extended model are ones with which we are already familiar – simply the tools of indifference curves and budget constraints. To simplify our introduction to this topic, we will use a “one-shot” model in which the individual has no savings decision to make – that is, there is no future, so that the only economic decisions to be concern the present. Once we understand how the one-period consumption-leisure model works and we study the consumption-savings model to come, we will bring the two models together to complete our analysis of macroeconomic consumer theory.

In addition to considering the structure of the consumption-leisure model, we will embed within it from the beginning a consideration of government tax policy. We will have much more to say later about the role of macroeconomic tax policies, and it will turn out that one of the major schools of tax policy thought to have emerged in the past 30 years crucially hinges on the main features of the consumption-leisure model.

The Two “Goods”: Consumption and Leisure

In our initial look at consumer theory, we supposed that there existed two broad categories of consumption goods, “good 1” and “good 2.” We will now condense these two categories into just a single category called “consumption.” That is, consumption is any and all “stuff” that individuals might purchase in order to obtain utility (happiness). Thus, consumption, which we will denote by c (without any subscripts), is an argument to individuals’ utility functions.

Because we are interested in studying how consumers “choose” their income, we must specify how consumers in fact earn their incomes. One seemingly obvious way of proceeding is to suppose that consumers obtain their income by working. An individual can choose to work some number of hours (per day or per week or per month, etc – we will specify this more carefully below) for which he receives **pre-tax pay** of W dollars per hour. That is, W dollars per hour is the individual’s gross wage rate, which in general is **not** what the individual actually gets to keep as the result of his efforts. In most countries, individuals are subject to a variety of government taxes – of the many

kinds of taxes which exist, the most common type (and certainly the type to which the greatest number of people are subject) are income taxes. Income taxes in the U.S. and essentially all other countries are specified as some percentage of an individual's total earnings. For example, if the labor tax rate in the U.S. were 30 percent and an individual earned \$50,000 in a given year, the amount of tax he would have to pay that year is $\$50,000 \times 0.30 = \$15,000$.¹³ Thus, it is as if the hourly wage W is subject to a 30 percent tax, making the individual's **after-tax wage rate** $0.70W$ dollars per hour. More generally, if we denote the tax rate on labor income by t (where $0 \leq t \leq 1$), then the after-tax wage rate is $(1-t)W$ dollars per hour. Because it is ultimately disposable income that individuals care about, $(1-t)W$ is the relevant wage rate for an individual's decision-making.

Presumably, working is a "bad" for individuals – that is, individuals dislike working because it reduces their total utility. In order to fit our model into standard consumer theory, we can easily recast the "bad" of working into a "good" by defining **leisure** to be the total number of hours an individual has available to him during some relevant period of time minus the total number of hours he spends working during that period. For example, suppose we were to consider each calendar week as a separate period of time. Denote by n the number of hours in a week the individual spends working. There are $24 \times 7 = 168$ total hours available to the individual during the week. Then, by our definition, the individual's leisure during the week, which we will denote by l , is $l = 168 - n$.¹⁴ Leisure is the opposite of working, so because working is a bad, leisure must be a good. We postulate leisure to be the second argument to the individual's utility function.

Indifference Map for Consumption and Leisure

The two objects in our model from which an individual obtains utility are thus consumption and leisure, giving rise to an abstract utility function $u(c, l)$. We will refer to both consumption and leisure from here on as "goods," even though clearly leisure is not a tangible object. Consumption and leisure as we have defined them are very broad categories of goods. As such, it is most useful to think of the general properties of the utility function $u(c, l)$ as being the same as those of the utility function $u(c_1, c_2)$ when we first studied consumer theory. Thus, we assume from now on the following properties:

1. Utility is always strictly increasing in consumption (that is, $\partial u / \partial c > 0$).
2. Utility is always strictly increasing in leisure (that is, $\partial u / \partial l > 0$).
3. Utility exhibits diminishing marginal utility in consumption (that is, $\partial^2 u / \partial c^2 < 0$).
4. Utility exhibits diminishing marginal utility in leisure (that is, $\partial^2 u / \partial l^2 < 0$).

¹³ The calculation of an individual's tax burden is not nearly so straightforward in reality due to a great many complicating features of tax laws. However, for our purposes this simple example will suffice.

¹⁴ Notice that because of our definition, leisure should not be thought of as time spent "having fun." Rather, it is time spent **not working**. Thus, time spent sleeping, for example, counts as leisure.

Indeed, these are exactly the same properties of the utility function that we have already studied. With these assumptions, we can construct an indifference map over consumption and leisure, as illustrated in Figure 7. Each indifference curve has all the usual properties we initially encountered in our study of consumer theory. Specifically, each indifference curve is downward-sloping, is bowed-in towards the origin, and crosses no other indifference curve.¹⁵

Although our two goods consumption and leisure are not both “market goods” (that is, one cannot really “purchase leisure” in a market), there is still a well-defined notion of a **marginal rate of substitution (MRS)** between the two. Again as usual, the MRS measures how many units of one good the consumer is willing to give up to get one more unit of the other good. Graphically, the MRS is the slope of the indifference curve.

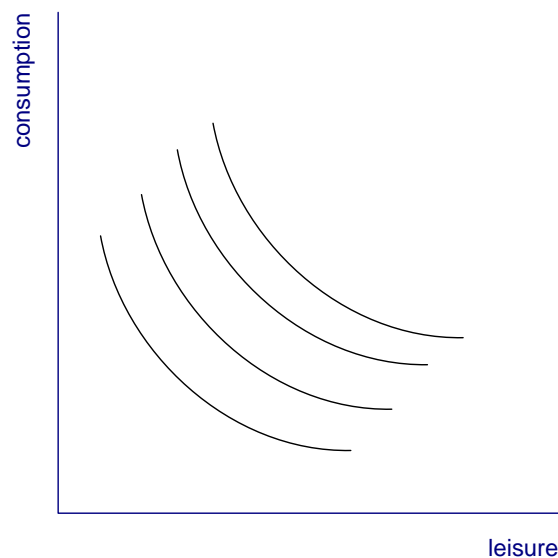


Figure 7. Indifference map defined over consumption and leisure.

Budget Constraint

Indifference maps alone are of course not enough to study an individual’s optimal choice. To study optimal decision-making, we need to consider the individual’s budget constraint, and it is here where our model of consumption and leisure most differs from the simple model of consumer theory we initially studied. In our simple model, a consumer simply had the income Y to spend on consumption (of good 1 and good 2). Here, the amount of income an individual has to spend on consumption (of the *one* market good) depends on how much he chooses to work. Let us now study formally the budget constraint in this model, supposing that the length of period of time in our model is one week (168 hours).

¹⁵ At this point, this should all be review. Recall especially that these three properties of indifference curves arise precisely because of our four assumptions on the utility function.

We assume the individual can work as few or as many hours as he wants. Regardless of how many hours he works, he gets paid the pre-tax wage W dollars per hour.¹⁶ As mentioned above, though, it is ultimately his after-tax income (his disposable income) that an individual cares about – his after-tax wage rate is $(1-t)W$ dollars per hour. Because he will choose to work n hours per week, his total disposable income is simply

$$Y = (1-t) \cdot W \cdot n .$$

The number of hours of work per week plus the number of hours of leisure (denoted by l) per week must equal 168, so that $n = 168 - l$. With this, we can write disposable income as a function of leisure,

$$Y = (1-t) \cdot W \cdot (168 - l) .$$

As in our earlier study, we make the simplifying assumption that the individual spends all of his income on consumption and saves nothing for the future. Each unit of consumption c can be purchased at the market price P (the individual is a price-taker). Thus, the individual's consumption each period (week) is

$$Pc = Y .$$

Combining the last two expressions yields the budget constraint in the consumption/leisure model,

$$Pc = (1-t) \cdot W \cdot (168 - l) .$$

In this budget constraint, the consumer takes as given the price P , the hourly wage rate W , and the tax rate t ,¹⁷ and he chooses his level of consumption c and hours of leisure l .

A useful rearrangement of this budget constraint is

$$Pc + (1-t)Wl = 168(1-t)W .$$

In this version with both the consumption good and leisure on the left-hand side, we see that the “price” of leisure is the after-tax wage rate $(1-t) \cdot W$. Of course, leisure (time off from work) is not directly bought and sold in markets. But the wage is the opportunity cost of leisure – every hour spent in leisure is an hour that could have been spent working. Thus, from an economic point of view, where we explicitly take account of

¹⁶ Clearly the assumption of being able to work as few or as many hours as one wants does not capture reality literally. Most workers have some semi-fixed schedule they must adhere to, at least in some relevant “short-run.” In a “longer-run,” workers are freer to move to jobs that better accommodate their lifestyles, etc. Thus, think of our consumption/leisure model as more of an attempt to capture this latter sense rather than the former sense.

¹⁷ That is, the individual is a price-taker in both the consumption-good market as well as the labor market.

opportunity costs, the after-tax wage is the price of leisure because it is what is being given up for every extra hour of leisure taken.¹⁸

As always, a budget constraint describes the set of choices which are available to the consumer but does not tell us anything about which point in that set he will choose. To graph this budget constraint in a diagram like Figure 7, we can rearrange again to get

$$c = \frac{168 \cdot (1-t) \cdot W}{P} - \frac{(1-t) \cdot W}{P} \cdot l.$$

This budget constraint is a straight line, as shown in Figure 8, with vertical intercept $\frac{168(1-t)W}{P}$ and slope $\left(\frac{-(1-t)W}{P}\right)$. By inserting the value $c=0$ in the budget constraint, we find that the horizontal intercept is at $l=168$, which simply states that if the individual wants no consumption, he can use all the hours in a week for leisure.

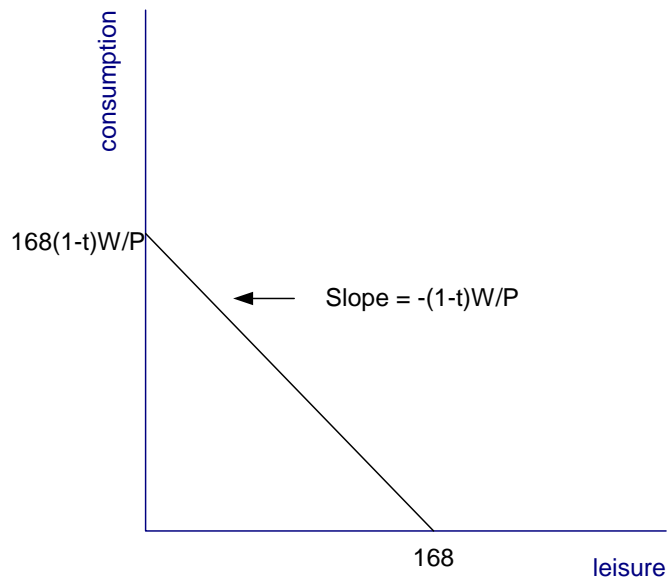


Figure 8. The budget line in the consumption-leisure model.

In our earlier model, in which consumers took their income Y as a constant, changes in income led to parallel shifts of the budget line. In this model, it is not income that individuals treat as a constant, but rather the after-tax wage rate $(1-t)W$. Notice how the after-tax wage rate enters the budget constraint here. Any change in the after-tax wage rate leads to a rotation of the budget constraint around the horizontal intercept (which recall is fixed at $l=168$) because $(1-t)W$ appears in both the vertical intercept and the slope of the budget line.

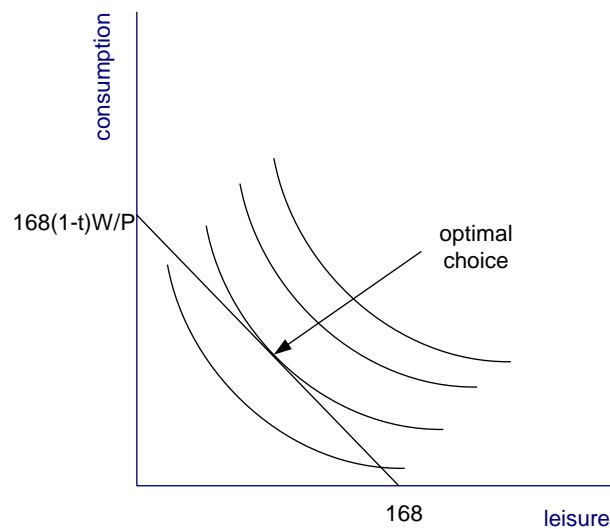
¹⁸ This is a very general notion of a “price.” A price is anything that must be given up in order to obtain something else.

Optimal Choice

As always, to consider optimal choice we must consider the interaction of the individual's preferences (indifference map) with his budget constraint. Superimposing the budget line and the indifference map, we have that the optimal choice of consumption and leisure is as shown in Figure 9.

Labor Supply Function

When the individual optimally chooses to spend l hours of his time in leisure, he is of course choosing to spend $n = 168 - l$ hours of his time working. He is thus supplying n hours of labor to the labor market. Clearly, the optimal choice of labor in Figure 9 depends on the after-tax wage $(1-t)W$. A definition is in order before proceeding: the **real after-tax wage** is the after-tax wage in money terms divided by the price of consumption in money terms. With our notation, the real after-tax wage is simply the ratio $\frac{(1-t)W}{P}$,¹⁹ and we see that it is only the real after-tax wage that matters for the vertical intercept and slope of the budget line. For the rest of this section, we will study how the optimal choice of labor varies as the real after-tax wage varies.



¹⁹ The term “real after-tax wage” comes from the units of measure associated with $(1-t)W/P$. Because the units of $(1-t)W$ is ($\$/\text{hour of work}$) and the units of P is ($\$/\text{unit of good}$), the units of $(1-t)W/P$ is (units of goods / hour of work), and hence the terminology: $(1-t)W/P$ measures the number of actual (real) goods a worker earns for each hour of work after he has paid his taxes. This is yet another example of how unit analysis helps us think about the relationship between economic variables.

Figure 9. At the optimal choice of consumption and leisure, the budget constraint is tangent to an indifference curve.

We begin our analysis by supposing that the initial real after-tax wage is quite low. Denote this initial wage by $((1-t)W/P)_1$. At this low initial real after-tax wage, the optimal choice is labeled point A in Figure 10. This initial optimal choice has associated with it n_1 hours of work (not shown of course because the axes contain c and l , not n). Now suppose that with the price P held constant, the nominal after-tax wage rises to $((1-t)W)_2$, so that the new after-tax real wage is $((1-t)W/P)_2$. **Notice that there are two ways the nominal after-tax wage rate can rise:** the gross wage W can rise while the tax rate remains constant or the tax rate can fall while the gross wage W remains constant. Regardless of which mechanism by which it occurs, the rise in the real after-tax wage causes the budget line to become steeper by pivoting around the horizontal intercept. With this higher real after-tax wage, the individual's optimal choice is point B. At point B, the individual has more consumption than at point A. The individual also enjoys less leisure at point B than at point A – which means that he now works n_2 hours, with $n_2 > n_1$. In other words, n has risen as the real after-tax wage has risen from $((1-t)W/P)_1$ to $((1-t)W/P)_2$.

An important note is in order here. You may be looking at Figure 10 and wondering why the optimal choice under the higher real after-tax wage did not feature more consumption **and** more leisure. In other words, you may be wondering why the indifference map is not such that the new optimal choice lies to the northeast of the original optimal choice, rather than northwest as is drawn. The answer is not at all a theoretical one, but rather one due to evidence about the real world. Much microeconomic and behavioral research has shown that when individuals currently have a low real after-tax wage, an increase in the real after-tax wage induces them to work more (presumably because, for example, they simply need the earnings to meet basic expenses). This suggests the partial indifference map in Figure 10.




Figure 10. As the real after-tax wage rises from $((1-t)W/P)_1$ to $((1-t)W/P)_2$, the individual optimally chooses more consumption and less leisure – the latter implying that he chooses to work more.

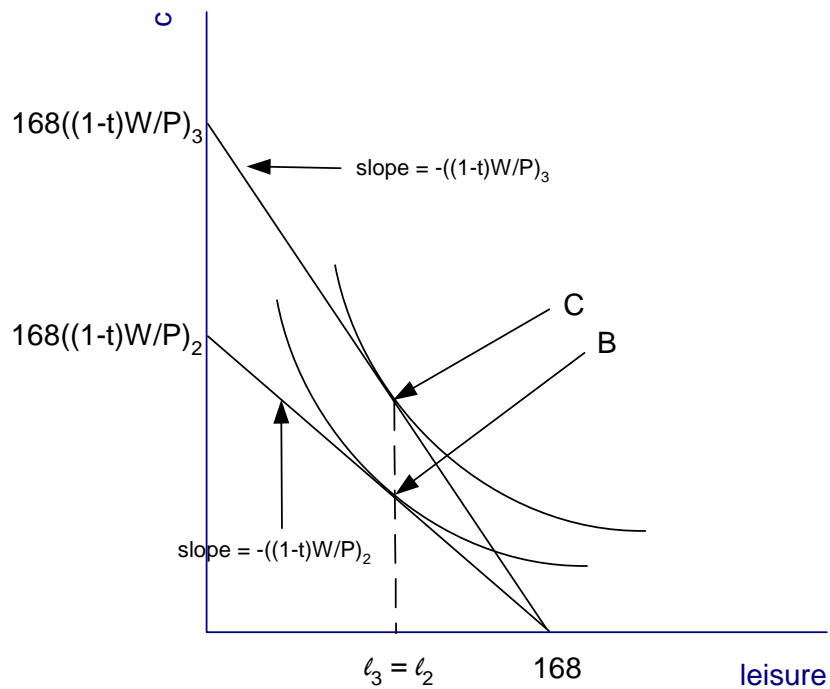


Figure 11. As the real after-tax wage rises from $((1-t)W/P)_2$ to $((1-t)W/P)_3$, the individual optimally chooses more consumption but an unchanged amount of leisure – the latter implying that he chooses to not adjust his hours worked.

Now suppose with the price P still held constant, the nominal after-tax wage rises again, to $((1-t)W)_3$. Thus, the real after-tax wage has now risen to $((1-t)W/P)_3$. The optimal choice at this new higher real after-tax wage is labeled point C in Figure 11. Comparing point C to point B, we see that the individual has chosen to not adjust the number of hours he works (and thus also not adjust the amount of leisure time he enjoys) when the real after-tax wage rose from $((1-t)W/P)_2$ to $((1-t)W/P)_3$. Thus, at this higher real after-tax wage, the individual is working n_3 hours, with $n_3 = n_2 > n_1$.

Consider yet another increase in the nominal after-tax wage, to $((1-t)W)_4$, which has associated with it the new real after-tax wage $((1-t)W/P)_4$. At this point, the real after-tax wage has gotten quite high and it may be that the individual simply does not need to spend more time working because his basic expenses (and perhaps even some luxuries) have already been met. At a very high after-tax wage, it may be reasonable to expect that the individual will now choose to spend **less** time working and spend more of his time in leisure. Such a situation is depicted in Figure 12, in which the rise in the real after-tax wage to $((1-t)W/P)_4$ induces the optimal choice to move from point C to point D. At point D, the individual is working fewer hours than at point C. That is, hours worked n_4 given the real after-tax wage $((1-t)W/P)_4$ is smaller than hours worked n_3 given the real after-tax wage $((1-t)W/P)_3$.




Figure 12. As the real wage rises $((1-t)W/P)_3$ to $((1-t)W/P)_4$, the individual optimally chooses more consumption and **more** leisure – the latter implying that he now chooses to work less.

To re-emphasize a point made above, notice there is no **theoretical** reason why the optimal choices of the individual should change in the way depicted in Figure 10 through Figure 12 when the real after-tax wage rate rises. Rather, such a description is justified on the grounds of evidence about how individuals do actually seem to respond to changes in their real after-tax wages.

Substitution Effect and Income Effect

We can decompose the effect of the change in the real after-tax wage on the optimal choice of leisure into two separate components: an effect called the substitution effect and an effect called the income effect. Both of these effects have very general meanings in economics and indeed can be applied to any optimal choice problem, not simply the consumption-leisure model. However, for our purposes, we will restrict our discussion of these effects to the consumption-leisure model.²⁰

Simply put, in the context of our consumption-leisure model, the **substitution effect of a higher real after-tax wage leads an individual to take less leisure (and hence work more)**. This is because as the real after-tax wage rises, the opportunity cost of leisure rises (because they are one and the same). As this “price” of leisure rises, an individual would tend to demand less leisure – simply because leisure has become more expensive!

Conversely, the **income effect of a higher real after-tax wage leads an individual to take more leisure (and hence work less)**. This is due to the higher income that a higher real after-tax wage tends to bring.²¹ **With a higher income, an individual would want to consume more of all normal goods.**²² So long as leisure is a normal good, an increase in income would lead an individual to want to take more leisure, and thus spend less time working.

The substitution effect and income effect are both always present. From the preceding discussion, it should be clear that they have opposing effects on an individual’s optimal choice of leisure (and hence opposing effects on an individual’s optimal choice of labor). For any given real after-tax wage and subsequent rise in the real after-tax wage, then, one of two things must occur. Either the substitution effect dominates (is stronger than) the income effect and the rise in the real after-tax wage leads the individual to choose to work more (take less leisure), or the income effect dominates (is stronger

²⁰ And defer a more general discussion of substitution effects and income effects to a more advanced course on microeconomic theory.

²¹ Note the distinction between “wage” and “income.” The wage is the hourly rate of pay, while income is the product of the wage and the actual number of hours worked.

²² Recall that this is in fact the definition of a normal (as opposed to an inferior) good.

than) the substitution effect and the rise in the real after-tax wage leads the individual to choose to work less (take more leisure).

With these notions of substitution and income effects, let us reconsider the events depicted in Figure 10 through Figure 12. The rise in the real after-tax wage from $((1-t)W/P)_1$ to $((1-t)W/P)_2$ led the individual to work more (take less leisure), as illustrated in the move of the optimal choice from point A to point B. Thus, it must be that over this range of the real after-tax wage, the substitution effect dominates (is stronger than) the income effect.

When the real after-tax wage rose again from $((1-t)W/P)_2$ to $((1-t)W/P)_3$, the individual decided to not adjust the amount of time he spent working, as illustrated in the move of the optimal choice from point B to point C. Thus, it must be that over this range of the real after-tax wage, the substitution effect exactly cancels with the income effect.

When the real after-tax wage rose yet again from $((1-t)W/P)_3$ to $((1-t)W/P)_4$, the individual decided to work less (take more leisure), as illustrated in the move of the optimal choice from point C to point D. Thus, it must be that over this range of the real after-tax wage, the income effect dominates (is stronger than) the substitution effect.

The Backward-Bending Labor Supply Curve

Let us now graph this individual's choice of number of hours worked as a function of the nominal after-tax wage, with the price of consumption held constant at some value P . The resulting graph is the individual's labor supply curve. The following table summarizes the labor supply schedule we found above in Figure 10 through Figure 12:

Nominal wage	Number of hours worked
$((1-t)W)_1$	n_1
$((1-t)W)_2$	n_2 , with $n_2 > n_1$
$((1-t)W)_3$	n_3 , with $n_3 = n_2$
$((1-t)W)_4$	n_4 , with $n_4 < n_3$

Graphing this data and passing a smooth curve through the points gives the individual's labor supply curve in Figure 13. The labor supply curve is said to be "backward-bending" because at high levels of the after-tax wage, the amount of hours worked declines as the after-tax wage rises.

The labor supply curve in Figure 13 is for a single individual. Every individual in an economy makes a similar labor supply decision, so in principle we have backward-bending labor supply curves for each individual. If the positions of every individual's labor supply curve is the same, then summing these individual labor supply curves horizontally yields an economy's labor supply curve (called the aggregate labor supply curve), which will also be backward-bending.

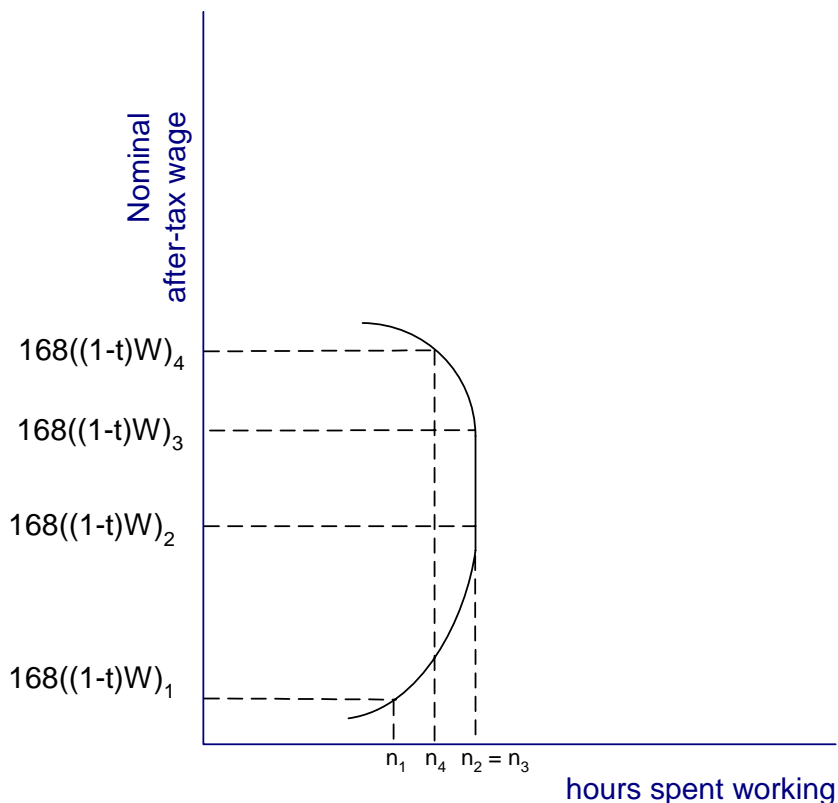


Figure 13. The backward-bending labor supply curve. In this diagram, the price of consumption is held constant at some value P . With this, at very low levels of the nominal after-tax wage, the substitution effect outweighs the income effect and thus the labor supply curve has a positive slope. At very high levels of the nominal after-tax wage, the income effect outweighs the substitution effect and thus the labor supply curve has a negative (“backward-bending”) slope. At intermediate levels of the nominal after-tax wage, the substitution effect roughly cancels out against the income effect, giving the labor supply curve its vertical region.

Aggregate Labor Supply Curve

Even if every individual in an economy has a backward-bending labor supply curve, though, the *aggregate* labor supply curve actually need not be backward-bending. This can occur if the exact positions of each individual’s labor supply functions are *not* identical. More precisely, if we are interested only in some “usual” range of macroeconomic outcomes *and* we wish to model events using the representative agent framework, then our representative agent should *not* have a labor supply curve that is backward-bending. This means that our analysis in Figure 10 through Figure 12 must be modified: the successive optimal choices traced out in the progression of these diagrams must feature always decreasing quantities of leisure, which equivalently means always increasing quantities of labor supply. In the terminology of substitution and income effects, we require

indifference curves (more fundamentally, a utility function) that features a substitution effect that is always stronger than the income effect with regards to leisure.²³

Because in macroeconomic data, there is no evidence of a “backward-bending” labor supply curve, the utility functions used in macroeconomic analysis feature just such a property. The particular functional forms for utility we encounter throughout our studies will exhibit this property.

Consumption Demand Function

Regardless of whether the aggregate labor supply curve is backward-bending or not, we derived it by considering how the optimal labor choice varied as the real after-tax wage varied. We can use the same analysis in Figure 10 through Figure 12 to consider how the optimal choice of consumption varies as the price P varies, **holding the nominal wage W and the tax rate t constant.** After all, a change in the real after-tax wage rate can be initiated by any one of P , W , or t changing with the other two held constant. In this section, we will suppose that it is the price of consumption which varies and examine how the optimal consumption demand varies.

Begin again at point A in Figure 10, and suppose that the price of consumption falls. This means that the budget line rotates to become steeper, just as shown in Figure 10. Point B then shows the new optimal choice of consumption, clearly larger than at point A. Turning to Figure 11, we see that if the price of consumption falls yet again, the budget line again becomes steeper, leading to a yet higher consumption choice at point C. If the price falls yet again, the budget line becomes even steeper and the optimal consumption choice increases again, as at point D in Figure 12. The point should by now be clear: a fall in the price leads to rise in optimal consumption, all else being held constant. Indeed, this is simply the law of demand that you learned in basic microeconomics. This analysis yields the downward sloping aggregate consumption function in Figure 14.

²³ You should be able to trace out for yourself the analogs of Figure 10 through Figure 12 for the labor supply curve to not be backward-bending.

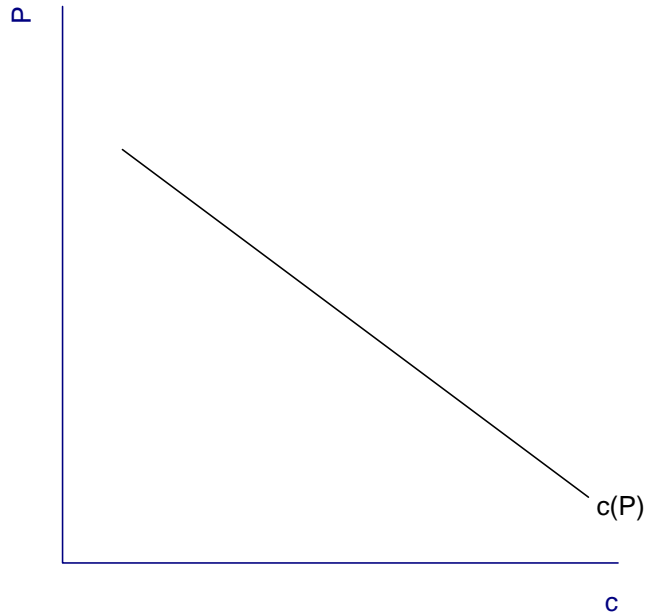


Figure 14. The downward-sloping aggregate consumption function, which is derived holding the labor tax rate t fixed and the nominal wage rate W fixed.

Lagrange Characterization – the Consumption-Leisure Optimality Condition

Let's now return to the decision problem of the representative agent (that is, before we aggregated things up to an aggregate labor supply function and an aggregate consumption demand function) and study the optimization problem our Lagrange tools.

To cast the problem into Lagrange form, we must first identify the objective function (i.e., the function that the consume seeks to maximize) – that is simply the utility function $u(c, l)$. Then, we must identify the constraint(s) on the maximization problem. The only constraint is the budget constraint; to cast it in our general Lagrange form, we write it as $g(c, l) = 168(1-t)W - Pc - (1-t)Wl = 0$. Proceeding as we have a couple of times already now, having identified the objective function and the constraint function(s), we now must construct the Lagrange function; in our problem here, the Lagrange function is

$$L(c, l, \lambda) = u(c, l) + \lambda [168(1-t)W - Pc - (1-t)Wl].$$

The first-order conditions we thus require are those with respect to c , l , and λ ; respectively, the are:

$$\frac{\partial u}{\partial c} - \lambda P = 0$$

$$\frac{\partial u}{\partial l} - \lambda(1-t)W = 0$$

$$168(1-t)W - Pc - (1-t)Wl = 0$$

The usual second step of the Lagrange method is to eliminate the Lagrange multiplier between the first-order conditions on the main variables of economic interest – here, c and l . Solve the first expression above for the multiplier gives us $\lambda = \frac{\partial u / \partial c}{P}$. Inserting

this in the second expression above gives us $\frac{\partial u}{\partial l} = \frac{\partial u}{\partial c} \frac{(1-t)W}{P}$. Rearranging in one more step gives us

$$\frac{\partial u / \partial l}{\partial u / \partial c} = \frac{(1-t)W}{P},$$

which is the representative consumer's **consumption-leisure optimality condition**. It states that when consumers are making their utility-maximizing consumption-leisure choices, they choose consumption and leisure such that their MRS between consumption and leisure (the left-hand-side of the above expression – recall that the ratio of marginal utilities **is** the MRS) is equated to the slope of the relevant budget constraint. As we saw graphically above, the slope of the relevant budget line here is the after-tax real wage $(1-t)W/P$.