Cost-Benefit Analysis

September 5, 2022

We have discussed models in a static sense. We now introduce a method to consider benefits and costs over time. Any economic agent (individual, firm, government organization, policy maker, etc.) should be aware that any project undertaken or policy implemented today has benefits as well as direct costs (the costs of the project) and opportunity costs (what else could be done with the resources allocated to the project) both today and likely in the future. The goal is to understand how to value these benefits and costs so as to make informed decisions about which projects and policies should be undertaken and which should not.

1 Present Value

The general concept of present value can be summed up with the line from the Popeye character Wimpy, who was known for saying "I will gladly pay you Tuesday for a hamburger today."[1] He realizes that a dollar today is more valuable than a dollar in the future, so he wants to hold on to his dollars today and pay in the future. There are a few reasons that dollars today are more valuable than dollars in the future. One is that those dollars today can be used to purchase goods today and, if the transactions are voluntary, the goods are going to provide at least as much utility today as the dollars themselves. Another is that the value of a dollar over time tends to decline due to inflation. The value of a dollar is in what it can buy,[2] and having $100,000 today will not buy as much as having $100,000 in 1950. A third reason is that a dollar today can earn interest if invested, and that invested dollar today will be worth more than a dollar in the future.

1.1 Projecting Present Dollars to the Future

To project present dollars to the future, we need a way to determine the value of those dollars in the future. One method of valuing current dollars in the future is to use a known interest rate. If an individual has $1,000 today and the annual interest rate is 10%, then the individual will have the original $1,000 plus an additional $100 in interest a year later, for a total of $1,100. If the $1,000 is invested for two years at 10% interest, then the individual will have $1,100 after the first year, which then pays an additional 10% ($110) the next year, for a total of $1,210. If invested for a third year, the individual would have $1,210 plus an additional $121, or $1,331, after three years. Let P be the original investment amount and r be the interest rate.[3] We can write the future value, which we will denote as FV_i for any length of time in years i, for T years in the future as follows:

$$
\begin{array}{cc}
\text{One year:} & FV_1 = P * (1+r) \\
\text{Two years:} & FV_2 = FV_1 * (1+r) \\
\text{Three years:} & FV_3 = FV_2 * (1+r) \\
\vdots & \vdots \\
\text{T years:} & FV_T = FV_{T-1} * (1+r)
\end{array}
$$

This equation simply takes the starting amount at the beginning of the year and finds the future value one year later. However, if we assume the interest rate is the same each year, we can project the value of

[1] This quote works less well when class meets on Tuesday than when it meets on other days.

[2] A dollar bill is a 6.14 inch x 2.61 inch piece of a cotton and linen blend. The physical item itself is not all that special, except that it can be exchanged for other goods.

[3] In future and present value calculations, r is also commonly called the discount rate.

any principal amount P any number of years into the future much more simply. Notice that we use FV_1 in determining FV_2, but we also know how to calculate FV_1 using the principal amount. We can then substitute in for FV_1:

$$\begin{aligned} FV_2 &= FV_1 * (1+r) \\ FV_2 &= P * (1+r) * (1+r) \\ FV_2 &= P * (1+r)^2 \end{aligned}$$

So if a principal amount is invested at the same interest rate for two years, we simply need to multiply the principal amount by $(1+r)$ two times to determine its value in the future. Now that we know FV_2 in terms of P we can find FV_3 in terms of P:

$$\begin{aligned} FV_3 &= FV_2 * (1+r) \\ FV_3 &= P * (1+r)^2 * (1+r) \\ FV_3 &= P * (1+r)^3 \end{aligned}$$

Hopefully the pattern is clear by now. For any principal amount P invested for a number of years T, the future value, T years later, will be:

$$FV_T = P * (1+r)^T .$$

1.2 Projecting Future Dollars to the Present

In addition to projecting current dollars to the future, we need a method of converting future dollars to the present. If an individual is promised $1,000 a year from now, how much would the individual be willing to pay today to receive that $1,000 in the future? Again, we will use the interest rate as a measure of how much society values future dollars.

We begin by examining the case with no inflation. As our starting point, we will use the already developed method for projecting present dollars into the future and work backwards. We know that $FV_1 = \$1,000$, so now we want to determine the present value of that future payment, which is the maximum amount of money an individual would be willing to pay for that future amount. In essence, we want to find P from our future value equation, though we will now call the principal amount the present value or PV:

$$\begin{aligned} FV_1 &= PV * (1+r) \\ \frac{FV_1}{(1+r)} &= PV \end{aligned}$$

If our future value is $1,000, and the interest rate is 10%, the the present value is $909.09. If an individual were to receive $1,000 two years from now, we can use the formula for FV_2 to calculate the present value of that money today:

$$\begin{aligned} FV_2 &= PV * (1+r)^2 \\ \frac{FV_2}{(1+r)^2} &= PV \end{aligned}$$

With $FV_2 = \$1,000$ and $r = 10\%$, we can calculate that the present value of that amount as $826.45. Note that this amount is less than the amount the individual would be willing to pay to receive $1,000 one year from now, as the individual is willing to pay less today for the same amount of money further into the future because they are waiting longer to receive the money. In general, if an individual receives a future value, FV_T, T years in the future, the amount the individual would be willing to pay today for FV_T is $PV = \frac{FV_T}{(1+r)^T}$.

Now suppose the individual receives payments annually for T years rather than one lump sum payment after T years. We will let R_0 be the payment the individual receives today, R_1 be the payment the individual receives one year from now, R_2 be the payment the individual receives two years from now, etc. for T years. How much is this stream of payments worth to the individual today?[4] The amount the individual actually

[4] These are the types of calculations being made by organizations who buy out streams of settlement payments or by lottery officials in determining how much the winner receives today if they choose to receive a lump sum today rather than a stream of payments over 20 years.

receives over time is $R_0 + R_1 + R_2 + ... + R_T$, but that money is all not received today. We know that an individual who receives an amount one year into the future values that amount as $\frac{R_1}{(1+r)}$ today,[5] and two years into the future as $\frac{R_2}{(1+r)^2}$ today. If we add up these discounted payments, we have:

$$PV = R_0 + \frac{R_1}{(1+r)} + \frac{R_2}{(1+r)^2} + ... + \frac{R_T}{(1+r)^T}$$

If an individual were to receive $1,000 each year for five years (which would make $T = 4$), if $r = 0.1$ that individual would be willing to accept $4,169.87 today in lieu of that stream of payments. If $r = 0.05$, then that individual would be willing to accept $4,545.95 today. If $r = 0.01$, then the individual would be willing to accept $4,901.97 today. As the interest rate decreases, future dollars are not discounted by as much, so the individual needs to be paid more money today to be willing to give up the future stream of payments.[6]

The interest rate is an important factor in these calculations. If a project today is expected to pay off $20 million thirty years from now, the present value of that $20 million is around $14.8 million if the interest rate is 1%, $4.63 million if the interest rate is 5%, and $1.15 million if the interest rate is 10%. If an individual is weighing the option of investing $5 million in that project today, whether the individual chooses to invest depends on the alternative rate they can receive.[7] At an alternative rate of 1%, investing $5 million seems worthwhile; at an alternative rate of 10%, investing $5 million does not seem worthwhile.

Inflation Case The previous analysis ignores inflation.[8] While we will not work through all the details, if inflation is anticipated then inflation should be automatically priced into the payments and interest rates and the present value calculation should not change. An important consideration is whether those payments and rates are adjusted for inflation. If the payments are not adjusted for inflation, and the individual is receiving a 5% interest rate but inflation is 2%, then it is similar to receiving a 3% interest rate because about two percentage points of the five percent return will go towards paying inflated prices in the following years.

2 Valuing Project Benefits and Costs

With the present value calculation, we now have a method for valuing a stream of payments over time. We can use this method to determine whether a project should be taken, as well as which of two projects should be undertaken. We begin with a discussion of private projects and rules and then turn to a discussion of public projects.

2.1 Private Projects

We begin with two criteria for evaluating private projects. The first examines a single project to determine whether it is *admissible*. A project is admissible if the benefits from the projects exceed the costs. Given our earlier discussions regarding whether an action should be undertaken this rule should be surprising. However, in earlier discussions we used marginal benefits and marginal costs; here we are using total benefits and total costs. The project has not yet been undertaken so before committing resources we want to evaluate whether there will be a positive return to the project. The second criterion examines two admissible projects to determine which is *preferable*. If two projects are admissible, the project that is preferable is the one with the higher net return. If two projects, M and N, have benefits and costs that all occur today, then we just calculate the difference between the benefits and costs for each project, $B^M - C^M$ and $B^N - C^N$, determine which project has the greater benefit (assuming both are admissible), and pursue that project.

[5] The term $\frac{1}{(1+r)}$ plays an important role in present value calculations. You may also see it referred to as the discount factor. Recall that r is the discount rate, so it is similar terminology (factor vs. rate) for these two terms.

[6] At the extreme, if the interest rate is 0%, then the individual would only be willing to accept an amount equal to the sum of the undiscounted payments.

[7] This alternative rate should be recognized as the opportunity cost of investing the funds in the project.

[8] Inflation is a measure of how much the general price level in an economy has risen. Having $1,000 in 1920 is not the same as having $1,000 in 2020, as the $1,000 in 1920 would buy more goods than $1,000 in 2020.

But the world is not that simple as benefits and costs accrue and occur over time. Fortunately we can value a stream of benefits and costs over time using present value. If there are T periods over which benefits accrue or costs occur,[9] we can use the present value calculation to determine the value of each project today:

$$\begin{aligned}
\text{Project } M \quad &: \quad PV^M = \left(B_0^M - C_0^M\right) + \frac{\left(B_1^M - C_1^M\right)}{(1+r)^1} + \frac{\left(B_2^M - C_2^M\right)}{(1+r)^2} + \ldots + \frac{\left(B_T^M - C_T^M\right)}{(1+r)^T} \\
\text{Project } N \quad &: \quad PV^N = \left(B_0^N - C_0^N\right) + \frac{\left(B_1^N - C_1^N\right)}{(1+r)^1} + \frac{\left(B_2^N - C_2^N\right)}{(1+r)^2} + \ldots + \frac{\left(B_T^N - C_T^N\right)}{(1+r)^T}
\end{aligned}$$

Using the present value framework, a project is admissible if $PV \geq 0$. A project is preferable if it is admissible and if $PV^i \geq PV^j$ for any two projects i and j.

The discount rate, r, is critical in these present value calculations and changes in the discount rate can cause changes in both project admissibility and preferability. Projects with returns in the distant future will be undervalued if the chosen discount rate is too high.

There are other criteria that are popular to use in cost benefit analysis. One is the benefit cost ratio, $\frac{B}{C}$, which admits projects if $\frac{B}{C} > 1$. For admissibility, the benefit cost ratio works well because $\frac{B}{C} > 1$ implies that $B > C$, which is the same rule as present value. But for preferability, individuals can strategically identify benefits as "negative costs" or costs as "negative benefits" to alter the numerator or denominator in a way that makes their own personal preferred policy or project seem more favorable. A second is internal rate of return, in which the firm determines the discount rate, ρ, that makes the present value of the project equal zero. The idea is that if a project has a rate of return greater than ρ, then the firm is exceeding the return of its opportunity cost so the project should be implemented. As with benefit cost ratio, internal rate of return works for admissibility – after all, if the firm can receive a higher rate of return on the project than investing the resources in its outside it would be better off completing the project. However, internal rate of return does not work as well when comparing projects of different scales. If an admissible project that uses \$5,000 returns 10% after a year then the firm has gained \$500. However, if using that \$5,000 prohibits the firm from engaging in a different project that requires \$11,000 but "only" returns the lower rate of 9%, then the firm could be making \$990 with the lower rate of return project.[10] When comparing projects of equal size the internal rate of return would provide an accurate answer to which project should be pursued, but when comparing projects of different sizes more caution should be taken when using this approach. The present value approach does not suffer from either of the problems of the benefit cost ratio or the internal rate of return valuation methods.

2.2 Public Projects

With private projects, firms and individuals may have a reasonable estimate of the rate of return they are forgoing when pursuing a project[11] as well as the costs and benefits. The question then is how to determine the rate of return on public projects as well as how to value costs and benefits.

2.2.1 Rate of Return

The economist's answer to determine the rate of return for funds invested into public projects uses the marginal rate of return for those dollars if they were invested into a private project. If there is \$90,000 that is invested into a public project that would have generated a before-tax 12% return if invested into a private project, then society only gains from the public if the rate of return is greater than 12%.

That calculation ignores that funds for public projects come from many sources, including individuals, who might use those funds for consumption. If an individual has a 12% rate of return on his funds, but

[9] Note that we are not limited to a finite number of periods T. We can calculate present value for an infinite number of periods. As long as $\frac{1}{(1+r)} < 1$, the present value will be finite. The Appendix of these notes contain a discussion.

[10] We do need to consider the return on the additional \$6,000 that the firm uses for the second project to complete the evaluation. If it will generate more than \$440 in returns then the \$5,000 project plus this additional return from the \$6,000 is better than pursuing the \$10,000 project.

[11] Historically a standard default for the opportunity cost of private funds was the U.S. Treasury Bill rate. However, those rates were less than 1 and for the most part (essentially) zero between September 2008 and March 2016. They had begun to climb higher (between 1.5%-2.5%) but fell back to (essentially) zero in March 2020 due to the pandemic.

must pay 25% in taxes, then the individual is only receiving a 9% rate of return. Another method for calculating the public discount rate would be to use a weighted average of the before-tax (for funds that come from private investment) and after-tax (for funds that come from consumption) rates. A difficulty with this approach is determining the weights on those two rates.

Another approach is to use a social discount rate, which measures how much society values the consumption that is lost today. Again, the economist's immediate response would be to argue that the social discount rate would be equal to the private discount rate. However, this argument may miss some key factors involved in public projects. One is that public projects should be putting weight on future generations, including those who are not yet born. However, private individuals and firms value future generations as well and it is difficult to claim the government will value the future more appropriately. A second argument is paternalism, in that society does not know what is best for itself collectively. Again, it is difficult to argue that the government, which is also comprised of people, will be any better at determining what is best for people. Finally, there may be market inefficiencies which the discount rates of private firms and individuals do not reflect.[12]

At this point all these potential methods of determining the appropriate discount rates for public projects seem lacking. One common approach is to consider all reasonable ranges of what the discount rate could be and then determine if the present value of the project is always positive. In essence, we would be conducting sensitivity analysis.

2.2.2 Benefits and Costs

In order to value public benefits and costs, an economist would immediately suggest that market prices be used as a starting point. After all, in a perfectly competitive market, the prices should reflect the marginal cost to society of the goods being traded. However, there are reasons why market prices may not reflect the true marginal cost to society. We have already seen that a monopoly produces deadweight loss, and that the price paid by consumers does not reflect the monopolist's marginal cost. Later in the course we will discuss how taxes have a similar distortionary effect on market prices. For these types of consideration, an adjusted valuation may be used. The extreme would be to use either the supplier's cost or the consumer's value, depending upon whether production is expected to increase (use the supplier's cost) or stay the same (use the consumer's value). As with the discount rate, a weighted average could also be used.

When market prices exist, we can use them, or an adjusted version, as a starting point for determining costs and benefits. However, what if market prices do not exist? In building our model of consumer choice, we focused on utility theory. An alternative approach is that of revealed preference, in which the choices individuals make, and not an inherent underlying utility function, are the foundation of consumer choice. The appeal of that approach is that preferences are based on actual choices – which toothpaste did the consumer choose when faced with multiple products at multiple price points? However, the approach is lacking if one wishes to consider goods/options that the consumer has never had to make a choice over. For valuing costs and benefits, we can use decisions made by individuals to estimate values for concepts like "time" or "life."

We can estimate the value of leisure time by examining the decisions the involuntarily unemployed make when they begin a new job. If we know the wage rate, then the value of leisure time should be less than the wage rate.[13] Additionally, if an individual can control how much time is spent working, then we know what the marginal value of leisure is if we know the wage rate of the individual's last hour of work. An alternative approach to valuing time could be to examine transportation choices. If one individual takes a slower, but less costly, method of commuting, than another, we can use that information to estimate how much an individual values time. Of course, either of those approaches makes assumptions that may not hold. Workers may not be able to control how many hours they work, and there are any number of reasons beyond the value of time why one individual may choose one method of commuting and a second may choose another.[14]

[12] We will discuss some market inefficiencies in the next set of notes.

[13] We will examine labor-leisure decisions more directly later in the course.

[14] Some may like driving in general; some may not drive at all. Some may live close to public transportation; others may live further away. Personally, I like driving, but after living in northern Virginia for two years I realized I really disliked commuting. One can resolve the commuting problem by choosing to live close to where one works.

Estimating the value of life is a much more difficult task, in part because it is a more difficult philosophical question to contemplate the value of a life. An emotional argument is generally made that life is priceless, but the reality is that society has made choices about tradeoffs between preserving life and other aspects of society. If we placed a regulator (or governor) on all automobiles that limited their speed to 10 miles per hour, we would greatly reduce, and probably nearly eliminate, deaths from traffic fatalities.[15] However, that would greatly increase travel time. That policy would also likely lead to substitution of alternative modes of transportation – bicycles and horses can move at speeds above 10 MPH – and some traffic fatalities would become fatalities from those other modes. But valuing a life is difficult. We could estimate the value of a life by using lost future earnings upon death, but that would suggest that anyone who would not be in the labor force in the future has no value, which would severely underestimate the value of those individuals. Instead of using lost future earnings we could examine choices that individuals make to reduce the probability of death. We know life is finite but we do not know the point at which we will die so we attempt to reduce the probability of death by taking certain actions. Which safety devices individuals use, how much they reduce the probability of death, and which prices they pay for those devices can help form an estimate of the value of life.

Finally, as you will hear on sports broadcasts, there are always intangibles, like the player who "holds the team together" or whose "value is more than you see in the statistics," which are difficult to value. For public projects, "national pride" is an intangible – how does one measure that? Because measurement is difficult, intangibles can be used to override the entire cost-benefit analysis. However, one can at least attempt to put a bound on the intangibles. There is a big difference if the proposed cost of a project is $10 million or $10 billion, and trying to determine where the supporters' break even cost point is can provide some evidence as to what they believe the value is. Also, as was mentioned in our discussion of firms, cost minimization should be a goal in a project if that project is being pursued, regardless of how much intangible benefit the project brings.

2.2.3 Evaluating Risk

Throughout this cost benefit analysis we have assumed that project benefits and costs are known. While that may be a reasonably accurate description for benefits and costs that accrue and are borne soon, that assumption is less likely to be true for benefits and costs well into the future. We have previously discussed how to incorporate risk for individuals, and a similar method can be used for public projects. The present value calculation for projects with known benefits and costs and known discount rates is:

$$PV = (B_0 - C_0) + \frac{(B_1 - C_1)}{(1+r)^1} + \frac{(B_2 - C_2)}{(1+r)^2} + ... + \frac{(B_T - C_T)}{(1+r)^T}$$

Let $E[\cdot]$ represent the expectations operator, so that $E[B_t]$ is the expected value of benefits t years from today, $E[C_t]$ is the expected value of costs t years from today, and $E[r]$ is the expected discount rate. Expected value is chosen rather than expected utility for two reasons. First, a utility function is needed if expected utility will be calculated, and determining which utility function to use adds another dimension of complexity to the problem. Second, one might expect that a government agency, like a firm, would be more likely to be risk neutral than risk averse. If we assume the government agency is risk neutrality then we can use the expected value calculation interchangeably with risk neutral expected utility. Our expected value of present value would then be:

$$E[PV] = (E[B_0] - E[C_0]) + \frac{(E[B_1] - E[C_1])}{(1+E[r])^1} + ... + \frac{(E[B_T] - E[C_T])}{(1+E[r])^T}$$

There is not much that changes in the case with risky outcomes, except now we replace the known outcomes with the expected outcomes each period. When considering projects now, we could add one more criterion, variability in outcomes. For two projects with the same $E[PV]$, the choice may be to implement the one with lower outcome variance. While a truly risk neutral maximizer would be indifferent over two projects with the same $E[PV]$, a prudent policy maker would likely want to narrow down the range of possible outcomes to reduce the likelihood of a very bad outcome from occurring.

[15] I would also predict that we would have many people who modified engines to circumvent the policy. We do, after all, live in Charlotte and the surrounding areas, which are home to many NASCAR teams.

3 Compensation Principle

Thus far we have examined projects from a present value perspective. We could also use the concept of Pareto optimality and only consider policies that yield Pareto improvements, meaning that all individuals are at least as well of as before and at least one individual is strictly better off. However, that is a very tough standard as any policy that applies to a group of individuals is likely to make some individuals better off and some worse off. But we can use the conceptual idea of Pareto improvement by considering policies under which those made better off could compensate those made worse off, and, as long as the (hypothetical) transfers leave all individuals at least as well off as before and one individual strictly better off, the policy can be implemented. It is important to note that the transfer are hypothetical – if they could be made they would actually lead to a Pareto improvement. The compensation principle is often referred to as the Hicks-Kaldor compensation principle.

At times I have alluded to this concept earlier in the course when mentioning that we could take surplus from individuals in some transactions and shift some of that surplus to others to make them better off. If an individual values an item at \$20, the individual still receives consumer surplus whether the price is \$8 or \$12, it is just that the individual receives less surplus if \$12 is paid. If the firm itself is charging the higher price then the surplus simply shifts from consumer surplus to producer surplus, but in this case we want to consider the possibility that the price paid by the consumer and the price received by the firm may differ, as when taxes are imposed on purchases.[16] If the consumer pays \$12 and the firm only receives \$11, then \$1 of the total surplus from that transaction goes elsewhere. The consumer and producer are still better off than they were before (otherwise they would not have made the voluntary transaction), but some other individual in society should now be benefitting from that \$1 of surplus created from the transaction that is captured by neither the consumer nor the producer.[17]

Using our model of individual consumer choice we will examine two types of compensation at the individual level.

3.1 Compensating Variation

Suppose there is a policy change that increases the price of one good. This price increase will reduce the consumer's budget set and lead to the consumer being worse off (i.e. on a lower indifference curve). A question we could ask is: how much money would we need to give the consumer, under this new set of prices, so that the consumer would be just as well off as before the price change went into effect? Figure 1 shows that under the initial budget constraint the consumer is on indifference curve I_2. After the price of Good 2 increases, the budget constraint shifts and the consumer is now choosing the bundle on I_1. If we keep the prices as they are after the price change, we can shift that budget constraint out to determine which point on I_2 would be parallel to a budget constraint with the new price ratio. The key is shifting that new budget constraint out in a parallel shift until it is tangent to the original indifference curve. We can then calculate the additional income needed to make the individual just as well off as before the price of Good 2 increased by looking at the new intercept for Good 1. If we know the price of Good 1 and how much of Good 1 the consumer buys under the initial budget constraint and under the hypothetical budget constraint, we can calculate the income the consumer would need to be indifferent to the original price ratio and the new price ratio. In essence, the compensating variation is a measure of a consumer's willingness to accept a new price change.

3.2 Equivalent Variation

An alternative question to ask is how much income we could take away from a consumer that would leave the consumer no worse off than if a price change occurred. Figure 2 shows a graphical representation of the equivalent variation. Under the initial budget constraint the consumer is on I_2, and after the price increase the consumer is on I_1. The hypothetical budget constraint in Figure 2 takes the original price ratio and finds the budget constraint under that price ratio that is tangent to the new indifference curve I_1. In essence, we

[16] We will discuss taxes more formally at a later date in the course.

[17] It is possible that the consumer or producer in the transaction is the individual who benefits from the \$1 of surplus that is not captured by either of them in that particular transaction.

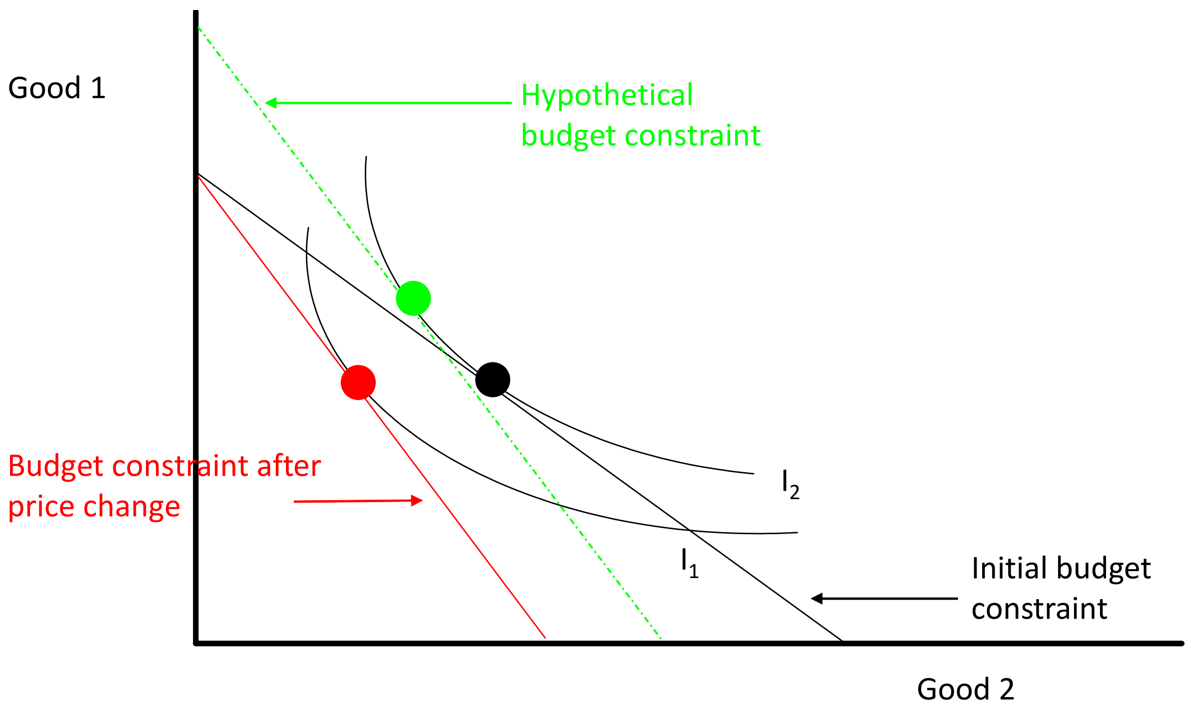

Figure 1: The compensating variation for a price increase of Good 2. The original budget constraint/bundle is in black, the new budget constraint/bundle is in red, and the hypothetical budget constraint/bundle is in green.

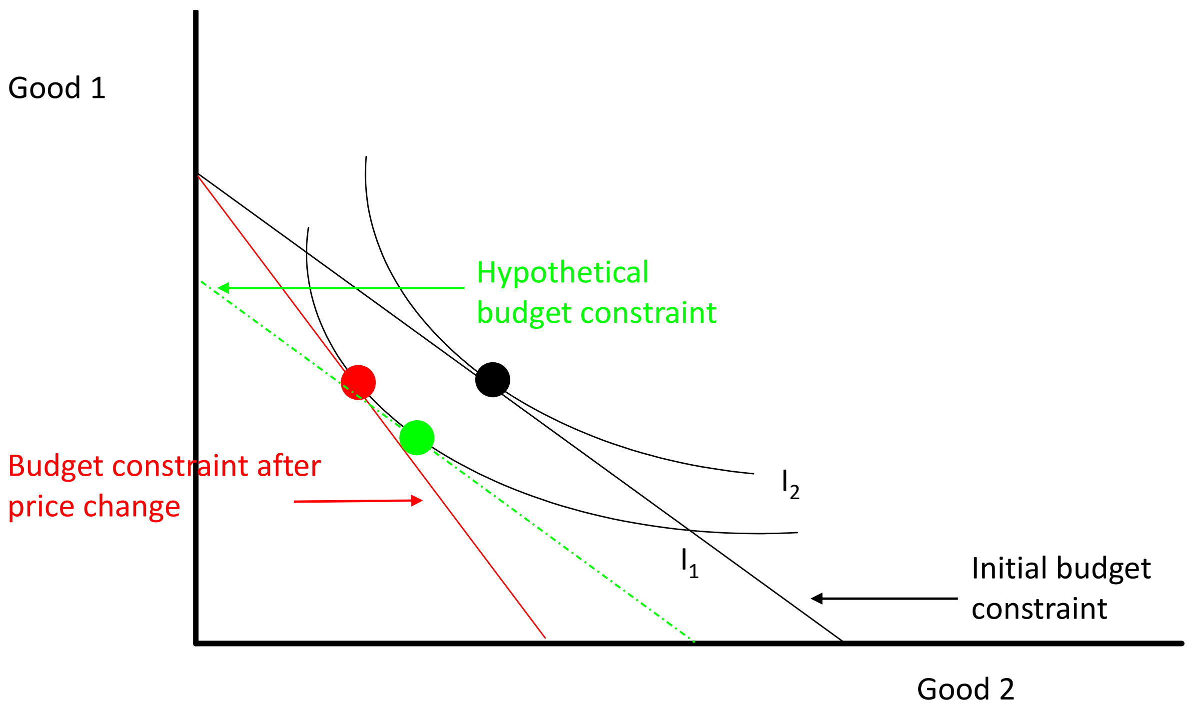

Figure 2: The equivalent variation for a price increase of Good 2. The original budget constraint/bundle is in black, the new budget constraint/bundle is in red, and the hypothetical budget constraint/bundle is in green.

could either keep prices as they were originally and take away income from the consumer that would leave them on I_1 (at the green bundle) or increase the price of Good 2 such that the consumer ends up on I_1 (but at the red bundle). Either way, the consumer is indifferent between those two policies. In essence, the equivalent variation is a measure of a consumer's willingness to pay to avoid a price increase.

When there is a price change, both substitution and income effects occur in the consumer choice model. A substitution effect occurs when the consumer substitutes some consumption to the now relatively less expensive good. The income effect occurs because the consumer, after a price change, does not have to shift all consumption to the now relatively less expensive good, but can recalculate the optimal bundle and reallocate income to both goods. The compensating variation and the equivalent variation would be the same if the income effect was zero, and they are essentially the same for negligible income effects. But theoretically, we should be able to determine how much income is needed or can be taken away to make a consumer just as well off before or after a price change.

One might ask why we do not just give the individual the same amount of money back (or taken away) so that they can buy the original bundle. If we did that, we would actually be making the consumer better off than they were initially. That result occurs because of the income effect – if you shift the hypothetical budget lines back to the original bundle, you will see that they intersect the indifference curve, but is not tangent to it. So the consumer would actually move to a higher indifference curve.

4 Criticisms and Difficulties

Throughout the discussion I have mentioned criticisms of some methods. The biggest difficulty is accurately calculating the costs and benefits of a policy proposal. It is also difficult to forecast what the appropriate discount rate will be in the future.

The compensation principle informs us to choose those policies for which Pareto improvements are hypothetically possible, even if they are not implemented in practice. However, if many policies are implemented over time and those transfers are not implemented in practice, it is possible that some consumers continually lose from the policy implementations. While the compensation principle guarantees efficiency, we now return back to the more difficult question of determining equity. If the goal is to transfer resources from a wealthier group to a less-wealthy group, then the compensation principle may not apply as we may not be seeking Pareto improvements.

5 Appendix

I mentioned that it is possible to calculate a stream of payments that are paid forever provided that the discount factor is less than zero, or $\frac{1}{1+r} < 0$. Assume a constant payoff each period of Π_t, for $t = 0, 1, 2, ..., \infty$. Rather than stopping payments at some fixed period T they continue forever. It is important to note that:

$$\left(\frac{1}{1+r}\right)^t = \frac{(1)^t}{(1+r)^t} = \frac{1}{(1+r)^t}$$

Mathematically those are equivalent because we can raise the numerator and denominator to a power separately and because raising one to any power is still one. Using that result, the present value can be written as follows:

$$PV = \Pi_0\left(\frac{1}{1+r}\right)^0 + \Pi_1\left(\frac{1}{1+r}\right)^1 + \Pi_2\left(\frac{1}{1+r}\right)^2 + ... + \Pi_\infty\left(\frac{1}{1+r}\right)^\infty$$

This equation looks slightly different than the earlier equation because the first term, for today's payoff, is written as $\Pi_0\left(\frac{1}{1+r}\right)^0$ rather than Π_0, but recall that when we raise a number to zero the result is one, so $\Pi_0\left(\frac{1}{1+r}\right)^0 = \Pi_0$. We also know that if $r > 0$ that $\frac{1}{1+r} < 1$. Let us define $\frac{1}{1+r} = \delta$ as the discount factor to make the equation a little less cumbersome. We now have:

$$PV = \Pi_0\delta^0 + \Pi_1\delta^1 + \Pi_2\delta^2 + ... + \Pi_\infty\delta^\infty$$

Because $\Pi_0 = \Pi_1 = \Pi_2 = ... = \Pi_\infty$ we can factor that constant payoff out (call it Π) to find:

$$PV = \Pi \left(\delta^0 + \delta^1 + \delta^2 + ... + \delta^\infty\right)$$

Focusing on the term with the δ, we know $\delta^0 = 1$, so we have:

$$\delta^0 + \delta^1 + \delta^2 + ... + \delta^\infty = 1 + \delta^1 + \delta^2 + ... + \delta^\infty$$

There is a mathematical result, which we will not prove, that as long as $\delta < 1$ this term sums to:

$$1 + \delta^1 + \delta^2 + ... + \delta^\infty = \frac{1}{1-\delta}$$

So if $\delta = \frac{1}{2}$, then the result is 2; if $\delta = \frac{2}{3}$ then the result is 3. For us, we have $\delta = \frac{1}{1+r}$, where r is likely to be close to zero. If $r = 0.05$, then $\delta \approx 0.95$, and the sum is 20; if $r = 0.1$, then $\delta \approx 0.91$ and the sum is 11. To calculate the present value we would simply take whatever that resulting sum is and multiply it by the constant payoff of Π.