The Myth of the Income Effect

Pascal Salin

Microeconomics traditionally splits the consequences of price changes into two different effects: the substitution effect and the income effect.[1] The effects are supposed to work in diverging directions. The precise reaction of demand or supply to a relative price change can operate either way. The supply curve is not necessarily upward-sloping and the demand curve is not necessarily downward-sloping.

The income effect plays a crucial role in mainstream micro-economics, as well as in applied economics, because it induces economists to think, for instance, that there are backward bending curves. Such curves are particularly assumed to exist for the supply of labor. Within a certain range, a decrease in the real wage would not bring about a decrease in the supply of labor, since people want to maintain their income and, therefore, increase their supply of labor. If this is true in tax theory, the "Laffer-effect"[2] would not exist, at least in certain conditions in which the supply curve for labor is supposed to be atypical. With the income effect, an increase in the tax rate on labor income—i.e., a decrease in the real after-tax wage—would be compensated for by working more. The "Laffer-effect" assumes that there is always an inverse relation between the tax rate and productive efforts. In a similar way, the income effect would

Pascal Salin is professor of economics at the Université Paris-Dauphine.

[2]The increase in a tax rate decreases the tax base and, beyond a certain tax rate, this effect is sufficiently important for the increase in the tax rate to induce a decrease in tax revenues instead of an increase. It is now usual to label this process as the "Laffer-effect." But it had been developed before Arthur Laffer gave it its popular expression. One can find it, for instance, in a refined form, in Geoffrey Brennan and James Buchanan, *The Power to Tax* (Cambridge University Press, 1980).

account for a possible increase in the rate of savings when the
return on savings is decreasing.

   The existence of the income effect assumes that the concept
of income can be defined in a non-ambiguous way. It can be
measured and quantified. In fact, this is not true. Income can only
be understood using a precise concept of utility which does not
allow room for measurement. In the present paper we show that
the income effect is not consistent with a purely logical theory
of utility. In reality the income effect can exist only if one adds
a specific assumption. Contrary to the belief that the income
effect is a general principle, it is in fact only a possible conse-
quence of a specific assumption. The substitution effect—consis-
tent with the general theory of utility—is the only general prin-
ciple.

   Let us consider an individual's supply curve for labor in a
world in which there are only two goods, leisure and wheat. Time
is the only scarce production factor. Land is not operationally
scarce, so that the physical productivity of one hour of labor in
terms of wheat is constant. There is a certain relative price
between leisure and labor (the production of wheat). In one hour,
one can produce one hour of leisure or one pound of wheat. For
this given relative price, an individual wants to share his avail-
able time between leisure and the production of wheat. If he does
not desire to produce and to consume more wheat, it is because
sacrificing one hour of leisure to get one more pound of wheat is
not worthwhile from his own point of view.

   Now, economists usually define the income of an individual
by the quantity of tradable resources (wheat) he is producing at
each period; his production of leisure services is excluded. But
such a notion of income is arbitrary in comparison with the more
general concept of satisfaction (or utility): The individual's total
satisfaction is equal to the "sum" of the satisfaction he draws from
leisure and the satisfaction he draws from consuming the wheat
he has produced. Saying that the desired level of income is equal
to y is in fact saying that a maximum satisfaction is obtained
when the production of wheat—named income—is equal to y and
the production of leisure services is equal to x, knowing that $x +
P_x y = t$, where P_x is the price of one unit of wheat in terms of hours
of leisure and t is the available quantity of time, for instance 24
hours a day.

   Let us assume now that a disease strikes the wheat crop and
it cuts the productivity of labor by one half or—the equivalent—a
tax has been introduced on the production of wheat, so that one

can keep only half a pound of wheat for one hour of work. The opportunity cost of producing and consuming wheat increases in terms of leisure. The logic of choice would imply that one does not produce as much wheat as before.

Referring to this logic of choice, what could be the meaning of the income-effect? It implies that the individual aims at maintaining the level of his income, at least partially. That is, he wants to maintain his production and consumption of wheat. But it is inconsistent to hold that the individual has two aims in mind. He would have an income target (a certain amount of y): leisure time would be nothing but an instrument which can be used to reach this income target through a process of transformation. On the other hand—and contradictorily—he would have a satisfaction target: he aims at maximizing his level of satisfaction by obtaining an optimal mix of x and y.

In the above example, it is traditional to say that the individual facing an increase in the cost of wheat may not decrease his production of wheat—although it is most costly in terms of sacrificing leisure—but may decide to work more. Therefore, the supply curve would be backward-sloping.[3] There are more work hours when the return of one hour of work is decreasing. The income effect would thus dominate the substitution effect. In fact, the belief in the existence of the income effect implicitly requires *specific assumptions* which are not absolutely and logically necessary. This is to say, the income effect does not exist as a general phenomenon; it is a mathematical illusion in a badly specified world.

What happens in reality? If one assumes that an individual increases his supply of labor when the return on work hours decreases, it is necessarily because one has added an additional assumption (if not, one ought to admit that the substitution effect—which necessarily exists—is the only one to be met). Very likely, it is implicitly assumed that wheat—contrary to leisure—is necessary for one's survival, so that, if the substitution effect was the only existing effect, there would be a risk for an individual of passing under the subsistence level, which he cannot accept. But that is really saying that there are three goods (or three sources of utility) and not two, namely the pleasure

---

[3]Hans-Hermann Hoppe raised an interesting question on a previous version of this paper: If the income effect does not exist, does it imply that a supply curve cannot be backward sloping? We explain in the appendix why, in our opinion, it is impossible.

drawn from leisure, the pleasure drawn from consuming wheat, and survival services obtained from wheat consumption. One also assumes implicitly that there is only one possible technology to get survival services, namely wheat production (which is called "income"). But why would there not be another one? As an example, it may happen that the concerned individual knows that he can survive by extracting roots from the soil, the taste of which he does not like very much, but the nutritive value of which is indisputable. So long as wheat is not expensive (in terms of time) in comparison with roots, he prefers to eat wheat. If the opportunity cost of wheat production increases too much, he will certainly consider possible substitutions between wheat, roots, and leisure.

It may happen that shifting to another consumption pattern cannot be done immediately, since one has to find the most productive field of roots, or to invent suitable instruments to extract the roots. Therefore, during some period of time, the individual may increase his production of wheat in order to survive and to prepare the change in technology. But this is also saying that the existence of the income effect is due to two specific assumptions: The existence of some subsistence level and the existence of transfer costs when shifting from one activity to another one. Now, if the individual could have forecast the relative price change between wheat and leisure, he would certainly not have waited until the last moment, which obliges him to increase his supply of work in order to produce wheat. He would probably have decreased his leisure time earlier in order to prepare the new survival technology. Because of the increasing price of wheat relative to roots, there is a substitution effect between both these goods which make it possible to meet in different proportions both targets (survival and gastronomic pleasure).

The apparent appearance of an income effect is due to a model which has been badly specified and surreptitiously adds assumptions which can be interpreted as the introduction of an exogenous event in a world in which there are rigidities of a technical or institutional nature. Choices are then reshuffled and, during a transition period, what is named the income effect may appear in certain circumstances, as the result of a search for minimizing adjustment costs until the process of change in activities has led to all desired substitutions.

The above remarks can be reconciled with traditional statements of microeconomics. Let us take the traditional figure which is used to illustrate the existence of both the substitution and the
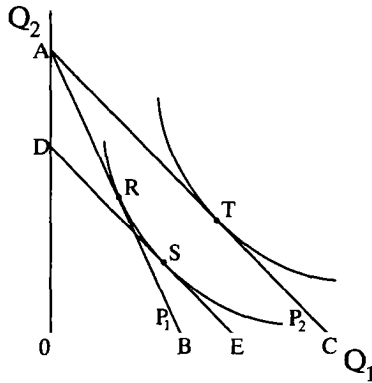
Figure 1

income effects (figure 1).[4] The initial equilibrium position of a "consumer"[5]—who has to choose between two goods, Q_1 and Q_2, for a relative price P_1—is point R. If the price of Q_1 decreases, the new budget constraint is AC and the equilibrium point becomes T. Now, one can obtain this later point by a formal splitting of the path from R to T into two paths: The path from R to S represents the substitution effect (pure relative price change, the utility level being maintained constant), and the shift from S to T represents the income effect (the decrease in the price of Q_1 corresponding to a gain in the purchasing power of income). Such a presentation does not help to clarify an important aspect of the problem, namely that the rotation of the budget constraint from AB to AC implicitly assumes that the consumer holds a stock of goods Q_2 equal to OA, so that the relative price change between Q_1 and Q_2 does not affect the position of point A: For a stock of goods Q_2 equal to OA, the individual can obtain a different (higher) quantity of Q_1. On the contrary, if the individual was holding (or producing) good Q_1, it would be point B which would

---

[4]Figure 1 is extracted from the classic textbook by James M. Henderson and Richard E. Quandt, *Microeconomic Theory*. But a similar graph can be found in most microeconomics textbooks.

[5]Let us use this traditional concept, although it induces the danger of splitting a human being into different, non-coordinated, parts (the consumer, the producer, the taxpayer, etc.).
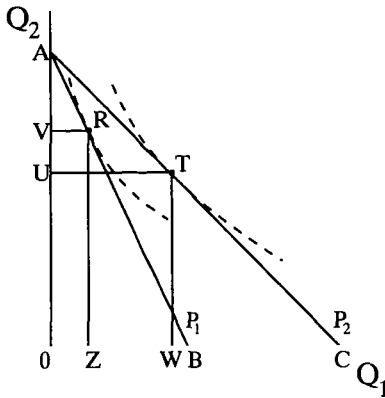
Figure 2

stay at the same place and the budget constraint would then cut the Q_2 axis at a point lower than D. Total utility would thus diminish. The result is perfectly logical: If the individual supplies Q_2 and demands Q_1, a decrease in the relative price of Q_1 (increase in the relative price of Q_2) leads to the transformation of more Q_2 in goods Q_1. The graph can then be interpreted from three different points of view.

First Interpretation

The relative price change reflects a technological change. Thus, with the same quantity of work, OA, an individual can obtain a greater quantity of wheat (OC instead of OB). The figure then does not represent the choices of the abstract "consumer" of textbooks, but the choices of the individual who *acts* (for instance Robinson alone on his island). If the productivity of labor is increasing, the individual will produce more wheat. On figure 2, OA represents the number of available hours per day, namely 24 hours. When the relative price is P1, the amount of work is equal to AV, the quantity of leisure to OV and the quantity of wheat produced to OZ. When the price of wheat in terms of work hours shifts to P2, work increases (AU), leisure decreases (OU) and the quantity of wheat produced increases (OW). The marginal sacrifice which is necessary to obtain one more unit of wheat having diminished, one substitutes wheat to leisure. There is a pure substitution effect.

Second Interpretation

Figure 2 now represents an exchange economy. The individual can obtain wheat (Q_1) by selling tomatoes (Q_2) which he is producing. Given the existing techniques and the limitation of his available time, the maximum quantity of tomatoes he can produce is OA. If the wheat price decreases from P_1 to P_2 (the price of tomatoes increases), he will sell more tomatoes to get wheat: On figure 2 his sales of tomatoes increase from AV to AU and his consumption of his own tomatoes decreases from OV to OU. His income does not change if it is measured in terms of tomatoes; it increases in terms of wheat. But there is only a process of substitution between both commodities and no income effect.

Third Interpretation

Figure 1 could be interpreted as representing the *partial* behavior of an individual *only* in his role as a consumer. He is a producer of another commodity, Q_3, which is not represented on the figure, and he demands Q_1 and Q_2. When the budget constraint AB is rotating to the position AC, one implicitly assumes that the relative price decrease of Q_1 in terms of Q_2 does not affect that valuation of Q_3—the supplied good—in terms of Q_2: The supply of commodity Q_3 in exchange for Q_1 and Q_2 would remain equal to OA (in terms of Q_2), which is debatable. In fact, there is a decrease in the price of Q_1 relatively to Q_3, which ought to affect the supply of Q_3. The new budget constraint is very likely situated between AC and DE (figure 1).[6]

These various interpretations are useful since they make it possible to clarify what is precisely the behavior of the individual who acts. But they may also help to demonstrate why the mainstream distinction between the substitution effect and the income effect is misleading. From a purely formal point of view, one can distinguish, as we have seen, the shift from R to S and the shift from S to T (figure 1). But these shifts do not correspond to any *actual* human action.

Let us take the first or the second interpretation.[7] It is meaningless to say that point S corresponds to a pure substitution

---

[6]As an example, one may imagine that the budget constraint is rotating around point R.

[7]Both interpretations are very close to one another. They imply that, from a given stock of resources—for instance time resources—an individual can obtain a good either via a technological process of transformation he is doing by himself, or via a process of exchange involving his labor services or the outcome of his labor.

effect (with a constant level of utility), since point D has no meaning for the individual. In fact, he *has* resources equal to OA and there is only a substitution effect, expressed by the rotation of the AB line to AC: Given the time resources he has, substitutions are possible along the budget constraint and these substitutions depend on the relative price between Q_1 and Q_2.

In the third interpretation, one cannot determine the value of the so-called income effect as long as one has not made clear what are the actual substitutions between the quantity of good Q_3, on the one hand, and the quantities of goods Q_1 and Q_2 on the other hand. But this is precisely saying that there is no income effect, and that there are only substitution effects between several possible uses of resources. Let us take the case of an individual who is choosing, at some point of time, a given allocation of his resources, specially as regards his time resources, to obtain a certain basket of goods, among which leisure has to be considered. To accept sacrificing an additional unit of leisure, he necessarily has to get compensation by obtaining a greater quantity of other goods. If ever there is an increase in the market price of time (a decrease in the price of other goods), it creates purely and simply a substitution effect.

The general belief in the existence of an income effect partially stems from an artificial separation between consumption activities and production activities. This distinction is usual in microeconomics: The theory of consumption and the theory of production are founded on different grounds and different behavioral assumptions. In fact one does not supply a good without demanding another good at the same time, so that the decision to produce and to supply cannot be separated from the decision to demand and to consume. By cutting, for instance, decisions related to consumption or decisions related to work from other decisions, the traditional theory implicitly maintains as constant certain variables which cannot be maintained as constant. It considers, for instance, the choice between two consumption goods, as if the whole set of activities was not implied. A specific assumption concerning choices made about these other activities is then called the income effect. As we have seen, a change in the relative price between two commodities has consequences which depend not only on the choices made by an individual about these commodities, but about the whole set of possibilities open to him.

In mainstream microeconomics, the income effect and the substitution effect are presented as similarly general effects. In fact, the substitution effect is the only one that actually exists. It
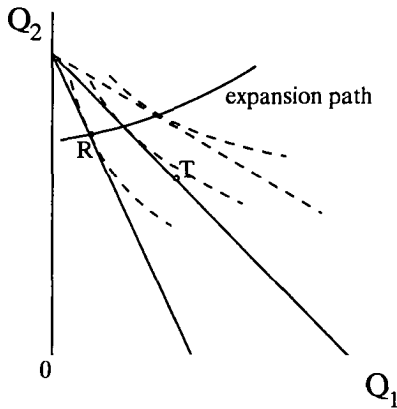
Figure 3

may appear that something close to the so-called income effect appears under specific assumptions when one adjusts from one equilibrium point to another one. However, this effect, which describes a specific adjustment path, is only transitory and it does not necessarily exist. This difference between both effects has important consequences. It means, for instance, that the Laffer effect necessarily exists: An individual shifts from an activity which is more heavily taxed to other activities. However, the process of adjustment may not be instantaneous, so that one may have the feeling that people react contrary to what the substitution effect would lead them to expect. However, this effect, if it ever exists, is necessarily a transitory and secondary phenomenon.

From a formal point of view, one can draw an expansion path so that the quantity of labor is decreasing as long as its price is increasing (figure 3), but this is nothing else than a graphical representation. Now, a graphical assumption does not give any information on the real motives of human behavior. The latter has to be explained by the law of decreasing marginal utility which is in reality a theory of human choices founded on rational logic. If one believes—as we should—that all reasoning in the social sciences has to be coherent with rational logic, it just means that one must get rid of the graphical or mathematical appearances which are inconsistent with this logic.
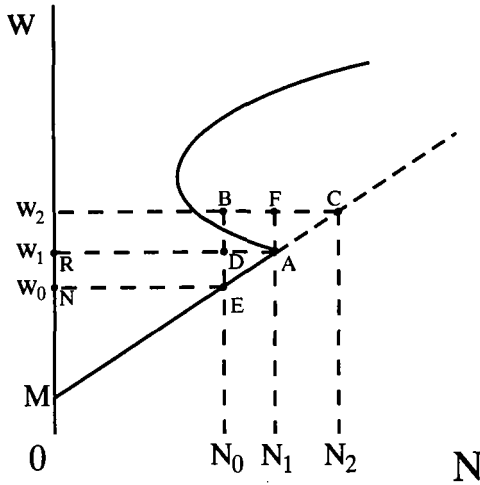
Figure 4

Appendix

Is a Backward-Bending Supply Curve Possible?

As we mentioned in footnote 3, an interesting question can be raised: Is it consistent to admit that both the income effect does not exist and that, however, the supply curve (for instance for labor) can be locally backward-sloping? To try to solve this question, let us take the case of an individual who is going up his supply curve when his real wage increases (figure 4). His supply of labor then shifts from E to A when his real wage shifts from w_0 to w_1. This increase in supply obviously corresponds to the law of diminishing marginal utility: In order to give up one additional unit of leisure, the individual wants a higher real wage as a compensation for his loss of utility. This is exactly the meaning of the supply curve.

However, one has to clarify the distinction between the marginal return and the average return. Let us assume, as an example, that the supply curve on figure 4 represents the daily supply of labor of the individual. At point A, he works during 10 hours for an hourly wage of $1, so that his daily return is $10. He considers that it is worthwhile working one more hour only if this additional hour is paid at least $2, which means that he wants at least an income of $12 for 11 hours of work. From his point of view,

it is exactly equivalent to get $2 for the marginal hour, the other hours being paid $1, or to get a higher average hourly wage—in our example $1.09—under the condition that he works 11 hours and not 10. If his job is totally specific—there is no possible substitution with other works—the price of successive hours can be differentiated and the employer may be ready to pay more for the additional hour.[8]

Now, if we consider a market situation with substitutable workers, such a differentiation of wages according to the number of hours worked within a day is not possible, so that the supply curve of figure 4 represents the amount of work hours offered for any hourly (average) real wage. Let us anyhow consider the behavior of one given worker. Let us assume that he is at point E when the real wage is equal to w_0. The precise meaning of this point is the following: Any wage lower than w_0 would be insufficient to compensate the individual for the disutility of an amount of labor equal to N_0. He needs at least w_0 or more. However, if the market wage is w_0, he cannot get more than w_0. Now, if the market wage rate increases from w_0 to w_1, the preferred point is A. As there is no wage differentiation, the new wage is paid not only on the marginal hours (N_1-N_0), but even on those hours which were already worked when the real wage was w_0, namely N_0. The change from w_0 to w_1 thus represents the minimum increase in the *average* wage rate which is necessary for the *marginal* return thus obtained to be considered as a sufficient compensation for the increased disutility of work.

Under the traditional neo-classical approach it would be said that the individual got a surplus equal to MEN and that the increase in the market wage brought him an additional surplus equal to NEAR. In fact, the individual was ready to offer N_0 for a wage equal to w_0. Now, he can obtain w_1 for this amount of work, so that he gets a sort of windfall profit equal to DENR; and he also gets a marginal surplus equal to EDA. For reasons linked with the distinction made above between the average and the marginal wage rate, the existence of this surplus can be disputed: One might say that the increase in the average wage rate (w_1-w_0) is just sufficient to allow the individual to get the minimum marginal wage rate he wants in order to offer additional work

[8]There is always a certain degree of specificity in the employment of a worker. For instance, as transaction costs do exist, an employer will prefer paying more for extra hours by a worker who is already employed in his firm than hiring an additional worker.

time (from N0 to N1). Therefore, we cannot quantify something like a surplus. But we can be certain that point A is preferred to point E. For a wage rate equal to w_0 the individual did not want to work more than N0. An increase in the average wage rate necessarily means an increase in the marginal rate, thus inducing him to work more: He prefers A to D. This is the rationale for the upward-sloping supply curve.

This discussion on the very meaning of the supply curve may help us to decide whether this curve can become backward-sloping at some point, for instance at point A. If such a situation exists, it means that, if ever the wage increases from w_1 to w_2, the supply of labor decreases from N1 to N0 (instead of increasing to N2). This is exactly what is assumed by those who believe in the existence of an income effect: As the individual gets more money for any work hour, he works less hours. Is that possible?

Saying that the curve is backward sloping, as represented on figure 4, is saying that point B is preferred to A.[9] As we do know, point B is preferred to point D and to point E (the higher is the wage rate, for a given amount of work, the more satisfied is the individual). We also know that D was not preferred to A, but there is no reason for B not to be preferred to A. However, the real question is not knowing whether B is preferred to A, but knowing whether it makes sense saying that B is on the supply curve. We explained above why A is preferred to D (or, at least, not less desired). Exactly the same reasoning has to be made as regards point B: If the wage rate increases from w_1 to w_2, the individual desires a point which is not on the left of point F (he prefers for instance point C). Therefore, although it may be true that B is preferred to A, C is preferred to B. It means that point B is not a point that the individual may *actually* choose.

Another way to look at this problem is as follows. As we saw above, the supply curve is nothing but a frontier between points which are not desired (those under the curve) and those which are accepted (those which are on or above the curve). At point E the individual declines all wage rates lower than w_0 and accepts all higher wage rates. At point B, the individual is supposed to decline all wage rates lower than w_2 and to accept all higher wage rates. But it does not make sense to assume both that wage rates higher than w_0 are accepted for an amount of work equal to N0, and that wage rates under w_2 are declined for the same amount of work.

[9]If B was not preferred to A, it would just mean that the supply curve cannot go through B.