Utility and Indifference Curves

Any point on the same ***indifference curve*** $U_{A(n)}$ represents an equal amount of ***utility*** U_n for person A. Utility is defined as a measure of well-being. We use utility to relate the value of different amounts and combinations of things ("***allocations***"). The ***convex*** shape of this individual's utility curves results in the most utility being obtained when there is a nice balance between work and leisure. In the following images, U_3 would be orange in the 3D utility "elevation map" and U_2 would be green and U_1 would be blue, with a lower degree of overall well-being. Note that we can compare different utility curves for the same person but cannot compare utility for two different people. The trade-off of income and leisure is given by the ***marginal rate of substitution*** and we can in fact compare an MRS across people because it is not measured in utility units. Someone who relatively prefers income more will have differently shaped indifference curves. This utility function is ***monotonic*** (more is always better) over both income and leisure because more income and more leisure (without reducing the other one) always result in higher utility. The red line represents the ***budget constraint*** (which is the limitation of what is feasible or affordable) and in this case is defined by the fact that there are only 24 hours per day to use. This red line represents every combination of income hours and leisure hours which sum to a total of 24. The red point on the right graph corresponds to the intersection of the budget constraint with the highest attainable indifference curve corresponding with the optimal balance of income and leisure for this person's convex preferences. That optimal (best attainable possibility) indifference curve in this example is U_2 and (with wage=1) is only feasible with exactly 12 hours allocated towards income and 12 hours allocated towards leisure:

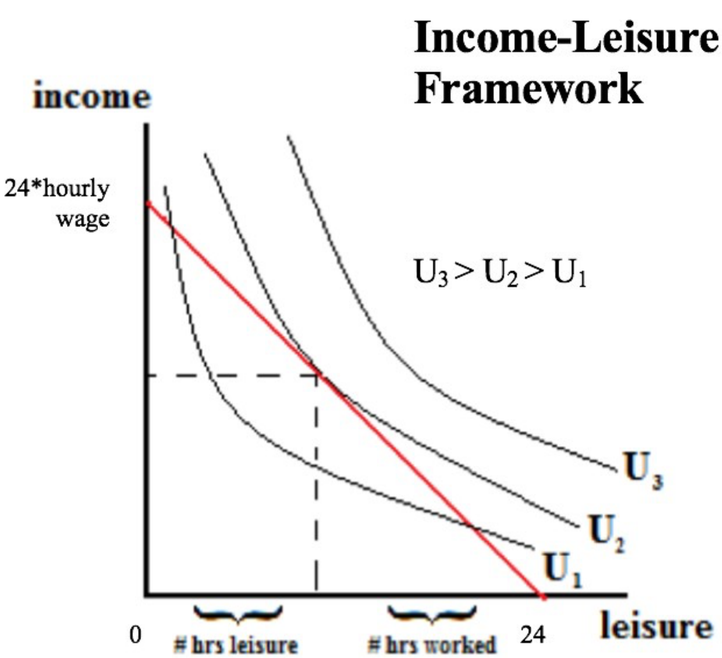

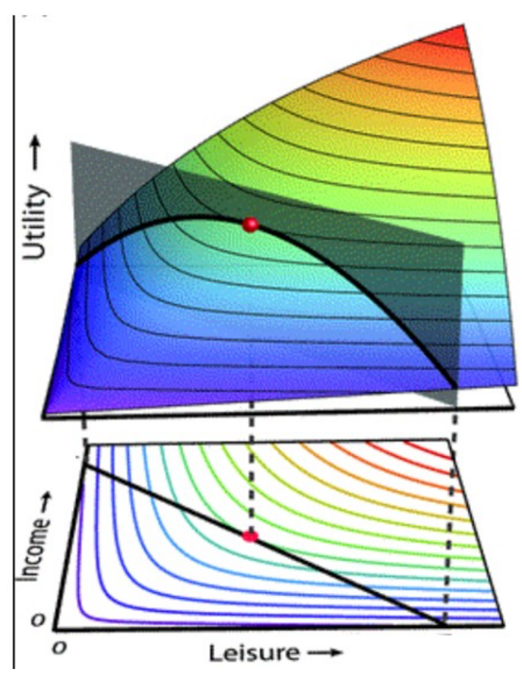

If, for simplicity, we had perfectly symmetrical indifference curves described by Andrew's utility equation U_A = [leisure * income] and normalized hourly wage to be equal to 1, we could solve this maximization problem to obtain optimal allocations of 12 for income and 12 for leisure with total utility maximized at U=144. Allocating one more hour to either option would have an "opportunity cost" greater than the marginal benefit, and thus result in negative marginal utility. To see this, note that 13*11 or 11*13 would each yield only 143 for utility. Every other combination would obtain even less utility. More generally, this optimization problem equates the ***marginal utility*** (change in utility per unit increase in the source of utility) from each of the two options: the best choice is where the utility increase from one more unit of income ("marginal utility of income") is equal to the increase in utility from one more unit of leisure ("marginal utility of leisure") … and of course this point must also be feasible based on the budget constraint. Also note that the fact that income and leisure are multiplied means that having more of either one increases the overall benefit of the other. Intuitively this should make sense with income and leisure time: more income increases the potential for what you can afford to do with leisure hours, and more leisure time means you have more hours to enjoy whatever you want to spend your income on.

This framework can operate to model other trade-offs, such as the utility from consumption through allocating money towards two different goods, *x* and *y*. In the following example, our utility function is again monotonic and convex with ***diminishing marginal returns*** (decreasing marginal utility) for both goods. There is an equal level of utility along each indifference curve, where each point on that curve represents a different combination of things used to obtain the same level of utility. For example, look at points A and B and C along utility curve U_4:

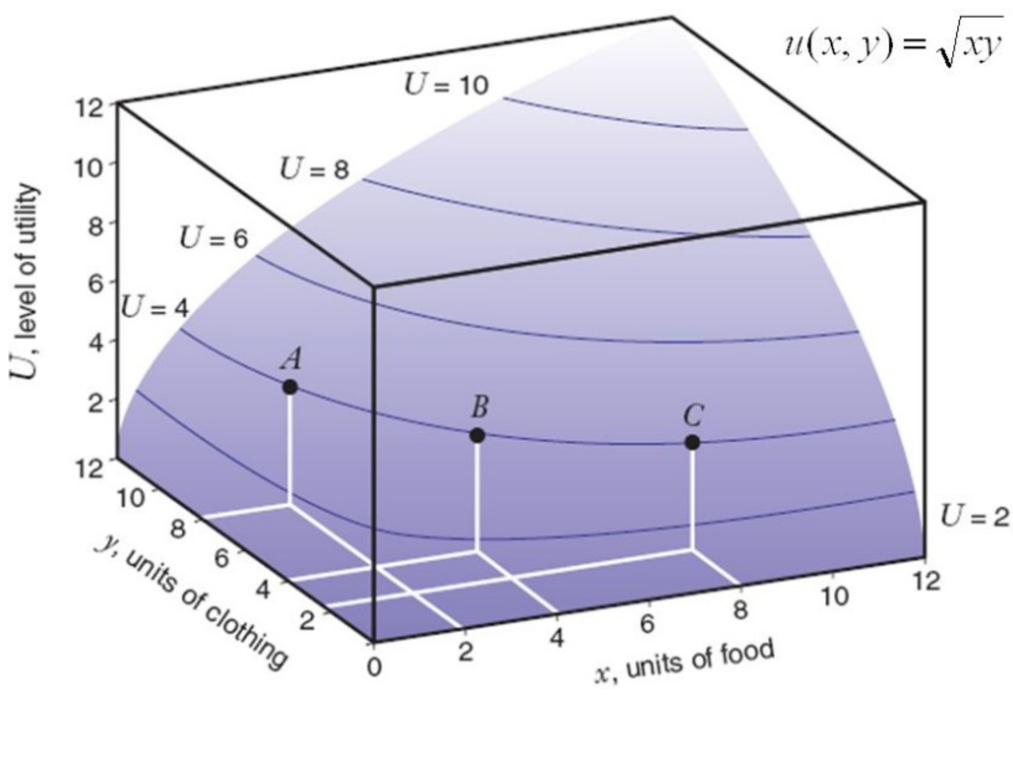

Income/Leisure Framework: Optimizing Utility Trade-Offs

Consider an investment banker with non-symmetric utility function $\mathbf{U_B}(I,L) = I^2 * L$ where I represents *income* (income = hours worked * wage) and L represents *leisure*. We are still normalizing wage here (w=1) for simplicity.[1] Using this, a convenient way to simplify the framework is to incorporate the *budget constraint* (I+L=24) to express hours worked as 24 minus leisure hours: equivalently, we could re-write our utility function as $\mathbf{U_B}(L) = (24-L)^2 * L$ and optimize with only one variable now using the derivative approach or testing values in a table. Below is a table of all integer values which sum to 24:

utility = income ^2 * leisure	hours worked	income utility component	leisure utility component	Marginal utility
0	**0**	0	24	
23	**1**	1	23	23
88	**2**	4	22	65
189	**3**	9	21	101
320	**4**	16	20	131
475	**5**	25	19	155
648	**6**	36	18	173
833	**7**	49	17	185
1024	**8**	64	16	191
1215	**9**	81	15	191
1400	**10**	100	14	185
1573	**11**	121	13	173
1728	**12**	144	12	155
1859	**13**	169	11	131
1960	**14**	196	10	101
2025	**15**	225	9	65
2048	**16**	256	8	23
2023	**17**	289	7	-25
1944	**18**	324	6	-79
1805	**19**	361	5	-139
1600	**20**	400	4	-205
1323	**21**	441	3	-277
968	**22**	484	2	-355
529	**23**	529	1	-439
0	**24**	576	0	-529

We can see that our optimal allocation is 16 hours working for income and 8 hours of leisure. This obtains the highest utility (U* = 2048 when I = 16) but there is an important point to make about that: the number itself for utility is totally meaningless! ***Utility*** is a *relative measure of value* or a "unit of happiness" which represents someone's overall individual well-being. Utility is useful in that it compares levels of wellness across different situations for one person to describe how relatively beneficial or harmful things are for that person, but the concept of utility is limited in that <u>we can never compare one person's utility to another person's utility</u>. We can, however, utilize the utility functions of two different people to determine each of their optimal behaviors and then compare their implied optimal choices (utility-maximizing allocations) since the decision of how many hours to work is in appropriately comparable units. ***Marginal utility*** is useful because most people can observe something change (such as more hours working or more food consumed, etc.) and conclusively determine whether they are more happy or less happy as a result, but attempting to measure their "actual level of exact happiness" is not possible.

1. We are still normalizing wage to always equal 1 here, but if that were to change then we would see a shift in the vertical axis intercept: increasing the wage would raise the intercept on the income axis. This would affect the optimal choice of working hours since the highest accessible indifference curve would then intersect the budget constraint at a point with a different set of allocations towards income and leisure. To understand this better, see the page on income effect and substitution effect.

Using the calculus approach instead, we can take the partial derivative of the utility function with respect to each choice variable, set each of those equal to zero ("first order conditions"), and solve the system with algebra by substituting with the budget constraint:

$$\frac{\partial U}{\partial I} = 2(I)L = 0 \qquad\qquad \frac{\partial U}{\partial L} = I^2 = 0$$

$$2(I)L = I^2$$
$$2L = I^*$$

Using the budget constraint:
L + I = 24
L + 2L = 24, so we conclude that $L^* = 8$ and $I^* = 16$ are the optimal allocations of time.

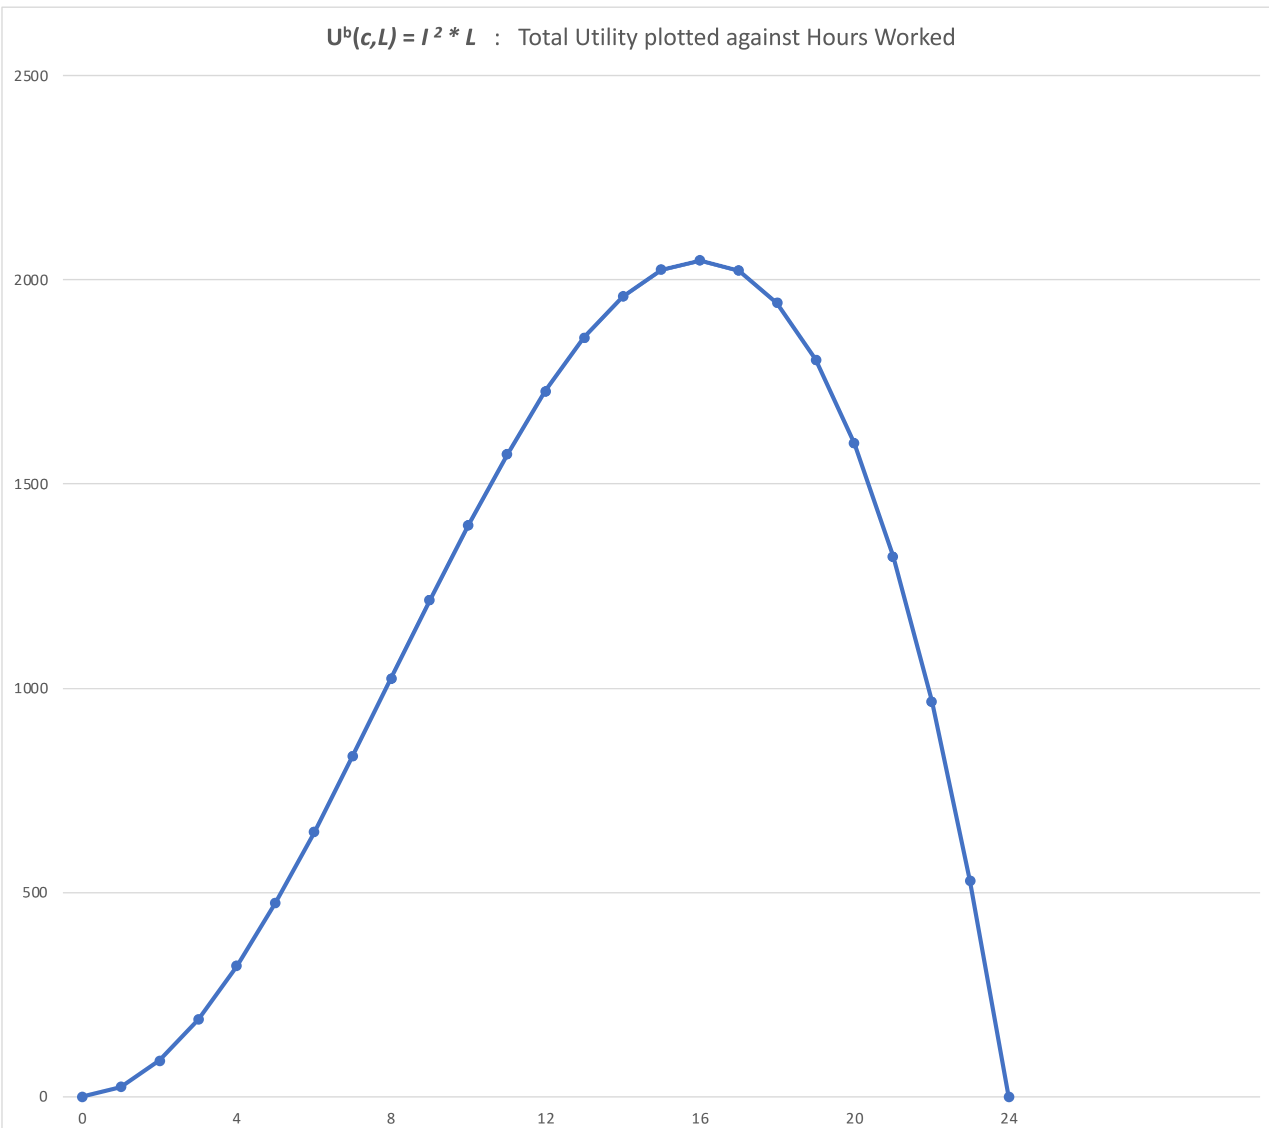

Lagrangian Optimization Approach:

Income-Leisure Framework: Utility-maximizing time choices - Investment Banker

Investment banker B's utility function over hours spent working (**I**) and leisure (**L**) is $\mathbf{U_B(I, L) = I^2 \cdot L}$ and she has 24 hours in a day which she can allocate as needed. Find this banker's optimal allocations of time which maximize her utility and graph her utility over hours spent working.

SOLUTION

The budget constraint is in terms of time here, and expressed mathematically it is $I + L = 24$. This also means that the utility equation can be re-written in terms of only one variable, as either $\mathbf{U_B(I) = (24 - L)^2 \cdot L}$ in terms of leisure hours or equivalently as $\mathbf{U_B(L) = I^2 \cdot (24 - I)}$ in terms of income hours. Graphing either one of these obtains an inverse parabolic function that is concave on both sides of the peak, which corresponds with the optimal allocations, and shows diminishing returns and then decreasing utility beyond the maximum.

To solve using a Lagrangian:

$$\begin{aligned}\mathbf{L}(I, L, \lambda) &= U_B(I, L) + \lambda \cdot g(I, L) \\ &= I^2 \cdot L + \lambda \cdot (24 - I - L)\end{aligned}$$

The three "first order conditions" describing the optimal allocations are obtained by taking the partial derivative of the Lagrangian with respect to the three variables:

$$\frac{\partial \mathbf{L}}{\partial I} = 2I \cdot L - \lambda = 0$$

$$\frac{\partial \mathbf{L}}{\partial L} = I^2 - \lambda = 0$$

$$\frac{\partial \mathbf{L}}{\partial \lambda} = 24 - I - L = 0$$

Equating the first two FOCs we have $2I \cdot L = \lambda = I^2$, which we can solve to obtain $2L = I$ to find our optimal ratio of time usage. Substituting this back into the budget constraint, which is also our third FOC after differentiating with respect to lambda, we get $24 - L - (2L) = 0$ which we can solve to obtain $I^* = 16$ and $L^* = 8$ as the optimal allocations of time to maximize utility.

Income/Leisure Framework Extensions & Applications

Now let's analyze the most basic changes to the income-leisure model: suppose we were to reduce wages and/or increase taxes, thus lowering disposable income per hour. The benefit of one hour of leisure is unchanged but the benefit of one hour of working has decreased. This moves the intercept down along the income axis as our budget constraint shifts from red to blue, so with the same utility function (and therefore the same indifference curves) the optimal combination of income and leisure is going to change. This change depends on the shape of the indifference curves:

In case A, the shape of the indifference curves given by A's utility function results in more leisure being consumed: the marginal benefit of leisure is unchanged but the marginal benefit of working has decreased.

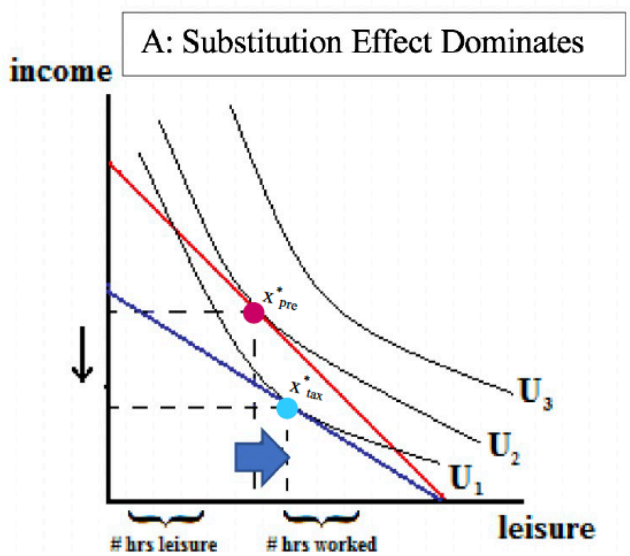

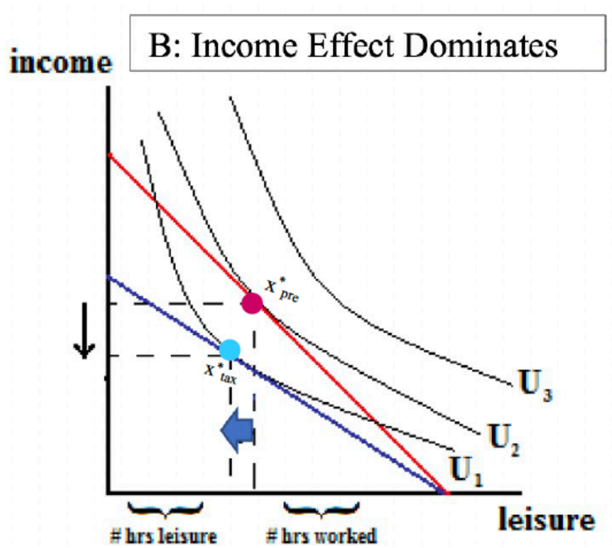

In case B, the indifference curves are shaped differently because person B has different preferences (perhaps less existing wealth, for example) so B actually responds to the taxation increase by working more to compensate for the disposable income reduction.

Aggregate Effects of Individual Labor Supply Choices

The Laffer Curve shows how a government's tax revenue varies based on the tax rate it chooses. (Another word for this is "taxable income elasticity".) For simplicity we can assume all workers pay the same tax rate on all levels of income. Tax revenue will always initially increase as the tax rate increases, obviously starting at zero for both. As the tax rate approaches its revenue-maximizing point t*, the rate of increase in revenue decreases because of an increasingly strong substitution effect exhibited by the average taxpayer, resulting in less work being done in the overall economy. In other words, there are diminishing returns to tax revenue as we increase the tax rate. Taxation can also interfere with the efficiency of capital allocation, and high levels of taxation can create various inefficiencies in and across markets, potentially reducing output and economic activity.

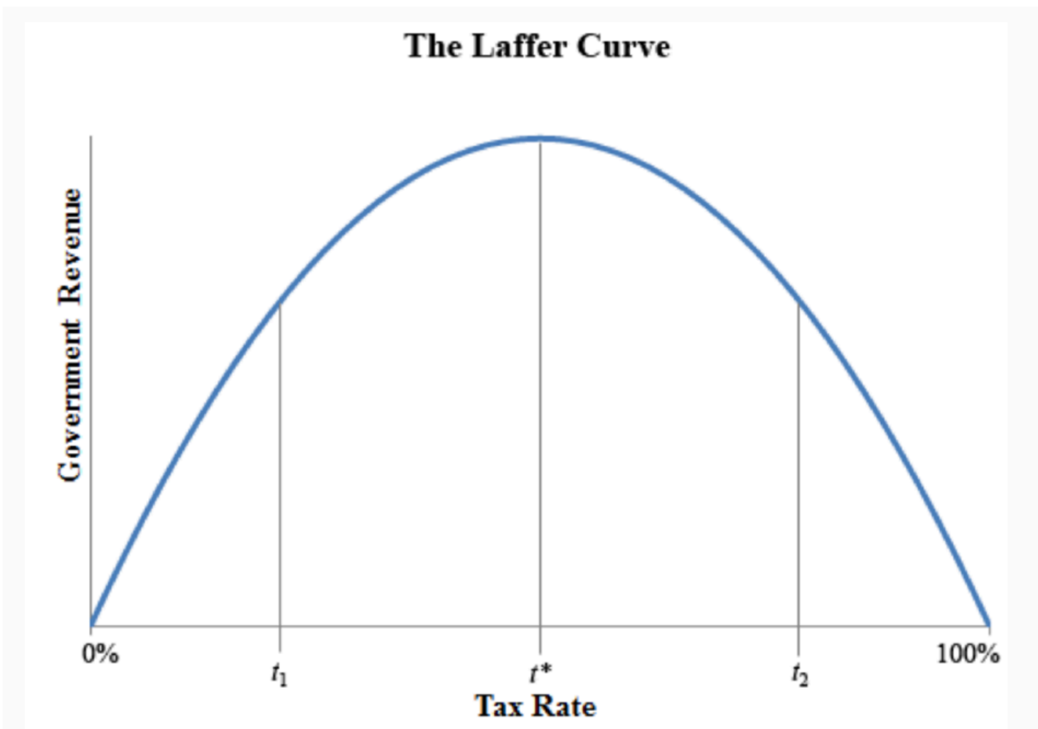

The revenue-maximizing point, ***t****, is where the inefficiency of capital allocation in society and the dominant substitution effect decreases the incentive to work and thus overall output exactly enough to maintain the same total tax revenue as taxes increase by one more unit. Beyond this point ***t****, the loss in economic output from "less work being done" actually reduces tax revenue more than the increased tax rate increases revenue. Conceptually this is equivalent to taking an increasingly large share of a pie that is getting smaller: if you decrease the total size of the pie enough, then a larger portion of it will actually get you a smaller slice.

The positive correlation that we generally observe between tax revenue and tax rates in reality strongly suggests that most economically advanced democracies have tax rates situated to the left of their Laffer maximum, wherever that may be for each country. Many economically successful countries have taxes as high as 40-50% and some, including the US in the past, have had top marginal rates of 70% or more. A different consideration is where and how severely taxation might negatively affect economic growth and business investment.

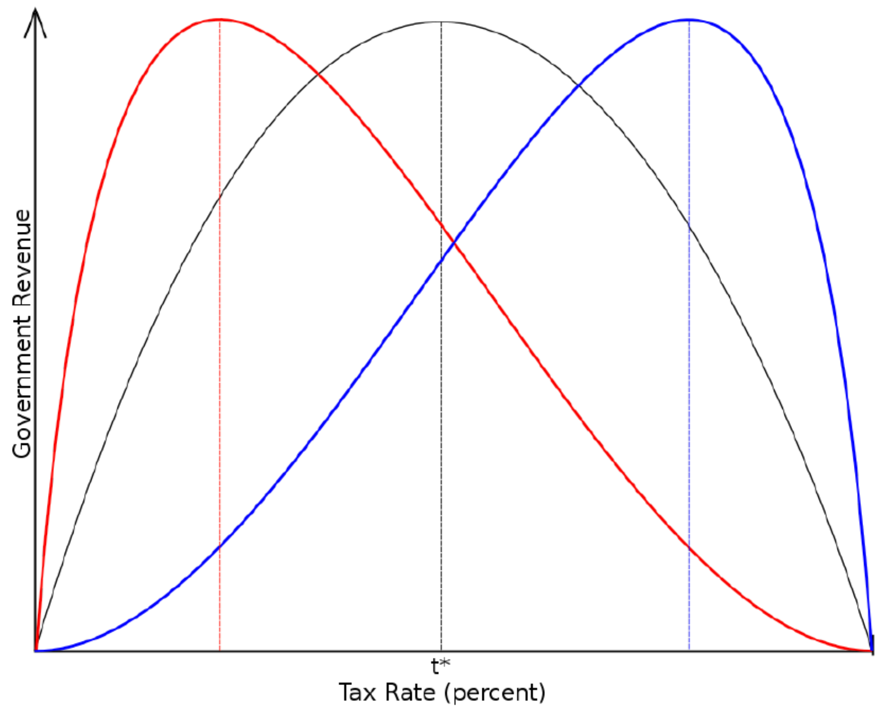

Note that the three different curves here represent separate estimations of the Laffer Curve. The way in which workers respond to tax changes will vary from person to person. If we assume some level of cultural differences exist across countries (this would be assuming different average or median utility functions over income and leisure) there may be different revenue-maximizing peak points in different nations. All are inverse parabolas, but the concavity and convexity and resulting peak locations vary significantly. Thinking about what the red and blue colors might represent, we can see how one of the largest debates in US politics is actually a dispute about the income/leisure framework and curvature of this graph. The differences in policy proposals are based on the use of tax money, the morality of redistribution itself, and of course on what will happen regarding workers' income/substitution effects aggregated to a societal level. Someone with fully socialist beliefs, for example, might believe that the peak is at 100% and the substitution effect would never become dominant… so people would never choose to work less even if all of their income was taken via taxes and everyone was guaranteed identical wealth outcomes regardless of their behaviors.