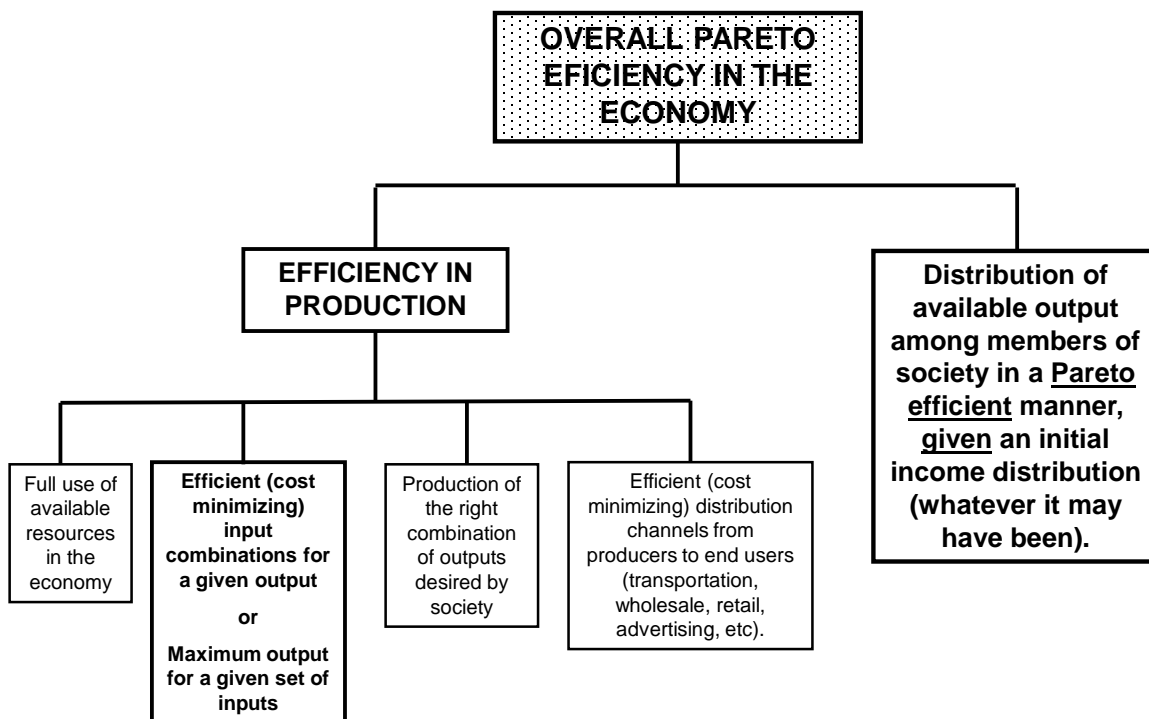


THE CONCEPT OF "EFFICIENCY" IN ECONOMICS

The concept of "efficiency" as used in economics is multi-faceted, as is shown in the chart below.

First, a distinction is made between (a) *efficiency in the production* of goods and services and (b) *efficiency in the distribution* of services from producers to end users.

FIGURE 1 -- FACETS OF "ECONOMIC EFFICIENCY" IN AN ECONOMY



Efficiency in production requires

- a) That available resources are fully used (which means among other things that an economic with involuntary unemployment is *ipso facto* inefficient)
- b) Real resources are used so as to maximize the total social value of the output to be had from any bundle of real resources or, which is the flip side, that any level of output with a given value to society be produced by the combination of real-resource inputs which minimizes the opportunity costs of those real resources.
- c) Producers collectively produce the right quantity and combination of outputs.
- d) that goods and services are carried from their producers to end users through cost-minimizing distribution channels.

Efficiency in distribution requires

- a) That the distribution of goods and services among end users be **Pareto efficient**.

In this lecture, we shall be concerned exclusively with Efficiency in Distribution – the right-most box in bold frame in the chart above.

There is not much controversy over the criteria for **efficiency in production** or in distribution channels (indeed, one could style the latter as part of production). These criteria should be well known to any student in economics.

There is, however, considerably controversy over the concept of **Pareto efficiency** regarding the distribution of output among individual members of society. It is especially so with respect to certain basic commodities, such as health care, education and justice – commodities that most modern societies do not wish to distribute strictly on the basis of price and ability to pay.

The rest of this write-up explores the nature of this controversy.

I. PARETO EFFICIENCY

About a century ago, the Italian economist Vilfredo Pareto offered the world a proposition that can be stated in words as follows:

"An allocation of resources in the economy is economically efficient (now called Pareto efficient) if it is impossible to reallocate the resources so as to make at least one person feel better off without making someone else feel worse off. "

Consider now a two-person economy has a given set of resources (inputs) that could be allocated to the production of this or that set of commodities, each of which could then be distributed to the two persons in a particular way.

We shall think of a particular allocation of input resources to production, coupled with a particular distribution of outputs among these two persons, as an "allocation." What we would like to know is this: when is a particular allocation efficient, and when it is not.

Figure 2 on the next page illustrates the trade-off this hypothetical, two-person economy faces.

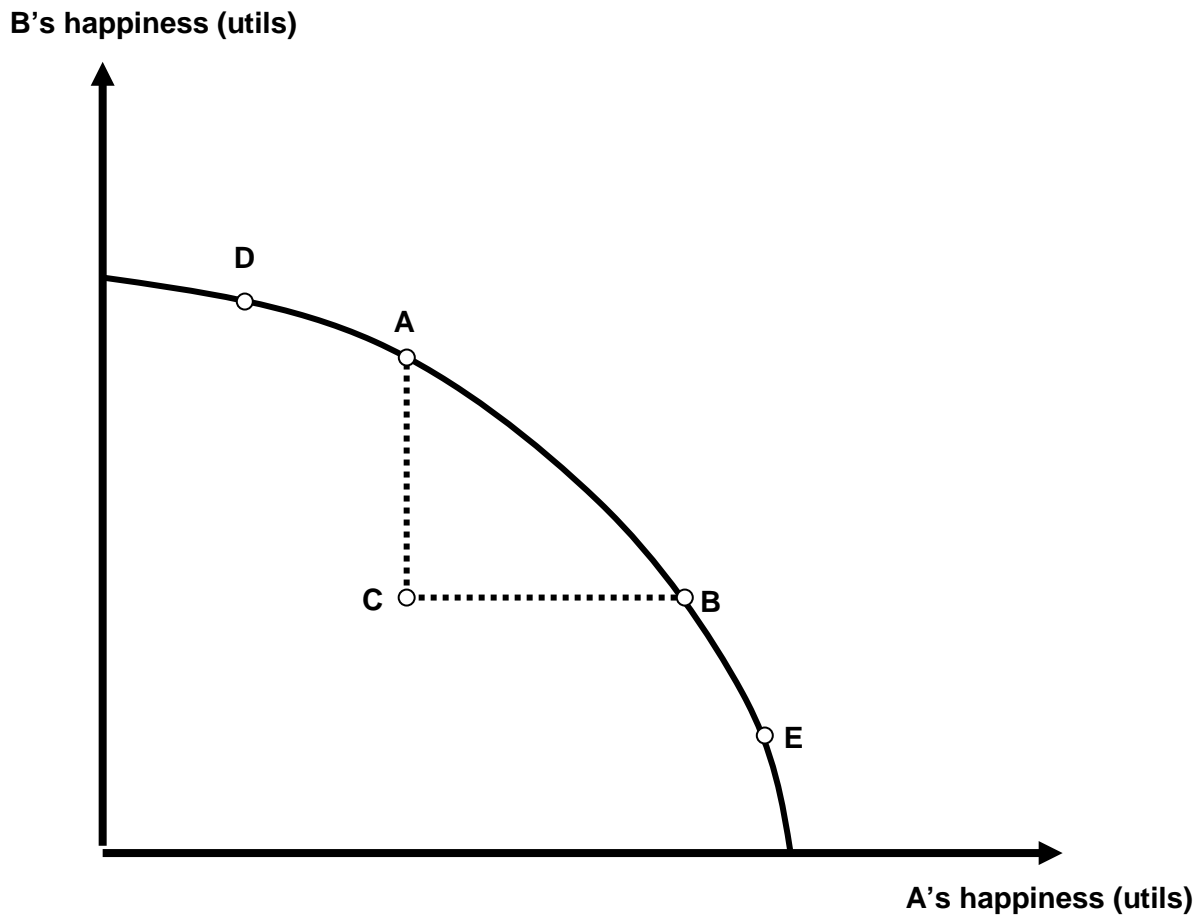


FIGURE 2

Inefficient Allocations: Point C, in the interior of the feasible set of allocations, cannot be judged efficient, because starting with that allocation, one could rearrange the allocation of resources within this two person economy so as to make one or the other of them better off without making the other one worse off, or one could make both better off. In fact, all of the points lying on line segment AB on the efficient frontier are unambiguously **Pareto superior** to point C (more efficient than allocation C), and moving from C to any point on line segment AB is unambiguously a **Pareto improvement**. Because on segment AB both persons are happier, or at least not less happy, we can feel safe in calling a **Pareto improvement** an enhancement of “social welfare.”

Policy Relevance of Pareto Efficiency: What are we to make of the Pareto criterion of efficiency, one so celebrated in economics?

To be sure, Pareto’s proposition does make perfect sense; but as a guide to concrete policy decisions, how far beyond a tautology does it really go? After all, rare are the real-world policy applications that can use Pareto’s criterion to advantage. Most public policies do make some people worse off as

others are made better off, and here this criterion cannot help us. Most public policies create winners and losers – e.g., a move from the inefficient point C to an efficient point D. Relative to a no-trade policy, for example, foreign trade creates winners and losers. Health policy – e.g., the Affordable Care Act of 2010 – almost always creates winners and losers.

Is “Efficient” also “Optimal”?: Let us note in passing that careless economists have fallen into the habit of referring to a Pareto-efficient allocation of resources as "*Pareto optimal*."

That usage does violence to the Latin language. In Latin "optimum" means "best." A *Pareto optimal* allocation, however, can be anything but "best."

Many a situation that civilized people would find abhorrent can be judged "Pareto efficient" without warranting the label "optimal." As noted above, on Pareto's criterion an economy in which some folks (e.g., person A in Figure 1) are literally drowning in resources while others (person B in Figure 1) are starving to death would nevertheless be judged "Pareto efficient" by economists, as long as the diversion of resources from the opulent have's (person A) to the starving have not's (person B) would make even one of the opulent have's feel worse off. Let economists rate the abhorrent status quo "efficient;" but would any *real Mensch* judge it "optimal"? In my view, the term "optimal" should be applied only to situations that reasonable people actually would call "best," in plain English. Don't you agree?

II. POTENTIAL PARETO IMPROVEMENTS

Matters become more complicated – and more like the real world in which we live – as we contemplate moves from, say, point the Pareto inefficient point C in Figure 2 to one of the Pareto efficiency points D or E. Can we economists, as objective scientists, say anything about the social merits of such moves from an inefficient to an efficient allocation of resources, but one at which at least one person is worse of than (s)he was at the inefficient point C? Once again, think of free intrnational trade.

Would such moves towards greater efficiency be a **Pareto improvement**, that is, unambiguously good—an enhancement of “social welfare”? This is an age old question that has plagued the economics profession for over a century and for which there does not exist a satisfactory answer.

One could, of course, try to convert such cases into unambiguous Pareto improvements (welfare enhancements) through a system of side payments (bribes). Thus, one could imagine a rearrangement of the economy such that it is initially moved from the inefficient point C to, say, the efficient point D, at which person B is the winner and person A is the loser. Now if one could arrange it so that person B bribes person A into accepting that change, one could eventually have the economy settle, after payment of the bribe, in the line segment AB at which both are better off or at least neither is worse off.

Many public policies that move the economy to greater efficiency but redistribute economic privilege in the process could be made politically more acceptable if such bribes were actually paid.

Let us define such situations carefully.

A policy (i.e. rearrangement of the economy) under which the winners could, in principle, bribe the losers into accepting that policy is called a potential Pareto improvement.

If the bribe in a situation of a **potential Pareto improvement** is actually paid, then one has thereby been converted it into an **actual Pareto improvement**. Very often free markets make such arrangements by themselves, without even the hand of government. At other times government can set up the arrangement.

The point to note is that converting **potential Pareto improvement** into **actual Pareto improvement** sometimes is possible, and economists are at their professional best when they figure out such arrangements.

But what if, for one reason or other these bribes (economists call them “side payments”) simply cannot be made? Are we stuck then, or is it possible nevertheless to make a scientific statement about the desirability of a move from an inefficient allocation such as C to an efficient one, such as D or E?

In fact, I believe we are stuck here, period. Many economists, however, have been reluctant to give up so easily. They have tried to solve this problem with a bit of shady hand waving—the so-called **Kaldorian welfare criterion**. It works in practice, because many clients buy it. But it does not work in theory. It lacks scientific integrity.

III. THE KALDORIAN CRITERION FOR SOCIAL WELFARE

Consider once again Figure 2 above. Remember that the solid curve represent the maximum happiness this society's resources could afford person B for any degree of happiness accorded person A. It is an interpersonal happiness-trade-off curve.

Suppose now that the policy makers presiding over this two-person society were contemplating a proposal that would move that society from the **Pareto-inefficient** point C to the **Pareto-efficient** point E. (Point E is Pareto-efficient, because at that point one could not make one person better off without making another person worse off). The move from C to E would make person A much, much happier. Alas, it would make person B less happy. Can economists say anything objectively about the economic merits of

such a move--that is, about what such a move would do to the overall "social welfare" in this two-person society?

The late British economist Nicholas Kaldor thought one could. In a nutshell, he proposed the following criterion for judging the impact of policy changes on overall social welfare:

"A reallocation of resources in an economy is an improvement in social welfare if those who gain from the reallocation evaluate their gains at a higher figure than the value which the losers set upon their losses."¹

What do you think of this proposition? To see what a mouthful it is, let us rephrase it slightly:

"A reallocation of resources in an economy is an improvement in social welfare if those who gain from the reallocation (whoever these gainers may be, whatever their tastes may be, and whatever their income and wealth may be) evaluate their gains at a higher figure than the value which the losers set upon their losses (whoever these losers may be, whatever their tastes may be, and whatever their income and wealth may be).

Notice that nothing in this *credo* requires that the gainers bribe the losers into accepting the change. It merely states that the gainers must evaluate their gain at a higher cash equivalent (or other common yardstick) than the losers evaluate their loss. Presumably an objective arbiter – typically an economist – will estimate what the cash equivalent gains and losses are and render the judgment regarding overall social welfare. We therefore can restate the Kaldorian criterion one more time as follows:

"A reallocation of resources in an economy is an improvement in social welfare if the gainers from the reallocation gain enough to be able to bribe the losers from the reallocation into accepting it, EVEN IF THAT BRIBE IS NOT ACTUALLY PAID.

The Policy relevance of the Kaldorian Criterion:

What do you think of this magic wand to create the Good Society?

After all, the Kaldorian *credo* is in the nature of a magic wand by which those who subscribe to this *credo* *believe* to be able to convert any putative **potential Pareto improvement** into a putative **actual Pareto improvement**, even if those who gain from the change in policy never actually do compensate those who lose from the change for their loss? A better name for it, of course, would be a **virtual Pareto**

improvement, because there is nothing real about it.

Do you buy this seemingly scientific magic? If you do, that is, if you do accept the **normative** judgments economists base on that magic, then you elevate the economics profession to the vaunted status of a **collectivist, Platonian high-priesthood** that is empowered somehow to issue normative *dicta* on complex public policy, thereby to direct the allocation of scarce resources among members of society. Are you willing to delegate that task to economists without even a whimper?

You may ask why I use the "collectivist." Think about it! The Kaldorian *credo* embodies the following, highly dubious social ethic:

\$1 for me and \$2 for you = \$2 for me and \$1 for you.

The rule abstracts entirely from who the "me" and "you" is. There's no distinction among individuals here at all. Members of society are just all one big, grey mass. Economists believe they can, legitimately, redistribute economic privilege within the masses, according to their Kaldorian *credo*, individual rights be damned. It reminds me a bit of a cattlefarm in which hay and other food is distributed among the cattle so as to maximize the salable meat on the hoof. Normative economics is cattle-farm economics.

One wonders whether this strange ethical tenet would have appealed even to ardent Communists on a collective farm. Do you find it suitable for an individualist, democratic society? Suppose, for example, you walked into a meeting of the economics faculty, grabbed one econ prof's wallet, took out, say, 20 or \$50 from that wallet and gave the money to another professor at the meeting, telling the assembly that "social welfare" has not changed. Do you suppose these econ profs would buy your argument?

One even wonders whether a card-carrying economists beholden to the normative ethical doctrine of welfare economics would consider it welfare neutral if you just took a book of his or her shelf and kept it.

Measuring "Happiness" (utility) by Money: The Kaldorian criterion itself is still expressed in terms of utility. It merely says that if the winners from a policy could, in principle, bribe the losers into accepting that policy (where the bribe could be paid in money or other favors), then such a policy enhances social welfare even if the bribe is not paid.

In practical applications, however, economists typically and blithely measure the utility of individuals by money – e.g., their willingness to pay money to prevent a policy from being implemented or the money they would have to be paid to accept that policy. Furthermore, in making the translation of human happiness into monetary equivalents under the Kaldor trick, economists tacitly assume that a given change in a person's wealth represents the same change in happiness to rich and poor or, that is, **that the marginal utility of wealth is constant at all levels of wealth**. Consider, for example, the following statements in a textbook on intermediate micro-economics:

According to the [Pareto] efficiency criterion, any change in policy that makes George \$ 2 richer and Martha only \$ 1 poorer is a good thing. Any change in policy that makes George \$ 1 richer and Martha \$ 2 poorer is a bad one. ... Many economists regard [this] efficiency criterion as a good rough guide to policy changes¹.

Usually, the writers of textbooks in economics who write passages like that do cover themselves somewhere, in passing, by observing that this should not be the only criterion for judging a policy, although in the world a practical affairs, the word "efficient" does carry an undue weight. As the famous economist Kenneth Arrow has observed candidly on this point:

"A definition is just a definition, but when the *definiendum* is a word already in common use with a highly favorable connotation, it is clear that we [economists] are really trying to be persuasive; we are implicitly recommending the achievements of optimal states²

Before accepting that a change that makes George \$ 2 richer and Martha \$ 1 poorer is a "good thing," would you not want to know who Martha is and who George is? For example, if you knew George were a billionaire and Martha a waitress paid the minimum wage, would you judge *ipso facto* as "a good thing" a policy that makes George \$ 2 richer and Mary \$1 poorer?

Whatever you may think of the Kaldor trick, however, the economics profession has lapped it up eagerly, like a bunch of long-starved puppies, and for a very simple reason: it has allowed economists to pretend to the uninitiated that they could say something scientifically objective about public-policy changes whose evaluation is, in fact, an inherently subjective matter and, hence, a political matter. After all, such a presumption does enhance the profession's power and thereby the earnings potential of economists as legal- or business consultants.

IV. THE KALDORIAN CRITERION AND COST–BENEFIT ANALYSIS (CBA)

It is important always to recognize that most benefit-cost analyses produced by economists today are based on the Kaldorian ethical credo, as are most other *normative* statements proffered by modern economists.

Such analyses almost always rest on the following ethical platform, which seems preposterous on its face:

¹ Steven E. Landsburg, *Price Theory and Applications*, 3rd. ed., 1995; p. 258-9.

² Kenneth Arrow, "Uncertainty and the welfare economics of medical care," *American Economic Review*, 1963; p. 942.

If the monetary equivalents of the benefits flowing from a proposed policy (to whomever these benefits may accrue, and evaluated in terms of the gainers' actual or imputed willingness and ability to pay for those benefits, given their preferences and their budgetary circumstances) exceed the monetary equivalent of the costs occasioned by the proposed policy (on whomever these costs may fall, and evaluated by the willingness and ability of the losers to pay to prevent their losses from the policy being implemented, given their preferences and budgetary circumstances) then that policy enhances social welfare overall and should be implemented.

When economists use the "willingness-to-pay" criterion in evaluating policy changes, for example, they implicitly use the Kaldor trick as well. Thus, if waitress Martha is willing to pay \$ 1,000 to prevent a policy change for which billionaire George is willing to pay \$ 1,500, then making the policy change is the "efficient" thing to do, say economists, as it "enhances social welfare" by \$ 500.

Economists in the Service of the Well-to-Do: Note how wondrously helpful the Kaldor criterion is in judging the social merits of macro-economic changes – e.g., free trade, tax cuts, corporate restructuring, and so on.

The massive corporate restructuring in America during the past two decades, is known to have reallocated the division of GDP away from labor and to the owners of financial capital, has served vastly to spread the distribution of income and wealth in America, resulting in a significant fall in real hourly wages for the bottom third or so of wage earners. The series of tax cuts enacted in 2001-2005 have yielded far higher dollar savings in taxes for high-income individuals than for low-income individuals.

It is claimed, often by economists, that these measures have helped to increase average real GDP per capita in the U.S. relative to what it would otherwise have been. Relying implicitly on the Kaldor trick, economists may happily score this average economic growth as an increase in social welfare. In your view, is that necessarily so? Do the voters necessarily think so?

The Welfare Economics of Health Insurance Revisited: Let us next explore the application of the Kaldorian criterion to the economic welfare analysis of health insurance as it is so often presented in textbooks. We shall use Figure 3 on the next page to explain the analysis.

In Figure 3 we assume that there exists a perfectly competitive market for some standard health service illustrated in Figure 3. Initially no buyer in this market has health insurance.

We assume, as welfare economists do, that the market demand curve represents the **marginal social value (MSV) curve**. This is easiest to understand if we assume that every individual buys only one unit of health care per period in this market, so that a point on the market demand curve represents a particular individual's maximum bid price for that unit of health care. It is the marginal value that individual assigns to that unit of health care. If we array individuals from left to right along the curve in decreasing magnitude of their marginal bid prices, we get the market demand curve. For the market as a whole it then makes sense to call it the **marginal social value (MSV) curve**.

$P = X$, price per unit of health care

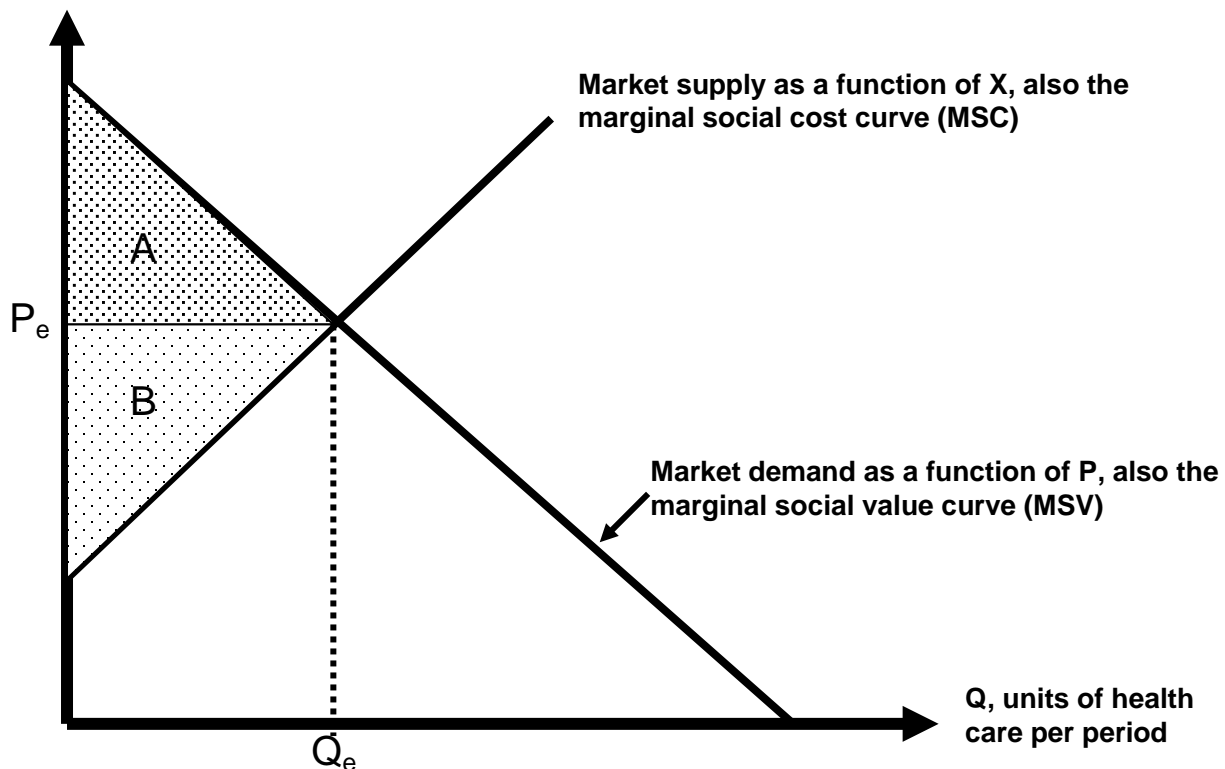


FIGURE 3 – THE MARKET FOR HEALTH CARE, NO HEALTH INSURANCE

Similar reasoning makes us view the upward-sloping supply curve as the **marginal social cost (MSC) curve**. Here we think that providers of care in this market are arrayed along the supply curve in terms of increasing minimum ask prices for units of care. These ask prices increase as we move along the supply curve to the right, because the marginal cost of producing services experienced by these providers increases. On the left segment of the curve are low-cost providers. On the right are high cost providers.

In equilibrium, this market will clear at a price P_e , which is paid by the buyers of care and received by the providers. All but the marginal buyer at the intersection of demand and supply would have been willing to pay more than P_e . They reap a so-called consumer surplus, defined as their maximum bid price

minus P_e . If we add that across all buyers, we get the **total consumers' surplus** in this market, equal to area A in Figure 3.

Similarly, at P_e , only the marginal seller at the intersection of the demand and supply curve just breaks even. All suppliers to the less get paid more (P_e) than their marginal cost of producing care. Their marginal profit per unit is thus P_e minus their marginal costs. If we add this up across all suppliers actually selling in the market, the **total producers' surplus** is area B.

Finally, the **total social surplus** in this market, in equilibrium, is the sum of areas A and B. Maximizing the total social surplus to be had from a market is what economists call "**maximizing the social welfare**" yielded by this market. It is the overarching normative goal of what is called "welfare economics."

With this set up as the baseline, let us now assume that the government introduces into this market a tax-financed, public health insurance program under which it pays 60% of the cost of health care for some patients (e.g., the poor or the elderly). Assume the providers of care collect 40% of their bills from patients and are paid the other 60% by the government program.

Let

X = the total price (40% paid by patients, 60% paid by the government) received by the providers of health care per unit of care and ,

$P = (0.4)X$ the out-of-pocket price patients pay per unit of care.

The impact of this policy on this market is shown in Figure 4. In that graph, the vertical axis shows both, the total price X producers are paid (half by patients, half by the government), and the price patients pay producers, equal to $P = 0.4X$. The market demand curve depicts how many units of care patients demand at alternative out-of-pocket prices they pay. That curve is still the consumers' **marginal social value (MSV) curve**.




FIGURE 4 – THE MARKET FOR HEALTH CARE, WITH HEALTH INSURANCE

The solid, upward-sloping supply curve depicts the number of units of care that providers supply at different total prices X they are paid (40% by patients, 60% by government). That curve is still the producers' **marginal social cost (MSC) curve**.

Finally, the dashed upward sloping supply curve shows the number of units of care providers supply to patients at the out-of-pocket price P patients pay. It is important to recognize that this curve does **not** represent the **marginal social cost curve**. It is simply a reaction curve – the suppliers' reaction to the out-of-pocket price $P = 0.4X$ patients pay. The providers know, of course, that when patients pay them P , they will actually receive $X = P/(0.40)$ in toto per unit of care.

We see in Figure 4 that, after the introduction of this health insurance scheme, the out-of-pocket price P patients pay has fallen below the original equilibrium price P_e , while the total price X providers receive (40% from patients, 60% from government) has increased above the original P_e . The total quantity of health care used has increased markedly by $\Delta Q = Q_h - Q_e$.

Textbooks in health economics invariably use a graph such as Figure 4 to point out that health insurance entails a welfare loss, shown in Figure 3 by the shaded triangle C. The argument might proceed

as follows:

1. The total value (measured by willingness to pay) that patients assign to the added units of health care $\Delta Q = Q_h - Q_e$ consumed by patients as a result of the “moral hazard” of health insurance is the sum of areas D and E.
2. The total additional social cost of producing these extra units of care is the sum of areas C, D and E.
3. Therefore, the total incremental social costs associated by the introduction of health insurance exceed the incremental social value associated with that policy by area C, which is a so-called “deadweight loss” or decline in “social welfare.”
4. Health insurance is “inefficient” in this sense.

The question now is what *normative* significance a policy maker should attach to these conclusion. It is true that society as a whole will incur additional money costs equal to areas C+D+E to bestow on some members of society added health care for which these recipients would have paid at most a money amount equal to areas D+E. But what does that suggest about the merits of the proposed policy?

Note that the statement about *money* tells us nothing about *utilities*, that is, the net change in *human happiness* (well being) brought about by that policy. For all we know, those who paid for the added care lose less happiness collectively than the happiness gained collectively by the recipients of the added care. Simply to assume that amounts of money measures this change in happiness accurately is dubious on its face, especially when those who are made to pay money for the extra care are in a different income class than those who get the extra care. To assume that the deadweight loss C expressed in money terms also reflects a net social loss in human happiness (utilities) tacitly assumes that utility can be measured in money terms and, moreover, that the marginal utility of money wealth is constant across income classes.

The preceding analysis speaks about net costs and gains. A somewhat richer welfare analysis identifies specific groups within society that are affected by the income redistribution inherent in the introduction of health insurance. We use Figure 5 and the standard welfare accounting system in the table below to perform that analysis, using three groups in society: (1) patients who receive the extra care, (2) the providers of health care who produce the extra care and are paid for it, and (3) taxpayers who have to pay 60% of the total cost of health care used by patients in this market.

WELFARE ANALYSIS OF HEALTH INSURANCE

INTEREST GROUP	BEFORE POLICY	AFTER POLICY	CHANGE
Buyers (Patients)	$F + A$	$F+A+B+D$	+B + D
Providers of Care	$H + B$	$H + B + A + G$	+A + G
Government (Taxpayers)	0	$-A - G - C - B - D$	-A - G - C - B - D
TOTAL (SOCIETY)	$F + A + H + B$	$F + A + H + B - C$	- C




FIGURE 5 – THE MARKET FOR HEALTH CARE, WITH HEALTH INSURANCE

Figure 4 and the accounting table show that, relative to the original situation without health insurance, the introduction of the public program under which government picks up 60% of the cost of

health care redistributes money-equivalent surplus from taxpayers to both patients and the providers of health care, as certainly Medicare did when it was introduced in 1965.

In our illustration, patients reap additional consumers' surplus equal to areas B and D. Providers reap additional producers' surplus equal to areas A + G. On the other hand, taxpayers, who previously had not paid anything toward health care, now pay actual money equal to areas A + G + C + B + D. Of this outflow of money, (A + G + B + D) flows as additional surplus to patients and providers.

The remainder, area C, the "deadweight loss" so dreaded by economists, goes up in smoke, so to speak. It accrues to no one in society. It is the "welfare loss" economist attribute to health insurance, which causes the production of additional health care at a cost of E + K + D + C for which patients would have been willing (and able) to pay collectively only E + K + D.

Once again, however, this redistribution of money equivalents tells us nothing about the associated redistribution in utility. To assume that it does, we must assume that utility is well measured by willingness to pay, which abstracts from ability to pay and assumes that the marginal utility of money wealth is constant across all income classes.

Cash Transfer instead of Tax-Financed Health Insurance: It may be argued, and economists routinely do argue, that all patients collectively could be made happier if they were simply given cash in the amount taxpayers now pay to subsidize health care.

The argument is that, instead of receiving additional consumer surplus of B + D through the health insurance program, if patients were given cash equal to A + G + B + D + C they could create far more happiness for themselves than is yielded them by the added consumer surplus B + D.

This is undoubtedly true, because patients could then still recreate precisely the same outcome that health insurance would, but they might prefer to spend some of the cash transfer on other pressing items – e.g., a better car.

This argument, however, has a major flaw. It is doubtful that voters would approve of so large a transfer of unrestricted cash. Can one really believe that voters would be indifferent between funding, ay, the current Medicaid program and simply giving the poor the same amount now spent on Medicaid in the form of unrestricted cash transfers, letting the poor spend that cash as they see fit? Where have economists been even to dream of such an outcome?

It is highly unlikely that the typical taxpayer wishes to see the poor maximize their own happiness for every dollar of taxes paid by taxpayers. Instead, taxpayers want the poor to behave in certain ways and not in others. Thus they want the poor to use tax-financed health care for themselves and their children, but taxpayers do not want the poor to enjoy tax-financed football games or alcohol. Why economists find that so hard to grasp has long been a mystery to me.

VI. NORMATIVE ECONOMICS AS FAITH-BASED ANALYSIS

If economists wish to follow the Kaldor *credo* within the analytical worlds they create for themselves--and a pure *credo* it is--that is their professional prerogative. It allows economists to produce papers that they and their peers enjoy.

The question, however, is: Why should anyone else buy this odd social ethic. Why should they obey what they are told by what I have called earlier the collectivist, self-styled *Platonian high priests of normative economics*? After all, *believing* that the faithful application of the Kaldorian criterion to public policy will lead to the good society really is *faith-based analysis*.

Can we really be surprised, for example, that the economist's definition of "economic efficiency," and the benefit-cost analyses erected on it, sometimes caused serious communication problems between economists and politicians? Does it mean that politicians are intellectually challenged, or does it mean quite the opposite?

Distributive Weights: One could, of course, blend the economist's notion of efficiency with the political construct of **normative, distributive weights**. For example, through a political process it might be decided to give Mary, the waitress, a weight of 2 and George, the billionaire, a weight of 0.9 in the benefit cost calculus for a particular policy change. If that policy change made George \$ 2 richer and Mary \$ 1 poorer, it might then be scored as $(.9)\$2 - 2(\$1) = - \$0.20$, a social net loss. Consequently, it would not be judged a "good thing" and, therefore, it would be rejected. Some economists have proposed that approach. Note, however, that it is not the economist's prerogative to set the distributive weights. In a democratic society, they can only be the product of a purely **political** process.

Honorable Economics: Honorable economists will always be quick to apprise their audiences and clients of the distributional impact of proposed policy changes. Alas, the real world is full of "seasoned economists" who do not, either because they wish to push their own ideology, or because they are paid to abstract conveniently from the distributional consequences of proposed policies--e.g., proposed health reforms or trade policies. "This policy is efficient," they may loudly proclaim, adding inaudibly under their breath "abstracting from distributional effects, of course." Sadly, some policy makers buy this stuff, either because they innocently believe the economist, or because they find it profitable or ideologically soothing to do so.

Caveat Emptor: As a potential future consumer of economic analysis you may wish to remember the following GPL (*Guide for Prudent Living*):

When economists or other policy analysts lapse into their *normative* mode – that is, when they pretend to be using scientific methods to suggest what *ought* to be done and what is or is not "efficient" or "welfare enhancing" -- a red warning light should go on in your mind.

There is always the chance that you are being addressed by someone knowingly and cleverly playing politics in the guise of science, or by someone commissioned by an interest group to structure information felicitously, toward a desired end – that is, by someone insufficiently respectful of the limits of economics as a science.

Finally, there is always the chance you are face-to-face with an economist who has never been properly taught or, in any event, has never fully grasped the meaning of the word “efficiency” within economic analysis.

I admit that this rule comes across as a bit cynical. Indeed its application may sometimes be unjust, for some economists may sincerely *believe* that, by *faithfully* (really: "*faithfully*" like in "*full of faith*") following the Kaldorian *credo*, they are doing God's work on earth: maximizing human happiness.

For example, one of my colleagues, whom I respect greatly, once told me:

Even if the application of the Kaldor criterion were unjust in each instance, if it is consistently applied over time, and to a large portfolio of public policies, then in the long run its use will beget a giant positive-sum gain from which everyone probably gains.

The idea here is that the direction of individual instances of injustice will wash out--that an individual will lose from some public policies and gain from others, but that over time and over the entire portfolio of public policies, everyone will book a net gain. This may be so; but it cannot be analytically proven to be so. Nor, to my knowledge, has this hypothesis been convincingly empirically supported, at least not to the point that we would not argue about it. In other words, in my view this hypothesis, too, remains in the nature of a *credo*. Thus, I repeat the *Caveat Emptor!*

A Code of Ethics for Economists: Curiously, unlike many other professions, the economics profession has never seen fit to articulate an explicit code of professional ethics to which all economics profess allegiance, although economists routinely structure information for important policy decisions that affect the well being of human beings. One would hope that, one day, the economics profession might feel moved to mimic the many other professionals who do profess allegiance to an explicit code of ethics (even

if sometimes the codes are observed in the breach). Such a code, which might be developed by the *American Economic Association (AEA)*, would frown upon careless or duplicitous normative economic analysis. In the meantime, the lack of an explicit code of ethics for applied economics makes "creative economics" even more dangerous than is "creative accounting."³

³ Such a code of ethics would be accompanied by a more formal process of auditing empirical research than is customary today. Economists routinely audit one another's methodology through rigorous peer review, but rarely do they insist in that peer review on replicating empirical results with the original data tape used by the authors of empirical research. It is a major shortcoming. I have written about this lack of auditing in U. E. Reinhardt, "Making economic evaluations more respectable," *Social Science and Medicine*, vol. 45, No. 4, 1997; pp. 555-62.


APPENDIX A

EFFICIENCY IN PRODUCTION

So far in our analysis we have dealt mainly with efficiency in the *distribution of output* among members of society, taking it for granted that this output has been efficiently produced. That may or may not be so.

As you can imagine, however, the concept of Pareto efficiency can easily be adapted from the context of happiness production to the *production of goods and services*. This jump from one context to another is illustrated in the graph below.

OUTPUT OF PRODUCT "Y" PER PERIOD



Here we have on the axes not human happiness, but volumes of output per period of two distinct commodities, X and Y. Continue to assume that these two commodities are produced for the two-person society we have explored so far.

The question now is this: Can we assume that a reallocation of resources that moves this two-person society from output (not happiness!) combination C to any output combination on the Pareto efficient line segment AB is *ipso facto* a *Pareto improvement*, an *enhancement of social welfare*, because the two-person society we are modeling collectively now has more of one output without needing to sacrifice any of the other output?

That conclusion has great intuitive appeal if you do not think too deeply about it—that is, if like a computer your cerebrum is on a stand-by mode. For example, economists almost instinctively believe that any increase in average GDP per capita represents an increase in social welfare. Indeed, an increase in real GDP per capita is the deity economists worship.

But, remember, the concept of “social welfare” refers to *human happiness*, not tons of physical output. While the consumption of goods and services surely is one (but only one) means of producing human happiness, commodities by themselves are not the same thing as human happiness.

To certify honestly that a move in the commodity space from output combination C to line segment AB is, in fact, a genuine Pareto improvement in terms of human happiness (“social welfare”), we really should know exactly how the extra output is distributed among the two folks in this two-person society. Might it not be that at point A, or B, or any other point on line segment AB in this commodity space, there is actually less human happiness in society than at point C, because one or the other human being was made worse off in the happiness space?

Put another way, is it not possible that a move from, say, output combination C to A in the commodity space actually makes one person in this two-person society **feel** worse off than he or she did at point C? It might happen, for example, if all of the extra output achieved through greater efficiency were given to only one of the two people, kindling in the other person a sense of relative deprivation. After all, there is such a thing as social envy. It can cause enormous social tension, crime, riots and even full-fledged revolutions.

To think about this problem, beam yourself back into your own childhood. Imagine yourself part of a household with two children who form our two-person society. If your parents gave each of you one Snickers Bar on Day 1 and, on the next day, gave two Snickers Bars to your brother or sister, but only one to you, is it your recollection from your childhood that, on Day 2, there would then be more bliss all around in your household than there had been on Day 1, because your sibling got more Snickers Bars on Day 2 than on Day 1, and you got no fewer than you did on Day 1? If you answer “Yes,” I’d like to know more about your unique family. You children would have been either, like, totally weird, or you might have been nascent neo-classical economists.

You may argue that children are, well, childish -- that social envy is a childish trait, totally unknown among seasoned adults, such as you now are. Oh yeah? Have you ever heard the term “keeping-up-with-the-Joneses,” which is one of the major driving forces behind economic growth? Have you ever read any book by Thorstein Veblen (e.g., his tome on the theory of conspicuous consumption)? If, say, a costly new medical miracle were made available only to some Americans who need it and can afford to pay for it, but not to others who need it just as much, would the American public really judge such a development as a genuine *Pareto improvement*, an enhancement in American social welfare? If not, do economists have the right simply to overrule with their Kaldorian *credo* the infidel, unruly, socially-envious plebs on this point? What do you think?

Consider again an economy that has experienced rapid economic growth over time, but in a way that has bestowed most of the benefits of that growth on the high-income classes, while the real (inflation-adjusted) incomes of families in lower income-classes have remained merely constant. Do we score this as a genuine Pareto improvement?

It is by now well known that the bulk of the real growth in U.S. GDP during the past 20 years has accrued to families in the top 10th percentile of the family-income distribution and, within that class, to the top 1 percent. Have American voters been indifferent to this fact? Are we a happier nation because average GDP per capita went up?

Economists might say so. *The Wall Street Journal's* editors might say so too. These people seem to believe that if average real GDP per capita in a society has risen, social welfare in that society has *ipso facto* been enhanced.

Would you think so? Would everyone else in our society say so?

APPENDIX B

HOW ECONOMISTS BASTARDIZED BENTHAMITE UTILITARIANISM AND BECAME SHILLS FOR THE WELL-TO-DO

A. BENTHAMITE UTILITARIANISM

Jeremy Bentham (1748-1832), an English philosopher, entered Oxford University at the ripe age 12 and graduated from there at age 15.

He and John Stuart Mill (1806-1873) are widely regarded as the founders of the controversial philosophical school called *Consequentialists*.

Consequentialists hold that the merits or demerits of human acts – business decisions, legislation, crime and punishment – should be judged strictly by the pleasure or pain, or both, that these acts visit on human beings (or animals), rather than by some other intrinsic merit or demerit of those acts (e.g., some religious stricture).

Bentham called the attempt to value the total consequences of pleasure and pain of an act as the *felicific calculus*. You have had occasion to engage in *felicific calculus* on some your homework assignments, and some day you will hire economists who will perform it for you. You used it in what we have called “welfare analysis,” when we added up triangles under demand curves and above supply curves to calculate what we have called “social surplus.” In fashioning their normative *dicta*, economists seek to maximize this mysterious something called “social surplus.”



Jeremy Bentham

Bentham, Mill and their disciples believed that the pleasure and/or pain begotten by an act could be quantified and measured as “utility” or “disutility,” which is why they are also known by the more common label of “utilitarians.” They argued that public policy – in local commerce, in intern-ational trade and in the law -- should be conducted so as to *maximize the sum of utility (happiness) in society* which, you will quickly recognize, is roughly the same idea that drives modern welfare economics as well. In fact, 19th century utilitarianism can be regarded as the intellectual foundation for what we now know as “welfare economics” or “benefit-cost analysis. Modern welfare economists are *Consequentialists par excellence*.

The practical problem with utilitarianism – and with its modern successor, welfare economics -- is how to quantify and measure utility and disutility (happiness and unhappiness), especially in a manner that allows one to make *inter-personal utility comparisons* and to sum utility across different individuals into what economists call “social welfare.”

If you read the writings of the 19th century utilitarians, especially of Bentham, you will discover that they were not at all cavalier about the *cardinality* of utility in practice (where by “*cardinality*” is meant that utility or disutility could be quantified and measured so that, say, 150 utils represented 50% more pleasure than 100 utils, and so on.)⁴

Furthermore, they recognized that there existed inter-personal dependencies of utility – what we have called externalities in consumption – and that it would be extraordinarily difficult in practice to make inter-personal utility comparisons, or to sum utilities over different people, without attaching additive weights to individuals. Such weights, for example, might count the utils of some person as more than those of others.

Finally, and very importantly, the utilitarians assumed that the *marginal utility of wealth was strongly diminishing*. In Jeremy Bentham’s words:

Of two people having unequal fortunes, he who has most wealth must by a legislator be regarded as having most happiness. But the quantity of happiness will not go on increasing in anything near the same proportion as the quantity of wealth: ten thousand times the quantity of wealth will not bring with it ten thousand times the quantity of happiness.

*It will even be matter of doubt, whether ten thousand times the wealth will in general bring with it twice the happiness. The effect of wealth in the production of happiness goes on diminishing, as the quantity by which the wealth of one man exceeds that of another goes on increasing: In other words, the quantity of happiness produced by a particle of wealth (each particle being of the same magnitude) will be less at every particle; the second will produce less than the first, the third than the second, and so on.*⁵

The practical implication of the hypothesis of *diminishing marginal utility of wealth* in the context of Benthamite utilitarianism is momentous.

It implies, *inter alia*, that public policy should achieve a more equal distribution of wealth than would be begotten by a market system left to its own devices. The hypothesis has been used by Benthamite philosophers to advocate progressive taxation, on the “equal sacrifice” principle—equal in terms of disutility of taxes, that is, not equal in terms of sacrifices of money.

B. UTILITARIANISM AND MODERN WELFARE ECONOMICS

⁴ In modern consumer choice theory, we deal with *ordinal* utility. It means that we merely can say that a person prefers a thing to something else. Thus, if a person assigns 150 utils to good A and only 100 utils to good B, we can merely say that good A is preferred by the person to B, but we cannot say that the person derives 50% more happiness from good A than from good B.

⁵ http://highered.mcgraw-hill.com/sites/0072875577/student_view0/chapter3/origin_of_the_idea.html#2. By the way, as is mentioned in this source “Bentham donated his own body for dissection, largely to promote acceptance of the practice. This, however, was not the end of Jeremy Bentham. He also left his estate to University College, London, but under the condition that his remains be present at all board meetings. His padded and dressed skeleton still sits (in a glass case) at the college. A wax head sits atop the body, and his actual head was preserved using South American headhunting techniques. Bentham’s real head used to sit on a plate between his feet, but in recent years has been relocated to a safe at the college, presumably to protect it from would-be pranksters using it on the soccer pitch or in other unspeakable ways.”

Modern welfare economics, as already noted, can be viewed as an offshoot of Benthamite utilitarianism, but modified in ways some might call dubious. Let us see why this is so.

In a nutshell, in their economic welfare analyses, economists measure the satisfaction (utility) people derive from a thing by the maximum money price they would bid for the thing. A market demand curve, you will recall, in effect is an array of bid prices offered by different potential buyers of a thing, arrayed from left to right in terms of descending bid prices. The fundamental question you must ask yourself is this: given that the bid prices expressed in a market reflect not only the different tastes of the bidders for a thing, but also their ability to pay (income and wealth), how legitimate is it to use bid prices as a measure of the utility the bidder expects to derive from the thing? Concretely, if Bill Gates bids \$10 million for a painting and you only \$50,000, does it mean Bill Gates expects to derive 20 times as much satisfaction from the thing than you expect?

In any event, to apply Bentham's utilitarian construct to their analyses, as they do, modern welfare economists must make these two assumptions:

1. **First**, the pleasure (or "economic welfare") that a person derives from a good thing can be measured *cardinally* by the money price he or she is willing to bid for the thing.
2. **Second**, an additional, say, \$100 will yield the same pleasure to a billionaire as it would to a pauper, and that a \$100 loss will visit on the billionaire the same degree and intensity of pain as it would on the pauper.

The second of these assumptions are crucial to welfare economics. They allow economists to make statements such as "*on the efficiency criterion any change in policy that makes George \$2 richer and Martha only \$1 poorer is a good thing,*"⁶ or, "*If, as a result of implementing Policy A, Jack gains \$10 and Jill loses \$5, Policy A yields a social welfare gain of \$5*"⁷ without having to give a thought to the relative incomes enjoyed by George and Martha or Jack and Jill.

Implicitly, most benefit-cost analyses offered by economists proceed on this assumption, aside from the rare instances in which economists use distributive weights, that is, weight equal dollar benefits or costs accruing to different people differently. Such weights, however, would be purely subjective. Implicitly, your textbook also proceeds on these assumptions when it covers consumers' and producers' surplus and the definitions of "efficiency" and "social welfare."

C. MAKING THE SOCIAL VALUE OF THINGS (INCLUDING HEALTH CARE) DEPENDENT ON THE WEALTH OF THEIR RECIPIENTS.

To illustrate what has just been said, consider the market demand curve for a certain thing, shown on the next page. Points E, F and B on that market demand curve can be thought to represent to maximum bid prices offered by individual buyers of the thing. Although each individual may be represented by more than one point on the demand curve, it helps us here to make the harmless assumption that each individual buys at most one unit of the thing per period.

The relative position of individuals along the market demand curve may reflect differences in taste


⁶ Steven E. Landsburg, *Price Theory and Applications*, 3rd ed., (1995): 258.

⁷ See <http://www.thebigquestions.com/2010/08/30/efficiency-experts/>, a blog exchange between Steven Landsburg and Uwe Reinhardt, exchange toward the end of the commentaries.

rooted the individuals' age, gender, education, prior experience with the product and so on.

But it is reasonable to assume also that, say, individual E is willing **and able** to bid more for a unit of the thing than are B or G because E is wealthier than B or G. Even if tastes were the same for all individuals in the graph, they might offer different maximum bid prices simply because they have different incomes.

THE MARKET DEMAND CURVE FOR A THING



OCBX it is simply the sum of the different *maximum bid prices* that the 100 customers who bought a unit of the thing being traded here would have been willing to pay per unit of the thing, if push had come to shove, that is, if the seller charge each customer a different price for the thing. Area OCBX would then be the **maximum sales revenue** that the seller of the units could have extracted from the customers under that condition of perfect price discrimination.

But is it legitimate to view this “*maximum potentially extractable revenue*” as the same thing as “*social value*”?

The author of your textbook – indeed virtually all textbooks in economics – would routinely assume so. But can we really be sure at all that the last customer to deal in this market, whose maximum bid price was B, “valued” or “appreciated” a unit of the thing less – that is, derived less happiness from it – than did the customer willing to bid the higher price A?

Think about it? The person bidding B might be relatively poor, for example, and the person bidding A might be quite rich.

In equating revenue with social value, are we not in effect saying that allocating a unit of the thing to the person willing to bid A creates greater “social value” (utility) than allocating a unit to the person willing to bid only B? But that is what many economists assume implicitly and others explicitly. In a reading assigned earlier in the course, the author asserted that

Rationing through markets and prices is usually an effective, [efficient](#) method of allocating commodities. It works automatically and under the right circumstances, it

achieves an efficient allocation of resources. That is, commodities are allocated to those buyers willing to pay the highest price because they receive the greatest satisfaction.

(See http://www.amosweb.com/cgi-bin/awb_nav.pl?s=wpd&c=dsp&k=price+rationing .)

To put this another way, are we not saying that if a unit of output were taken away from someone willing to bid maximally only \$10 per unit and allocated that unit to someone willing to bid \$25, that the net social value begotten by that unit increases by \$15? Do you buy that proposition?

Finally, in yet other words, are we willing to say that the “social value” of a thing depends on the wealth of its recipient which, in turn, drives the maximum bid price (s)he is able to offer for the thing?

To make this question very concrete, how would you answer this question if the “thing” in question were a dose of a pain killer or a drug-eluting stent implanted in a person or a cancer treatment, are we prepared to say it creates greater “social value” or “economic welfare” if given to a wealthy person willing to bid a high price for the treatment than it would if given to a poorer person unable to match the maximum bid price the wealthy person can bid?

Remarkably, textbooks in economics rarely if ever raise these questions when they trot out welfare economics to you. But don't you think that these are entirely reasonable questions to raise in a social science?

D. CONSTANT MARGINAL UTILITY

Our problem here would, of course, go away if the marginal utility of wealth were indeed constant over all levels of wealth for each individual and, moreover, an additional \$1 of wealth triggered the same intensity of happiness (utility) for all citizens.

If that were the case, then one could legitimately argue that the different maximum bid prices different people would offer for a unit of the thing do represent equally sized differences in the satisfaction (or happiness, utility or “economic welfare”) that the different bidders derive from the thing.

Furthermore, one could then with a straight face argue that because individual E was willing to bid a much higher maximum bid price for a unit of the thing than was, say, individual F, greater human happiness or economic welfare or social value is created by the unit going to E than is created if the unit that goes to F.

Finally, one could then pretend that the shaded area in the graph above truly does represents a meaningful measure of the “total economic welfare” created by 100 units of the thing per period, prior to the deduction of production costs.

Economic welfare analysis then would make perfect sense, as would the benefit-cost analyses based on it.

Alas, the assumption of a constant marginal utility of wealth is “unrealistic” for two reasons.

First, it does not conform to common sense, as Bentham noted.

Second, the very economists who implicitly (rarely explicitly) posit **constant marginal utility of wealth** for their normative economic welfare analyses routinely posit **diminishing marginal utility of wealth** when they explain, for example, why people buy insurance.

E. CONCLUSION

Modern applied welfare economics may look like objective science, because it typically is cloaked in mathematical symbols or graphs. At its core, however, economic welfare analysis is but one particular **distributive ethic. It is just one of many different moral doctrines.**

You may or may not buy that particular distribute ethic. Whether or not you do may possibly depend on – you guessed it – your wealth and ability to offer high bid prices for things.

APPENDIX C

**THE FIRST AND SECOND THEOREM OF OPTIMALITY
IN WELFARE ECONOMICS⁸**

Many economists seem to believe that if a nation's distribution of income and wealth is "right" (i.e., ethically acceptable), then the allocation of goods and services among individuals can be left to the private market. This *credo* – and a *credo* it is, as is much of economics – is based on the First and Second Optimality Theorems.

THE FIRST OPTIMALITY THEOREM

If a market that meets all of the stringent conditions for qualifying as a "perfectly competitive market" has an equilibrium, and if all goods and services relevant to utilities (human well being) and costs are, in fact, priced in that market, then that equilibrium is necessarily optimal in the Paretian sense, which means:

There is no other allocation of resources (inputs and output) which will make some participants in the market feel better off without making someone else feel worse off.

For a market to be perfectly competitive, at least the following stringent conditions must be met:

1. Both buyers and sellers can enter and exit the market freely, without restrictions imposed by high entry or exit costs, by other competitors or by government.
2. Both buyers and sellers must be fully informed about and understand all of the relevant attributes of the good or service being traded in the market.
3. Neither buyers nor sellers in the market individually have enough market power to influence with their decisions the market-clearing prices of the goods or services being traded in that market.

THE SECOND OPTIMALITY THEOREM

Every equilibrium position in a perfectly competitive market in which both, the happiness-production functions (utility functions) of individuals and the production functions for goods and services, are **convex** throughout corresponds to one **unique** initial distribution of purchasing power among the participants in that market.

Policy Implication: In an economy characterized by such markets, distributive social policy can confine itself simply to achieving a fair initial distribution of purchasing power. Market forces will then automatically drive the economy to a Pareto optimal equilibrium. By changing the initial distribution of purchasing power, any desired socially optimal equilibrium can be obtained through an iterative process, if that redistribution of purchasing power can be achieved without distorting the economy (a huge IF).

⁸ See Kenneth J. Arrow, "Uncertainty and the Welfare Economics of Medical Care, *American Economic Review*, December, 1963, pp. 942-3.