WORKING PAPER SERIES

Alternative Measures of Welfare in Macroeconomic Models

Vipin Arora

December 2013

Summary

The impacts and costs and benefits of different policies and scenarios can be calculated in several different ways. The measure chosen often depends on the class of model employed and the purposes of the policy and/or study. The sum total of costs and benefits, or changes in costs and benefits, is termed welfare. Traditionally, EIA has used measures such as GDP, consumption, and unemployment (among others) as ways to describe the overall economic impacts of policies, mainly highlighting changes in consumption as a proxy for welfare.

Using a variety of different measures of welfare to evaluate policy changes is desirable. This is particularly true in the case of utility, which is unique because it can incorporate the direct and indirect costs and benefits of different policies. Given the assumptions and complications required to make welfare calculations using the current NEMS setup, using a CGE model to make such calculations is a good option. The feasibility of using a CGE model with NEMS and/or WEPS+ should be further explored.

The mechanics of welfare measures in macroeconomic models

Basics

Macroeconomic models commonly use utility, GDP, consumption, and unemployment (among other measures) as proxies for welfare. The utility measure is more popular in optimization-based macroeconomic models, while the other proxies are more common in non-optimization based macroeconomic models. In general, optimization-based models begin at the level of consumers and firms and build up to the macro-economy. In these models equilibrium is defined as a situation where consumers and firms do the best they can subject to their constraints, and all markets clear. Popular examples include both Computable General Equilibrium (CGE) and Dynamic Stochastic General Equilibrium (DSGE) models. Non-optimization based models may or may not build up from consumers and firms, but they do not clearly build on the optimization of agents. A popular example of a non-optimization based model is the Global Insight U.S. Macroeconomic model.

Measuring utility

One welfare measure which has not generally been used by EIA in conjunction with policy scenarios is utility. Utility is the usefulness or satisfaction an individual derives from an activity. In optimization-based models, the consumer's problem is framed as one of maximizing utility. The results of this optimization problem are demand functions for each good, which specify the quantity of each good demanded based on prices and income. The aggregation of individual demand curves for a particular good leads to a market demand curve. A particular market is in equilibrium when the market demand curve intersects with a corresponding market supply curve to yield a price where the quantity demanded equals the quantity supplied.

In optimization-based models, aggregate utility is based on these demand functions, in that it depends at least in part on the total consumption of various goods and services by each consumer. It is common for modelers to measure the changes in aggregate utility in the face of some policy change. While this gives a comprehensive look at the impacts of a policy, aggregate utility is an ambiguous concept.

It is ambiguous because the units of utility measures, utils, are difficult to interpret and compare with other measures. To overcome this, most impacts of policy changes in optimization models are based on the concepts of compensating and equivalent variation. These are two different ways of translating the effects of policy changes from welfare terms into income: they convert utils to dollars.

A policy change leads to a rise or fall in consumer utility which works through price changes. Compensating variation calculates the income required to "compensate" an individual for price changes resulting from a policy change. Using the prices after the policy change, compensating variation is the amount of income which changes utility by the same amount as the policy.[1] Equivalent variation is similar, but used more frequently because it is based upon the original prices. It calculates the income change which would make an individual "equivalent" to what they would be after price changes resulting from a policy change. Using the prices before the policy change, equivalent variation is the amount of income which changes utility by the same amount as the policy.[2]

Both compensating and equivalent variation are useful ways to capture the total welfare costs and benefits of a policy change in terms of income. However, they require the demand functions derived from utility maximization described above. Unfortunately, these are generally not available in non-optimization based models, although modelers can still calculate changes in utility due to policy changes in such models using the method of Harberger triangles. This requires assuming that the observed market prices and quantities are based on market supply and demand curves, and those demand curves are generated from utility maximization.

If this is the case, the area under the market demand curve can be interpreted as total utility for the individuals who consume that good. Aggregating over all markets gives the total utility for the current prices and quantities. The prices and quantities will change with the introduction of new policies. Total utility in each market and for all markets can again be calculated at these new prices and quantities due to the policy change. The difference in the two levels of utility gives the total welfare impact of the policy.

Strengths, weaknesses, and alternatives

Using utility as a welfare measure fully accounts for the direct and indirect costs (and benefits) of policies. This is primarily because utility captures the preferences of consumers, and allows calculation of how a policy impacts desired allocations as compared with those preferences. Utility also captures the costs and benefits of non-market activities, and can be aggregated into one number for large economies. There are some heroic assumptions, however, that are necessary to aggregate utility over individuals.

---

[1] It is similar to asking: at the new prices, how much income would an individual need to return to the level of utility achieved under the old prices?

[2] It is similar to asking: at the old prices, how much income would an individual need to get to the level of utility under the new prices? Another way to think of equivalent variation is that it is the income which the consumer would be indifferent about accepting/paying in lieu of the price change.

Because individual demand curves are based on income, they cannot easily be aggregated, as there is a distribution of individual demand curves for each income level. One way to overcome these aggregation issues is to assume that people are similar enough that they can be grouped under a ``representative'' agent, or multiple representative agents. This similarity requires specifying a social welfare function for the representative agent, which can be controversial. Additionally, it is also unclear exactly what utility is or is not. Finally, the utility maximization procedure may not be a good description of actual choice outcomes.

Non-optimization models often use GDP, consumption, and unemployment as measures of policy impacts. The benefit of using such metrics is that they are simple aggregates which most people are familiar with. In fact, these are often the aggregates which a policy is designed to impact. In general, the movements of these measures above or below their trend are correlated strongly, indicating they are strongly related. GDP is also strongly correlated with commonly accepted measures of well-being, such as health outcomes or literacy.

The major drawback of using these measures is that they do not capture the full costs or benefits of policy changes. This is because they cannot account for non-market activities or costs and benefits related to preferences (indirect costs and benefits). As with utility, aggregation of these measures presents complications as well. Table 1 summarizes the advantages and disadvantages of each measure or proxy for welfare.

Table 1. Comparison of common welfare measures and proxies

Aggregate utility	
Advantages	*Disadvantages*
1. Includes non-market activities.	1. What is it?
2. Comprehensive: Incorporates all costs and benefits.	2. Aggregation requires strong assumptions.
3. Summarized in one number.	3. How to interpret?
	4. Cannot be easily integrated into existing NEMS framework.
Compensating and equivalent variation	
Advantages	*Disadvantages*
1. Utility-based, but in terms of income.	1. Require demand functions.
2. Calculations are straightforward.	2. Cannot be easily integrated into existing NEMS framework.
Harberger Triangles	
Advantages	*Disadvantages*
1. Can be used with any model.	1. Calculations are complicated.
2. May be possible to integrate with existing NEMS framework.	2. Need to assume that market demand curve comes from utility maximization.
	3. Need full information on all preexisting distortions in each market.
	4. Assumes supply and demand functions are linear.
GDP and other proxies	
Advantages	*Disadvantages*
1. Simple and understandable.	1. Do not capture non-market activities.
2. Summarized in one number.	2. Cannot capture costs and benefits in relation to preferences.
3. May be the metric of interest.	3. Aggregation may not be straightforward.
4. Already used in NEMS.	

Welfare calculations and NEMS

There are two ways in which welfare calculations can be made in conjunction with NEMS. The first are sometimes called Harberger triangles. This method uses the existing setup and infers utility changes in each individual market as outlined above.[3] The benefit of this approach is that it does not require any substantial changes to the current NEMS framework. The drawback is that the calculations can be difficult, and they depend on two key assumptions.

First, linear supply and demand curves must be assumed in each market. This requirement may be fine for small policy changes, but is likely inappropriate in analyzing larger ones. Second, all pre-existing distortions in the model must be identified. In a model as complex as NEMS this may be a difficult task. Here, a distortion refers to some factor (such as a tax) which causes the price of a good to be above or below marginal cost.

The second approach is to build a CGE model to integrate with NEMS. The major benefit of this approach is that it is designed specifically for welfare calculations, and can provide estimates of both equivalent and compensating variation. The model also has theoretical consistency, possibly making interpretation of policy changes easier. Unfortunately, the drawbacks are large.

A CGE model cannot provide all of the variables currently found in the Global Insight model. In particular, this class of models is weak in representing the financial sector. Money is usually not an important component of CGE models either. While this may not be an issue if working exclusively in real terms, the impact of monetary policy cannot be adequately captured. The closure of a CGE model specifies the split between exogenous and endogenous variables. In larger CGE models varying this closure may give substantially different results, which is problematic because there is no commonly accepted rule for picking a specific closure.[4] Finally, CGE models are not known for their forecasting, but are used for policy analysis. To the extent that policy changes require forecasting, a CGE model may not be the appropriate tool for NEMS integration.

The use of different welfare measures in macroeconomic models at EIA

The preceding analysis makes clear than any approach has its strengths and weaknesses with respect to measuring and evaluating welfare in the face of policy changes. For this reason using a variety of different measures of welfare to evaluate these policy changes is desirable. This is particularly true in the case of utility, which is unique because it can incorporate the direct and indirect costs and benefits of different policies. Given the assumptions and complications required to make welfare calculations using the current NEMS setup, using a CGE model to make such calculations is a good option. The feasibility of using a CGE model with NEMS and/or WEPS+ should be further explored.

---

[3]There is no utility in NEMS, but the procedure gives a way to infer changes in utility if utility is assumed to underlie market demand, and if both supply and demand curves can be approximated as linear.

[4]For example, in an open economy model, one could choose to make the trade deficit exogenous and let interest rates change. But the reverse closure could also be used.

Appendix A: The Mechanics of Alternative Measures of Welfare in Macroeconomic Models

Overview

This document is the appendix to *Alternative Measures of Welfare in Macroeconomic Models*. The main text broadly surveys the use of utility as a welfare metric, and makes some comparisons with GDP and consumption-based approaches. It also explains the challenges of using utility as a welfare measure in the current EIA framework, the National Energy Modeling System (NEMS). Technical details are provided in the appendix.

The current EIA approach

The impacts and costs and benefits of different policies and scenarios can be calculated in several different ways. The measure chosen often depends on the class of model employed and the purposes of the policy and or study. The sum total of costs and benefits, or changes in costs and benefits, is termed welfare. Traditionally, EIA has used measures such as GDP, consumption, and unemployment (among others) as ways to describe the overall economic impacts of policies, mainly highlighting changes in consumption as a proxy for welfare. One welfare measure which has not generally been used by EIA in conjunction with policy scenarios is utility. This remainder of this document broadly surveys the use of utility as a welfare metric, and makes some comparisons with GDP and consumption-based approaches.

Utility

Overview

In the context of a macroeconomic model, utility is generally interpreted as the usefulness or satisfaction that an individual derives from an activity. Often this activity is consumption of one or more goods, but may include leisure as well. The calculation of utility or changes in utility can differ substantially between optimization-based and non-optimization based models. In most optimization-based models consumers are assumed to be utility maximizers, so utility theory is a key building block of such models. In solving this constrained optimization problem, consumers try to obtain as much utility as possible by their actions subject to any constraints that may be present (budgetary, informational, etc.).

It is common for modelers to aggregate many different types of goods and services into one or more composite consumption goods. For example, a small model might have two such goods: a service good and a manufacturing good. The consumer's optimization problem is to choose how much of each good to consume today and in the future given their constraints.[1] A solution to this constrained optimization problem is a demand curve for the consumer, which gives the quantity of a good demanded as a function of its price. The demand curve holds other variables, such as income or the price of other goods, fixed. If those variables change then the optimal decision of the consumer may do so as well.

The area under the consumer's demand curve can be interpreted as the total utility which the consumer gains from acquiring a certain amount of the good.[2] This interpretation follows directly from utility maximization, which shows that a condition for such a maximization is that the price of a good equals its marginal utility (sometimes called the marginal benefit) if all other variables are held fixed. Marginal utility is the utility gained from consumption of an additional unit of a good, and summing up n marginal utilities (i.e. adding up the area under the demand curve) gives the total utility for consumption of n units of the good.

These results are for one consumer, but macroeconomic models need to deal with entire populations. As a result, optimization-based macroeconomic models must aggregate up over every individual's demand to generate a market demand. This turns out to be problematic because the derivation of an individual demand curve assumes that all else is equal. In particular, each individual demand curve is derived based on a given level of income. These individual demands cannot simply be added up because individuals may have different income. That is, there is a distribution of market demand curves based on individual income. This is conceptually and computationally difficult -imagine tracking the different incomes of the 300 million people who live in the U.S. The most common work-around is to assume one or more representative agents in a model. These are consumers who have an income that is a sum of many (or all) individual incomes. For example, a representative consumer might stand in for all consumers, or a model might have two such consumers, one standing in for low-income individuals and the other for high income individuals. One can then take as given the total income of this representative consumer and derive a market demand curve, as the demand from this one consumer is the total of all demand.

Unfortunately, there are difficulties in assuming a representative agent as well. If the individual's demand curve is derived from solving a utility maximization problem, then the representative agent's

[1] See the appendix for more on optimization in the context of these models.

[2] This is not the consumer surplus, which is the utility the consumer gets less the price paid for the good.

market demand curve should be the result of such a maximization as well. This will require the agent have a utility function, which means the modeler must specify a function that incorporates the utility of all individuals. This is sometimes termed a social welfare function. But how to go from individual utility functions, which may be very different, to such an aggregate social welfare function?[3] The simplest way is to restrict the utility functions of all agents who will comprise the representative agent. The restriction amounts to forcing agents to have homothetic preferences: if the consumer is indifferent between 2 applies and 3 pears, they must also be indifferent between 4 applies and 6 pears.[4] In other words income does not change consumer preferences with respect to proportions, only aggregate quantities. This is obviously a restrictive assumption, but allows the derivation of a social welfare function and the corresponding aggregate demand curve.[5]

Measures

In optimization-based models the aggregate level of utility may be used as a measure of welfare. This number, however, does not provide much information that is easy to conceptualize. Because modelers are often interested in the impacts of a policy change, the change in aggregate utility can be used as well. This may give a better idea of how welfare varies with a policy change than absolute utility, but is still not easily understandable. A common approach is to couch changes in utility in terms of income by using the concepts of compensating variation or equivalent variation. The ideas behind each of these two concepts are similar, with each focusing on the welfare effect of a price change.

Suppose the modeler conducts a policy experiment, such as changing the tax on gasoline. This will result in price changes across a variety of goods and services (gas, cars, car parts, etc.) and likely change total utility for consumers. How much money would it take to compensate a consumer for these price changes? This is the idea behind compensating variation: the income needed to return the consumer to the old level of utility, given the new prices. In our example the compensating variation will be positive, but it can in general move either way. One can think of this as the amount that the consumer needs to be compensated for a price change.

Equivalent variation is the amount which the consumer would be indifferent about accepting/paying in lieu of the price change. How much would the consumer pay to avoid the tax on gasoline? They would pay an amount that would be equivalent to the price change in terms of its welfare impact. Equivalent variation is negative if the price change would make the consumer worse off (as with the gasoline tax),

[3]This question ignores the fact that comparing utilities across individual is meaningless. Because we cannot directly measure/observe utility, 10 utils to one person might be much different than 10 utils to another.

[4]Technically, a function f defined on a set K is homothetic if given any two elements of that set $\widehat{x}$ and $\widehat{y}$ with $f(\widehat{x}) = f(\widehat{y})$ then $f(t\widehat{x}) = f(t\widehat{y})$, $t > 0$.

[5]Several justifications exist for using this assumption. Some believe that many of the changes between individuals net out when aggregating up, others posit that for some issues income distribution is not important (i.e. international finance/trade), others concede the assumption is absurd but a necessary simplification to get at other important issues.

and positive if the consumer is made better off. Equivalent variation is used in cases where current prices are the best place to make a comparison.

This is often the case in CGE models, where the base year prices are consistent with actual observed data, which has been modified to the level of aggregation of the model. This makes equivalent variation a popular way to measure the welfare impact of policy changes in such models. Examples of its use include Paltsev et al. (2009) with MIT's Emissions Prediction and Policy Analysis (EPPA) model, various reports for the U.S. EPA using the Economic Model for Environmental Policy Analysis-Computable General Equilibrium version (EMPAX-CGE, see RTI (2008) for details), and various reports from Resource for the Future (RFF) (see for example Fischer and Fox (2009)) among others.

Calculation of compensating and equivalent variation require knowledge of consumer and market demand functions. The foundation of non-optimization based models is generally not utility maximization, so consumer demand functions based on such a procedure may not be available. However, every model will have information on demand, supply, and prices in each relevant market. If a modeler is willing to assume that these are the result of the intersection of supply and demand curves, and these demand curves are derived from utility maximization, non-optimization models can still provide an estimate of how utility changes in the face of different policy options (Krupnick et al., 2010). These changes in utility can be inferred from changes in equilibrium prices and quantities in the face of a policy change.

The basic idea uses the fact that the area under a demand curve can be interpreted as total utility. Then any set of equilibrium prices and quantities in each market has a given level of utility. The same is true of the new prices and quantities after a policy change. The difference in this utility in each market can be approximated by the area of a triangle. And the area of the triangle can be calculated if both the old and new prices are known.[6] The changes in utility can then be summed up over all markets to give the aggregate utility change due to a policy change. There are some key assumptions made here in order to generate this result. First, the market demand and supply curves which are implicit in the analysis are assumed to be linear. This may be inappropriate for large policy changes. Additionally, any previous policy distortions must be known, as this can impact the shape of the area representing the utility change.

Advantages and disadvantages

As a measure of welfare utility has many advantages. The most important is that it fully captures the impacts of any policy change on consumers. This includes the direct costs/benefits, as well as those costs/benefits which are more difficult to measure. For example, consider again the tax imposed on

[6] There are some caveats to this point, please see the appendix.

gasoline. A direct cost of this to consumers is the money they pay for that tax, and there may be some direct benefit if the tax revenue is used to improve roads, for example. However, a change in utility would also capture the less-obvious cost to consumers of driving less than they might otherwise prefer, or the benefit of spending less time in traffic because fewer other people are driving. Calculation of such less-obvious costs/benefits requires a way to quantity individual preferences, and utility provides such a mechanism. Another benefit of using utility is that it captures the costs and benefits of non-market activities. Utility also gives one number which aggregates all of the costs and benefits of different policies for consumers.

The concept, however, is not without problems. First, it is unclear exactly what utility is or is not. Is it satisfaction? Happiness? This ambiguity leaves many modelers uncomfortable. Also problematic is the representative consumer assumption necessary to aggregate from an individual utility function to a social utility function (or from an individual demand curve to a market demand curve) in macroeconomic models. Does a social welfare function exist? Another issue is that aggregating to a representative consumer abstracts from many important questions about income distribution. A final point is that utility maximization may not accurately model the outcome of consumer decisions. And the assumptions underlying this procedure may also seem unreasonable. The appendix has additional details on these assumptions.

Integration with NEMS

There are two ways in which welfare calculations can be made in conjunction with NEMS. The first, an approach taken by (Krupnick et al., 2010), is to use the existing setup and infer utility changes in each individual market as outlined above. The benefit of this approach is that it does not require any substantial changes to the current NEMS framework. The drawback is that the calculations can be difficult, and they depend on two key assumptions. First, linear supply and demand curves must be assumed in each market. This requirement may be fine for small policy changes, but is likely inappropriate in analyzing larger ones. Additionally, it is unclear if anything like a market demand curve actually exists in NEMS. That is, it may be impossible to find a utility function from which the relationship between price and quantity demanded is derived in NEMS. This is the problem of integrability, see the appendix for additional details. The supply curves may not be smooth, as some likely have steps or possibly discontinuities. Second, all pre-existing distortions in the model must be identified. In a model as complex as NEMS this may be a difficult task. Here a distortion refers to some factor (such as a tax) which causes the price of a good to be above or below marginal cost.

The second approach is the build a CGE model to integrate with NEMS. The major benefit of this approach is that it is designed specifically for welfare calculations, and can provide estimates of both equivalent and compensating variation. The model also has theoretical consistency, meaning that interpretation of policy changes may be easier. Unfortunately, the drawbacks are large. A CGE model cannot fully replicate all of the variables currently provided by the Global Insight model. In particular, this class of models is weak in representing the financial sector. Money is usually not an important component of CGE models either. While this may not be an issue if working exclusively in real terms, the impact of monetary policy cannot be adequately captured. Finally, CGE models are not known for their forecasting, but are used for policy analysis. To the extent that policy changes require forecasting, a CGE model may not be the appropriate tool for NEMS integration.

GDP

Overview

GDP is an aggregate measure of macroeconomic performance usually output from policy simulations.[7] Economic models which use GDP as a proxy for welfare primarily focus on real GDP in order to make comparisons across time. Nominal GDP is the sum of the dollar value of consumption, investment, government spending, and net exports (the value of exports less imports).[8] Specifically it is the price of goods and services used in each of these categories times the quantity.

Changes in nominal GDP can be due either to changes in price or to changes in quantities. Comparing the productive capacity of an economy across time requires looking at how quantities change. Real GDP does this by fixing prices so the changes in production can be isolated. Because there are many prices (both in reality and in many models), price indices are constructed to aggregate into one number, and nominal GDP can be converted to real GDP using such indices.[9]

Advantages and disadvantages

Using real GDP as a proxy for welfare has some advantages. Most importantly, it is a simple statistic understood by many people. This is a compelling argument, as other measures such as utility may not be so clear. Furthermore, movements in GDP above or below its trend growth rate (i.e. the business cycle) are strongly positively correlated with movements in employment, consumption, investment, and labor productivity. Using GDP as a welfare proxy may also shed light on these variables. Other measures of well being such as positive health outcomes or literacy, which are difficult to model, also have positive correlations with GDP. Using GDP as a proxy for welfare may also provide some information on these other variables.

[7] See for example many reports from the Congressional Budget Office (CBO) on policy proposals: `http://www.cbo.gov/publications/`.

[8] This is called the expenditure method of calculation. Nominal GDP can also be calculated using the income method (the dollar value of income paid to all factors of production plus profits), or the value added method (the changes in the value of a goods over different stages of production). Theoretically, these measures should all equate, though they often fail to do so due to measurement error.

[9] Examples of price indices include the Consumer Price Index (CPI) and variants, the Personal Consumption Expenditures (PCE) index, and the Producer Price Index (PPI), among others.

There are several problems with using GDP as a measure of welfare. Importantly, GDP does not measure non-market activities, it only accounts for paid work. This excludes important activities such as unpaid child care or unpaid care for elderly parents, which are generally grouped as "leisure" (Krupnick and McLaughlin, 2011). Using GDP as a proxy for welfare also abstracts from how income is distributed. To the extent that this is important to the modeler, GDP is unable to provide any information on the distribution of income. In a multi-country model, comparisons of GDP across the different countries can be problematic. To make such a comparison, at least one of the GDP calculations must be converted at market exchange rates, which may not capture actual purchasing power.[10] Finally, some also argue that GDP does not consider environmental sustainability, and may not capture environmental damage.

Consumption

Overview

As with GDP, macroeconomic models which use consumption as a proxy for welfare generally use real consumption. There are two ways in which consumption can be used as a welfare metric. The most common approach is to consider how different policy changes impact the aggregate value of real consumption. The level of disaggregation a modeler uses depends on the question of interest. For example, the consumption changes of agents in different income levels may be used, or consumption of various types of goods and services may be more appropriate. However it is used, using consumption as a proxy for welfare in this way is closely related to GDP, but with more of a focus on the consumer. This approach can be used in optimization and non-optimization based models.

The other option is used in small optimization-based models and relies on utility (Lucas, 1987). This approach asks how much of a percentage increase in consumption a representative agent requires as compensation for volatility in their consumption. In other words, how much does this consumer need to be paid (in terms of consumption) so that their utility is the same between a volatile or non-volatile consumption stream? This can be modified and applied to policy changes as well. Take for example a tax on gasoline. How much additional consumption would a representative consumer require to compensate them for the tax? This would not generally be the same as the value of the tax. This is because the tax has changed consumer behavior from what was previously optimal, so there must also be compensation for this as well. Notice the similarity with compensating variation, but focused on consumption volatility instead of price changes.

[10] Some models correct for this by assuming that the law of one price holds. This states that a good trades for the same price, irrespective of the currency of denomination. For example, if the dollar buys 100 yen, and an apple in the U.S. costs one dollar, the law of one price states that the same apple must cost 100 yen in Japan.

Advantages and disadvantages

Using the level of consumption or changes in consumption as a welfare proxy has similar advantages as using GDP. It is a simple statistic that correlates well with other measures of well-being (and GDP). An advantage of using consumption is that it is something that is directly important to consumers, whereas GDP may or may not have that feature. This is because GDP also considers investment, government spending, and net exports. The main problem, as with GDP, is that consumption does not capture all costs/benefits associated with policies. Consumption also may not easily allow the modeler to distinguish between individuals of different income levels. Finally, a more subtle issue is that using an absolute level of consumption does not account for diminishing marginal value. It is likely the case that the millionth unit of consumption does not have as much value as the 100th to a consumer, but an aggregate level of consumption cannot distinguish between the two.

Calculating the consumption required to compensate an individual for volatility in their consumption stream does account for different costs and benefits which are not directly observable. It may also be a simpler statistic to understand when comparing two policies, as one is talking in terms of consumption instead of utility. However, the calculation itself is not simple and may be difficult to implement in larger models.

Other

There are many other proxies for welfare used, and these often depend on the questions of interest to a modeler (Krupnick and McLaughlin, 2011). Examples include the level of employment/unemployment, investment, the level of emissions, miles driven, the level of interest rates, and even government spending. In general, each of these measures is unable to account for indirect costs and benefits of a policy other than what is associated with that particular measure. For example, consider the impact of a policy change on the unemployment rate. The modeler can clearly assess how the policy change may increase/decrease this rate. But how to asses the impact on average hours worked per person? Did the policy change lead the structure of other industries to change? And how did it impact hours in those sectors? Are consumers as a whole better or worse off? Questions such as these can only be answered with a comprehensive measure of welfare such as utility. A substantial benefit of these other measures of welfare, however, is that they are often the metric of interest. Policy makers are understandably concerned with how a policy change impacts measures such as unemployment, and likely not as much with other indirect impacts.

References

Conlisk, John, "Why Bounded Rationality?," *Journal of Economic Literature*, 1996, *34*, 669–700.

Fischer, Carolyn and Alan K. Fox, "Combining Rebates with Carbon Taxes: Optimal Strategies for Coping with Emissions Leakage and Tax Interactions," Discussion Paper RFF DP 09-12, Resources for the Future 2009.

Harberger, Arnold C., "The Measurement of Waste," *American Economic Review*, 1964, *54* (3), 58–76.

Krupnick, Alan J. and David McLaughlin, "GDP, Consumption, Jobs, and Costs: Tracking the Effects of Energy Policy," Issue Brief 11-08, Resources for the Future 2011.

_ , Ian W.H. Parry, Margaret Walls, Tony Knowles, and Kristen Hayes, "Toward a New National Energy Policy: Assessing the Options," Full Report, Resources for the Future 2010.

Lucas, Robert E., *Models of Business Cycles*, Wiley-Blackwell, 1987.

Paltsev, S., J.M. Reilly, H.D. Jacoby, and J.F. Morris, "The Cost of Climate Policy in the United States," Joint Program Report Series 173, MIT Joint Program on the Science and Policy of Climate Change 2009.

RTI, "EMPAX-CGE Model Documentation," Interim Report, Prepared for the U.S. EPA 2008.

Varian, Hal R., *Microeconomic Analysis*, third ed., W.W. Norton and Company, 1992.

Further technical details

These technical details summarize some aspects of welfare measurement in macroeconomic models. The first section differentiates specifically between general types of macroeconomic models. This can be important because different classes of models may use differing measures for welfare (or as a proxy for welfare). The next section introduces the rational decision making process which underlies utility maximization in optimization-based models. To the extent that the demand curve in non-optimization based models is treated as a marginal utility curve, this procedure implicitly underlies that construction as well. The following section uses the results of the rational decision making procedure to derive the calculation of compensating and equivalent variation in detail. The final section details the specifics of how to calculate utility in non-optimization based models using the method of Harberger (1964).

Optimization and non-optimization based models

Macroeconomic models can be broadly differentiated between those which are based on consumer/firm optimization and those which are not. Optimization-based models have been popularly used since the late-1970s for policy analysis and include both partial equilibrium and general equilibrium models. Partial equilibrium models simulate certain markets or segments of an economy in isolation from other markets or sectors. For example, a model of the gasoline industry may be national in scope but exclude other industries in its structure. The general equilibrium concept incorporates all markets and segments of an economy. Specifically, these models fully describe the economic environment: individual consumers, firms, government, and specific factors that influence the decisions of these agents, such taxes or subsidies. Consumers and firms have both objectives and constraints, and they also have access to markets. An equilibrium is a situation where consumers and firms do the best they can (they optimize), and their decisions are coordinated through markets, subject to the actions of the government and any specific factors which influence their decisions. Popular examples of general equilibrium models used in policy analysis are computable general equilibrium (CGE) and dynamic stochastic general equilibrium (DSGE) models. These models are also often used for historical decompositions, although they are less popular for forecasting.

One implication of equilibrium is that markets clear (supply is equal to demand) at a given price. Most general equilibrium models go further and assume that the economy is continually in such an equilibrium and instantaneously adjusts (through changes in prices or quantities) to maintain the equilibrium given any changes in the environment.[11] For example, if the government increases spending on goods and services, the prices and quantities of those goods and services will instantaneously adjust to accommodate the increase.[12] The majority of general equilibrium studies

[11] These changes are often referred to as shocks.

[12] This will generally impact all other agents in the model as well, but the changes work through variation in the prices of goods and services.

which consider welfare analysis use utility as the underlying measure. This is because utility is central to the construction of the model: the goal of consumers is to maximize their individual utilities. All consumer decisions are either explicitly or implicity based on this optimization.

Models which are not based on consumer/firm optimization include macroeconometric models, reduced form models (often based on vector autoregressions (VARs)), and input-output models. Given their focus, these models may or may not specify the objectives of consumers or firms, but always rely on various theory or hypothesized relations in deriving equations. For example, the Global Insight Macroeconomic Model of the U.S. economy used at EIA incorporates Keynesian, Monetarist, and Classical economic theory in specification of the model. The output of a policy change in such models often considers a proxy for welfare such as GDP, consumption, or unemployment.

Rationality

In order to build their models many macroeconomists begin at the level of individual decision making. Such models may be referred to as "microfounded" and are based on the theory of consumer choice. A first approximation to consumer choice that many use is rationality, or the rational decision making process.[13] This is a three-step process the consumer uses when making decisions:

1. What is feasible?
2. What is desirable?
3. Choose the most desirable from the feasible.

Each of these steps are outlined in detail below, with the final step analogous to the consumer maximizing utility.

What is feasible?

To characterize feasibility, commodities, commodity bundles, the consumption set, and the budget constraint are used. Commodities are the primitive object in consumer choice. These are the goods and services available for purchase in the market. In the simplest case, it is assumed that there are a finite number equal to L, with each commodity indexed by $l = 1, 2, ..., L$. A commodity bundle (or consumption bundle) is a list of the amounts of the different commodities, denoted by the vector $\hat{x}$. The lth entry of this vector represents the amount of commodity l consumed. For example, suppose an economy has rice, wheat, and corn as its three commodities. One possible commodity vector is $\hat{x}_1 = [1 \quad 0 \quad 5]'$. This represents one unit of rice, no units of wheat, and five units of corn.

[13] This approach has been heavily criticized, but is still the predominant method for modeling consumer choice in optimization-based macroeconomic models. See Conlisk (1996) for an overview of the arguments.

Commodities differ based on their physical characteristics, time, and location.[14]

The different consumption bundles can be grouped into a consumption set, X. This finite consumption set has as its elements the consumption bundles which the individual can consume given any physical constraints, $X = \{\hat{x}_1, \hat{x}_2, ...\hat{x}_j\}$, where j can be large but is finite. For example, an individual cannot consume a hot dog in Chicago and Paris at the same time, so a consumption bundle with this combination could not be in the consumption set. The set itself is limited by physical constraints, but the consumer also has an income constraint. This is summarized by a budget constraint, which restricts the value of purchases by the consumer to be no larger than their wealth (w). Quantities consumed are contained in the individual consumption bundles, but calculating a value requires prices. A key assumption usually made is that the consumer takes the price for good l (p_l) as given.

The combination of the consumption set and the budget constraint yields a new set, called the (Walrasian) budget set, B. This new set has as its elements all consumption bundles for the consumer which are feasible and affordable.[15] The prices of these goods are assumed to be given. Figure 1 shows the budget constraint (the diagonal solid line) and the associated budget set (the area under the constraint) in the two-good case, with income denoted by M.

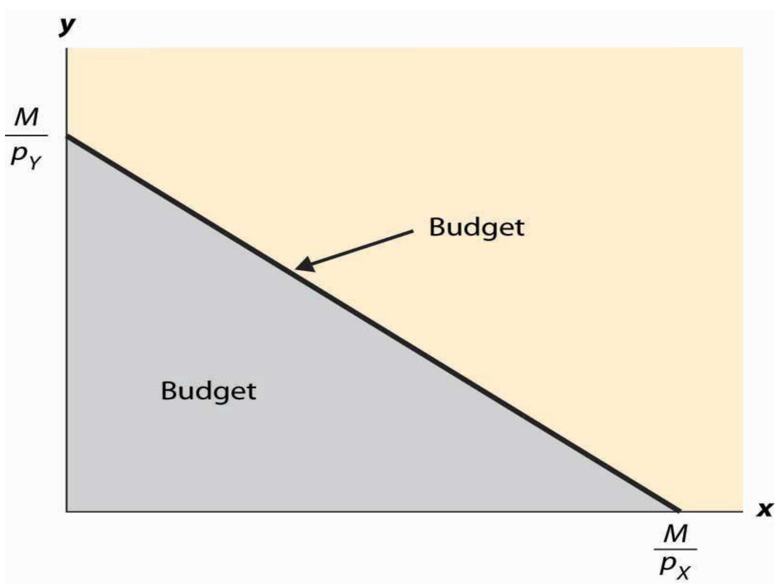

Figure 1: The consumer's budget constraint and budget set, from Varian (1992)

[14]As an example consider the same brand of hot dog made in Georgia. If one of these hot dogs is consumed on Sunday and another on Monday, they are considered different commodities. Furthermore, if one is consumed in Chicago and another in New York, they are also different commodities. These possibilities expand the number of commodities substantially.

[15]These are consumption bundles which already lie in X, and when associated with prices have a value less than the consumer's wealth.

What is desirable?

How to choose the most desirable bundles in the consumption set? In the classical approach to consumer choice, one first specifies the consumer's preferences over these bundles in order to compare between the bundles. The preferences are summarized by a preference relation ($\succeq$, read "at least as good as") which is assumed to be rational.[16] Rationality of the preference relation means that it has the properties of completeness and transitivity.[17]

A consumer's preferences are complete if they are able to compare any two consumption bundles. This means they either prefer one to the other, or are indifferent between the two. This assumption is stronger than it may seem at first glance, as consumption bundles can be very different, making comparisons difficult.[18] Completeness allows the consumer to compare and rank all bundles in the consumption set. Transitivity ensures consistency in these comparisons. This assumptions states that given any three consumption bundles $\hat{x}, \hat{y}, \hat{z}$ in X, if $\hat{x}$ is preferred to $\hat{y}$ and $\hat{y}$ is preferred to $\hat{z}$, then $\hat{x}$ must be preferred to $\hat{z}$.[19] Taken together, completeness and transitivity assume that the consumer can consistently rank the bundles in their consumption set according to preferences.

Two additional assumptions are usually made to ease computation in the third step of the process. The first is called non-satiation, which is the assumption that more is always preferred to less. This will allow the consumer to differentiate between bundles which are separated by only a small difference in commodity quantities. The other assumption is convexity of the preference relation. This amounts to stating that the consumer has a preference for diversity in their consumption bundle. It will allow a solution to the problem presented below.

The four assumptions give a complete characterization of what is desirable for the consumer. It would be much more useful, however, if the preference relation could somehow be quantified. This is done by introducing a familiar function, the utility function, which is a numerical representation of the consumer's preference relation.[20] To ensure that such a representation exists, continuity of the preference relation is also assumed. This means that the utility function will be continuous as well, and ensures there are no sudden jumps or breaks in preferences. One important note is that the utility function representing a preference relation is not necessarily unique. Any positive monotonic transformation of the utility function also represents the same preference relation. The implication is that the absolute level of utility output from any utility function is meaningless in and of itself. Interpretation requires comparison of other alternatives which come from the same utility function.

A useful way to summarize preferences graphically is by using indifference curves, as shown in Figure

[16] Technically, this is a binary relation on X.
[17] To many economists, these properties are what make a consumer "rational".
[18] Completeness: $\forall \hat{x}, \hat{y} \in X$, $\hat{x} \succeq \hat{y}$ or $\hat{y} \succeq \hat{x}$ or both.
[19] Transitivity: $\forall \hat{x}, \hat{y}, \hat{z} \in X$, if $\hat{x} \succeq \hat{y}$ and $\hat{y} \succeq \hat{z}$ then $\hat{x} \succeq \hat{z}$.
[20] The utility function is mapping from the consumption set to the real line, $U : X \rightarrow \Re$.

2. These curves plot (in the two-good case) the different combinations of commodities which give the same level of utility. In Figure 2, all bundles of the two goods along the curve U_1 give the same level of utility. And because more is preferred to less, $U_3 > U_2 > U_1$. Given the utility function, both feasibility and desirability can now be characterized analytically (and graphically in simple cases), leading to the final step in the rational decision making process.

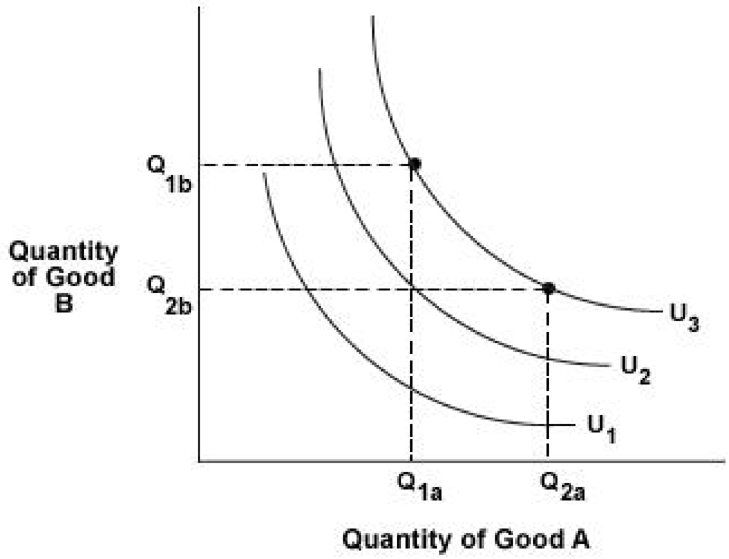

Figure 2: The consumer's indifference curves, from Varian (1992)

Choose the most desirable from the feasible

In this final step of the rational decision making process, the consumer chooses their preferred consumption bundle from the budget set. Given the steps above, this is the same as the consumer maximizing utility (choosing the most desirable) subject to their budget constraint (from the feasible). In the two good case, this can be written:

$$\max_{x_1,x_2} \quad U(x_1, x_2)$$

subject to:

$$p_1x_1 + p_2x_2 = w$$

The consumer chooses consumption of each commodity (x_1, x_2) to maximize utility $(U(x_1, x_2))$ given their budget constraint. Under appropriate assumptions outlined above, the unique solution to this problem will be a consumption bundle, represented by a pair of ordinary (or Marshallian) demand functions. These functions give the optimal demand for each commodity as a function of prices (p) and wealth. If the consumer has Cobb-Douglas utility ($U(x_1, x_2) = x_1^\alpha x_2^{1-\alpha}$) then the demand functions from this problem are:

$$x_1^*(p_1, p_2, w) = \frac{\alpha w}{p_1}$$

$$x_2^*(p_1, p_2, w) = \frac{(1-\alpha)w}{p_2}$$

In the two-good case this optimal bundle (x_1^*, x_2^*) can be seen graphically as shown in Figure 3. This will be at the point where the indifference curve is tangent to the budget constraint.

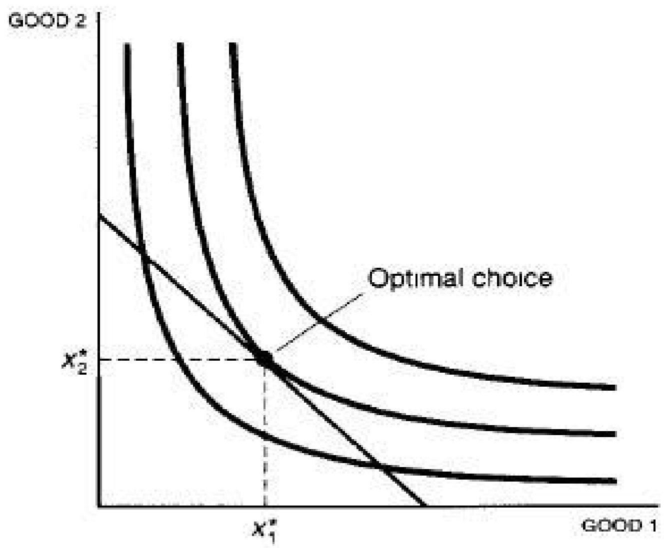

Figure 3: The consumer's optimal bundle, from Varian (1992)

The discussion has shown how a utility function representing a preference relation, under suitable conditions, can be used to derive individual demand curves. The main text described how these could be aggregated to yield market demand curves. Is it possible to move in the opposite direction? That is, given a market demand curve is it always possible to find a utility function (and so a rational preference relation) which would give such market demand? Not always. This is the integrability problem. Various authors have shown that if the market demand curve has certain restrictive properties, then it could have been generated by some rational preference relation.[21]

[21] The demand curve must be homogenous of degree zero, satisfy Walras' Law, and have a symmetric and negative semi-definite substitution matrix.

Calculating welfare changes

As outlined in the main text, many macroeconomists use a utility-based approach to assessing welfare. This can entail comparisons of aggregate utility changes in the face of a policy change. In optimization-based models, these can be found by calculating the new level of utility given the definition of the utility function, as is shown in the previous section. The alternative is to couch the impacts of the policy in terms of income, using compensating or equivalent variation, which are detailed in the next section. In non-optimization based models, changes in utility due to a policy change can be calculated using the method of Harberger (1964), which is detailed in the final section.

Compensating and equivalent variation

Calculation of either compensating or equivalent variation requires the ordinary demand functions from utility maximization, as derived in the previous section. These functions can be inserted into the utility function to give the indirect utility function, v. This function gives the maximum level of utility possible subject to income and prices. In the two-good Cobb-Douglas case from above, the indirect utility function is (where $\hat{p} = (p_1, p_2)$):

$$v(\hat{p}, w) = \left(\frac{\alpha w}{p_1}\right)^{\alpha} \left(\frac{[1-\alpha]w}{p_2}\right)^{1-\alpha}$$

While this function depends on prices and income, its units are in terms of utility. The higher the value of this function, the better off is the consumer. For example, consider a price change in either one or more of the goods from $\hat{p}^0$ to $\hat{p}^1$. The consumer is better off if $v(\hat{p}^1, w) - v(\hat{p}^0, w) > 0$, assuming that wealth does not change.

As before, this may be difficult to interpret because it is in terms of utility. To convert this to income terms (for use in compensating and equivalent variation) the expenditure function can be used. This function gives the optimal expenditure of a consumer in terms of prices and utility, $e(\hat{\hat{p}}, v(\hat{p}, w))$.[22] Here, $\hat{\hat{p}}$ is a vector of prices, which are not necessarily the same as the prices in the indirect utility function. If $\hat{\hat{p}}$ is chosen to be the same as $\hat{p}$, this function gives the dollar value of optimal consumption at the given prices (and utility). Any changes in utility will also change e, from which a difference in utility in terms of dollars can be calculated. This is the basic idea behind compensating and equivalent variation.

Consider a two-good example, where the price of one good falls, while the price of the other good and wealth do not change. This is also depicted in Figure 4 as a pivoting up of the budget constraint. The

[22]The expenditure function comes from the consumer's cost minimization problem, where the consumer minimizes their total costs in purchasing commodities subject to their utility. Duality theory shows that this problem is equivalent to the utility maximization described here.

expenditure function before the price change, if the general prices are chosen to be the same as the indirect utility function, is $e(\hat{p}^0, v(\hat{p}^0, w))$, or $e(\hat{p}^0, v^0)$. When the price of the good falls, this becomes $e(\hat{p}^1, v(\hat{p}^1, w))$, or $e(\hat{p}^1, v^1)$. One implication of static utility maximization (or cost minimization) is that a consumer will expend all of their income, so that expenditure will be equal to wealth. Because more is preferred to less, and there is no consideration for tomorrow, consumers will leave no income unspent. In our example, this means that before and after the price change all wealth is spent, or $e(\hat{p}^0, v^0) = e(\hat{p}^1, v^1) = w$.

As explained in the main text, the compensating variation is the dollar amount which compensates the consumer for the price change. This uses the prices after a policy has been implemented. In our current notation, it is equal to $e(\hat{p}^1, v^1) - e(\hat{p}^1, v^0) = w - e(\hat{p}^1, v^0)$. This is the expenditure amount at the new prices and the new utility less the expenditure amount with the new prices and old utility. In other words, the amount of income required to give the old utility (v^0) at the new prices (p^1). In Figure 4 this is shown in the second graph.

Calculation of equivalent variation uses the old prices, and is the dollar amount the consumer would be indifferent about accepting in lieu of the price change. This can be written as $e(\hat{p}^0, v^1) - e(\hat{p}^0, v^0) = e(\hat{p}^0, v^1) - w$. This is the expenditure amount at the old prices and new utility less the expenditure amount at the old prices and old utility. This is the first plot in Figure 4.

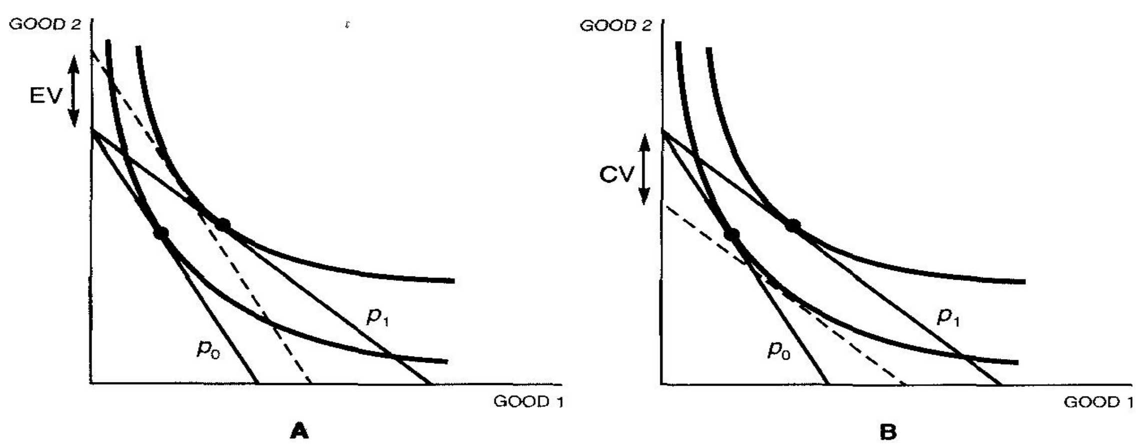

Equivalent variation and compensating variation. In this diagram $p_2 = 1$ and the price of good 1 decreases from p_0 to p_1. Panel A depicts the equivalent variation in income—how much additional money is needed at the original price p_0 to make the consumer as well off as she would be facing p_1. Panel B depicts the compensating variation in income—how much money should be taken away from the consumer to leave him as well off as he was facing price p_0.

Figure 4: Compensating and equivalent variation, from Varian (1992)

Harberger triangles

Different methods must be applied to calculate utility changes in models where utility maximization is not explicit. A common approach is based on the work of Harberger (1964), and approximates the change in utility given a change in policy using individual market prices and quantities. The discussion here is based on Appendix C of Krupnick et al. (2010). The basic idea is to consider each market individually and calculate the utility change in that market after a policy change. Utility can be inferred from the prices and quantities in each market if one assumes that the prices and quantities come from the intersection of supply and demand curves. Furthermore, the demand curve must be assumed to be a result of the aggregation of individual demand curves based on utility maximization, as described in the main text. Summing up the utility changes in all markets yields the total utility change for the policy. A key implicit assumption here is that the market demand curve from NEMS has the properties required by integrability described above. Without these properties it may be impossible for a utility function representing a rational preference relation to have generated the demand curve.

As an example, consider a tax on oil products, whose market supply and demand curves are shown in Figure 5. Point b in the figure is the starting equilibrium quantity (q^*) and the starting equilibrium price (p^*). The initial consumer surplus, which is the difference between the consumer's willingness to pay and what is actually paid, is represented by the triangle abc. The producer surplus, the difference between the marginal cost of production and the price received, is represented by the triangle cbh. Total surplus in this market is given by the sum of the consumer and producer surplus, or the triangle abh.

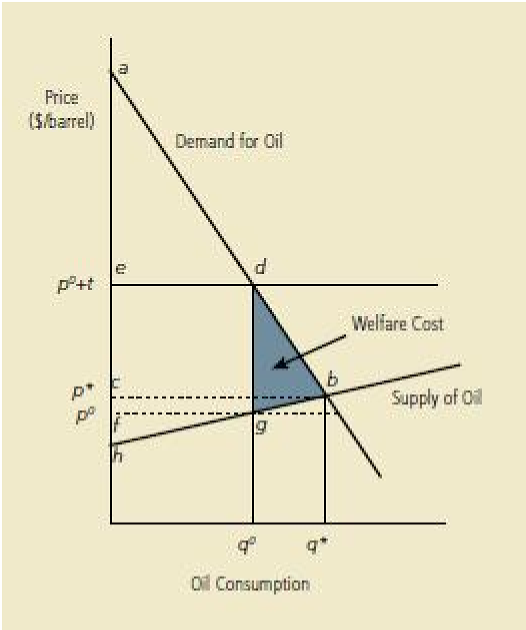

Figure 5: An example of a Harberger triangle, from Krupnick et al. (2010)

Suppose a tax of t is placed on sale of each of these units. As shown in Figure 5, this will create a wedge between the price paid by the consumer $(p^0 + t)$ and the price received by the seller (p^0), the size of the wedge is equal to the amount of the tax, t. Notice the equilibrium quantity also falls to q^0. The tax will also change consumer, producer, and total surplus in this market. Consumer surplus shrinks to ade, producer surplus shrinks to fgh, the government now receives a surplus equivalent to $edgf$ (the tax times the quantity sold). This gives a total surplus represented by $adgh$, which is less than the original surplus of abh. The difference in the two measures, called the deadweight loss, is represented by the shaded area dbg. This area is also the change in utility for the policy change if it is assumed the market demand curves are based on utility maximization. The idea behind Harberger triangles is to calculate this area.

To use this method, one first assumes that both the demand and supply curves are linear. This approximation is good for small policy changes, but can be inaccurate for large ones. In NEMS it may be problematic because supply curves may not be linear, and the shape of demand curves is unclear as well. One can then apply the formula for the area of a triangle to find dbg:
$A_{triangle} = \frac{1}{2} * Base * Height$. Notice in this case that the base of the triangle will be the difference between the old and new equilibrium quantities $(q^* - q^0)$. Similarly, the height is the difference between the old and new prices. This implies that calculation of the area requires only the old equilibrium prices and quantities along with the new equilibrium prices and quantities.

Figure 5 shows a case where there are no previous distortions before the new tax is implemented (i.e. other taxes, price ceilings, etc.). This may or may not be true in any particular market, but needs to be considered when calculating the change in utility. Similarly, the market structure must also be considered. It may be the case that supply (or demand) are perfectly inelastic, in which case there will be calculations of other areas in addition to the triangle in Figure 5. This caveat does not change the basic procedure. The benefit of this method is that it can be used for any type of model if the prices and quantities before and after a policy change are available. It is, however, an approximation depending on the linearity of the supply and demand curves.