# Chapter 18: Economics and the Good Society

> The preceding discussion makes it plain that any rigid inference from effects on economic welfare to effects on total welfare is out of the question. In some fields, the divergence between the two effects will be insignificant, but in others it will be very wide.
>
> Nevertheless, I submit that, in the absence of special knowledge, there is room for a judgment of probability. When we have ascertained the effect of any cause on economic welfare, we may, unless, of course, there is specific evidence to the contrary, regard this effect as probably equivalent in direction, though not in magnitude, to the effect on total welfare; (AC Pigou (1920). *The Economics of Welfare*, KL: 438-443).

## I. Introduction: The Utilitarian Foundations of Welfare Economics

To this point, we have explored the implications of rational decision-making in the sense of optimization in various choice settings and used those implications to understand how consumers allocate their funds among products, why firms produce and sell what they do, the prices at which such goods are sold, and also analyzed the factors that can induce changes in all those activities. The latter provided us with a basis for understanding how such factors affect the extent and growth rates of market-based economies. The analysis has been "positive" in that no claims were made about whether the outcomes that emerge from the great networks of voluntary exchange are "good" or "bad," desirable or not, "moral" or not, or "ideal" or not. Such conclusions were left to the reader's own internalized beliefs about such things.

This chapter explores some of the normative ideas that economists have used to assess whether markets are "good" or not and whether the markets that emerge are as "good as possible" or not. This area of study was given the name "welfare economics" or "the economics of welfare" by Arthur Pigou in 1920. That name is a bit confusing for American readers, where the term "welfare" is not a description of the quality of one's life, but instead used to describe a particular set of public policies designed (at least in principle) to improve the lives of the poorest persons in their society.

What Pigou had in mind was the utilitarian project, the efforts of an important group of philosophers and public policy reformers to determine the nature of a "good society" and how to move towards such a society. Pigou did not create the utilitarian project, which dates back at least as far as Aristotle's analysis of the relative merits of different forms of governments and societies in his book the *Politics*. The term utilitarian originated with Jeremy Bentham's work in the late eighteenth and early nineteenth centuries, who founded that school of philosophy and also, like Aristotle, attempted to characterize a good society. Bentham also had an interest in economics and wrote a textbook for the field. He is the first to clearly state the possibility of mutual gains from trade and to insist on what would later be called methodological individualism—the approach to social science that insists that all social phenomena should be explained in terms of individual decisions and interests. That methodology has characterized microeconomics since its beginning.

This is not surprising once one realizes that many of the most prominent microeconomists of the late nineteenth and early twentieth centuries were utilitarians of one variety or another. Utilitarians interested in markets wrote many of the first economic textbooks, and they were often pioneers in economic analysis, including the marginal revolution of the late nineteenth century.

In a nutshell, the utilitarian normative perspective is that an activity, rule, or public policy is "good" or "morally proper" if and only if it tends to increase the sum of utility in the community of interest (where utility is interpreted as the extent of personal happiness or satisfaction). If an activity, rule, or public policy tends to increase the sum of utility in a community, it is a moral activity, a good rule, or a welfare-increasing public policy. If it does the opposite, then it should be avoided by all virtuous persons. Voluntary exchange, from this perspective, is a good activity because the parties involved must each benefit from every exchange—unless fraud or coercion is involved or the trade significantly harms other persons not directly involved in the exchange.

The latter argument, together with Adam Smith's conclusion that market activities tend to be self-regulating and to make societies wealthier in the long run, helped persuade politicians and voters throughout the West to gradually deregulate markets, to stop selling monopoly privileges, and to open occupations up to all persons qualified (e.g., with the required skills). These reforms of the old medieval order significantly increased the openness and competitiveness of markets, which, as parts I and II of the text demonstrate, tend to expand the scope of trade—other things being equal. Supporting such reforms are, thus, virtuous activities from the perspective of utilitarian normative theory.

For utilitarians of the nineteenth century, market activities—exchange, production, innovation—were inherently "good" rather than morally neutral or to be avoided because they undermined virtue. This was a significant shift from the Western medieval perspective where commerce was considered at best a morally neutral activity, and at worst one to be avoided to the extent possible.

As neoclassical economics developed and utilitarian ideals were sharpened, utilitarians began to ponder whether markets produced the very best outcomes for society. It was possible, after all, that markets improved society without truly maximizing the sum of utilities. To explore these issues, Pigou and others used what might be called monetized indices of utility to approximate aggregate utility. The use of aggregate net benefits (measured in dollars or other currencies) as a proxy for aggregate utility allowed stronger claims to be made about whether the sum of utilities could be increased through various public policies such as economic regulation or the provision of various kinds of social insurance. A subset of these have been incorporated into public policies. For example, some government agencies are required to use cost-benefit analysis to inform their policy choices.

Pigou believed that public policies could improve welfare in this sense, although other utilitarian economists disagreed. As it turned out, there were assumptions about markets that suggested that competitive markets tend to maximize social net benefits and others that argued that market outcomes could be improved upon. Scholars who agreed about the latter cases might nonetheless disagree about the political feasibility of the policies that could do so. Neither politicians nor voters are all utilitarians. Thus, there were many disagreements among utilitarian economists about whether particular policies increased social welfare or not—and these continue through to the present.

Another source of disagreement concerns the use of the utilitarian norms themselves. Utilitarian normative theory is not the only one employed by economists, although it and its variations are the most common. Several economists have used other normative principles to appraise the relative merits of market outcomes and public policies. These include the Pareto criteria and the contractarian approach.

This chapter provides overviews of the normative theories used most frequently by contemporary economists. As true of the other topics covered in this book, the aim is to provide an overview of the main or core theories and conclusions, rather than a comprehensive review of the literature—which would require a full-length book or two to do reasonably well.

## II. The Pareto Norms and Voluntary Exchange

We'll start with one of the simplest normative ideas used by economists, one that is attributed to the Italian socio-political economist Vilfredo Pareto, whose work was developed at roughly the same time as that of Alfred Marshall. A change from one social state (A) to another (B) is said to be a Pareto superior move if and only if at least one person is better off at B than at A, and no one is worse off. In utility terms, this occurs when at least one person realizes additional utility at B relative to A, and no one realizes more utility at A than at B. A social state (C) is said to be Pareto optimal or Pareto efficient if no Pareto superior moves from C are possible. Any move from C would either fail to make anyone better off than they were at C, or at least one person worse off than they were at C.

Thus, most moves away from a Pareto optimal state would violate the "do no harm" principle. If there is a unique Pareto optimal state (or policy), all moves away from that state violate the "do no harm principle."

### The Pareto Norms in a Simple Trading Game

A simple game-matrix representation of trade without transaction costs can be used to illustrate the Pareto normative criteria. Suppose that Al is a seller, Bob is a buyer, and that there are unrealized potential gains to trade between Al and Bob. Trade requires offers to be made and accepted. The payoffs are in utility gained or lost relative to their original position before this potential trading opportunity emerged.

**Table 18. 1 Gains from Trade without Transaction costs**

| | | | |
|---|---|---|---|
| | | Bob (a potential buyer) | |
| | | Accept Offer | Do Not |
| Al (a potential seller) | Make Offer | $\underline{3}$, $\underline{3}$ | $\underline{0}$, 0 |
| | Do Not | 0, $\underline{0}$ | $\underline{0}$, $\underline{0}$ |

Notice that the (3,3) cell is Pareto superior to the other three cells. Both traders realize higher utility in that cell than in the others, so both would be better off shifting from any of the other cells to the upper left-hand cell. Once the (3,3,) cell is reached, no other Pareto superior move is feasible (because of the assumptions of the trading game in the choice setting characterized). Thus, the upper left-hand cell is Pareto optimal. A move away from that cell to any of the others would make each worse off.

Nonetheless, there are two Nash equilibria in this game matrix, one at the upper left-hand corner and one at the lower right-hand corner. However, only one of the equilibriums is Pareto optimal. According to the Pareto criteria, the upper left-hand corner is better than the lower right-hand one.[1]

### The Pareto Norms and the Edgeworth Box

An Edgeworth box can be used to develop a somewhat more nuanced illustration of the Pareto principles. Figure 18.1 is a slightly more detailed version of the Edgeworth box used in Chapter 5. It begins with the same initial endowment and includes the same market-clearing price line. For the present chapter, however, we are more interested in the areas that characterize the possibilities for realizing mutual gains to trade and the points where such gains from trade do not exist (the sets of Pareto superior moves and Pareto optimal outcomes) than in characterizing an equilibrium price.

Figure 18.1 The Pareto Criteria in an Edgeworth Box

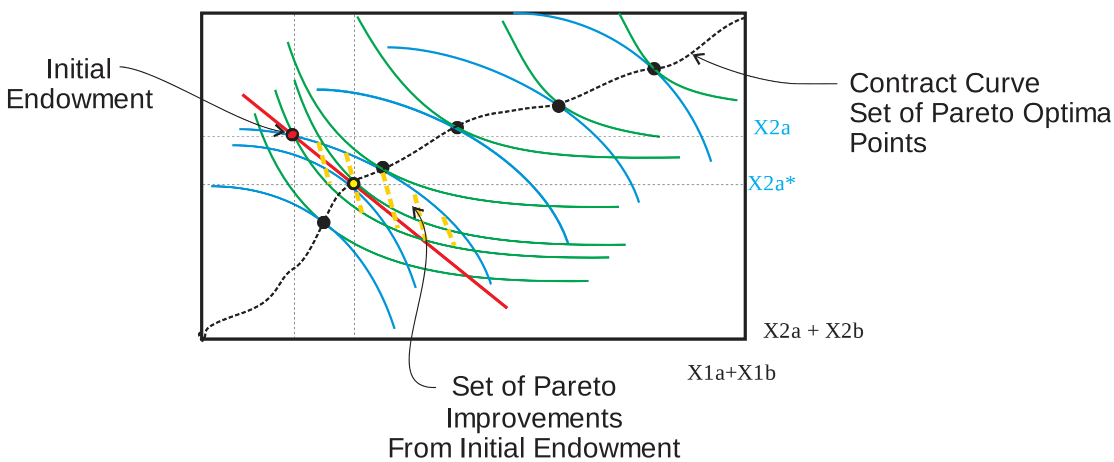


Given the initial endowment, the two indifference curves that pass through the initial endowment point create a lens-shaped area that represents the points in the Edgeworth Box that are Pareto

[1] Whether the latter equilibrium would be stable or not is debatable. It could be argued that the strategy of making and accepting offers is a weakly dominant strategy for each, because that strategy always yields a payoff that is at least as great as those associated with the no-trade outcomes. If that weakly dominant strategy is adopted by both traders, the (3,3) cell emerges as the equilibrium. However, it is nonetheless true that at the lower righthand equilibrium, neither can change their strategy and improve their condition.

Superior to the initial endowment. They are all the combinations of the two goods that could make at least one person better off without making the other worse off.

The points that are Pareto optimal are points where Al and Bob's indifference curves are tangent to one another. Any move from such a point will make at least one of them worse off. The set of all Pareto optimal points is sometimes referred to as the Pareto set or, in the case of an Edgeworth box, the contract curve. That curve is traced out and labeled in the box. Note that only part of the contract curve lies within the set of allocations that are Pareto superior to the initial endowment. Those are the subset of Pareto optimal outcomes that can be reached through a series of Pareto superior moves. Note that the equilibrium price produces just one of these possibilities. If more than one equilibrium price vector exists, each would produce a point on the segment of the contract curve within the set of Pareto superior points.

The Pareto optimal points in a given choice setting are not always unique, although there are cases in there is only a single Pareto optimal possibility, as in Table 18.1. The fact that there are Pareto optimal points that cannot be reached by a series of Pareto superior moves is also of interest. Pareto superior moves limit one to a subset of the feasible Pareto optimal outcomes—namely, those that make at least one person better off and no one worse off. These are the only points that could be reached via voluntary exchange.

In a trading situation, bargaining between the two players is likely to reach a Pareto optimal point, even if there are no posted prices, but which one will depend on the bargaining that takes place. If, for example, Al is the better bargainer, the result is likely to be a point on the contract curve within the set of Pareto superior points that is fairly close to Bob's indifference curve though the original endowment. If they are equally good bargainers, then points in the middle would be more likely to be the final result.

There are no unrealized gains to trade at the points that make up the contract curve. They are all Pareto optimal points. No changes from those combinations of goods are possible that would make one person better off without making the other worse off.

## III. Maximizing Social Net Benefits

Chapter 2 began with a model of demand based on the net-benefit maximizing model of individual decision-making. Individuals maximized a function from the family of $N = B(Q) - C(Q)$ functions, which implied that they chose a quantity, $Q^*$, where $dB/dQ = dC/dQ$. That is to say that individuals chose to engage in activities (including the purchase of goods and services that were regarded as consumption) up to the point where their marginal benefits equaled their marginal costs. This was shown both with geometry and with calculus—and only required the assumption that market participants knew their benefits and costs and that net benefit functions were strictly concave.

In cases where the marginal benefit functions are monotone decreasing in the quantity of the goods or services purchased, individual demand curves are inverse functions of their marginal benefit curves. Demand curves went through exactly the same points as the marginal benefit curves, but rather than being functions from quantities of the good (Q) into dollars per unit (marginal benefits), they were functions from prices (dollars per unit) into quantities purchased.

That individuals may maximize net benefits in a series of separate decisions rather than as a single decision using a lifetime utility function is suggested by the fact that most students find the net-benefit maximizing model to be more intuitive than the geometry and calculus of utility functions. Something like lifetime utility may roughly characterize all the tradeoffs among goods that one might purchase, but the particulars—as with which automobile or which cell phone to purchase—are often decided one good at a time, as assumed in the single-time-period net-benefit-maximizing models of demand and supply.

It turns out that the net benefit maximizing model can also serve as the basis for a quantifiable version of utilitarianism—for reasons briefly summarized by Pigou at the beginning of the chapter. The total benefits realized (in dollars) can be used to approximate the total utility gained by consumers. The total cost of production can be used to approximate the utility sacrificed by input providers and owners that produced the good of interest. Social net benefits can thus be used to approximate the net aggregate utility gained (or lost) from producing and consuming the goods of interest. Given these assumptions, utilitarian logic suggests that one should maximize social net benefits because aggregate utility is likely to be highly correlated with social net benefits.

The maximize social net-benefit norm is widely used in industrial organization and public economics as a method for evaluating the relative merits of different forms of markets and alternative public policies. When applied to evaluate specific policies, it is known as cost-benefit analysis. In most cases, this quantifiable normative approach employs a partial-equilibrium framework, focusing on one market and one public policy at a time.

The geometry of social net benefit maximization is based on results from Chapter 2. Recall that if $b_i(Q)$ is individual i's marginal benefit function for a good or service of interest, then $b^{-1}(P)$, the inverse of the marginal benefit function (if one exists) is individual i's demand function for the good of interest. Market demand ($Q^d$) is the sum of the individual demand functions or demand curves for all the people purchasing the good in the relevant community, state, nation, or world, thus $Q^d = \sum_{i=1}^{N} b_i^{-1}(P)$, where $b_i^{-1}$is the inverse of individual i's marginal benefit curve.

In principle, the social marginal benefit generated by successive units of a good can be determined from a market demand curve by gradually reducing the price of the good of interest from one that is higher than anyone is willing to pay to lower prices at which at least one consumer is willing to purchase one or more units. In this manner, the highest price at which successive units of a good can be sold can be determined. Those successive prices characterize the marginal benefit generated by successive units of the good of interest. It is determined by the highest demander for each of the successive units purchased.

Fortunately, this tedious method is not the only way to determine the social marginal benefit function. The inverse of a market demand function is the social marginal benefit curve for the individuals in the market of interest. In cases in which the inverse of a sum of functions is exactly the same as the sum of the inverses, the inverse of a market demand curve will exactly characterize the marginal benefits realized from successive units of a good, $inv(Q^d) = SMB(Q) = \sum_{i=1}^{N} b_i(Q)$. This is, for example, the case when the individual marginal benefit functions are linear.[2]

The same logic can be applied to the Ricardian supply curves derived in chapter 3 from individual firm cost functions. The social marginal cost function is the inverse of a market supply curve (in the Ricardian case). Recall that market supply is simply the sum of the individual firm supply functions in the market of interest, $Q^s = \sum_{j=1}^{M} c_j^{-1}(P)$ where $c_j^{-1}(P)$ is the inverse of the marginal cost function for firm *j*. In cases in which there are no production externalities, the social marginal cost curve associated with producing and distributing the good of interest is approximately the inverse of the market supply function, $inv(Q^s(P)) \approx SMC(Q) \approx \sum_{j=1}^{M} c_j(Q)$.

[2] Several proofs of this proposition can be found with a brief web search. Using the net-benefit maximizing model directly avoids the differences between demand curves derived from utility functions and their marginal benefit functions. In those cases, one should use the Hicksian compensated demand curves instead of their actual demand functions.

These two approximations allow us to use estimates of demand and supply functions to approximate the social marginal benefits and social marginal costs for the good or service of interest. Most estimates assume linearity and so inverses always exist. Geometrically, social net benefits are simply the areas between the demand (social marginal benefit) and supply curves (industry marginal cost) if we ignore fixed costs or assume that they are zero in the long run, $SNB(Q) = \int_0^Q SMB(Q) - SMC(Q)\, dQ$ .[3]

**Illustration: The Social Net Benefits and Static Deadweight Losses from Monopoly Power**

The static normative case in favor of competitive markets over monopolistic markets can be developed using this approach. Figures 18.2a and 18.2b below illustrate the geometry of social net benefit calculations for two types of markets. Figure 18.2a characterizes a perfectly competitive market with price-taking firms. Figure 18.2b characterizes a price-making (monopoly's) firm's output decision.

Figure 18.2: The Static DWL of Monopoly

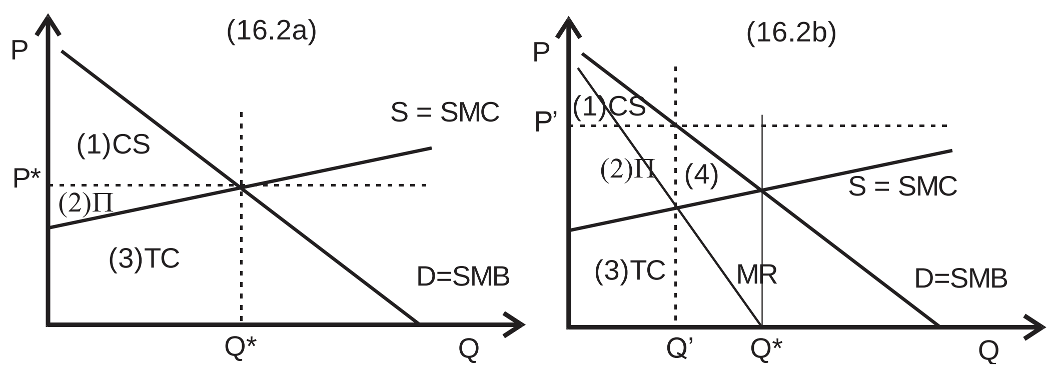


We'll first analyze figure 18.2a. The social benefit is the area under the demand curve from zero to Q* (Areas 1+2+3). The social cost is the area under the supply curve from zero to Q* (area 3). Subtracting social costs from social benefits yields social net benefits, ([1+2+3]-[3]=[1+2]). Areas 2

[3] The analysis is similar for pure public goods, but in this case, the SMB curve is the "vertical" sum of the individual marginal benefit functions rather than a horizontal sum (in the Q dimension). However, in this case there are no market demands to estimate, so approximating such sums is more difficult and normally requires many additional assumptions. Samuelson (xxxx) characterizes the notion of a pure public good and provides a more direct utilitarian analysis of the optimal provision of such goods.

and 3, thus represent social net benefits. In this case and others similar to it, social net benefits are the sum of consumer surplus (CS) and profit (Π).

Elementary calculus implies that social net benefits are maximized where SMB(Q) = SMC(Q), which occurs at exactly Q*, the quantity produced and sold in a competitive market (in the absence of externalities). This output level is also the unique Pareto optimal output of the good being modeled. No change in that output (e.g., away from Q*) can be made without reducing someone's net benefits and, thus, making someone worse off. This is the partial equilibrium version of what is often referred to as the "First Welfare Theorem." Competitive markets are Pareto efficient. At a market-wide (general) equilibrium, every market has this property, and thus, this result extends to the market as a whole. (The general case is characterized and demonstrated to be Pareto efficient in the appendix of this chapter.)

This is true of both the Marshallian and Ricardian variations of competitive markets, although the supply curve would normally be horizontal in the Marshallian version, and long-run profits would equal zero, because every firm realizes the "ordinary" risk-adjusted rate of return on its investment. In the Ricardian version depicted, some firms earn higher rates of return than the ordinary one, possibly because of their location or good fortune with their personnel.

At the competitive equilibrium price equals the marginal cost of production and the marginal benefit of consumption. In the absence of externalities (a topic taken up later in this chapter), the competitive price, thus, equals the social marginal cost of production and the social marginal benefit of consuming or otherwise benefiting from the goods and services produced. Insofar as market prices tend to converge to market clearing ones, this implies that competitive markets are self-regulating in a manner that is consistent with the social-net-benefit-maximizing norm (again, in the absence of externalities). Note that this implies that the social net benefit maximizing level of every good or service varies with input prices, the prices of substitutes, and consumer willingness to pay for the good or service of interest. It is not a single unique vector of outputs, but varies with market conditions.

Next, examine figure 18.2b. That diagram provides a geometric representation of the choice made by a "price-making" firm facing the same demand curve and having the same marginal cost curve as the industry depicted in Figure 18.2a. Recall that a price-making firm faces a downward-sloping demand curve and takes into account how its output affects prices and, thereby, profits. Such firms will produce where their marginal revenue (calculated from the demand curve) equals their marginal cost, which is denoted as Q' in Figure 18.2b. This calculation implies a significantly higher price than that in the competitive market and higher profits for the firm (or firms, if the monopoly is the result of a cartel rather than a single firm's unique product).

The social net benefits again consist of profits and consumer surplus (areas 1 and 2). However, the total net social benefits are smaller than they were in the competitive case. Area 4 is no longer included in the area that represents social net benefits. The missing area (area 4) is sometimes called the deadweight loss of monopoly. It characterizes the extent to which cartelization or monopolization of a competitive market tends to reduce social net benefits.

Because all potential gains from trade have not been realized, in principle, there is a Pareto superior move from Q' to Q*. For example, a Pareto superior move could be realized by selling the first Q' units at price P' and additional units at price P*. Notice that that combination makes both consumers and firms in this market better off. (Profits and consumer surplus are both increased by the two-price solution.)

The extent of deadweight loss, if any, depends on how the monopoly or cartel was established. If an innovative firm creates a new product that is temporarily protected by a patent or its own unique production methods, then the social net benefits of figure 18.2b are at least partly new. Social net benefits are generally increased by successful innovation. A successful new product must generate more consumer surplus than previous purchases, or it would not have been purchased. However, their previous purchases doubtless produced consumer surplus as well, the reduction of which should be subtracted from the consumer surplus in Figure 18.2b to determine the net benefit of the particular innovation of interest for consumers.

If Schumpeter is correct in his assessment that large firms with significant price-setting ability are generally more innovative than the smaller firms of competitive markets. In that case, there may be dynamic gains from monopoly that offset some or all of the static losses of monopoly in terms of social net benefits. Whether this is so or not is an empirical question, and it is likely to vary by market and innovation. (Note that social net benefits are, in principle, estimatable—an objective phenomenon rather than an entirely subjective phenomenon as aggregate utility is.

## IV. Social Net Benefits and Externality Problems

Partial equilibrium analysis of markets focuses on one or two interconnected markets. It makes no effort to characterize entire market networks, nor does it take account of effects on people other than consumers and producers of the product of interest. However, there are many cases in which market activities affect people that are neither producers nor consumers of the products or services being modeled.

Many folks decorate their houses for holidays, maintain fine gardens in front of their houses, and take steps to ensure that their sidewalks are clean, neat, and flat. Such "consumption" activities produce benefits for others walking past their houses that are not part of a market transaction. Such fastidious homeowners produce "positive externalities" (spillover benefits) that increase the welfare of all or most of those walking past their houses. On the other hand, other neighbors may have parties outdoors that create noise, smoke, and parking problems that reduce the welfare of their neighbors. Such "negative externalities" (spillover costs) may not be ignored by those throwing the parties. They would ignore them if they are pragmatists, as assumed in most economic models, but would take them into account if they have internalized utilitarian norms or were altruists.

Similar externalities are often associated with the production activities of firms. In the late nineteenth century, as large-scale production in industrial facilities became commonplace, firms would typically use the air and water to dispose of their waste products. As that usage increased, their waste products produced noticeable effects. Some were simply unpleasant sounds or odors. Others produced a variety of health problems. Such spillover costs and benefits are not part of market transactions and tend to be ignored by those generating them, unless they are induced by liability, nuisance, or other laws to do so. In the absence of such laws, externality-generating activities are rarely undertaken at levels that maximize social net benefits.

Figure 18.3 shows why this tends to be true. It illustrates an externality problem that might emerge from the production activities of firms, such as air and water pollution. As is usually the case in public-economic and environmental-economic textbooks, we'll assume that the externality generators are "pragmatists" and that no laws inhibit their externality-generating activities. As a consequence, the firm owners ignore the effects of their waste disposal methods on others. The previous chapter and the chapter on law and economics imply that this would not always be the case, but the normative analysis of externality problems nearly always assumes this.

Figure 18.3: Externality Problems

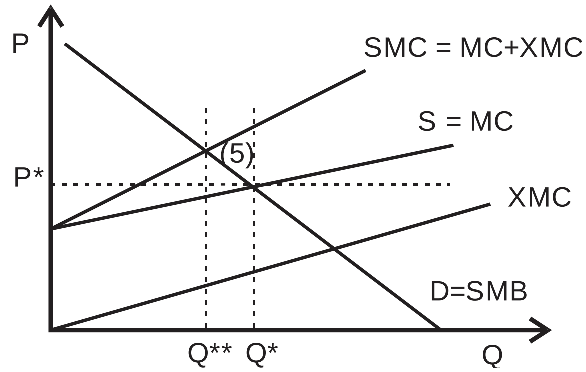


The externality illustrated is assumed to be an external spillover cost that varies directly with the extent of the product sold in this market. The spillover marginal cost is labelled XMC, for eXternal Marginal Cost. The output produced by a competitive market is, as usual, where supply equals demand, which is labeled Q*. The output that maximizes social net benefit occurs where social marginal benefits equal social marginal cost. As drawn, the demand curve is assumed to be a good proxy for the social marginal benefits generated by this market. The supply curve is assumed to be a good proxy for the industry's marginal cost of production. Were these the only costs and benefits, the market outcome would have maximized social net benefits and been Pareto efficient, for reasons discussed earlier in the chapter.

However, because of the external costs generated, the social marginal cost curve differs from the direct marginal cost of production. Social marginal costs (SMC) now include both the direct marginal cost of production and the spillover marginal costs associated with the "free disposal" of waste products using the air and water systems. The output where SMC equals SMB is labeled Q**. Note that the market produces a greater quantity than that which maximizes social net benefits. The area labeled with (5) characterizes the extent to which social net benefits could be increased if output and/or purchases of the goods of interest were reduced to Q**.

Economists often conclude that solutions require government action in such cases, but Chapter 14 suggests that such actions may not always be forthcoming and would not necessarily produce a regulation or tax that would induce output Q** or otherwise increase social net benefits net of the cost of the policies adopted. Moreover, the cost of implementing the policies that would ameliorate the problem should also be considered when making such recommendations. For "small" externalities, the cost of implementing corrective policies is often greater than any net benefits recovered—indeed, this property could be used as a definition of “small” externalities. In those cases, inaction maximizes social net benefits.

Nonetheless, there are cases where governmental or other actions can potentially increase social net benefits, as argued by Pigou. For example, a disposal fee or fine equal to the external cost generated at Q** could be imposed on producers.

## A Coasian Solution to Externality Problems

A possible non-governmental solution was suggested by Coase (1962). He argued that if transaction costs are low, the individuals affected by an externality could organize and contract with firms and consumers in the relevant market to reduce the sales and production of the goods that generate the externality. There are unrealized gains from exchange that exist between persons harmed and those who benefit from the production of such goods.

Note that “at the margin,” those affected by the externality would be willing to pay up to the vertical distance from the horizontal axis to the XMC curve to have producers and consumers reduce market sales, whereas consumers and producers require little or nothing (initially) to reduce production and consumption of the good. As the bargaining process continues, those harmed are willing to pay less and less for reductions (the height of the XMC curve falls as one moves from Q* to the left), while the amount required as compensation increases (the distance between the demand curve and supply curve increases). The potential gains from such Coasian contracts run out exactly at Q**, where the distance from the horizontal axis to the XMC curve exactly equals the distance between the supply and demand curves.

Such contractual solutions clearly become more difficult as the number of persons affected increases, although it bears noting that international environmental treaties closely resemble the agreements that Coase imagines being consummated. Environmental treaties are voluntary agreements among the parties to the treaty—e.g., they are contracts among sovereign nations.

One of the peculiarities of the single market approach is that the social net maximizing output level tends to be unique. Thus, there is only a single Pareto optimal output. Coase made use of this property when he developed what is now known as the Coase theorem—that the bargaining solution to externality problems tends to converge to the same output level regardless of whether the market participants have the “right” to engage in voluntary exchange, or if those affected by the externality have veto power over whether they can trade or not. This result holds as long as transaction costs are low (approaching zero) and the bargains negotiated do not themselves change and any of the curves drawn. The latter, in Coase's terms, implies that there are no significant income or liquidity effects from the various side payments used to consummate the contract negotiated.

### The Mathematics Behind the Externality Diagrams

The mathematics behind the externality diagram are straightforward and closely mirrored by Figure 18.3. Let B(Q) be the social benefit function, C(Q) be the industry cost function, and X(Q) be the external cost function. Net benefits (N) are $\mathrm{N} = \mathrm{B(Q)} - \mathrm{C(Q)} - \mathrm{X(Q)}$. Assuming that function N is strictly concave implies that N is maximized at the quantity that sets social marginal benefits equal to social marginal costs,

$$\mathrm{N}_Q = \mathrm{B}_Q - \mathrm{C}_Q - \mathrm{X}_Q = 0 \text{ at } Q^{**} \qquad (18.2)$$

Q** is the social net benefit maximizing output of the negative externality-generating activity. Note that this quantity is not usually zero because benefits as well as costs are associated with it.

To characterize this as a function of parameters of the choice setting requires a bit more descriptive characterization of the three component functions, as with $B = \mathrm{B}(\mathrm{Q}, \mathrm{P}^0, \mathrm{Y})$, $C = \mathrm{C}(\mathrm{Q}, \mathrm{w}, \mathrm{r}, \mathrm{T})$, and $\mathrm{X} = \mathrm{X}(\mathrm{Q}, \mathrm{D}, \mathrm{A})$, where the parameters of the demand and supply functions are the usual ones, and those of the external cost function include population density (D) and various probabilities of rain, wind directions, and air speed (A). Given that richer characterization of the B, C, and X functions, the implicit function theorem implies that $Q^{**} = q(\mathrm{P}^0, \mathrm{Y}, \mathrm{w}, \mathrm{r}, \mathrm{T}, \mathrm{D}, \mathrm{A})$.

The implicit function theorem can, in turn, be used in the usual way to characterize how the social net benefit maximizing output in this market varies as those parameters change.

### Ameliorating Externalities with Economic Regulation(s)

If the social net-benefit maximizing outcome, Q**, does not arise from the behavior of the individuals involved, then government regulations of various kinds can potentially increase social net benefits.

Assume that regulation R can reduce output or the extent of the externality in this market and is undertaken to maximize social net benefits. Suppose there are costs associated with implementing the regulation(s) of interest that can be characterized as, E=e(R). In the above case, the social net benefit maximizing stringency of regulation R can be determined by maximizing:

$$N = \mathrm{B}(\mathrm{Q}(\mathrm{R}), \mathrm{P}^0, \mathrm{Y}) - \mathrm{C}(\mathrm{Q}(\mathrm{R}), \mathrm{w}, \mathrm{r}, \mathrm{T}) - \mathrm{X}(\mathrm{Q}(\mathrm{R}), \mathrm{D}, \mathrm{A}) - \mathrm{e}(\mathrm{R}) \tag{18.3}$$

The ideal regulation, in this case, satisfies the following first-order condition.

$$N_R = \mathrm{B}_Q \mathrm{Q}_R - \mathrm{C}_Q \mathrm{Q}_R - \mathrm{X}_Q \mathrm{Q}_R - \mathrm{E}_R = 0 \text{ at } \mathrm{R}^* \tag{18.4a}$$

or

$$N_R = [\mathrm{B}_Q - \mathrm{C}_Q - \mathrm{X}_Q]\mathrm{Q}_R - \mathrm{E}_R = 0 \text{ at } \mathrm{R}^* \tag{18.4b}$$

Note that the term inside the brackets represents the first-order condition for maximizing social net benefits with costless policies. The implementation and administrative costs imply that the regulation should aim for a somewhat smaller reduction in output than the diagram suggests, because the cost of implementation must be taken into account.[4]

The implicit function theorem can, in turn, be used in the usual way to characterize how the social net benefit-maximizing regulations (if any) vary as market and regulatory parameters change.

## V. Net Benefits and Time in Cost-Benefit Analysis

There are both theoretical and applied uses for the social net-benefit maximization norm because this norm, if accepted, allows one to determine whether a change in public policy improves or worsens the existing situation using various combinations of economic theories and estimates of

---

[4] Environmental regulations often mandate that producers use specific emission reducing devices (E) and then monitor both the installation and the performance of the devices installed. Calculating the ideal stringency of such devices, E**, would be undertaken in a similar, but not identical, way to that used to characterize equation 16.2. In this case, X would be a function of the rigor of the emissions inhibiting device, E. Annual social net benefits would be $\mathrm{N} = \mathrm{B}(\mathrm{E}) - \mathrm{C}(\mathrm{E}) - \mathrm{X}(\mathrm{E}) - \mathrm{e}(\mathrm{E})$. If the emission inhibiting device somewhat reduces the quality of the product sold (to the average consumer) as some emissions inhibiting devices reduce auto fuel economy or power, then $B_E<0$. The marginal cost of using such devices is normally greater than zero as are the administrative costs of the regulation.

Social net benefits would be maximized when $X_E=-B_E + C_E + E_E$. The optimal rigor of the emission reducing device, E**, sets the reduced marginal loss from external costs (the marginal benefits from the emissions device) equal to the marginal cost of the emission reducing device in terms of increase production costs, $C_R$, reduced benefits from the good of interests, $B_R$, and the marginal cost of administering the regulation mandating the use of the emissions device $g_E$.

From the perspective of cost benefit analysis, such devices should be required (e.g., E**>0) only if they increase social net benefits.

economic relationships. These usually involve probabilistic assessments of net benefits over a significant timespan, and so require the use of expected value and present values.

To illustrate the process of cost-benefit analysis, suppose that robust and unbiased estimates of external marginal costs, supply, and demand have been undertaken for the functions depicted in Figure 18.3. Suppose also that the regulation to be imposed on the industry that costs F to establish and E=e(R*) per year to administer. The costs of regulation imply that the ideal output target is a bit less than the Q** of Figure 6.3, as shown by equation 18.4b, which implies that Q**=Q(R*).

Determining whether regulation R* should be adopted or not requires estimates of the various costs and benefits of that regulation through time. Consumer benefits and costs can be estimated using estimates of market demand (D) and supply (S), as previously developed.

$$\mathrm{B(Q)} \approx \int_o^Q D^{-1}(Q)dQ, and\ C(Q) \approx \int_0^Q [S^{-1}(Q) + X(Q)]dQ$$

where $D^{-1}(Q)$ and $S^{-1}(Q)$ are the inverse functions of the demand function and supply function. The annual social net benefits realized by moving from Q* to Q** are:

$$N \approx \int_o^{Q^{**}} [D^{-1} - (S^{-1} + X)]dQ - \int_o^{Q^{*}} [D^{-1} - (S^{-1} + X)]dQ - e(R^*) \tag{18.5a}$$

Which can be written as:

$$N \approx -\int_{Q^{**}}^{Q^{*}} [D^{-1} - (S^{-1} + X)]dQ - \mathrm{e(R^*)} \tag{18.5b}$$

The first term corresponds roughly to the triangle labeled (5) in figure 18.3. It represents the maximal annual net benefit achieved by the regulation if it is costless, which can be denoted as $N^R$. The second term, the one outside the integral, characterizes the annual cost of administering the regulation.

Even an ideal policy increases social net benefits only if the present value of the costs of the regulatory apparatus—including the fixed costs of establishing the regulatory regime for implementing the policy is less than the present value of the gain achieved in moving from Q* to Q**. Recall that Chapter 6 demonstrated that the present value of an essentially infinite series of benefits or costs of constant value (A) is PV(A) = A/r, where r is the annual interest rate, and A is the annual amount of the cost or benefit of interest. This implies that an ideal regulation increases the present value of social net benefits only if

$$\frac{N^R}{r} > F + e(R^*)/r \tag{18.6}$$

Whether regulation $R^*$ increases social net benefits or not depends on the specific values for administrative costs, e(R), administrative fixed costs, F, and the annual social net benefits, $N^R$, that can potentially be realized via regulation. Note that if the N of equation 18.5b is greater than zero, it is the fixed costs of regulation that will determine whether R* should be adopted or not.

If the planning horizon of the regulators is finite and relatively short, or the various costs and benefits are expected to vary through time (and perhaps be stochastic rather than deterministic constants), then the present value formula and expected value formula would have to be used to characterize some of the benefits and costs. However, essentially the same logic would apply.

## VI. Social Welfare Functions as a Generalization of Bentham's Aggregate Utility Norm

> A thing is said to promote the interest, or to be for the interest, of an individual, when it tends to add to the sum total of his pleasures: or, what comes to the same thing, to diminish the sum total of his pains. …
> An action then may be said to be conformable to the principle of utility, or, for shortness sake, to utility (meaning with respect to the community at large), when the tendency it has to augment the happiness of the community is greater than any it has to diminish it. (Jeremy Bentham,(1780). *An Introduction to the Principles of Morals and Legislation,* Kindle Edition.)

In the nineteenth century, there were many utilitarian contributors to economics theory (including Bentham) and many applications of utilitarian analysis to justify particular public policies. Pigou's quantitative approach to welfare economics is essentially a fusion of neoclassical economic arguments and utilitarianism ideas. His suggestion that money values can be used to approximate degrees of satisfaction or happiness is widely accepted among economists (and many others). Thus, social net benefit calculations and modes of reasoning are widely used by applied welfare economists.

The advantage of money measures is that they are "objective" and "operational." Net benefits can be estimated using economic models and statistical methods. And, those estimates can be used to determine whether particular public policies actually increase "welfare" on average or not within the polity of interest. Given this norm, it is possible determine whether one policy is better than another or one country better organized than another.

Nonetheless, even Pigou recognized that the monetization of utility was an approximation—albeit a very useful one—of what one really wanted to maximize, namely aggregate utility—the sum of happiness in a community. Thus, a good deal of theoretical work in welfare economics use utility functions and various social welfare functions to characterize aggregate utility.

The simplest measure of aggregate utility is one mentioned several times in Bentham's work. It is simply the sum of the utility levels of all the individuals living in the community of interest, $W = \sum_i^N u_i(Q_1, Q_2, Q_3, Q_4, \ldots)$. Many other functional forms were subsequently suggested, such as multiplicative functions. In the mid twentieth century, to avoid controversies over functional form, general welfare functions such as: $W = w(U_1, U_2, U_3, U_4, \ldots U_N)$ largely replaced such concrete functional forms for aggregate utility or social welfare. The only restriction on such functions is that $W_{Ui} > 0$ for all i, which is to say that everyone in the community counts. Thus, an increase in any single person's utility, holding the others constant, adds to social welfare (W). Some scholars also require all the derivatives of W with respect to $U_i$ to be the same, but many theorists—perhaps most—do not.

The proper aim for public policy makers is maximizing social welfare, according to proponents of the contemporary form of the utilitarian approach to policy analysis. Although few such scholars believe that real governments actively attempt to maximize a social welfare function—those employing this methodology for identifying ideal policies clearly wish that they did so.

That utilitarian approach has been used in most applied microeconomic fields to analyze the merits of alternative public policies and areas where markets and political systems may fail to realize ideal outcomes or to adopt ideal policies. There are problems with the maximize social welfare norm that will be discussed towards the end of this section, but first we'll review one of the most impactful of the applications of that approach.

## Samuelson's Theory of the Optimal Provision of a Pure Public Good

Before 1954, neoclassical economics had regarded all goods and services to be "private goods" that were consumed by one person at a time. There may be externalities, but in a sense "goods were goods" there were all essentially the same sort of divisible objects. If two persons equally share a candy bar, their share of the candy bar falls from 1 to one half. Paul Samuelson (1954) argued not all goods have that property and, perhaps surprisingly, that some goods are perfectly shareable. The consumption of such goods (termed pure public goods) does not decline at all as more people consume them. Examples include broadcast radio shows, broadcast television shows, national defense, and environmental quality. A few years later, Buchanan (1965) suggested that in between pure public goods and pure private goods were other types of goods that were more sharable than pure private goods, although not as sharable as pure public goods. Buchanan called these good "club goods." Club goods were "congestible." The services that they provide users decrease with the number of persons using them, but not proportionately to their numbers. Examples include swimming pools, public parks, playgrounds, bicycle paths, highways, lectures, and so forth. It turns out that there is a spectrum of types of goods with varying different degrees of "shareability."

However, for the purposes of this chapter, our main interest is Samuelson's characterization of the ideal service level of a pure public good. His argument places a generalized welfare function at the center of his analysis.

The first step in reviewing Samuelson's model and results is to develop some notation. Let G be the level of a pure public good, let Xi be the level of a pure private good received by individual *i*, let Ui = u(G, Xi) be the utility of individual *i* associated with a particular combination of the public good G and private good Xi received by *i*. Second, we need some macro-choice notation. Let W be a social welfare function and let T(G,X) = 0 be the technological frontier of combinations of the public good and private goods, with the sum of the private goods being equal to their total production, $X = \sum_{i-1}^{N} X_i$. The constraint function is in its "zero form," T(G,X) = 0, as required for the Lagrange method. Assume that there are N persons in the society of interest.

The problem of maximizing social welfare can be undertaken with the approach developed by Lagrange, e.g., first create a Lagrange function, then differentiate with respect to the control variables and set the result equal to zero:

$$\mathrm{L} = w(U_1, U_2, U_3, U_4, \ldots U_N) - \lambda(\mathrm{T}(\mathrm{G}, \mathrm{X})) \qquad (18.7)$$

Differentiating the Lagrangian with respect to G, $X_1$, $X_2$, $X_3$ .... $X_N$, and λ yields the first-order conditions for the combination of services and private incomes that maximizes social welfare:

$$\sum_{i=1}^{N} W_{U_i} U_{iG} - \lambda \mathrm{T}_G = 0 \qquad (18.8a)$$

$$W_{U_i} U_{iX} - \lambda \mathrm{T}_X = 0 \qquad (18.8b)$$

for all i = 1 ... N (Equation 18.6b thus represents N first order equations)

$$\mathrm{T(G, X)} = 0 \tag{18.9c}$$

After obtaining the Lagrangian first order conditions, the next step is to manipulate the first order conditions into a form that is both economically interesting and useful. Samuelson uses a rather clever series of steps to do so.

First, as usual, shift the lambda terms to the right by adding their positive equivalents to each side. Then divide the first such equation and one of the members of the second group of first order conditions to eliminate the lambda.

$$[\sum_{i=1}^{N} W_{U_i} U_{iG}]/W_{U_j} U_{jX} = \frac{\mathrm{T}_G}{\mathrm{T}_X} \tag{18.10}$$

(I have used the j-th of the private good first order conditions to avoid confusion with "i" the counter for the summation in the public goods terms in the numerator)

Since the denominator does not change with "i" it can be brought inside the brackets--because it is essentially a constant as far as this fraction is concerned.

$$\sum_{i=1}^{N} [W_{U_i} U_{iG}/W_{U_j} U_{jX}] = \frac{\mathrm{T}_G}{\mathrm{T}_{jX}}$$

Now note that the first order conditions for the private goods imply that the marginal social welfare generated by the private goods allocated to each individual are equal to one another at the margin:

$$W_{U_j} U_{jX} = W_{U_i} U_{iX} = \lambda \mathrm{T}_X \tag{18.11}$$

This condition holds for all "i" and "j" (for every person's private good). This equivalence means that you can substitute each one of the $W_{U_j} U_{jX}$ with its equivalent $W_{U_i} U_{iX}$ term and rewrite the equation under part b as:

$$\sum_{i=1}^{N} [W_{U_i} U_{iG}/W_{U_i} U_{iX}] = \frac{\mathrm{T}_G}{\mathrm{T}_{jX}}$$

Note that the denominator now has a counter that moves with the sum. (This is the clever part of the derivation. The rest is pretty straight forward.)

This allows us to simplify quite a bit—the marginal weights $W_{Ui}$ in the numerator and denominator are the same and can be divided out of each. This implies that at G*:

$$\sum_{i=1}^{N} [U_{iG}/U_{iX}] = \frac{\mathrm{T}_G}{\mathrm{T}_X} \quad \text{at G*} \tag{18.12}$$

The ideal level of a pure public good sets the sum of the marginal rates of substitution between the private and public goods equal to the technological rate of transformation between private and public goods. (In benefit/cost terms, G* occurs at a point where the sum of the marginal benefits equals the marginal cost of the public good in terms of reductions in the private good**.)** **Surprisingly, the optimal level of a pure public good is completely independent of the social welfare function used!**

However, the distribution of the private good will vary with the specific for of the social welfare function assumed. So, the complete solution cannot be fully characterized without knowing more about the social welfare function assumed and the individual utility functions.

Note also that this a pure "benevolent central planner" model in which the imaginary planner can make both the supply of pure public good decisions and the has complete control over the allocation of private income or consumption decision without input from the individuals in the society. Moreover, the central planner him or herself is costless—there are not administrative costs.

### Weaknesses of the Maximizing Social Welfare Approach

All normative approaches can be thought of as a particular scholar's or group of scholar's methodology for thinking about ideal policies, who may hope that policy makers will take their ideas into account when actually adopting public policies.

However, there are several problems with the maximizing social welfare approach to normative analysis as it has generally been applied. Perhaps the most important is the assumption that a theorist can know individual utility functions—not simply some of their general characteristics. Actual numerical estimates of every person's utility functions have to be worked out and, similarly, the economic environment in which policies are to be implemented has to be fully characterized. For example, in Samuelson's characterization of the ideal level of a pure public good, we do not have to argue about the form of social welfare functions, but we do need to know the values and functional forms of $U_{iG}/U_{iX}$ for every individual and $\frac{T_G}{T_X}$ for the economy as a whole in order to determine the ideal quantity of the pure public good of interest, G*. Notice that ordinal utility functions are not sufficient for this. Cardinal ones (e.g., utility functions that can be used for arithmetic, algebra, and calculus) are required.

This approach also implies that individuals truly have utility functions, rather than simply act as if they had one. This stronger assumption is not required to undertake a positive analysis based on the existence of such functions—because, for the most part, it uses such models to characterize behavior rather than utility per se. As long as individuals generally behave as if they had utility functions, utility-based models will predict responses to changes in circumstances that will be consistent with their behavioral responses—albeit not perfectly so.

A third problem is the lack of political context in most analyses, which implicitly or explicitly assume an all-powerful, benevolent dictator or central planner, rather than a democracy in which normative ideas must be shared by at least a majority of the pivotal voters to systematically influence public policy. A majority of voters, or at least pivotal voters, would have to be utilitarians for utilitarian policy analysis to affect election-driven public policies. Moreover, by neglecting more political settings, their analyses neglect possible indirect impacts that public policies may have on aggregate utility because of their political consequences.

Lastly, when pushed to its limits, utilitarian analysis can reach conclusions that conflict with other normative theories, as with rights-based frameworks. For example, if harvesting the organs of a single person can prolong the lives of two dozen people, utilitarian analysis would tend to conclude that such policies—randomly killing healthy persons for their body parts—is a completely moral activity. Individual rights never trump interests in maximizing aggregate utility from a utilitarian standpoint.[5]

There is arguably a domain in which utilitarianism generates morally plausible results, but that domain does not include all morally relevant questions—despite the arguments made by utilitarians.

## VII. Contractarian Norms—a Possible Alternative to Utilitarianism

The contractarian line of reasoning preceded Bentham's works on utilitarianism by more than a century. Contractarian normative theory has its roots in Hobbes' (1651) and Locke's (1690) theories of legitimate governments, both of whom refer to social contracts and governments created by consensus as commonwealths. Governments grounded in the consent of those governed are

[5] It should be noted that "rule-utilitarians" exist and often develop arguments that provide utilitarian defenses of some rights. Nonetheless, such lines of argument are rarely used in the utilitarian strand of welfare economics. See, for example, Harsanyi (1982).

formed to advance the shared interests of their citizens. In Hobbes' case, a commonwealth is created to escape from "the war of every man against every other."

> The final Cause, End, or Design of men, (who naturally love Liberty, and Dominion over others,) in the introduction of that restraint upon themselves, (in which we see them live in Commonwealths,) is the foresight of their own preservation, and of a more contented life thereby; that is to say, of getting themselves out from that miserable condition of War, which is necessarily consequent (as hath been Shown) to the Natural Passions of men… (Hobbes, 1651, p. 93)

> [W]hen **men agree** amongst themselves, to submit to some Man, or Assembly of men, voluntarily, on confidence to be protected by him against all others. This later, may be called a Political Commonwealth, or Commonwealth by Institution… (Hobbes, 1651, p. 96).

Locke's approach is similar but, rather than assuming that a war of every man against every other exists, he assumes that individuals create a commonwealth to reinforce natural laws, which in the absence of governmental support, provide a less perfect constraint on narrowly self-interested behavior. In both cases, the governments of commonwealths are grounded in the consent of the governed, rather than conquest by an organization able to impose laws on others.

The use of consent as the primary process for determining ethical principles and ranking political and legal procedures and states of the world is the main point of agreement among contractarians. Many of their conclusions differ from one another—as true of utilitarians—because of differences in assumptions, but voluntary agreement lies at the heart of all of their normative analyses. With respect to good governance, contractarians argue that in a good society, every person's interests should be advanced through a government's policies, not just an elite or favored group. Contractarian reasoning was employed, for example, by such well-known scholars as Montesquieu (1748), Rousseau (1762), and Blackstone (1765-70).

Contractarian analysis was supplanted by utilitarianism during the nineteenth and early twentieth centuries. It was re-established in the mid-twentieth century by two eminent scholars: John Rawls (1921-2002) and James Buchanan (1919-2013), both of whom were dissatisfied with various aspects of utilitarianism. Rawls' approach was similar in spirit to Locke's, and Buchanan's resembled Hobbes' approach—although each provided new foundations for and methods of contractarian analysis. Rawls' approach is widely regarded as one of the major innovations in philosophy in the twentieth century. Buchanan won the Nobel Prize in economics in 1986 for his contributions to constitutional theory, most of which were grounded in his contractarian approach.

Rawls (1971) argues that what might be called moral intuitions tend to converge on two principles of justice and that these are compatible with many, but not all, institutional designs:

> First: each person is to have an equal right to the most extensive basic liberty compatible with a similar liberty for others. Second: social and economic inequalities are to be arranged so that they are both (a) reasonably expected to be to everyone's advantage, and (b) attached to positions and offices open to all. (Rawls, 1971, pp. 60-61).

His analytical device for finding that consensus—the veil of ignorance—resembles Adam Smith's notion of the impartial spectator in that it requires individuals to imagine a choice setting in which one is ignorant of one's place in society and to choose among societies given that he or she might conceptually wind up in any of the positions in that society. Such a perspective tends to make each person an impartial observer of an entire social order.

Buchanan (1975) develops a Hobbesian perspective on the emergence of governance. He suggests that bargaining possibilities at the Nash equilibrium that emerges in a state of anarchy allow mutually beneficial rules to be gradually adopted. In that book and several other papers, he argues that consensus is the ultimate foundation both for legitimate governance and legitimate changes in public policy.

> In the earlier book [Buchanan and Tullock, 1962], we argued that the criterion of acceptability or efficiency lay in agreement, in unanimity. Further, we argued that insofar as participants remain uncertain as to their own specific roles in subsequent operation under the rules chosen, they would tend to reach agreement on reasonably "fair" and "efficient" working rules. We did not postulate initial equality among individuals in property rights or in capacities, but our presumption of uncertainty served to generate a plausible basis for agreement on rules for collective action. (Buchanan, 1975, p. 81).

Notice that both Buchanan and Rawls argue that uncertainty tends to increase prospects for agreement by inducing individuals to take a broader view of the benefits and costs of governing institutions and grounding laws.

However, the thought processes that they have in mind differ significantly. In Rawl's case, the veil of ignorance is a thought process that individuals should adopt when evaluating the relative merits of alternative legal, political, and economic systems. This requires individuals to accept the idea that "the veil" is the only way to discern just or fair institutional arrangements. In contrast, Buchanan—a student of Frank Knight—suggests that uncertainty is inherent in all long-term decisions about political and legal institutions, because one cannot perfectly imagine how those institutions will affect the future or, therefore, an individual's own long-term interests. One can imagine a wide range of possible outcomes for oneself under every political and legal institution, which requires every individual to take a wide range of possibilities into account when assessing their relative merits.

For example, one could wind up rich or poor in the future under the rules agreed to, and that tends to induce each individual to favor rules that are advantageous both to rich and poor persons. Taking account of a range of such possibilities is for Buchanan, in contrast to Rawls, inherent in making long-term commitments about "the rules of the game." It requires no agreement, implicit or otherwise, about the relative merits of "the veil" as a method for discerning just or fair social systems. The veil of uncertainty is associated with every long-term decision.

Rawl's veil of ignorance and Buchanan's veil of uncertainty are sufficiently clear ideas that one can characterize both their analytical approaches and some of their conclusions using conventional rational choice models—as was also possible for utilitarianism after the mathematics of utility functions were worked out.

## Behind the Rawlsian Veil

Rawl's normative analysis rests on two principles: an equal liberty principle and a method for characterizing distributive justice, both of which are arguably implications of his veil of ignorance approach to analyzing alternative social systems. Rawl's notion of equal liberty constrains the types of decisions that can be just or fair in every social system. Rights associated with equal liberty constrain the institutions of every just society, as in Locke. However, in contrast with Locke, fundamental rights are developed from behind the veil of ignorance, rather than having divine origins. His just distribution principle requires the income or welfare of the least well-off person in the society of interest to be maximized. In principle, one may argue that neither is necessarily an implication of the veil of ignorance approach, but for the purposes of this section of the chapter, we'll accept Rawls' conclusion about this.

To simplify a bit, we'll assume that both the equal liberty principle (L) and some rule, possibly a political decision rule (R) and government service (G) have effects on the future distribution of wealth and thereby utility. We'll also assume that the effects tend to vary by individual because of differences in native talent, family, and location, $W_i=w_i(G, R, L)$. Each person's utility function or perceived quality of life may also differ for the same reasons, $U_i=u_i(W_i)$. The rules defining equal liberty are prior to the decision behind the veil and so L is held constant for deliberations behind the veil, while the procedural rules and some key government services are decided from behind the veil. Let $F_i$ characterize the probability of being a person of type i. If the typical individual is an expected utility maximizer, their choice from behind the veil would be:

$$U_i^e = \sum_{i=1}^{N} F_i u_i(w_i( G, R, L)) \tag{18.13}$$

Note that, as written, everyone maximizes the same fairness objective function, which implies that a perfect consensus would emerge among the individuals that use the Rawlsian approach. Assuming that each utility function (after the substitution) is strictly concave, we can use our normal optimization methods to characterize the ideal level of government services and type (or system of rules) that would be adopted from behind the Rawlsian veil. Differentiating with respect to G and R yields two first-order conditions that characterize the unanimous decision that would be made from behind the Rawlsian veil.

$$U_{iG}^e = \sum_{i=1}^{N} F_i u_{iW} w_{iG} = 0 \tag{18.14a}$$

$$U_{iR}^e = \sum_{i=1}^{N} F_i u_{iW} w_{iR} = 0 \tag{18.14b}$$

Both first order conditions would be simultaneously satisfied at this Rawlsian ideal institutional arrangement. Both ideals in turn can be characterized as function of the equal liberties adopted in stage 1, G*=g(L) and R*= r(L).

Notice that if everyone adopted the expected utility version of decision making from behind the veil, that each person's justice objective function turns out to be identical and so agreement from behind the veil is essentially automatic unless personal assessments of other people's utility functions or of the W functions vary. Nonetheless, conclusions reached about the best procedures (R*) and core government services (G*) would vary with the system of equal liberties in place and with the degree of risk aversion of the persons in the community of interest. Rawls assumes a very risk averse population that adopts a maximin strategy rather than a maximize expected utility strategy in as represented in equation 18.14b. He argues that people would agree to maximize the utility of the poorest person in society instead of average or risk adjusted utility. (This is one of many points where those adopting the Rawlsian approach might disagree.)

Note also that equation 6.13 closely resembles a Benthamite aggregate utility function. If Fi were interpreted as the number of persons of type i in the community of interest, it would be identical to a Benthamite social welfare function.[6]

Although, the model developed is an abstract one that simultaneously considers grounding institutions and likely public policies. The same approach could be used to characterize ideal fiscal systems, education systems, environmental programs, and demogrant programs, and so forth, given preexisting institutional arrangements. What differs fundamentally from the utilitarian approach is the term L, which limits the types of policies that can be adopted to those that are consistent with the equal liberty rules or rights adopted in the first round. Such constraints are ignored in the usual utilitarian approach.

## Behind Buchanan's Veil of Uncertainty

With three modifications, the model used to characterize Rawlsian reasoning can also be used to characterize the conclusions that tend to emerge from the Buchanan approach to contractarianism. In Buchanan's analysis, one remains "oneself," and so the relevant uncertainty concerns one's future wealth, rather than one's identity. Also, the probability that one realizes a particular wealth in situation "k" would be affected by the rules adopted. The same rules also affect the extent of one's future wealth associated with situation "k", as in the Rawlsian model. In addition, in Buchanan's analysis, there is always a *status quo ante*. The status quo ante implies that individuals have an original or pre-constitutional endowment of wealth, Wi, which may affect their future wealth and the probability of particular wealth outcomes, $F_i=f_i(G, R, L, W_i)$ and $W_i=w_i(G, R, L, W_i)$. To simplify a bit, we'll neglect present value calculations and focus on a future steady state, and also assume that there are K possible wealth distributions.

In the choice setting described, an individual's expected utility would be:

$$U_i^e = \sum_{k=1}^{K} F_k(G,R,L,W_i^0)u_i(w_{ik}(G,R,L,W_i^0)) \qquad (18.15)$$

[6] Rawls adopts an extremely risk averse version of the objective function that maximizes the worst possible payoff rather than the average payoff. His "maxi-min" approach yields a focus on the worst-off person in society and tends to promote a good deal of redistribution—more so than the expected utility approach used above. Of course, if persons behind the veil differed in their degree of risk aversion, agreement from behind the veil would be unlikely to occur.

Buchanan's approach would include liberties as a decision to be reached at the constitutional convention, so there are three first-order conditions in this case.

$$U_{iG}^{e} = \sum_{k=1}^{K} F_{kG} u_i(w_{ik}(G, R, L, W_i^0) + F_k U_W W_G) = 0 \quad (18.19\text{ a})$$

$$U_{iR}^{e} = \sum_{k=1}^{K} F_{kR} u_i(w_{ik}(G, R, L, W_i^0) + F_k U_W W_K) = 0 \quad (18.16\text{b})$$

$$U_{iL}^{e} = \sum_{k=1}^{K} F_{kL} u_i(w_{ik}(G, R, L, W_i^0) + F_k U_W W_L) = 0 \quad (18.16\text{c})$$

All three first-order conditions are satisfied at an individual's ideal system of services, rules, and rights defining liberties.

Notice that a consensus is not guaranteed in Buchanan's system. This is largely because no presumption that people all initially agree about how to rank alternative outcomes is assumed and also because of differences in initial endowments, wealth functions, and utility functions. For agreement to be more certain, individuals would have to be fairly similar in their utility functions (even more so than assumed to this point) and also would have to have similar expectations about how the alternative services, rules, and liberties affect personal wealth. Even in that case, the effects of uncertainty must largely swamp differences in initial endowments.

In Buchanan's defense, such commonalities are not entirely unrealistic for persons initially living in the same community, and hence, there are prospects for agreement behind Buchanan's veil as well. Alternatively, rather than a perfect convergence, ideals for the three variables ($G_i$*, $R_i$*, $L_i$*) may be sufficiently close to one another that a compromise can easily be worked out—because of the anticipated net benefits (an increase in expected utility from reductions in uncertainty and better protection of rights relative to the pre-contractual setting).

And, as true of the Rawlsian theory, Buchanan's approach can also be applied to analyze the relative merits of policy changes.

In contrast to the utilitarian approach, under the contractarian approach everyone in the relevant community has to anticipate improvements in their quality of life for a consensus to emerge. The requirement for consensus implies that it matters how any new social net benefits are distributed in the post-agreement society. It is not sufficient for the winners to win more than the losers lose, as it would be under the usual utilitarian norms. The losers have to be compensated for their losses.

## VIII. Some Conclusions—Economics and Normative Analysis

In Chapter 15, we demonstrated that a positive analysis of the effects of various ethical dispositions on the size and scope of markets could be undertaken. In this chapter, we shifted to a normative analysis of the positive consequences associated with laws, institutions, and public policies. Economists usually have used the utilitarian approach for normative analysis, although others, such as the contractarian approach, have also been used. Which normative framework should be used is beyond the scope of this chapter. Instead, the chapter has provided an overview of some of the implications of the utilitarian and contractarian approaches.

The normative theories used by economists are all "consequentialist." It is the consequences of economic, institutional, and policy decisions that matter rather than the inherent goodness of particular choices, institutions, or policies themselves. To a consequentialist, good policies have good consequences. For utilitarian economists, good policies increase aggregate utility or "social welfare." The correlation between income, wealth, and utility allows changes in aggregate economic output, income, and social net benefits to be used as proxies for aggregate utility. For contractarians, unanimous or near-unanimous support for a policy is the only reliable way to assure that a policy or institutional reform is expected to benefit everyone affected by it. If the affected individuals all support a policy or constitutional amendment, each must anticipate additional net benefits or additional utility from the change to be implemented.

An advantage of the consequentialist approach to normative analysis is that it is systematic. The consequences of alternative policies have to be imagined and estimated in some detail. The probabilities of different possible consequences have to be approximated, and net benefits associated with each possibility assessed. Such steps are useful even if one does not use one of the economic approaches to determine the relative merits of alternative policies. Although none of these steps is likely to be as precise in practice as textbooks imply, significant differences in net benefits and risks can often be identified, and sensible decisions reached that are likely to broadly increase net benefits within the community of interest.

The domain of welfare economics encompasses a range of changes, from relatively minor adjustments in local public policies to major policy and institutional reforms, and everything in between. It is a large field but is ultimately dependent on the quality of the positive models used to forecast the consequences of the changes under consideration. Major reforms may have consequences beyond the scope of microeconomics, but insofar as economic consequences are taken into account, the positive core of microeconomic analysis remains a significant factor in the normative conclusions reached.

Welfare economics helps both economists and others better understand some of the broad moral issues associated with market-based societies—both their costs and benefits. By doing so, welfare economics and its various precursors have provided justifications for the institutions that are prerequisites for extended market networks to exist. Without the customs and institutions that facilitate voluntary exchange, markets would be far smaller, and it is unlikely that the field of microeconomics would have emerged.

## References


Arrow, K. J. (1973). Some ordinalist-utilitarian notes on Rawls's Theory of Justice. *Journal of Philosophy* 70(9): 245-63.

Buchanan, J. M. (1965), An Economic Theory of Clubs," *Economica*, 32, (1):1-14.

Buchanan, J. M. (1975). *The Limits of Liberty: Between Anarchy and Leviathan* (No. 714). Chicago IL: University of Chicago Press.

Buchanan, J. M. (1988). Contractarian political economy and constitutional interpretation. The *American Economic Review*, 78(2), 135-139.

Buchanan, J. M., & Stubblebine, W. C. (1962). Externality. Economica, 29(116), 371.

Coase, R. H. (2013). The problem of social cost. *Journal of Law and Economics*, 56(4), 837-877.

Coleman, J. S. (1966). The possibility of a social welfare function. *American Economic Review*, 56(5), 1105-1122.

Hamilton, J. H. (1986). The flypaper effect and the deadweight loss from taxation. *Journal of Urban Economics*, 19(2), 148-155.

Harsanyi, J. C. (1982). *Rule utilitarianism, rights, obligations and the theory of rational behavior*. Chapter 12 of Harsanyi, J. C. (1982). *Papers in Game Theory*. Springer, Heidelberg, DE.

Hicks, J. R. (1943). The four consumer's surpluses. *Review of Economic Studies*, 11(1), 31-41.

Irvine, I. J., & Sims, W. A. (1998). Measuring consumer surplus with unknown Hicksian demands. *American Economic Review*, 88(1), 314-322.

Lipsey, R. G., & Lancaster, K. (1956). The general theory of second best. *Review of Economic Studies*, 24(1), 11-32.

Martinez-Vazquez, J., & Winer, S. L. (Eds.). (2014). *Coercion and Social Welfare in Public Finance: Economic and Political Perspectives*. Cambridge University Press.

Mueller, D. C. (1967). The possibility of a social welfare function: Comment. The American Economic Review, 57(5), 1304-1311.

Rawls, J. (1971). *A Theory of Justice*. Cambridge MA: Harvard University Press.

Samuelson, P. A. (1954). The pure theory of public expenditure. *The Review of Economics and Statistics*, 36(4), 387-389.

Samuelson, P. A. (1977). Reaffirming the existence of" reasonable" Bergson-Samuelson social welfare functions. *Economica*, 81-88.

Pigou, A. C. (1951). Some aspects of welfare economics. *American Economic Review*, *41*(3), 287-302.

Pigou, A. (2017). *The Economics of Welfare*. Oxford UK: Routledge.

Sørensen, P. B. (2014). Measuring the deadweight loss from taxation in a small open economy: A general method with an application to Sweden. *Journal of Public Economics*, 117, 115-124.