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, "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, whether specific public policies improved things 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 is rather 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 influential group of philosophers and public policy reformers to determine the nature of a "good society" and how various reforms would move us towards or away from such a society. Pigou did not create the utilitarian project, which goes back at least as far as Aristotle's analysis of the relative merits of different forms of governments and societies in the *Politics*. The term utilitarian is much newer than the project itself. It was first used in Bentham's *An Introduction to the Principles of Morals and Legislation* (1789). He and his fellow travelers established utilitarianism as an influential strand of moral philosophy. Utilitarian ideas were used as a basis for supporting and opposing policy and institutional reforms from then on.

Its influence on economics is nearly as old. Bentham also had an interest in economics and wrote a textbook for the field. He was 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.

The connection between utilitarianism and economics is unsurprising if one realizes that many of the prominent academic economists prior to the twentieth century held chairs in moral philosophy. This was true, for example, of Adam Smith and Alfred Marshall. Utilitarians who were 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 qualified persons (e.g. with the required skills). These reforms of the old medieval order greatly increased the openness and competitiveness of markets, which, as parts I and II of the text demonstrate, tend to increase the scope of trade—other things being equal. Supporting such reforms were, 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 in perspective from the Western medieval one where commerce was always 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 wonder whether markets produced the very best outcomes for society. It was possible, after all, that markets generally increased aggregate utility without actually maximizing the sum of utilities (or, for some utilitarians, the products of utilities).

Efforts to assess whether this was the case or not led Pigou and others to use what might be called monetized indices of utility to approximate aggregate utility. This allowed stronger claims to be made about whether the sum of individual utilities (welfare) was or could be increased through particular public policies—including forms of economic regulation—insofar as a nation's national income or gross output could be used as a proxy for aggregate utility. (See, for example, the quote from Alfred Pigou at the beginning of this chapter.)

Pigou and many other utilitarians of his time believe that revised or additional public policies could improve welfare in this sense. Nonetheless, there were many disagreements about whether particular policies increased social welfare or not among utilitarians—and such discussions and disagreements among utilitarians continue to the present.

Utilitarian normative theory is not the only one used by economists. Many economists have used other normative ideas to appraise the relative merits of activities, rules, and public policies including the Pareto criteria and contractarian approaches, but the utilitarian one is by far the most commonplace. For example, cost-benefit analysis has been explicitly required for a subset of public policies and is often one of the factors taken into account when it is not mandated.

This chapter provides overviews of the normative theories used most frequently by contemporary economists. The utilitarian one is the dominant one, and so it takes up most of the chapter. Other approaches including the Pareto criteria and contractarian models are also reviewed. As true of most other topics covered in this book, theories of "the good society" are manifold, and the chapter aims to provide an overview of the main or core theories and conclusions, rather than a comprehensive overview of the literature—which would take a full-length book or two to do reasonably well. Readings at the end of the chapter provide students with additional material that would broaden and deepen their understanding of results and controversies in this area of research.

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, Pareto superior move occurs when at least one person realizes additional utility at B relative to A, and no one realizes less utility at B than at A. 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 transactions costs can be used to illustrate the Pareto normative criteria. Suppose that Al is a seller and 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 net-utility gained, or lost relative to their original position, before this potential trading opportunity emerged.

Table 18.1 Gains from Trade without Transactions Costs

		Bob (a potential buyer)	
		Accept Offer	Do Not
Al (a potential seller)	Make Offer	3, 3	0, 0
	Do Not	0, 0	0, 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 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 in an Edgeworth Box

An Edgeworth box can be used to develop a somewhat more nuanced illustration of the Pareto principles. Figure 18.1 provides a somewhat more detailed version of the Edgeworth box used in the appendix of 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. These characterize the sets of Pareto superior moves from the initial endowment, and the set of Pareto optimal outcomes. Note that the shift from the initial endowment to the priced induced division of the resources is only one of the possible Pareto superior moves and only one of the possible Pareto optimum states.

---

[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 right-hand equilibrium, neither can change their strategy and improve their condition.

Figure 18.1 The Pareto Criteria in an Edgeworth Box

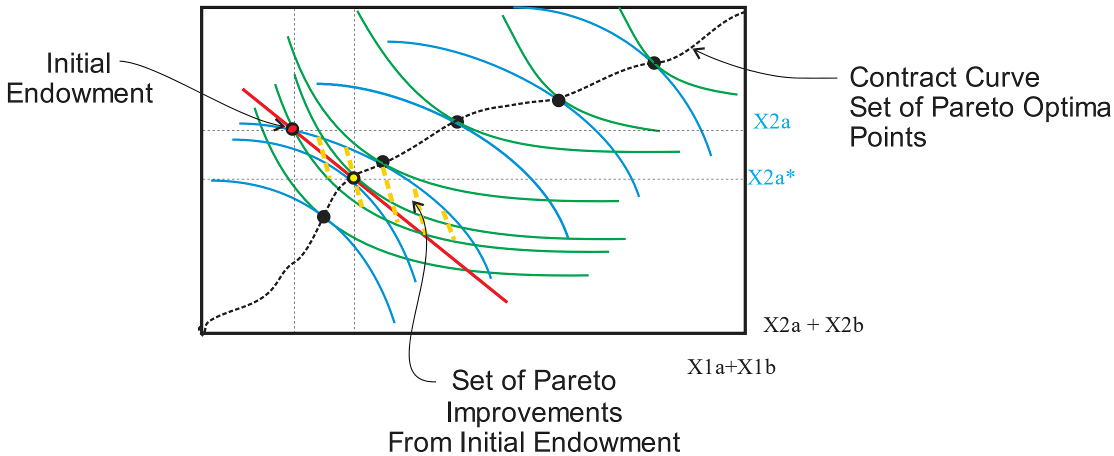

Given the initial endowment, the two indifference curves passing through that endowment characterizes a lens-shaped area between the two indifference curves. That lens-shaped areas is the set of Pareto superior moves at the initial endowment. A move from the initial endowment to any point in that lens makes at least one person better off and no one worse off. (Thus, the initial endowment is not Pareto optimal.) Moves from the initial endowment to all the points on the parts of the two indifference curves that form the outer edge of the lens are also Pareto superior moves, except for the point at the opposite side of the lens where the two indifference curves intersect. They make one person better off without making the other worse off. (The intersection at the other side makes neither better off and so is not a Pareto superior move.) The points on the same two indifference curves that are beyond the lens make someone worse off.

The Pareto optimal points inside the lens 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. It includes distributions of the two goods that are Pareto superior to the initial endowment (those inside the lens, which includes the Walrasian equilibrium) and ones that are not (those outside the lens). That curve is traced out and labelled in the box.

Those that are within the lens can all be reached through a series of Pareto superior moves, as might be undertaken in a sequence of trades between the two individuals whose endowments and utility levels are characterized by the box and its indifference curves. Note also that the equilibrium price outcome is just one of these possibilities.

The Pareto optimal points in a given choice setting are not always unique, although there are cases in which 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 can be reached via voluntary exchange (in the absence of mistakes).

In a trading situation where there are no posted prices, informal bargaining between the two players is likely to reach a Pareto optimal point, 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 through 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 voluntary moves away from those combinations of goods are possible because every move would make one or the other or both worse off. One can apply the same logic to determine the Pareto superior moves and Pareto optimal outcomes in all the game matrices used in the previous chapters.

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 realize benefits, B, and bear costs, C, from every action that they choose to take. A net-benefit maximizer maximizes a function from the family of $N(Q) = B(Q) - C(Q)$ functions. This implies that each consumer purchases the quantity, Q*, that sets dB/dQ = dC/dQ. 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 marginal benefit functions are initially higher than their associated marginal cost functions and that the intersect in the positive domain (within the feasible set). This could be assured by strict concavity for goods whose marginal benefits were greater than marginal costs at 0 with marginal benefits greater than marginal costs, but strict concavity was not a necessary condition for the geometric demonstration that purchasing quantities that equate marginal benefits with marginal costs maximize net benefits. That assumption did, however, make the calculus much simpler.

In cases where the marginal benefit curve or functions are monotone decreasing in the price of the goods or services purchased, individual demand curves are simply 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 (Q) into dollars per unit (marginal benefits), they were functions from prices (dollars per unit) into quantities.

This allows the net benefit maximizing model to 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 or another currency) can be used to approximate the total utility gained by consumers and the total cost of production can be used to approximate the utility sacrificed in other uses by input providers and owners. Social net benefits can thus be used to approximate the net aggregate utility gained (or lost) from producing and consuming the goods of interest.

If $b_i(Q)$ is individual i's marginal benefit function for a good or service of interest, then $b_i^{-1}(P)$, 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 $Q^D = \sum_{i=1}^{N} b_i^{-1}(P)$.

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 determine the highest price at which successive units of a good can be sold. Those successive prices characterize the marginal benefit generated by successive units of the good of interest. It is determined by the individual willing to pay the most for each successive unit of the good of interest—e.g. by the highest demander for each of the successive units that might be purchased.

The inverse of a market demand function derived from net-benefit maximizing choices thus is approximately the inverse of the social marginal benefit curve for the individuals in the market of interest. In cases in which the inverse of a market demand curve is the sum of the inverses of each individual's demand function, the inverse of a market demand curve will exactly characterize the marginal benefits realized from successive units of a good. In this case, $inv(Q^{dD}) = \sum_{i=1}^{N} b_i(Q) = SMB(Q)$. This is, for example, the case when the individual marginal benefit functions are linear. In other cases, it is a reasonably good approximation for the social marginal benefit function (in settings without externalities).[2]

Essentially the same logic can be applied to the Ricardian supply curves that were derived in chapter 3 from individual firm marginal cost functions. The industry supply curve can be used to determine the lowest price at which successive units of a good are produced by gradually raising price above zero until a unit is provided, then raising it until the second unit is provided and so forth. This marginal cost is determined by the least cost producer of successive units of the good. Thus the social marginal cost function is approximately 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, $inv(Q^s(P)) \approx SMC(Q) \approx \sum_{j=1}^{M} c_j(Q)$.

These two approximations allow us to use estimates of the inverses of the demand and supply functions to approximate the social marginal benefits and social marginal costs for the good or service of interest.[3] Social net benefits are simply the difference in the areas under the associated social marginal benefit and marginal cost curves (ignoring fixed costs, or by assuming they are zero in the long run),

$$SNB(Q) = \int_0^Q [SMB(Q) - SMC(Q)]\, dQ. \qquad (18.1)$$

[2] There are several proofs of this that can be found with a brief web-search. Using the net-benefit maximizing model directly avoids differences between demand curves derived from utility functions and marginal benefits and the necessity of using Hicksian compensated demand curves instead.

[3] The geometry is similar to that used to analyze pure public goods. In this case, the SMB curve is the "vertical" sum of the individual marginal benefit functions. Samuelson (1954) characterizes the notion of a pure public good and provides a more direct utilitarian analysis of the optimal provision of such goods.

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 16.2: The Static DWL of Monopoly

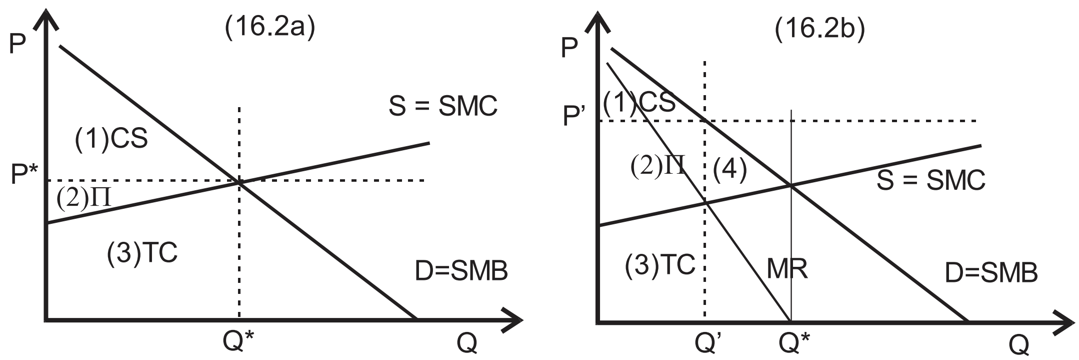

We'll first analyze figure 18.2a. The social benefit is the area under the demand curve from zero to Q* (Areas 1+2+3). The social cost is the area under the supply curve from zero to Q* (area 3). Subtracting social costs from social benefit yields social net benefits. $([1+2+3]-[3]=[1+2])$ Areas 1 and 2, thus represent social net benefits. In this case, social net benefits consist of 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. This output level is also the unique Pareto optimal output of the good being modelled—other things being equal. 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," that competitive markets are Pareto efficient.[4]

[4] I had originally planned to add an appendix with a general proof of the first welfare theorem. But the chapter turned out to be longer than anticipated. For a general calculus-based demonstration see Lange (1942). For a more general proof, see Arrow (1951).

This is true of both the Marshallian and Ricardian variations of competitive markets, although the long run supply curve would normally be horizontal in Marshallian case and profits would equal zero, e.g. every firm makes the "ordinary" 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.

Notice that at the market equilibrium of competitive markets, price equals the marginal cost of production and the marginal benefit of consumption, but also the social marginal cost of production and the social marginal benefit of consuming or otherwise benefiting from the goods and services produced. Thus, competitive markets are self-regulating in a manner that is consistent with the social net benefit maximizing norm (again in the absence of externalities). The social net benefit maximizing level of a good or service also varies with input prices, the prices of substitutes, and consumer willingness to pay for the good or service of interest.

Next focus on figure 18.2b. That is the case of a "price-setting" firm facing exactly the same demand curve and with the same marginal cost curve as the industry's marginal cost curve in figure 18.2a. The 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), but 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 how cartelization or monopolization of a competitive market tends to reduce social net benefits.

In principle, there is a Pareto superior move associated with a shift from Q' to Q*, which might be accomplished by selling the first Q' units at price P' and the last units at price P*. Notice that that combination makes both consumers and firms in this market better off.

Of course, the existence of such deadweight losses depends on how the monopoly or cartel was established. If an innovative firm creates a new product that is temporarily protected by a patent, then the social net benefits of figure 18.2b are new and part of the consumer surplus represents the increase in social net benefits associated with successful innovation.

(It may be presumed that the new product produces more consumer surplus than their previous purchases, but their previous purchases doubtless produced consumer surplus as well.) If Schumpeter is correct in his assessment that large firms are generally more innovative than the smaller firms of competitive markets, 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.)

IV. Social Net Benefits and Externality Problems

Partial equilibrium analyses of markets focus on one or at most a few interconnected markets. It makes no effort to characterize entire market networks, nor does it take account of effects on other non-market activities. However, there are many cases in which market activities affect the residents of a community in ways that are not intermediated by markets.

Many folks decorate their houses for holidays, maintain fine gardens in front of their houses, and take steps to assure 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 in the usual sense. Such fastidious homeowners produce positive "externalities" (spillover benefits) that at least slightly 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 noises, smoke, and parking problems that increase the welfare of some neighbors (who enjoy the music or smells) but reduce the welfare of other neighbors who dislike the noise and odors. Such positive and negative externalities (spillover costs) may not be taken into account by those throwing the parties. (They would not do so if they are pragmatists, but might if they have strongly internalized utilitarian norms or are altruists.)

Similar externalities are often associated with the production activities by firms. In the late nineteenth century, as large-scale production in industrial facilities began being commonplace, and firms would normally use the air and water to dispose of waste products at levels that produced noticeable effects—odors and in some cases health problems. Such spillover costs and benefits are not part of market transactions and may not be taken into account at those generating them—whether firms or consumers. Because of that neglect, such externality generating activities are rarely undertaken at levels that maximize social net benefits.

Figure 18.3 illustrates an externality problem that might emerge from production activities of firms such as air and water pollution or of the uses to which consumers of the goods and services purchased by them. As is usually the case, we'll assume that the externality generators are "pragmatists" that take account of only their own narrow interests (consumer surplus or firm profits). The previous chapter implies that this would not always be the case, but the normative analysis of externality problems nearly always assumes this.

Figure 18.3: Externality Problems

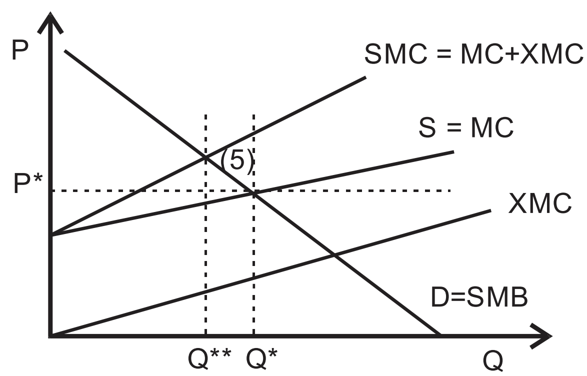

The externality illustrated is assumed to be an external spillover cost that varies directly with the extent of the product sold in this market. The spillover marginal cost is labelled XMC, for eXternal marginal cost. The output produced by a competitive market is, as usual, where supply equals demand, 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 again assumed to be a good proxy for the social marginal benefits generated by this market. The supply curve again 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.

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 either with "free disposal" of waste products using the air and water system or the manner in which consumers use the good purchased. The output where SMC equal 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 chapters 14 and 15 imply that such actions may not always be forthcoming and, if undertaken, 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.

The cost of the policies that would ameliorate the problem should also be taken into account, when making such recommendations, but is often neglected. 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 that case, inaction maximizes social net benefits.

Nonetheless, it is clear that there are cases where governmental or other actions can potentially increase social net benefits, as argued by Pigou.

A Coasian Solution to Externality Problems

A possible non-governmental solution was suggested by Coase (1960), where he argued that if transaction costs are low, the persons affected by an externality could organize and contract with the firms and consumers in the relevant market to reduce sales and production of the goods generating the externality.

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 affected 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, and 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 do so or not. This result holds as long as transactions costs are low (approaching zero) and that the bargains negotiated do not themselves alter any of the curves drawn. This requires that the various side-payments associated with a Coasian contract have only insignificant income and liquidity effects. Otherwise, the curves would shift as side payments were made.

The Mathematics Behind the Externality Diagrams

The mathematics behind the externality diagrams 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 quantity output of the externality generating activity. Note that this quantity is not normally 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 $B = \mathrm{B(Q, P^0, Y)}$, $C = \mathrm{C(Q, w, r, T)}$, and $\mathrm{X} = \mathrm{X(Q, D, A)}$, 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 wind or air speed (A). Given that richer characterization of the B, C, and X functions, the implicit function theorem implies that

$$Q^{**} = q(\mathrm{P^0, Y, w, r, T, D, A}). \qquad (18.3)$$

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

Ameliorating Externalities with Economic Regulation(s)

When the social net benefit maximizing outcomes such Q** do not arise from the behavior of the individuals involved, government regulations of various kinds can potentially increase social net benefits.

To model this, we'll assume that regulation R can reduce output in this market and is undertaken to maximize social net benefits. However, we'll also assume creating and implementing regulations is a costly activity and those costs should be taken into account when evaluating the relative merits of regulatory solutions.

The costs associated with implementing the regulation(s) of interest 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:

$$N = \mathrm{B}(\mathrm{Q}(\mathrm{R}), \mathrm{P}^0, \mathrm{Y}) - \mathrm{C}(\mathrm{Q}(\mathrm{R}), \mathrm{w}, \mathrm{r}, \mathrm{T}) - \mathrm{X}(\mathrm{Q}(\mathrm{R}), \mathrm{D}, \mathrm{A}) - \mathrm{e}(\mathrm{R}) \qquad (18.4)$$

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

$$N_R = \mathrm{B}_Q\mathrm{Q}_R - \mathrm{C}_Q\mathrm{Q}_R - \mathrm{X}_Q\mathrm{Q}_R - \mathrm{e}_R = 0 \text{ at } \mathrm{R}^* \qquad (18.5a)$$

or

$$N_R = [\mathrm{B}_Q - \mathrm{C}_Q - \mathrm{X}_Q]\mathrm{Q}_R - \mathrm{e}_R = 0 \text{ at } \mathrm{R}^* \qquad (18.5b)$$

Note that the term inside the brackets is the first order condition for maximizing social net benefits with an ideal costless policy. The implementation and administrative costs imply that the regulation should aim for a bit smaller reduction in output than the diagram suggests—because the costs of implementation should also be taken into account.[5]

---

[5] 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. This type of policy can be modeled in a similar way. The ideal stringency of such devices, E**, would also maximize social net benefits which can be characterized in a manner similar to equation 18.4.

In this case, X would be a function of the rigor of the emissions inhibiting device, E. Annual social net benefits can be characterized as $\mathrm{N} = \mathrm{B}(\mathrm{E}) - \mathrm{C}(\mathrm{E}) - \mathrm{X}(\mathrm{E}) - \mathrm{e}(\mathrm{E})$. The ideal policy maximizes net benefits. Social net benefits are maximized when $X_E = B_E - C_E - e_E$. The optimal rigor of the emission reducing device, E**, again sets the marginal social costs, including marginal regulatory and environmental costs equal to marginal benefits. (In this

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

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 economic relationships. Normally these 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 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 18.3 as shown by equation 18.5b.

To determine 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 for market demand (D) and supply (S) for reasons already developed.

$$B(Q) \approx \int_0^Q D^{-1}(Q)dQ$$

And

$$C(Q) \approx \int_0^Q [S^{-1}(Q) + X(Q)]\mathrm{dq}$$

---

case, the social marginal benefit is the diminution of the externality and the cost is the reduction in the consumer surplus and profit and increase in taxes to fund the policy's administration.)

From the perspective of cost benefit analysis, such devices should be required (e.g., E**>0) only if they increase social net benefits. This will be the case if marginal regulatory costs are small relative to the marginal external costs.

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:

$$N \approx \int_0^{Q^{**}} [D^{-1} - (S^{-1} + X)]dQ - \int_0^{Q^*} [D^{-1} - (S^{-1} + X)]dQ - e(R^*) \qquad (18.6a)$$

Which can be written as:

$$N \approx -\int_{Q^{**}}^{Q^*} [D^{-1} - (S^{-1} + X)]dQ - \mathrm{e}(\mathrm{R}^*) \qquad (18.6b)$$

The first term, the one inside the integral, corresponds roughly to the triangle labeled (5) in figure 18.3. It represents the maximal annual net benefit achieved by the regulation if it can be costlessly implemented and enforced, which can be denoted as N^0. 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:

$$\frac{\mathrm{N}^0}{\mathrm{r}} > F + e(R^*)/r \qquad (18.7a)$$

This assumes that the benefits and administrative costs begin immediately. However, if they are delayed by T years, then that delay would have to be taken into account. In this case, timing would have a more obvious effect on the calculations and the initial fixed costs of creating the regulations and purchasing the required equipment would play a larger part in the assessment.

$$[\frac{1}{(1+r)^T}][\frac{\mathrm{N}^0}{\mathrm{r}}] > F + [\frac{1}{(1+r)^T}][\frac{e(R^*)}{r}] \qquad (18.7b)$$

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^0, that can potentially be realized via regulation. Note that if the N of equation 18.6b 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 another of the present value formulas should be used for this calculation and expected values would also have to be used to characterize some of the benefits and costs. However, the same logic applies.

V. Social Welfare Functions, Generalizations of Bentham's Additive Utility Norm

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 utilitarian 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 the Pigovian approximation, it is possible to determine whether one policy is better than another or one country better organized than another using various currency-based measures as in the first half of this chapter.

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, most theoretical work in welfare economics uses 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, $W = \sum_i^N u_i(Q_1, Q_2, Q_3, Q_4, \dots)$. 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: $W = w(U_1, U_2, U_3, U_4, \dots U_N)$ largely replaced such concrete functional forms for aggregate utility or social welfare. The only restriction on such functions is that $W_{U_i} > 0$ for all i, which is to say that everyone in the community counts, although not necessarily equally. 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 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.

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

Prior to 1954, neoclassical economics had regarded all goods and services to be "private goods" that were consumed 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 individual consumption of the candy bar falls from 1 to one half.

Paul Samuelson (1954) argued that 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 goods "club goods." Club goods were "congestible." The services that they provide users decrease with the number of persons using them, but not in proportion to their numbers. Examples include swimming pools, public parks, playgrounds, bicycle paths, highways, lectures, and so forth. It turns out that there is actually a spectrum of types of goods with varying different degrees of "shareability."

However, for the purposes of this section of Chapter 18, our main interest is Samuelson's characterization of the ideal service level of a pure public good. His analysis is calculus based and 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 $U_i = u(G, X_i)$ be the utility of individual *i* associated with a particular combination of the public good G and private good Xi received by *i*. Note that each person gets the same level of pure public good G, because it's perfectly sharable, but particular amounts of the pure private goods, which may be the same or different.

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 maximal combinations of public good and private goods that can be produced. The quantity of the pure private good X is distributed among N persons, $X = \sum_{i=1}^{N} X_i$. The constraint function is in its "zero form," $T(G, X) = 0$, as required for the Lagrange method. 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 results equal to zero:

$$\mathrm{L} = w(U_1, U_2, U_3, U_4, \ldots U_N) - \lambda[\mathrm{T(G, X)}] \qquad (18.8)$$

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

$$[\sum_{i=1}^{N} W_{U_i} U_{iG}] - \lambda \mathrm{T}_G = 0 \qquad (18.9a)$$

$$W_{U_i} U_{iX} - \lambda \mathrm{T}_X = 0 \qquad (18.9b)$$

for all $i = 1 \ldots N$ (Equation 18.9b thus represents N first order equations)

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

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 clever series of steps to develop a surprising result.

First, as usual, shift the lambda terms to the right by adding their positive equivalents to each side. Then divide the first such equation by the j^{th} member of the second group of first order conditions to eliminate the lambda.

$$[\sum_{i=1}^{N} W_{U_i} U_{iG}]/W_{U_j} U_{jX} = \frac{T_G}{T_X} \quad (18.10)$$

(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.

$$\sum_{i=1}^{N} [W_{U_i} U_{iG}/W_{U_j} U_{jX}] = \frac{T_G}{T_X}$$

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 is equal to one another at the margin:

$$W_{U_j} U_{jX} = W_{U_i} U_{iX} = \lambda T_X \quad (18.11)$$

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

$$\sum_{i=1}^{N} [W_{U_i} U_{iG}/W_{U_i} U_{iX}] = \frac{T_G}{T_X}$$

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 straightforward.)

This allows us to simplify quite a bit. The marginal welfare weights, $[W_{U_i}]$, in the numerator and denominator are the same; thus, they can be divided out of each term in the sum. This implies that at G*:

$$\sum_{i=1}^{N} [U_{iG}/U_{iX}] = \frac{T_G}{T_X} \text{ at } G^* \quad (18.12)$$

The ideal level of a pure public good sets the sum of the marginal rates of substitution between the private and public goods for each individual in society equal to the technological rate of transformation between private and public goods. **Surprisingly, the**

optimal level of a pure public good is completely independent of the social welfare function used!

(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**.)**

However, the distribution of the private good will vary with the specific form 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 that this is a pure "benevolent central planner" model. An imaginary idealistic planner determines the supplies of both pure public and pure private goods. The central planner also has complete control over the allocation of private goods among the individuals in the society characterized. Moreover, the central planner him or herself is evidently a devoted utilitarian—there are no principal-agent or administrative costs.

Weaknesses of the Maximizing Social Welfare Approach

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

There are several problems with the maximizing social welfare approach to normative analysis. Perhaps, the most important is the assumption that individual utility functions can be known by a theorist—not simply some of their general characteristics. Numerical estimates have to be developed for every person's utility function and for the technological-economic environment in which policies are to be implemented. Samuelson's characterization of the ideal level of a pure public good requires both estimates of U_{iG} and U_{iX} for every individual and of T_G and T_X for the economy as a whole in order to determine the ideal quantity of the pure public good of interest, G*. Ordinal utility functions are not sufficient for the former, cardinal ones (e.g., utility functions that one 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 positive analysis based on the existence of such functions—because it uses utility functions to characterize behavior rather than utility per se. As long as individuals generally behave as if they have utility functions, utility-based models will accurately predict how individuals respond to changes in circumstances.

A third problem has to do with the lack of political context in most analyses—which implicitly or explicitly assume an all-powerful benevolent dictator or central planner, or a population of devout utilitarians, rather than a democracy in which differences in normative ideas, tastes, and income cause voters to disagree about the nature of the best policies. A majority of voters, or at least pivotal voters, would have to be devout utilitarians for utilitarian policy analysis to affect election-driven public policies.

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 a policy—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.[6] There is a domain in which utilitarianism generates morally plausible results, but this does not include all morally relevant questions—in spite of the arguments made by utilitarians.

VI. Contractarian Norms—a Possible Alternative to Utilitarianism

An alternative to the utilitarian approach is the contractarian perspective. 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. More recent versions of contractarian thought were developed by Rawls (1971) and Buchanan (1975). Governments grounded in the consent of those governed are formed to advance the shared interests of their citizens.

In Hobbes' analysis, 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

[6] 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.

> 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 assumes that, rather than a war of every man against every other, individuals create a commonwealth to reinforce natural laws, which in the absence of governmental support provide a less perfect constraint on 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. It is the use of consent as the principal process for determining ethical ideas and for ranking political and legal procedures and states of the world that is the main point of agreement among contractarians. Many of their conclusions differ from one another—as is true of utilitarians. The contractarian line of reasoning preceded utilitarian thought by a century and a half. With respect to good governance, it argued 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). And, it continues to play a role in ideas about legitimate government and the bounds of the authority of such governments today.

The twentieth century revival of contractarian analysis was launched more or less independently by two 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 analysis, and Buchanan's is more in the spirit of Hobbes'—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 completely 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 work (and in most of his subsequent work), he argues that consensus is the ultimate foundation for legitimate governance and also for 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 Rawls' case, the veil of ignorance is a thought process that individuals should use to think about legal, political and economic systems when evaluating them. This requires them to accept his reasoning about it being an effective 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 political and legal institutions, which requires each individual to take a wide range of possibilities into account. 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, inherent in making long term commitments about "the rules of the game" and requires no agreement, implicit or otherwise, about the relative merits of "the veil" as a method for discerning just or fair social systems.

Rawls' 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.

Modelling Choices from Behind the Rawlsian Veil of Ignorance

Rawls' notion of equal liberty in principle constrains the types of decisions that can be properly made from behind the veil of ignorance. To simplify a bit, we'll assume that both the equal liberty principle (L) and a government characterized by a political decision rule (R) and expected government services (G) have effects on the future distribution of wealth and thereby utility. We'll also assume that the effects tend to vary among individuals 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, $L = L^0$,while the procedural rules and some key government services are decided from behind the veil. Let F_k characterize the probability of being a person of type *k*. If the typical individual, *i*, is an expected utility maximizer, choices from behind the veil can be characterized as maximizing:

$$U_i^e = \sum_{k=1}^{N} F_k u_k(w_k(G, R, L^0)) \qquad (18.13)$$

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.

$$U_{iG}^{e} = \sum_{k=1}^{N} F_k u_{kW} w_{kG} = 0 \quad (18.14a)$$

$$U_{iR}^{e} = \sum_{k=1}^{N} F_k u_{kW} w_{kR} = 0 \quad (18.14b)$$

Both first order conditions would be simultaneously satisfied at the Rawlsian ideal institutional arrangement. The implicit function theorem implies that the characteristics of the ideal government (the one unanimously supported) can be characterized as a function of the equal liberties adopted in stage 1, the probabilities of being different persons, and the number of persons in the community:

$$G^* = g(L^0, N, F_1, F_2, \ldots F_N) \quad (18.15a)$$

and

$$R^* = r(L^0, N, F_1, F_2, \ldots F_N). \quad (18.15b)$$

If everyone adopted the expected utility version of decision making from behind the veil, each person's objective function is identical, and thus agreement from behind the veil is assured, unless personal assessments of other people's utility functions or of the W functions vary.

Conclusions reached about the best procedures (R^*) and core government services (G^*) would vary with the system of equal liberties in place and the probabilities of being each type of person.

Note also that equation 18.13 closely resembles a Benthamite aggregate utility function. If F_i were interpreted as the number of persons of type i in the community of interest, it would be identical to such welfare function.[7] Thus, utilitarians might reach similar conclusions.

What differs from the utilitarian approach is the term L. The equal liberties assumption limits the types of policies that can be adopted. They have to be consistent with the equal

[7] 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.

liberty rules, the system of rights, adopted in the first round—about which Rawls is not as clear.

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 for oneself, rather than one's identity. Also, the probability that one would realize a particular wealth in situation "k" is affected by the rules adopted. Those 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, W_i^0, which is likely to affect the probability of particular wealth outcomes, $F_{ik} = f_{ik}(G, R, L, W_i^0)$ and also the extent of wealth in outcome k, $W_{ik} = w_i(G, R, L, W_i^0)$.

To simplify a bit, we'll neglect present value calculations and focus on a future steady state, and also we'll assume that there are K possible wealth outcomes for each person in the future society.

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

$$U_i^e = \sum_{k=1}^{K} F_k(G, R, L, W_i^0) u_i(w_{ik}(G, R, L, W_i^0)) \qquad (18.16)$$

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

$$U_{iG}^e = \sum_{k=1}^{K} [F_{kG} u_i(w_{ik}(G, R, L, W_i^0) + F_k U_W W_G)] = 0 \qquad (18.17a)$$

$$U_{iR}^e = \sum_{k=1}^{K} [F_{kR} u_i(w_{ik}(G, R, L, W_i^0) + F_k U_W W_R)] = 0 \qquad (18.17b)$$

$$U_{iL}^e = \sum_{k=1}^{K} [F_{kL} u_i(w_{ik}(G, R, L, W_i^0) + F_k U_W W_L)] = 0 \qquad (18.17c)$$

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

Notice that a consensus about an ideal system is not guaranteed in Buchanan's system. Disagreement may occur because of differences in circumstance. Individuals differ in their utility functions, precontract wealth and may disagree about the probabilities of the K possible outcomes.

Agreements of this variety are more likely if individuals have similar utility functions and similar expectations about how alternative combinations of services, rules, and liberties affect the probabilities and individual wealth outcomes under scenarios 1...K. Even in that case, the effects of uncertainty must largely swamp the effects of differences in initial endowments.

However, ideal forms of unanimity are unnecessary. What is necessary is that there be at least one combination of G, R, and L that each individual expects to make themselves better off than at the status quo. In Hobbesian circumstances, there are many combinations of G, R, and L that are likely to do so—even if few or none are ideal from the perspective of all individuals. Given this, a constitutional decision rule may be adopted as a reasonable way to choose among the combinations of G, R, and L that are unanimously expected to produce outcomes that are superior to the status quo.

In somewhat better conditions, the set of institutions that are unanimously regarded to be better than the status quo tends to shrink, but not necessarily to the null set.

Note that 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 new social contract or revisions to the old one are, in this sense, Pareto superior to the initial conditions, the status quo ante. It is not sufficient for the winners to win more than the losers lose, as it would be under the usual utilitarian norms.

VII. 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 normative analysis grounded in particular normative ideas that have attracted the most attention from economists interested in assessing the relative merits of alternative policies as well as their consequences.

The various welfare economics approaches are all "consequentialist" in that it is consequences of policy decisions that matter rather than the inherent goodness of particular policies themselves. To a consequentialist, good policies have good consequences. Welfare economics naturally focuses mainly on economic consequences and uses utilitarian based analysis to determine whether "social welfare" has increased or not—e.g. whether aggregate utility has increased or not. The correlation between income and wealth and utility allows changes in aggregate economic output or income to be used as a proxy for aggregate utility. That inference also implies that monetized calculations of benefits and costs can be used to assess whether aggregate utility tends to increase when relatively narrow policies are adopted (or repealed).

An advantage of the economic approach to normative analysis is that it is systematic. The consequences of alternative policies have to be imagined and estimated in some detail, 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, large 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.

Broader issues and even systemwide redesigns have also been analyzed by welfare theorists in ways that shed light on larger economic-social-political-legal system issues. Many of these address concerns beyond the main foci of microeconomics. However, insofar as economic consequences are taken into account, welfare economics belongs to the field of economics—if along one of its farthest boundaries. It has more natural homes in philosophy and public administration.

Careful examinations of the issues, including economic ones, can help policy makers avoid mistakes and determine policies that are likely to improve life for residents of their communities. Avoiding such mistakes clearly requires political and economic analysis.

References and Suggested Readings

Arrow, K. J. (1951). An extension of the basic theorems of classical welfare economics. In J. Neyman (Ed.), *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability* (pp. 507–532). University of California Press.

Arrow, K. J. (1973). Some ordinalist-utilitarian notes on Rawls's Theory of Justice. *Journal of Philosophy*, 70(9), 245-263.

Blackstone, W. (1765–1770). *Commentaries on the laws of England* (4 vols.). Clarendon Press.

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). University of Chicago Press.

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

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

Buchanan, J. M., & Tullock, G. (1962). *The Calculus of Consent: Logical Foundations of Constitutional Democracy*. University of Michigan Press.

Coase, R. H. (1960). The problem of social cost. *Journal of Law and Economics*, 3, 1-44.

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

Congleton, R. D. (2024). Optimal taxation for democracies with less than perfect voters: A public choice perspective. *Kyklos*, *77*(1), 3-21.

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

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

Hobbes, T. (1651). *Leviathan*. Andrew Crooke.

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

Lange, O. (1942). The foundations of welfare economics. *Econometrica*, *10*(3/4), 215–228.

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

Locke, J. (1690). *Two Treatises of Government*. Awnsham Churchill.

Montesquieu, C. de S. (1748). *The Spirit of the Laws* (T. Nugent, Trans.). J. Nourse & P. Vaillant.

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

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

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

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

Rousseau, J.-J. (1762). *The Social Contract* (G. D. H. Cole, Trans.).

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*, 44(173), 81-88.

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.