



INTERNATIONAL

Overview of Cost-Benefit Analysis and its Applications in Public Policy Decisions

Prepared By:

Michael Trebilcock

Adonis Yatchew

Andy Baziliauskas*

June 2007

* Trebilcock is university professor and professor of law and economics at the University of Toronto, and is a Senior Consultant with CRA International; Yatchew is professor of economics at the University of Toronto, editor-in-chief of *The Energy Journal*, and a Senior Consultant with CRA International; Baziliauskas is a Principal with CRA International. This paper was commissioned by the Market Evolution Analysis and Research Group, IESO..

TABLE OF CONTENTS

EXECUTIVE SUMMARY	1
1. INTRODUCTION TO COST-BENEFIT ANALYSIS.....	2
1.1. WHAT IS COST-BENEFIT ANALYSIS?.....	3
1.2. DECISION RULES.....	3
1.3. IMPLEMENTING CBA.....	4
1.4. ESTIMATING GAINS AND LOSSES	6
1.4.1. Market Goods and Services	6
1.4.2. Non-Traded Goods.....	7
1.5. DISCOUNTING AND PRESENT VALUE.....	8
1.6. EFFICIENCY CONCEPTS AND COST-BENEFIT ANALYSIS	8
1.6.1. Allocative efficiency	9
1.6.2. Productive efficiency	10
1.6.3. Dynamic efficiency	11
1.6.4. Winners and losers.....	11
1.7. WEIGHTING EFFECTS ON DIFFERENT GROUPS DIFFERENTLY	12
2. EXAMPLES OF THE USE OF COST-BENEFIT ANALYSIS	15
2.1. ONTARIO ENERGY BOARD	16
2.2. CANADIAN COMPETITION BUREAU	16
2.3. THE U.S. ENVIRONMENTAL PROTECTION AGENCY	17
2.4. THE REGULATORY IMPACT ASSESSMENT REQUIREMENT IN THE UK	18
3. DEALING WITH ‘WINNERS’ AND ‘LOSERS’	19
3.1. MECHANISMS TO EASE TRANSITION FOR WORKERS	21
3.2. MECHANISMS TO EASE TRANSITION FOR FIRMS	22
3.2.1. Adjusting the Timing of Implementation and Compensation	23
3.2.2. ‘Grandfathering’	24
3.2.3. LMP Example: Compensation through Allocation of Financial Transmission Rights.....	25
4. SUMMARY AND CONCLUSIONS.....	25

EXECUTIVE SUMMARY

Cost-benefit analysis (“CBA”) is a decision standard that is commonly used by policymakers to assist in determining whether a policy or project should be implemented. CBA monetizes the effects of a policy on individuals or groups in order to facilitate comparisons with the *status quo* or with other policies. In its basic form, CBA places primary weight on economic efficiency, but it can also be modified to account for adverse wealth distribution effects by appropriately weighting the costs and benefits to individuals or groups.

Governments and regulators in Canada and other countries make use of CBA to choose among policy options or at least to inform their policy decisions. Examples of government bodies that use CBA are the Ontario Energy Board, the Canadian Competition Bureau, the U.S. Environmental Protection Agency, and the government of the UK.

Changes in policy or regulation almost always create both ‘winners’ and ‘losers’. Compensation is often paid to those harmed by a policy to promote fairness, to efficiently ease the transition to the post-change world, and to solidify political support for the policy. Governments have implemented a wide range of compensatory mechanisms including worker retraining and direct financial payment to offset the harm inflicted by policies on individuals or groups.

Compensation to firms should, whenever possible, be made in lump-sum form rather than as quantity-contingent subsidies, to avoid distorting incentives to make economically efficient decisions. Examples of compensation and transition mechanisms for firms include delaying implementation and ‘grandfathering’. In electricity markets that have shifted from uniform pricing to locational marginal pricing, financial transmission rights can also be used to compensate ‘losers’.

1. INTRODUCTION TO COST-BENEFIT ANALYSIS

Cost-benefit analysis (“CBA”) is a *decision standard* that is used to decide whether a project (which can include the introduction of a new set of regulations or a modification in existing regulations) should be implemented. In simple terms, a proposed project ‘passes’ a CBA if the benefits to individuals exceed the losses.

The advantages of CBA include its transparency, which encourages accountability by revealing where decisions are at variance with the analysis. CBA also provides a coherent framework for data collection and for identifying information gaps, and it allows for the aggregation of effects that are dissimilar, through the use of money as a metric, to provide a simple one-dimensional measure of the overall effects of a proposed project.

CBA does *not* take into account all factors that bear on the social benefits of a project. An important difference between CBA and other decision rules is that, strictly speaking, CBA considers only the implications of a project for economic efficiency, while ignoring the effects of the project on wealth distribution.¹ Most changes to regulations will harm some individuals while benefiting others, and depending on the identities of the winners and losers, this redistribution of wealth may be socially undesirable (for example if high income groups benefit at the expense of lower income groups). Although CBA does not account for wealth distribution effects, the analysis can be modified to give more weight to the effects of a project on certain socially protected or disadvantaged groups. Asymmetries in the distribution of gains and losses from a proposed project may also prevent efficient projects from being implemented, if the ‘losers’ have sufficient influence over the political process. As we discuss in this paper, there are mechanisms that can be used in the transition from existing to new regulations to correct for adverse distributional consequences and overcome opposition through the political process, to ensure that efficient regulatory changes are implemented.

¹ There are many other decision rules that are used to evaluate public sector projects, and a review of all of these other methodologies is beyond the scope of this paper. Another important difference between CBA and some other methodologies is that it measures effects through changes in individual welfare, which is assumed to depend on the satisfaction of individual preferences.

1.1. WHAT IS COST-BENEFIT ANALYSIS?

CBA is widely used by government agencies in Canada, the U.S., and throughout the world to evaluate proposed government projects and changes to regulations and policies. Although CBA was rarely used in the U.S. before the 1980's, it is now commonly applied by U.S. regulatory agencies, as a result of Executive Orders issued by the Reagan and Clinton administrations (in 1981 and 1993) and a number of Congressional statutes that require regulatory agencies to perform a CBA. The Competition Bureau and the Treasury Board in Canada, as well as a number of other Canadian government agencies and departments, routinely use some variant of CBA to at least inform their analyses of project proposals and changes to regulations. The Ontario Energy Board also appears to be adopting CBA, at least in some instances.

In general terms, CBA involves first evaluating both the costs and benefits to individuals of a proposed project or policy (such as change to regulations), and then deciding whether the change should be implemented according to a particular decision rule that compares costs and benefits. Loosely speaking, CBA justifies a decision, project, or regulatory change when the gains to individuals resulting from the change exceed the losses to individuals. The term 'cost benefit analysis' is also sometimes used to refer to any type analysis that trades off costs and benefits, even when it does not invoke a decision rule that approves a change when gains exceed losses.

1.2. DECISION RULES

A change is said to be a *Pareto improvement* if at least one individual is made better off by the change and no one is made worse off. This is a very restrictive criterion which is likely to prevent the implementation of many projects that most people agree would clearly be socially desirable. Changes to regulations typically affect large and diverse populations and consequently are likely to reduce the welfare of at least one person. Under the Pareto criterion many 'good' changes would not be implemented, no matter how many people benefit or how small the losses to the individuals who are harmed. For example, consider a project to produce a vaccine that would save

millions of lives, while taxing one (non-altruistic) wealthy person one dollar. Under the Pareto criterion, the *status quo* is efficient and the vaccine would not be produced, even though the project would be implemented under any reasonable decision standard.

A less restrictive decision rule uses the *Kaldor-Hicks* criterion. This criterion is usually stated as requiring that a change be approved if the winners from the change could *hypothetically* compensate the losers and still be better off. More accurately, the change is approved if the gains to ‘winners’ exceed the losses to ‘losers’, such that the change creates benefits that are sufficiently large to offset the losses. This criterion does not require that losers are actually compensated or that winners ‘bribe’ losers to accept a change; it only requires that winners or losers *could* provide compensation or ‘bribes’, in the sense that they could provide compensatory transfers to other parties to make them at least as well off as before while themselves remaining better off than under the *status quo*. Notice that the Kaldor-Hicks criterion is much more permissive than the Pareto criterion, as it approves projects that the Pareto criterion would not (because it requires only that winners could potentially compensate losers). It would approve the vaccine in the example discussed above, since the benefits accruing to vaccine recipients would be more than large enough to offset the loss to the single individual.

The Kaldor-Hicks criterion forms the basis of modern CBA. If winners gain enough from a project to compensate the losers while themselves remaining better off, we can say that gains exceed losses—in other words, the sum of benefits are greater than the sum of costs. This simply stated criterion can be difficult to apply, however, especially when there is a large number of individuals who are affected by the change in different ways. We turn next to implementation of CBA.

1.3. IMPLEMENTING CBA

The first stage of a CBA consists of calculating both the total costs and benefits to individuals that would result from the implementation of the proposed project. The focus on how a project affects individuals results from the fundamental assumption that the satisfaction of individual preferences gives rise to individual well-being. In

other words, individuals know what is good for them, their preferences reflect this knowledge, and they act in a manner consistent with their preferences.

There are two conceptual tools that are commonly used to estimate gains and losses to individuals, namely compensating variation (“CV”) and equivalent variation (“EV”). CV measures the amount of dollars that would have to be given to (or taken from) an individual to allow her to be as well off with the project as she was in the *status quo*. For example, if a project increases the price of widgets, then a consumer of widgets will be made worse off by the project; the CV is the amount that she would have to be compensated to make her as well off as before the change. The EV is the amount that a consumer would pay to avoid a change. In many circumstances, CV and EV differ only by a relatively small ‘income effect’,² and in what follows we discuss only CV in any detail.

The CV is *not* the amount of money that lets the consumer purchase the same number of widgets as before the price increase (in other words, the price increase multiplied by the quantity purchased before the change). This amount overstates the harm to the consumer because when the price of widgets increases she can avoid some of the harm resulting from the price increase by substituting towards other products. Consider an extreme case where a project increased the price of a can of Coke from \$1.00 to \$1.25. If the consumer is indifferent between Coke and Pepsi when the price of each is \$1 per can, then an increase in the price of Coke has *no* effect on the consumer (if the price of Pepsi remains at \$1), since the consumer can avoid the harm from the price increase by simply switching to Pepsi. Compensating the consumer with \$0.25 per can of Coke that she consumed before the change would clearly overcompensate her. In less extreme cases, where there are no perfect substitutes available, the consumer could (and would) avoid some of the harm from the price increase by substituting towards other products, and the harm that she could not avoid is the cost to her of the price increase.

² See, for example, Robert Willig, “Consumer’s surplus without apology”, *American Economic Review*, (66) 5, (1976). Others have argued that empirical and experimental studies have shown large differences between CV and EV in some cases; see Richard Smith, Jan Abel Olsen, and Anthony Harris, “A Review of Methodological Issues in the Conduct of Willingness-to-Pay Studies in Health Care III: Issues in the Analysis and Interpretation of WTP Data”, Center for Health Program Evaluation (1999).

1.4. ESTIMATING GAINS AND LOSSES

The measurement of the costs and benefits of a project can involve a wide range of tools and methods and, depending on the project, can require highly specialized skills and expertise. We provide just a brief overview of some of the methods that are commonly used.

1.4.1. Market Goods and Services

Changes in regulations can affect consumers through changes in prices and quantities of products and services. CV for traded goods (i.e. products and services that are traded in markets) can be estimated by using information about actual consumer behaviour. In practice, CV is often estimated using the closely related concepts of consumer and producer surplus.

The consumer surplus accruing to a consumer from the consumption of X units of a good is the difference between the maximum amount that she would pay for the X units and the amount that she actually pays. Consumer surplus can be measured using information from her demand curve. A consumer's demand curve for a good shows the maximum amount the consumer would pay for various units of the good. For example, a demand curve for apples could show that the consumer would pay \$1 for the first apple, \$0.90 for the second apple, \$0.80 for the third apple, \$0.70 for the fourth apple, and so on.³ If the price of apples is \$0.75, the consumer purchases three apples (the fourth apple is worth only \$0.70 to her) and her consumer surplus is calculated as her total willingness to pay for three apples ($\$1.00 + \$0.90 + \$0.80 = \2.70) less the amount that she actually pays ($3 \times \$0.75 = \2.25), which is equal to \$0.45. If the price increases to \$0.85, she buys two apples, and her consumer surplus is \$0.20 (her total willingness to pay is $\$1.00 + \$0.90 = \$1.90$ and she pays a total of $2 \times \$0.85 = \1.70 , so consumer surplus is $\$1.90 - \$1.70 = \$0.20$).

³ A reduction in the consumer's willingness to pay for an additional apple as the number of apples consumed increases implies that the demand curve is downward sloping. This reflects the assumption of diminishing marginal utility, which implies that as the consumer's consumption of apples increases, her increase in well-being from consuming an *additional* apple falls.

In this example, if a new regulation increased the price of apples from \$0.75 to \$0.85, the cost to this consumer of the regulation would be measured as the loss in consumer surplus. Since consumer surplus falls from \$0.45 to \$0.20, the cost to the consumer in the CBA is \$0.25.

In practice, consumer surplus is estimated by first estimating the parameters of the demand curve using econometric methods based on data on actual prices and consumption volumes. Consumer surplus is measured as the area under the demand curve but above the price.

Producer surplus is measured in an analogous manner, but using supply curves rather than demand curves. A supply curve shows the prices at which producers are willing to supply various units of output. The area above the supply curve but below the price is the producers' surplus. When the introduction of regulation changes the price received by producers, the effect (either a cost or a benefit, depending on whether price decreases or increases) is measured as the change in producers' surplus.

1.4.2. Non-Traded Goods

Some changes to regulations affect individuals through changes in the quantity of 'goods' that are not traded in markets, such as pollution. Changes to individuals' well-being from changes in the quantities of such goods cannot be measured using the price-based methods described above. The simple reason for this is that, since pollution is often not bought and sold,⁴ there are no prices for pollutants and no demand curves that can be used to estimate consumers' valuations for varying quantities of pollutants.

A variety of methods have been used to estimate the effects on consumers arising from changes in the quantities of non-market goods, including techniques that rely on stated preferences and techniques that use revealed preferences. Stated preference methods consist primarily of surveys that ask consumers to report their willingness to

⁴ There are many emissions trading markets where companies and countries can trade the right to emit certain pollutants.

pay (“WTP”) for the project at issue (for example, when the change in regulation reduces the amount of a pollutant, consumers would be asked to report how much they would be willing to pay to have the pollutant removed). There are many types of surveys, including structured surveys that attempt to elicit respondents’ WTP by confronting them with a choice, and conjoint analysis, which attempts to elicit respondents’ valuations for particular combinations of product attributes. Individuals have a clear incentive to overstate their willingness to pay for goods that they would like to have more of but don’t actually have to pay for, and these structured surveys are designed to limit this type of ‘gaming’. Revealed preference methods use information about choices made by individuals, and include hedonic pricing, the property-value approach, and the travel cost approach.

1.5. DISCOUNTING AND PRESENT VALUE

Both the costs and the benefits of proposed projects are typically experienced by individuals into the future. For example, a project might be expected to generate a certain dollar value of benefits to individuals every year for X years, and costs may be incurred annually for Y years. In CBA, the streams of costs and benefits are made comparable by calculating the present value of each stream. The present value of a benefit received in the future is today’s value of that benefit. It is calculated using a discount rate which results in the present value of a benefit being lower than its future value, reflecting the fact that people prefer to have a benefit today rather than a benefit of the same ‘face’ value in the future. The discount rate applied in a CBA ideally reflects the time preferences of the individuals affected by the policy or project being evaluated. As an example, the Treasury Board of Canada has, since 1976, required that benefit-cost analysts use a social discount rate of 10 per cent ‘real’ per annum; i.e. the discount rate is applied to constant, inflation-adjusted dollars.

1.6. EFFICIENCY CONCEPTS AND COST-BENEFIT ANALYSIS

There are three broad concepts that are commonly used by economists to evaluate whether an outcome is efficient, namely allocative efficiency, productive efficiency,

and dynamic efficiency. Below, we explain these concepts and discuss how they are related to cost-benefit analysis.

1.6.1. Allocative efficiency

An allocation of resources is *allocatively* efficient if all market participants produce and consume the 'efficient' amount of output. More precisely, an allocation is allocatively efficient if it is not possible to change the allocation of resources in an economy—for example, by increasing the output of a particular industry—in a way that satisfies the Kaldor-Hicks criterion. Recall that under this criterion, a change in the allocation of resources is efficient if the winners from the change could hypothetically compensate the losers and still remain better off. If it is not possible to make such a change, then the allocation must be efficient.

This concept can be illustrated with an example of an industry where demand for the industry's output is downward sloping (meaning that consumers demand more of the industry's output when its price is lower), and the industry marginal cost is increasing in output (or at least not decreasing). The efficient amount of industry output occurs where demand equals marginal cost.

To see this, recall that the quantity demanded by consumers at a given price corresponds to consumers' marginal benefit of consuming the industry's output at that price. For example, if consumers' marginal benefit from consuming the 10th unit of output is \$18, this is the amount that they are willing to pay for the 10th unit of output. In other words, consumers' demand at a price of \$18 is 10 units. To see that the allocatively efficient level of output is where marginal benefit is equal to marginal cost, suppose that the industry's marginal cost of producing the 10th unit of output is \$20, so that at 10 units of output, consumers' marginal benefit is less than producers' marginal cost. It is clear that it is not efficient to produce the 10th unit of output, since the cost of producing this unit exceeds the benefit of consuming it (in this case by \$2). Under this allocation of resources, a Kaldor-Hicks improvement is possible: if the industry produces one less unit of output, consumers lose \$18 in benefits, but the resources saved can be used to produce \$20 in benefits to consumers elsewhere in the economy. There is a net benefit to be realized—in particular, in the amount of \$2—and therefore the initial allocation would be allocatively inefficient.

Similarly, if the marginal cost of producing the 11th unit is \$16, and the marginal benefit is \$17, then efficiency would be enhanced if the 11th unit were produced. Resources can be taken away from another part of the economy, at a cost of \$16, and reallocated to this industry, for a benefit of \$17. There is a net gain of \$1, and therefore the re-allocation represents a Kaldor-Hicks improvement. Only when output is produced up to the point where the marginal benefit of the last unit produced is equal to the marginal cost, is there no scope for improving the efficiency with which resources are allocated. It is clear, then, that if a project improves allocative efficiency, then the benefits exceed the costs under a Kaldor-Hicks criterion. Such a project would be approved under a CBA (assuming of course that there were no other offsetting costs).

1.6.2. Productive efficiency

A market outcome is productively efficient if each firm is producing output at the lowest possible cost. If a producer is not minimizing its costs, it could, in theory reduce the amount of labour, capital, or some other input that it uses, and maintain the same level of output. The resources released as a result of more efficient production could be employed elsewhere in the economy to produce additional output. Alternatively, if costs are not being minimized, the firm could increase the amount of output it produces from the same amount of inputs, thus creating a benefit with no additional costs (except for any adjustment costs). Again, a reallocation of inputs when a firm is productively inefficient would satisfy the Kaldor-Hicks criterion, since someone in the economy would benefit from the increase in output.

The productive efficiency of a market outcome is sometimes also used to refer to the minimization of the industry's total cost of producing industry output. Even if all firms are minimizing the costs of producing their own output, if one firm is producing at a higher marginal cost than another firm, then output could be re-allocated across firms to reduce the industry's production costs. If it is not possible to re-allocate production from firms with higher costs to firms with lower costs, then the market is productively efficient.

A project that results in a reduction in firms' production costs, or a re-allocation of production from higher-cost firms to lower-cost firms, will be approved under a CBA.

This assumes, of course, that the resources that must be expended in reducing production costs do not exceed the gains.⁵

1.6.3. Dynamic efficiency

Allocative and productive efficiency are *static* concepts: they are used to assess the properties of a given allocation of resources at a point in time. It is also important to assess *dynamic* efficiency, which refers to the efficiency of the framework for decision-making over time. A market promotes dynamic efficiency if producers and consumers are guided by the appropriate incentives to take actions that promote efficient evolution of the market. For example, a market is dynamically efficient if producers have an incentive to make appropriate investments in the development of new products or production processes that will reduce costs or improve the quality of products or services. Such enhancements ultimately benefit consumers, either through increased utility or through lower prices.

1.6.4. Winners and losers

When economic resources are re-allocated as a result of a government project or change in regulations, some parties will gain, while others may lose. By definition, if the social benefits of the project exceed the losses, then the net gains exceed the net losses. This follows from satisfaction of the Kaldor-Hicks criterion in a CBA: if the criterion is satisfied, then the winners' gains are sufficient to fully compensate the 'losers' while retaining sufficient surplus so that the winners are better off than under the *status quo*. Recall that this criterion does not require that 'losers' actually be compensated, only that winners benefit sufficiently so that they would still be better off if they did compensate 'losers'.

⁵ For example, if the research resulting from a government-funded research program allows firms to reduce their production costs, then the costs of the program must be offset against the benefits of the cost reduction.

1.7. WEIGHTING EFFECTS ON DIFFERENT GROUPS DIFFERENTLY

Cost-benefit analysis focuses on economic efficiency while leaving aside social considerations, such as the effects of a project or policy on wealth distribution. To some extent, these other considerations can be dealt with using transition or compensation mechanisms, as we discuss in more detail below. However, cost-benefit analysis can itself be modified to account for other factors, by treating changes in welfare experienced by different groups differently. This can be done by assigning different weights to different individuals or groups of individuals; for example, the change in the welfare of low income individuals can be accorded more weight than changes in the welfare of higher-income individuals, or changes in the welfare of consumers as a group can be given more weight than changes in the welfare of producers (as we explain below, some competition enforcement agencies put more weight on consumer surplus than producer surplus when assessing the economic effects of mergers).

Consider an example where the effects of a policy on consumers are deemed to be more important than the effects of the policy on producers. When calculating the sum of costs and benefits across all groups, more weight can be put on the change in consumer surplus than on the change in producer surplus. Thus, if a policy results in a net reduction in consumer surplus of \$10 and a net increase in producer surplus of \$15, then it passes a strict cost-benefit analysis, but if the decision standard is modified so that changes in consumer surplus are given twice the weight of changes in producer surplus, then the policy is not approved.

To see this, notice that using a traditional (i.e. non-weighted) decision standard, a policy is 'approved' by a cost-benefit analysis if

$$(\text{Change in Consumer Surplus}) + (\text{Change in Producer Surplus}) > 0^6$$

⁶ In this example, we make the simplifying assumption that there are only two groups of individuals—namely, consumers and producers—and the surplus of every individual within each group is given the same weight, so that the change in surplus for each group is just the sum of changes in surplus in each group. Alternatively, we can assume that there are only two individuals: one consumer and one producer.

It is easy to see that, in the traditional case, the policy will be approved as long as the sum of changes in consumer and producer surplus is greater than zero. In the example where consumers surplus is reduced by \$10 and producer surplus increases by \$15, the policy is approved because $(-\$10) + \$15 = \$5 > 0$.

Now suppose that consumer surplus is given twice the weight of producer surplus. The decision rule is then to approve the policy if:

$$2/3 \times (\text{Change in Consumer Surplus}) + 1/3 \times (\text{Change in Producer Surplus}) > 0$$

Notice that the weight on consumer surplus, $2/3$, is twice the weight on producer surplus, $1/3$. This inequality can be rearranged to show that the policy will be approved only if the increase in producer surplus is more than twice the reduction in consumer surplus.

A closely related procedure, sometimes referred to as 'balancing weights' or 'relative weights', can be used to favour particular individuals or groups in a cost-benefit analysis.⁷ As an example, suppose that there are two groups of individuals to be considered, namely consumers and producers, and further suppose that implementing a particular project would reduce consumer surplus and increase producer surplus, and that the decision maker seeks to protect the welfare of consumers (perhaps because they are thought to have lower incomes than producers). Under the balancing weights procedure, the minimum weight on consumers' surplus that would be needed to tip the cost-benefit analysis in favour of consumers (that is, to prevent implementation of the policy) is first calculated. The decision-maker then determines whether there is evidence to support weighing consumer surplus at this minimum level.

⁷ This procedure was used by the Competition Tribunal in the *Superior Propane Redetermination Decision*, as discussed briefly below.

To make this more concrete, consider again the above example. The balancing weights approach would first calculate the value of the parameter α that satisfies the following equality:

$$\alpha (\text{Change in Consumer Surplus}) + (\text{Change in Producer Surplus}) = 0$$

or, using the values in the previous example, $\alpha (-\$10) + (\$15) = 0$. After rearranging, we obtain $\alpha = 1.5$. If the weight on consumer surplus is at least 1.5 times greater than the weight given to producer surplus, then the balancing weights procedure would tip the cost-benefit analysis in favour of consumers. The decision-maker then determines, using some procedure outside the cost-benefit analysis (for example by using some indicator of the socially tolerable level of income inequality), whether consumer surplus should in fact be weighted by at least 1.5 times as much as producer surplus.

When implementing a weighted cost-benefit analysis or using balancing weights, care must be taken to properly identify individuals as either being in the 'favoured' group or not. For example, consider a potential project that would increase the price of electricity in some areas, and would increase the profits of electricity producers. If the decision-maker seeks to favour consumers by increasing the weight on consumer surplus in a cost-benefit analysis, she must determine whether she seeks to favour all purchasers of electricity in the affected areas, or just final consumers. A significant proportion of direct purchases of electricity (i.e. purchases in the wholesale market) are accounted for by firms that use electricity as an input into the production of their final products, which they then sell to consumers. Some of the increase in the price of electricity will likely be passed on to final consumers in the form of higher output prices—potentially leading to a loss in *consumer* surplus—but some of the price increase will be absorbed by the firm, in the form of lower profits, which results in a loss of *producer* surplus. If the decision-maker seeks to favour only final consumers, then any reduction in producer surplus due to higher prices paid by direct purchasers of electricity should not be given the higher weight accorded the effect on the surplus of final consumers.

We note also that it is problematic to maintain a claim that the loss in producer surplus experienced by direct purchasers of electricity should be given the same

weight as the change in surplus to final consumers, and consequently more weight than the change in surplus experienced by electricity producers. While there may be a public policy rationale for placing more weight on consumer surplus, based on concerns relating to wealth distribution (i.e. increases in producer surplus benefit the shareholders of the benefiting firms, who are likely to be wealthier, on average, than the general population, from which the consumers of electricity are drawn), there does not appear to be a compelling public policy rationale for favouring changes in surplus to firms that are direct purchasers of electricity over changes in surplus to producers of electricity.

An additional complication arises when an affected producer is publicly owned. If, for example, OPG's producer surplus increases when a project is implemented, the beneficiaries are primarily Ontario taxpayers, because OPG is owned by the province. When a project increases electricity rates, thereby reducing consumer surplus, and also increases OPG's producer surplus, there is a significant overlap of winners and 'losers'. Furthermore, it becomes difficult to maintain a higher weighting for consumer surplus over producer surplus, as this would imply that ratepayers' benefits when they are received in their capacity as purchasers are given more weight than their benefits that are received in their capacity as citizens of Ontario.

2. EXAMPLES OF THE USE OF COST-BENEFIT ANALYSIS

Cost benefit analysis is widely used by government agencies in Canada and other countries to at least inform their analyses of regulatory change. Here we discuss cost benefit analysis as applied by the Ontario Energy Board, the Canadian Competition Bureau, the Treasury Board of Canada, the U.S. Environmental Protection Agency, and the government in the United Kingdom.

2.1. ONTARIO ENERGY BOARD

A recent Ontario Energy Board (“OEB”) Staff Discussion Paper⁸ proposes a test for evaluating transmission infrastructure investments, which is essentially the Kaldor-Hicks decision standard used in cost-benefit analysis. The test approves a proposed transmission investment (among a set of ‘alternative options in a majority of reasonable scenarios’) if it maximizes the (present value) of market benefit less costs. In the proposed test, only *direct* costs and market benefits to consumers, producers, and distributors of electricity are counted, while indirect costs and benefits, such as those related to the effects of an investment on the Ontario economy as a whole, are not included in the analysis. The discussion paper makes it clear that transfers among consumers and producers are not to be counted as either costs or benefits, which implies that the effects of an investment on consumers and producers are weighted equally (in other words, a ‘traditional’ non-weighted cost-benefit analysis is proposed). The paper also notes that California, Australia, and New Zealand also use cost-benefit analyses when evaluating proposals for transmission investments.

2.2. CANADIAN COMPETITION BUREAU

The Competition Bureau is the federal government agency responsible for enforcing the Competition Act. The Act’s merger provisions prohibit mergers that substantially lessen competition, which is interpreted to mean that a merger will be prohibited if it is likely to enhance the market power of the merging firms. However, the Act also includes an “efficiencies exception” that permits the consummation of otherwise anti-competitive mergers under certain conditions. In particular, a merger that is likely to result in significant cost savings may be approved, if the cost savings “are greater than, and offset” the likely anti-competitive effects.

The Competition Bureau has historically (although not consistently) interpreted the efficiency exception as mandating that a merger be prohibited only if the social costs

⁸ Ontario Energy Board, “Development of Filing Requirements for Transmission Infrastructure Investment (EB-2005-0470)”, *Staff Discussion Paper for Technical Advisory Team*.

of the merger, taking the form of higher prices to consumers, exceed the benefits of the merger, in the form of cost savings. The consumer harm from higher prices is calculated as the loss in consumer surplus, which is closely related to compensating variation. The cost saving to producers is calculated as producer surplus, which is the production-side analogue of compensating variation. A merger is prohibited if and only if the loss to consumers (measured as the loss in consumer surplus) exceeds the gain in producer surplus. This is the Kaldor-Hicks criterion, since it will challenge a merger if and only if the 'winners' from a merger (producers) gain enough to hypothetically compensate 'losers' (consumers), without actually requiring them to do so.

In an important contested merger case (Superior Propane), the Competition Bureau advocated the use of 'balancing weights', but the Competition Tribunal initially rejected this approach and allowed the merger--which was expected to generate substantial efficiencies--on the basis of a total surplus standard. The Bureau appealed the Tribunal's decision, and the Federal Court of Appeal granted the Bureau's appeal, remanding the case to the Tribunal and directing it to reconsider the efficiency defense, and in particular, to consider a broader range of anti-competitive effects arising from the merger. In its redetermination decision, the Tribunal increased the weight on the effects of the proposed merger on the lowest-income quintile of consumers, relative to the weights on the effects on producers and wealthier consumers. In effect, the Tribunal moved from a 'traditional' cost-benefit analysis in its initial decision, to a weighted cost-benefit analysis in its redetermination decision, in order to address income distribution concerns as it was directed to do by the Federal Court of Appeal.⁹

2.3. THE U.S. ENVIRONMENTAL PROTECTION AGENCY

The U.S. Environmental Protection Agency (EPA) is authorized by Congress to adopt regulations under environmental and public health protection laws. A preliminary analysis of a proposed regulation is published in the Federal Register to

⁹ See Michael Trebilcock and Ralph Winter, "The State of Efficiencies in Canadian Merger Policy", *Canadian Competition Record* (Winter 1999).

elicit comments from the public in accordance with the Administrative Procedure Act before being finalized.

The EPA is bound by several laws and executive orders that serve to guide the federal regulatory development process. One of these is Executive Order 12866 which calls for the EPA to execute economic analyses based upon a collection of principles that make the regulatory system operate well for the American people. Performing such analyses requires the consideration of benefits and costs of various alternatives to a particular form of regulation, including that of non-regulation.¹⁰

EPA evaluates regulations by addressing the following questions: Is it theoretically possible for the “gainers” from the policy to fully compensate the “losers” and still remain better off?; who are the gainers and losers from the policy and what are the associated economic changes? and; how did a particular group—especially a group that may be considered to be disadvantaged—fare as a result of the policy change?

These questions are addressed by employing cost-benefit analysis to determine the net social benefits, economic impacts analysis (EIA) to establish winners and losers and equity assessment to examine subsets of the population, particularly those classified as disadvantaged.

2.4. THE REGULATORY IMPACT ASSESSMENT REQUIREMENT IN THE UK

The Cabinet Office of the United Kingdom describes the Regulatory Impact Assessment (“RIA”) as “a framework for analysis of the likely impacts of a policy change and the range of options for implementing it”. It is considered by the government to be an all-encompassing and flexible analytical tool because it lends itself to all forms of regulations, takes into account all possible outcomes and considers the parties affected. Once completed, RIAs are made publicly available.

The RIA is considered to be a vital tool in providing improved regulation that is intended to be aligned with the government’s goal of regulating only if necessary and

¹⁰ United States Environmental Protection Agency, “Regulation: A Vital Tool for Protecting Public Health and the Environment”, EPA100F-03-001, July 2003.

to regulate proportionately to the inherent risk. All policy changes, whether European or domestic, that could potentially affect the public or private sectors, charities, the voluntary sectors or small businesses must be accompanied by an RIA. Five Principles of Good Regulation have been identified by the Better Regulation Task Force; regulation must be proportionate (to the risk); accountable (to ministers and the Parliament, to users and the public); consistent (predictable, so that people know where they stand); transparent (open, simple and user-friendly), and; targeted (focused on the problem, with minimal side effects).

Estimated costs and benefits are the chief analytical components of the RIA,¹¹ and will be critical in deciding if the benefits justify the additional costs of imposing regulation. Both direct and indirect effects are considered, as are distributive changes within and between groups. Policy effects are classified according to whether they are economic, social or environmental impacts.

3. DEALING WITH 'WINNERS' AND 'LOSERS'

Government projects and policies almost always benefit some members of society while harming others. For example, when price regulation is introduced, consumers typically benefit (at least in the short term) from lower prices, while regulated firms suffer a loss in profits. When price regulations are removed or relaxed, prices can increase, and consequently consumers are harmed and producers benefit. Note however, that if the project passes a cost-benefit test, the 'winners' in each case must benefit by more than the 'losers' are harmed.

As discussed above, projects that satisfy a cost-benefit standard may cause a redistribution of wealth that is deemed socially undesirable. We have explained how a cost-benefit analysis can be modified to account for wealth distribution concerns by weighting the effects of a project on favoured groups (e.g. lower-income households) more heavily than the effects on other groups. An alternative (or complementary) approach compensates individuals or groups harmed by a project. Government

¹¹ The Cabinet Office, "Regulatory Impact Assessment Guidance", obtained from the Cabinet Office website, www.cabinetoffice.gov.uk/regulation/ria/ria_guidance/index.asp, retrieved 14 November 2006.

policies to alleviate adverse distribution consequences may also be used to solidify political support from those harmed from a change. Re-training and compensation programs for workers harmed by trade liberalization policies are examples of government responses to significant changes that are aimed at addressing adverse distribution consequences and garnering political support.

Consider the effect of tariffs and quotas on consumers and workers. These policies restrict the volume of imports, and their main purpose is to protect the jobs of workers in economic sectors that are vulnerable to imports. Studies have concluded that, in many instances, the jobs 'saved' by import protection provide benefits for workers (in the form of continued employment in their current jobs) while harming consumers (in the form of higher prices for imports and domestically produced goods, and reduced quality and variety). In many cases the cost to consumers per job 'saved' exceeds--often by a substantial margin--the earning of workers in the protected industry. To cite just two examples, it has been estimated that voluntary export restraints (essentially a type of import quota) in the U.S. automobile sector in the 1980's cost U.S. consumers US\$105,000 to \$241,235 annually per job saved. Voluntary export restraints on automobiles in Canada have been estimated to have cost Canadian consumers C\$179,000 to \$207,166 per job saved.¹² In both cases, the costs to consumers clearly exceed workers' earnings (which are probably overestimates of the benefits of import protection, since many workers can presumably obtain employment in other economic sectors, although at lower wages). In instances such as these, there is considerable scope for enhancing economic efficiency by eliminating trade restrictions. At the same time there is an argument for compensating workers in some form, on distributional grounds and also to gain the support of workers (and their unions) to ensure that they do not use their political power to prevent the reforms. As discussed below, monetary payments and/or various forms of transitional assistance have been used in the past for these purposes.

¹² See Michael Trebilcock, Marsha Chandler, and Robert Howse, *Trade and transitions: A comparative analysis of adjustment policies*, London: Routledge (1990), pp. 55 and 67; and Edward Iacobucci, Michael Trebilcock, and Huma Haider, *Economic Shocks*, Toronto: D. Howe Institute (2001).

Governments respond to significant changes in the economic circumstances of individuals using both *ex ante* and *ex post* policies. *Ex ante* policies are put in place before a significant change occurs, to alleviate the negative impact of the change. *Ex post* policies are implemented in response to a particular change; examples include relief payments when natural disaster occur. It is important to note that government policies that are aimed at correcting undesirable wealth redistribution can often make matters worse, by creating inefficiencies that at least partly offset the efficiencies projected to occur from the project.

There are a variety of programs and mechanisms that can be used by governments which can both compensate 'losers' from government projects and also gain their political support. Below we discuss ways that can be used to compensate and/or facilitate transitions for workers and firms.

3.1. MECHANISMS TO EASE TRANSITION FOR WORKERS

There are several types of programs available that can be used to facilitate the transition from activities that are displaced by government projects/changes to regulations.

Job Search Assistance: Includes subsidizing job-search facilities that collect and post job listings, and counseling and advice on jobs, education, and skills development.

Relocation Assistance: Displaced workers can be provided with financial assistance to relocate to regions with more job vacancies.

Training Programs: These can include on-the-job or classroom programs that provide workers with the skills required for new jobs.

Also available are Government programs that compensate workers for economic changes, including employment insurance, which provides assistance to unemployed workers regardless of the reasons for job loss. Employment insurance can create an incentive for workers to inefficiently refrain from working, since it reduces the cost of leisure. But it also provides a temporary safety net for the period of transition between job loss and a new job, and it can also facilitate productive job search.

Special income benefits have sometimes been provided to workers dislocated by long term structural change, for example as a result of trade liberalization. Such benefits are often justified on the grounds that displaced workers should not have to bear the costs from structural adjustments that benefit the economy overall. However, it has also been argued that it is not clear that groups affected by particular events deserve special treatment, and that all workers unfortunate enough to be displaced should be assisted to an equal extent.

3.2. MECHANISMS TO EASE TRANSITION FOR FIRMS

Economically efficient compensation mechanisms incorporate lump-sum payments or transfers because, unlike per unit subsidies, they do not distort incentives for efficient decision-making. For example when a change in regulations or market rules in an electricity market results in higher prices for some consumers (this may occur, for instance, when locational marginal pricing (LMP) is introduced in a market that previously had uniform pricing), compensation payments should not depend on the amount of energy produced or consumed in the future. Prices in an electricity market (as in other markets) act as signals for efficient production, consumption, and investment. In a market that has adopted LMP, for example, the price in a given region equates the marginal benefit of consumption in the region and the marginal social costs of supplying energy to those consumers, including production costs, losses, and the effects of transmission congestion. Such a price ensures allocative efficiency because consumers consume up to the point where their marginal benefit is equal to marginal social costs of energy. It also promotes dynamic efficiency by providing the appropriate signal for investing in generation and consumption location: returns to generation investment are higher in high-price regions, where new generation is required; and consumers (in particular, electricity intensive industrial loads) have an incentive to locate in low-price regions.

A compensation scheme that lowers the prices paid by consumers that are 'losers' as a result of adoption of LMP or increases the price received by 'losing' generators drives a wedge between true marginal benefits and marginal costs, and thus creates inefficiency. Lump sum transfers, on the other hand, do not affect the prices paid or received by market participants (or more generally, do not depend on the quantities

consumed or produced) and therefore do not create incentives for inefficient behaviour.¹³ Note that compensatory payments that depend on *past* levels of consumption or production do not create inefficiencies because they do not affect prices facing market participants and therefore do not affect future consumption or production decisions.

In the following, we outline some compensation/transition mechanisms for consumers and producers in energy markets, which may become desirable because of changes in market rules. We do not discuss in any detail the forms of direct compensation except to note, as above, that if direct payments are made, then in order to promote efficiency, these payments should be lump-sum transfers or at least not tied to future levels of production or consumption.

3.2.1. Adjusting the Timing of Implementation and Compensation

One approach to mitigating the impacts of an efficiency-enhancing change in market rules on 'losing' market participants is to delay implementation to a fixed future date. If implementation is delayed (for example, for a period of three years), companies that are likely to suffer adverse impacts will be temporarily shielded, allowing them to investigate and implement loss mitigation strategies. If future losses are heavily discounted, then the present value of lost profits will be significantly lower than if implementation were to occur immediately. For publicly traded companies, this will reduce the impact on market capitalization. Private companies would also experience a lesser impact on asset values. Delay of implementation, however, would tend to reduce the present value of net social benefits of the change in market rules, since the stream of benefits would be pushed back in time.

In a similar vein, if it is contemplated that compensation to 'losers' is to be paid over time, then front-loading compensation by increasing payments in earlier years and

¹³ Any inefficiencies resulting from quantity-dependent compensation schemes would need to be taken into account in the cost-benefit analysis.

reducing them in future years, would increase the present value of the compensation stream to generators or consumers.

Delay in implementation may provide time for consumers and generators to adjust more effectively and efficiently to the change in pricing regime. For example, consumers may be able to adjust their operations so that they use less electricity and more of other forms of energy, or they may be able to substitute capital for energy. Such adjustments would generally require investment that would also require time to implement.

Indeed, the underlying cost-benefit analysis which assesses the viability of a change in market rules should take into account the timing of implementation since rapid adverse changes in a firm's business environment usually lead to higher adjustment costs than slower and predictable changes.

3.2.2. 'Grandfathering'

In some cases, it may be socially desirable to exempt current market participants from compliance with a new policy. The benefits of such 'grandfathering' can be significant when a firm previously made investments to comply with the existing policy, and where it must make significant additional investments to comply with the new policy. The costs of these additional investments to the firm (to which may be added various adjustment costs) may outweigh the incremental social benefits of the new investment, thus giving rise to an efficiency rationale for exempting the firm from the policy. At the same time, it may be efficient to require that new firms, and firms that have not made investments to comply with the previous policy, to comply with the new policy, so that firms may be treated differently, depending on whether they already have sunk investments in place.¹⁴

¹⁴ For a discussion of this rationale for grandfathering, see Steven Shavell, 'On Optimal Legal Change, Past Behaviour, and Grandfathering', Discussion Paper No. 57, John M. Olin Center for Law, Economics, and Business, Harvard University (2006).

3.2.3.LMP Example: Compensation through Allocation of Financial Transmission Rights

Compensation schemes may also be available that are suited to particular changes to market rules. As an example, consider the case where an electricity market operator adopts an LMP regime along with a system of Financial Transmission Rights (“FTRs”) as a hedging mechanism for market participants. In a uniform pricing system, participants are not exposed to price risks arising from transmission congestion (since transmission costs are folded into the uplift which is averaged over all consumers). However, under an LMP regime, congestion rents are included in LMP prices and a mechanism which allows hedging of these risks, is desirable. Many jurisdictions that have adopted LMP regimes have implemented FTR systems in part for this reason.

The owner of a share of an FTR associated with a particular “source-sink” pair has the right to a stream of revenues, which consists of the stream of congestion rents arising from the difference between prices at the two locations. FTRs thus have monetary value, and the allocation of FTRs, or the revenues arising from their allocation, can be used to compensate ‘losers’.

A number of options are available to the market operator for allocating FTRs, including direct allocation, and allocation through auction or through a competitive bidding process. Direct allocation of FTRs to ‘losers’ from LMP would provide a stream of revenues to fund LMP losses. Alternatively, losses could be offset by revenues from the sale of FTRs. The other two methods of allocation, auction and competitive bidding, generate revenues for the market operator which could be used to compensate losers directly.

4. SUMMARY AND CONCLUSIONS

This paper provides an overview of cost-benefit analysis, explains how it can be implemented, and briefly discusses some mechanisms for compensating individuals harmed by a policy and easing their transition to the post-policy world.

CBA is a decision standard with an associated range of methodologies for implementing the standard. Except in the simple (and very rare) case where a policy decision benefits all individuals equally and harms no one, policy makers are required to compare the benefits of the policy to “winners” with the costs borne by others, and then decide whether the benefits are worth the costs.

The first step is to calculate the benefits to individuals who gain from a policy and the costs to individuals who are harmed. Benefits and costs can be either monetary or non-monetary, but to facilitate comparisons all the effects of a policy are monetized, using various methodologies discussed in this paper.

Under the Kaldor-Hicks decision criterion that underlies CBA, a policy change merits approval if the monetized benefits to individuals exceed the monetized costs; that is, if those who benefit from the change could, hypothetically, compensate the individuals harmed by the change, while still remaining better off. It is in this sense that the benefits exceed the costs. CBA does not require that compensation actually be paid, only that there is enough of a benefit created by the change to offset the costs.

In its basic form, a CBA based on the Kaldor-Hicks decision criterion does not take into consideration any adverse wealth distribution effects. In order to address such concerns, the criterion can be modified to increase the weight placed on the costs or benefits to favoured groups, such as the less wealthy. Alternatively, adversely affected individuals can be compensated, either through existing mechanisms such as employment insurance and publicly funded retraining, or through specialized programs aimed at the compensating the ‘losers’ from a particular policy change. In this paper, we discussed delayed implementation, grandfathering, and allocation of financial transmission rights as compensation mechanisms for firms harmed by changes to the electricity market structure. Compensation mechanisms can also be used to promote political support among individuals and groups who are harmed by the policy change.