Formal Analysis: Utility, Optimization, and Welfare

Ben Ansell

Hilary Term 2017

1 Introduction

This lecture introduces students to the basic building blocks of formal analysis. Students should first read over my Math Prefresher guide on calculus and constrained optimization, which covers most of the necessary mathematics to follow this lecture. We will begin by thinking about the concept of a *utility function*, which is the bedrock of rational choice / economic analysis. This function is presumed to apply to a single agent (an individual or a group with identical preferences) though that does not mean that an individual is necessarily selfish - utility functions could include the utilities of other agents - that is, they can be altruistic as well as self-regarding.

We then move on to thinking about how a given agent *maximizes* their utility function given a set of constraints. These constraints are key in understanding political economy since we can think of political life as providing a series of constraints on any individual's ability to get what they want. Whereas microeconomists tend to think of constraints as deriving from having a limited income or limited resources, political economists view constraints as deriving from formal and informal institutions like legislative rules, electoral systems, or social norms.

We finish by talking about making normative judgments over social outcomes. Given agents' behavior under various institutional constraints, how can we judge whether one pattern of behavior or one institutional arrangement is 'better' than another? In general there is no single answer to this question. Different conceptions of social welfare are hotly debated in both welfare economics and political philosophy. We shall briefly discuss the concepts of utilitarianism, Rawlsianism, and Pareto optimality.

2 Utility Functions

The concept of a utility function is the bedrock of any formal analysis. Simply put, the structure of the utility function defines exactly what it is that agents are trying to achieve. They might be trying to become richer, or happier (not necessarily the same thing), they might be trying to avoid costly outcomes, they might be trying to avoid losses (again a slightly different thing), they might be trying to make everyone else happy, they might be trying to have a happier future, or they might be trying to have a happier present, damn the future. As long as outcomes can be ranked and as long as we can conceive of the things that agents are trying to maximize as variables that could be larger or smaller, then we can apply the utility function concept, *even if* the behavior of an agent does not correspond to an ideal typical form of rationality.

2.1 The Theory of Choice

The basic idea of a utility function comes from a *preference ordering* over different outcomes. Classical choice theory assumes only that agents can identify whether they prefer, weakly prefer, or are indifferent among two different outcomes x and y in the set of all outcomes X.

Following McCarty and Meirowitz (2007) I use the symbol R to mean 'weakly preferred', defined such that xRy implies that x is either preferred to, or as good as, y, or symmetrically, that y is definitely not preferred to x. The operator R is thus analogous to the mathematical expression $\geq$ to compare real numbers. The operator P is used to express strict preference such that xPy if and only if xRy but not yRx. This is analagous to the mathematical symbol $>$. Finally, the operator I expresses indifference between options, so that xIy is true if and only if xRy and yRx. This is analagous to the mathematical symbol $=$.

The idea of maximization using this basic setup is quite simple. The maximal set M is defined as $x \in X : xRy, \forall y \in X$. That is, those outcomes that are at least as good as all other possible outcomes. As McCarty and Meirowitz state 'the fundamental tenet of rationality is that agents choose outcomes from the maximal set'. 'Ideally' the maximal set would contain only one outcome. More critically, the maximal set must not be empty if the theory of choice is to make any sense.

There are two conditions that have to obtain (at least in the world of discrete choice) for the set to be nonempty and thus for an optimal choice (or set of choices) to exist. Firstly, outcomes must be comparable. A *complete and reflexive* relation is one where $\forall x, y \in X, xRy, yRx,$ or both. For choice theory to work agents cannot constantly change their minds between x and y. Furthermore, agents must be able to compare outcomes, they cannot refuse to decide.

Secondly, agents must have a *transitive* ordering of preferences. Formally this means, that if xRy and yRz, then xRz. That is, if one outcome is no worse than a second outcome, and the second is no worse than a third, then the first outcome must also be at least as good as the third outcome. This may seem logically obvious but in fact it often does not hold - think about sports teams - if Papua New Guinea beat France (luckily) in a game of soccer, and France beat Brazil in a second game, then transitivity (rather loosely) implies that Papua New Guinea are a better team than Brazil. Of course, this is not quite the same things as preferences but it is not hard to imagine that you prefer the music of U2 to that of Coldplay, Coldplay to that of Nirvana, but Nirvana to that of U2 if you were asked these questions separately. Most famously, Arrow's impossibility theorem shows that a transitive social ordering of preferences is not possible without a dictator (assuming the ordering is also Pareto optimal and independent of irrelevant alternatives).

Putting these conditions together provides the basic tenet of rational choice: for a finite set X, with a complete and transitive relation R, M is nonempty. So for a theory of choice to be useful all we need is agents to be able to compare all outcomes against one another and to have a transitive ordering among them.

In fact, things are a little more complicated with a infinite set of choices Y since we need the set Y to be nonempty and compact (this mathematical term means that the set of outcomes is closed and bounded like the set $[0, 1]$ but unlike the sets $(0, 1)$ or the set $[0, \infty)$) and the relation R to be complete, transitive and (lower) continuous. Continuity means that all points sufficiently close to y, where yPx, are also strictly preferred to x. Graphically, continuity is just that, a continuous, unbroken, smooth function.

2.2 From Choice to Utility

So that is basic choice theory. But this is not all that useful yet since we do not generally work with operators like R and P. What we need is a way of converting this to mathematical form so that we can use algebraic manipulations and perform calculus. A utility function is defined as $u(x) \geq u(y)$ if and only if xRy. From this we can derive $xPy \Rightarrow x > y$ and $xIy \Rightarrow x = y$. The optimal choice $M = \arg\max_{x \in X} u(x)$, which in

English means the maximal set equals those (that) x at which $u(x)$ reaches its maximum.

Provided that R is complete, transitive, and continuous then there exists a continuous utility function $u(x)$ that represents R. Similar results about the existence of a maximum apply as before: this time we need X to be compact and $u(x)$ to be continuous for M to be nonempty.

Since $u(x)$ can come from any complete, transitive, and continuous R this implies that utility functions, like the choice relations discussed before, or at base *ordinal*. The difference between $u(x)$ and $u(y)$ and that between $u(y)$ and $u(z)$ is unknown even if we do know that $u(x) > u(y) > u(z)$ because they simply came from the relationship $xPyPz$, that is, from an ordered relationship, without any cardinal values.

Because utility functions are ordinal we can theoretically apply any monotonic transformation to them and they will retain the same ordering. So $U(x) = 2u(x)$ carries the same information as $u(x)$ as does $u^*(x) = \ln(u(x))$. This has the important implication that it may not be possible to compare utility functions, and hence utilities across individuals. This then implies that coming up with normative judgments of different distributive outcomes is made very difficult since we cannot tell if the beggar or the miser would be made happier by giving them ten dollars if we are unable to compare their utilities. This seems rather problematic and is the focus of libraries of literature.

There is a further problem in that because any utility function can be transformed monotonically, we also do not know their shape. For example if x is income, then $u(x) = x$ implies utility is linear in income, $u(x) = \ln x$ implies that the marginal utility of income is declining as income increases (also a result of risk aversion) and finally $u(x) = x^2$ implies marginal utility is increasing in income (a result of risk acceptance). But all could represent the same preference ordering among outcomes - that is under all cases people prefer more money to less!

Since we prefer to distinguish these kinds of cases, in practice most political economists add more structure to utility functions, essentially treating them as cardinal. This is not without controversy. Elster notes the famous attack on interpersonal utility comparison by Becker and Stigler in an article entitled *De Gustibus non est Disputandum* (there's no accounting for taste). Two types of paramterization are particularly common: risk preferences over expected utility, and spatial preferences.

Risk preferences are associated with the famous concept of von Neumann - Morgenstern utility. These are preferences over *lotteries* or more simply over outcomes that occur with different probabilities. The basic idea of expected utility is that the expected utility of a lottery L between outcomes x and y where x occurs with a probability p and y with $1 - p$ is $EU(L) = pu(x) + (1 - p)u(y)$. If we had several outcomes this would require a probability associated with each, multiplied by the utility of each outcome, and then added together, or $EU(.) = \sum p_i u(x_i)$, where $\sum p_i = 1$. Now that we are multiplying utilities by probabilities and adding them, the expected utility function is no longer ordinal, and indeed monotonic changes to the utility functions themselves will alter the expected utilities of various lotteries. All said, once we enter into a world where uncertainty is pervasive we have to abandon ordinal utility and use cardinal utility.

In doing so we now care about the shape of the utility function. In particular we are interested in whether utility functions are linear, concave, or convex, which correspond to risk neutrality, risk aversion, and risk acceptance. Put simply, risk preferences determine how an individual values a bet between two extreme outcomes versus a moderate outcome for sure.

Let us use a simple example. Imagine comparing a bet between \$0 and \$100 with a fifty percent chance of getting each with a 'bet' where you received \$50 for sure. Under risk neutrality you would be indifferent - and we would get this outcome with a linear utility function, $u(x) = kx$. Here we compare the expected utility of bet A $EU(A) = 0.5 \cdot k0 + 0.5 \cdot k100 = k50$

with the expected utility of the for sure bet B $EU(B) = 1 \cdot k50 = k50$.

However, what if we are risk averse? Then we would prefer the sure thing over the bet. This means that our utility function at the sure bet is higher than a straight line between our utilities at the low and high end of the risk bet. Which mathematically means the utility function must be concave (or that the marginal utility of income is decreasing). Let us use the standard risk-averse utility function of $u(x) = \ln(x + 1)$ and compare lotteries A and B. Now we have $EU(A) = 0.5{\cdot}0{+}0.5{\cdot}ln(101) = 0.5ln(101) = 2.308$, compared to $EU(B) = ln(51) = 3.932$, so that the sure-thing is much preferred to the risk (and in fact, our agent is so risk averse that they would be approximately indifferent between the risky bet and $10).

Risk acceptance produces the opposite kind of result. Here we use a convex utility function $u(x) = x^2$, such that the agent prefers the bet to the sure thing. Again we compare expected utilities $EU(A) = 0.5 \cdot 0 + 0.5 \cdot 100000 = 5000$ versus $EU(B) = 2500$, showing that the agent would rather take the bet. Here we see that this agent would need a sure thing offer of $71 to induce them *not* to take the bet!

We finish our discussion of utility functions by briefly remaking on so-called *satiable* preferences. Up til now we have looked at nonsatiable utility - individuals always want more money or happiness. But what about cases when agents have a particular finite preference - for example, the amount of chocolate to eat in a day or their preferences for a politician along a one-dimensional scale like abortion rights? Here we need a utility function that reaches its maximum at some *ideal point* and then declines either side. Calling the ideal point x^*, the standard functions used are the linear absolute distance $u(x) = -|x - x^*|$ and the quadratic form $u(x) = -(x - x^*)^2$. Note that only the latter is differentiable at all points since the former has a spike at the ideal point. However, the latter also applies some form of risk aversion, which may be inappropriate depending on the matter at hand.

3 Constrained Maximization

Having established, at length, how utility functions are constructed we can now turn to more substantive analysis by considering how such functions are maximized under economic and political constraints. In the cases above, regardless of the shape of the utility function, when utility was nonsatiable individuals would always maximize their utility by getting as much as stuff as possible - to infinity and beyond. When utility was satiable then individuals would simply move to their ideal point. That's not particularly interesting or informative and barely moves beyond tautological. Our real interest in political economy is how different constraints affect agents' abilities to maximize their utility.

In this simple example we examine the case of an warlord seeking to maximize their income from a set of subjects. We can think of income as akin to utility. Note that presumably this is a nonsatiable utility function, since the ruler would like infinite income if possible. However, for a variety of reasons, some obvious (the income of the subjects is itself not infinite), some less obvious (institutional constraints on behavior), the ruler's maximal take from their subjects will be less than one hundred percent - this is known as an *interior solution* as opposed to a *corner solution*, wherein the optimal choice by the warlord would be either complete refrain or complete expropriation.

We being by constructing a general model of the warlord and the citizens. Let us assume that the citizens only earn income from the land L and split it according to population N: such that individual citizen income is $Y_{Ci} = \frac{L}{N}$. There is a single warlord who earns money only from appropriating the income of the population through 'taking' (or we could call this taxation) τ, and whose income is $Y_{Wi} = \tau N Y_{Ci} = \tau L$. Let us assume finally that the utilities of both the subjects and the warlords are linear in income - in fact, equal to income $U_{Wi} = Y_{Wi}, U_{Ci} = Y_{Ci}$ (note that this implies that both the warlord and citizens are risk-neutral).

Note that since we are restricting each citizen to the same income, we can simply express the income of the warlord as a function of land. We are also restricting taxation to be linear (i.e. a flat tax), which since all citizens earn the same income is equivalent to a lump-sum tax on each citizen. As we progress in the course we will alter these kinds of assumptions but for now our main interest is on the ruler so we retain these assumptions even though we know them to be false (see Morton, 1999 for more on the advantages and disadvantages of false assumptions).

Before turning to constraints let's examine the unconstrained effect of taxation on the utility of the warlord. To do this we take the derivative of warlord utility with respect to the tax level τ: $\partial U_{Wi}/\partial\tau = L > 0$. So a marginal increase in the rate of taxation adds extra land L to the warlord's income. Note that this effect is constant at L, which is greater than zero. This implies that the effect of increasing taxation on the utility of the landlord is *always* positive and thus a tax rate of infinity would be optimal. Obviously this is impossible. But why?

3.1 Budget Constraints

The first reason that the infinite tax rate is impossible is because tax rates are necessarily constrained to be between zero and one. There are in fact two ways of representing this information. Firstly, we could insert a mathematical constraint that $\tau \leq 1$ (or more broadly that $0 \leq \tau \leq 1$ but we already know that the warlord will not want a tax rate less than zero). Secondly, we could move further down the causal chain and ask *why* tax rates cannot be greater than one. To do so we argue that the maximum tax take that the warlord can take is equal to the amount of land in the economy and set this as a constraint, or $\tau L \leq L$. Of course, basic algebra tells us these conditions are exactly the same. Either way, since we know that the unconstrained effect of taxation on warlord utility is always positive $\partial U_{Wi}/\partial\tau = L > 0$ then we know that the warlord will choose the maximum possible value of τ, that is $\tau = 1$.

We also have another implicit budget constraint in the set-up above, $Y_{Wi} = \tau L$, where we assume that there is no *wastage* between taxed income of the subjects and the warlord's income. If we introduce wastage, then another constraint appears on the tax rate that the warlord prefers and now they might even prefer a tax rate less than one. Let us now assume instead that taxation is leaky and that as the tax rate increases, the leakage gets proportionally larger. What does this assumption mean? Well, a whole lot of things. It could be that as taxation gets higher, the subjects work the land less and produce less output (this is the famous Laffer curve assumption - as taxes approach one output approaches zero as no-one has the incentive to work). It could be that the subjects hide more of their output from taxation as taxes get higher (the peasant hoarding effect that frustrated Lenin). Or it could be that high tax rates encourage tax inspectors to cheat the warlord out of tax revenues and pocket some themselves. Note then that a mathematical assumption could have many substantive interpretations and if we are to distinguish among them we have to go further down the causal chain and develop more complicated models. It is useful nonetheless to know that a variety of substantively different factors could have the same outcome in the formal model - this helps clarify what kinds of assumptions the model supports and what it does not.

Let us mathematically reframe the warlord's utility as: $U_{Wi} = (\tau - \tau^2)L$, where we have added a quadratic loss - meaning that as taxes get larger, more and more leakage occurs. Now we can take the derivative once again of the warlord's utility with respect to the tax rate: $\partial U_{Wi}/\partial \tau = (1 - 2\tau)L$. To solve for the optimal level of taxation we set this expression to equal zero: the 'First Order Condition' (in fact, we also need to check that this is a maximum which means making sure the second derivative is negative, which it is $\partial^2 U_{Wi}/\partial \tau^2 = -2L$). This gives us the optimal tax rate of $\tau^* = \frac{1}{2}$.

Immediately we see that adding the leakage constraint means that taxes are no longer set to equal one but take the form of an 'interior solution' of one half. If taxes were raised any higher, the warlord would lose more in leakage than they would gain in additional revenue.

3.2 Institutional Constraints

So far the only constraints on the warlord's behavior have been economic. What if we allow some kind of political constraint? We can think of such constraints as legal, for example, a statutory limit on the warlord raising taxes above 0.25, which in all the cases above would prevent them from achieving their preferred outcome. Or we can change our warlord into a politician who faces a re-election constraint. In this latter case, their ability to set their preferred tax rate will depend on the information available to citizens, on whether they have to share the tax revenues back with the citizens in the form of public goods, on the election system that the warlord faces, and on the types of challenger they might face.

I will postpone specific discussion of particular political institutions until we reach later lectures but let us examine a general case. The warlord may face a potential challenger in some form of future 'election' - using that term very loosely. The 'election' could, for example, amount to a revolution that leads to their overthrow, or it could be a coup within the court. Borrowing from Przeworski Chapter 5 I assume that the warlord values not only their current consumption but also their future consumption V, which depends on whether they win this 'election'. Let us assume that they win election with probability $p(\tau) = 1 - \tau^2$. Before we go any further let us take the derivatives of this function: $p'(\tau) = -2\tau$ and $p''(\tau) = -2$. Note that this function is negative and convex - this implies that as taxes get higher the probability of getting elected gets lower and a given increase in taxes has a larger negative effect on election chances the higher are existing taxes (this comes from the second derivative). Well that might be true - citizens might be more willing to tolerate a change from ten to twenty percent taxation than from eighty to ninety percent.

The warlord's utility function is now $U_{Wi} = \tau L + Vp(\tau)$ and let us once more take the derivative of this expression with respect to τ: $\partial U_{Wi}/\partial\tau = L + Vp'(\tau)$. Since we know $p'(\tau)$ we can substitute to get $\partial U_{Wi}/\partial\tau = L - V2\tau$. Now notice that increasing taxes has positive effects *and* negative ones. Consequently, we can solve for the maximum utility attainable by setting this expression to equal zero and then rearranging to get $\tau^* = \frac{L}{2V}$. This is a very simple result but it has a clear intuition. Optimal taxes are increasing in the amount of land out there to be taken but decreasing in the value of getting re-elected. Let us assume that there is only one election, and hence in the next period the warlord will revert to taxing at one hundred percent, and assume also that the warlord values the future equally to the present - then $V = L$ and $\tau^* = 1/2$. Note that this is the same outcome as in our leakage example earlier but we got there down a slightly different path.

What instead happens if citizens react in a different way to taxation? Perhaps citizens are outraged when taxes increase from ten to twenty percent but are resigned to predation when they increase from eighty to ninety percent. In that case we could model $p(\tau) = (1-\tau)^2$. This function has derivatives $p'(\tau) = -2(1-\tau) = 2(\tau-1)$, which is negative when $\tau < 1$, and $p''(\tau) = 2 > 0$. We can reframe utility as $U_{Wi} = \tau L + Vp(\tau) = \tau L + V(1-\tau)^2$ and the first derivative is $\partial U_{Wi}/\partial\tau = L - V2(1-\tau)$. Again we set this equal to zero and solve for the optimal tax rate: $\tau^* = 1 - \frac{L}{2V}$.

Wait a second... that's strange we seem to be getting the reverse result because as V gets really big τ^* approaches one... Hang about, we forgot to check whether this was a maximum or a minimum so let's go back and check the second derivative $\partial^2 U_{Wi}/\partial\tau^2 = 2V > 0$. Ah hah! The problem is that this function is convex and so we actually found a minimum. Well, since we know that $p'(\tau) < 0$ when $\tau < 1$ then the maximum utility must occur when $\tau = 0$. This is called a corner solution since there is no internal optimum. Now this function doesn't make a huge amount of sense because we might assume that V is also a function of τ and if taxes are zero so is V. So then we would want to add that further level of complexity - I leave this for you as an exercise!

4 Social Welfare Functions

So far we have talked about constrained optimization in a 'value-free' fashion - that is, we have given agents assumed utility functions and then assumed a series of constraints that might exist and examined to what degree agents are able to fulfil their goals under such constraints. There are a variety of problems with this approach. The most obvious, and one we will wait a few weeks to deal with, is where are these preferences / utility functions coming from? Why don't warlords maximize the happiness of their citizens, or something completely different like the aesthetic beauty of their realm? Another set of problems come from the constraint side. Which constraints are most important? Can we include them all and keep the model tractable? Are these constraints really solid, material barriers, or are they socially constructed and hence alterable through ideational change? Again we will explore a whole variety of constraints throughout the course so I will punt on this for now but it's worth keeping in mind.

The other problem with the above analysis that it remains entirely unjudgmental. It is a pure *positive* analysis of political behavior without any *normative* judgment over what the 'best' outcome would be, or indeed among different institutional constraints what produces the 'best' behavior by the warlord(s). The vast discipline of welfare economics provides tools to attack this issue. It is not uncontroversial and much of modern liberal normative political theory, particularly the work of Rawls and his followers, has questioned many of the basic tenets of classical welfare economics. In this section, then, we examine these varied approaches to making welfare judgements over outcomes. The nice thing about this literature is that we can do so using the same mathematical tools with which we constructed the *positive* model in the first place.

I begin by covering the standard tools of classical welfare economics - the weak standard of Pareto optimality, which values efficiency over equality. We then turn to the concept of utilitarianism, which can promote either equality or efficiency depending on its formulation. I then move to the Nash and Rawlsian approachers that tend to promote higher levels of equality.

4.1 Pareto Optimality

The basic tool of welfare economists is the concept of Pareto optimality (after the 19th century Italian political economist Vilfredo Pareto). A Pareto optimal outcome is one where there is no way to make any individual better off without making one worse off. That is, there are no twenty-dollar notes on the sidewalk. Mathematically the way to think about a social welfare function that is Pareto optimal is to maximize the utility of one citizen with respect to the good / income to be distributed, *holding the utility of all other citizens constant*. In this manner we can be sure that no other citizen is worse off since we are holding constant their utility. We then repeat the procedure for all citizens. Hence we can directly connect the conception of constrained maximization explored in the previous section and the concept of utility functions developed in the first.

Note that we do not need cardinal utility functions to perform this task - all we need is ordinal utility functions that reach a (constrained) maximum. This makes Pareto optimality a useful tool since it involves no further assumptions about the comparability of utility functions across citizens. Nonetheless Pareto optimality is only satisfying from an efficiency perspective (in fact, the definition of efficiency from a welfare economics perspective is identical to Pareto optimality since no resources are being 'wasted'). From an equity perspective it is very problematic. For example it would not be Pareto optimal to take resources from a kleptocratic ruler and grant them to his impoverished citizens, since we would of course be making that ruler worse off. Hence we will need to move to more complicated analyses of welfare.

4.2 Utilitarianism

Unlike Pareto optimality, most ?welfarist' forms of ascertaining social welfare require a cardinal form of utility since we will explicitly apply mathematical operators to them. Of course, once we make the move to cardinal utility functions this means making interpersonal welfare comparisons, which opens us up to a wide variety of philosophical critiques about our ability to compare the value of ten dollars to a miser and a beggar, etc. Yet, this is precisely the point of the exercise - once we decide that value judgments that permits some to lose and others to gain are permitted then we necessarily have to ascertain how actions affect the utility of each person involved. Many welfare economists find that unacceptable. My view is that as political economists we are already far beyond worrying about this problem since we are making a whole host of other assumptions about expected utilities, risk aversion, etc that require us in any case to assume cardinal utility.

So once we have assumed cardinal utility, we can aggregate individual utilities in some manner to create a 'social welfare function' $W(u_i, ...u_n)$. The simplest way to do this is simply to add together all the individual utilities, a concept known as *utilitarianism*. The utilitarian social welfare function can be written as $W(.) = \sum_{i=1}^{N} u_i$. This basic utility function assumes that all individuals' welfare is valued equally.

We can adapt it slightly be applying a different 'weight' a_i to each individual, producing the *generalized utilitarian* social welfare function: $W'(.) =$

$\sum_{i=1}^{N} a_i u_i$. The advantage of the generalized function is that we can put more weight on the poor and less on the rich. But it is not intuitively obvious exactly what series of a_i one should use (the inverse of relative income perhaps).

Note that utilitarianism is in general not equity-producing unless we add the further assumption of declining marginal utility of income. If all individuals receive the same effect from a ten dollar gift then basic utilitarianism is indifferent between giving all the money to Bill Gates and giving it all to the poorest citizen. Generalized utilitarianism would give all the money to the person with the highest a_i. So we need to add more structure to the shape of the utility function for utilitarianism to produce equitable outcomes. Note though, that unlike Pareto optimality, utilitarians would be indifferent about taking $10 from Bill Gates and giving it to a beggar, even though this is *not* Pareto optimal, as Bill Gates is being made worse off.

4.3 Nash and Rawlsian Utility

To obtain a social welfare function with a higher valuation of equality than utilitarian ones we need to prevent the function from arbitrarily taking away the income of poorer citizens or leaving people without income at all. The Nash utility function does just that. Instead of being additive it is multiplicative: $W_N(.) = \prod_{i=1}^{N} w_i$ and its generalized form is $W'_N(.) = \prod_{i=1}^{N} w_i^{a_i}$. Notice that if we take away the utility of any individual these utility functions become zero (unless $a_i = 0$ in the generalized model).

Furthermore, the basic Nash utility function is entirely egalitarian in that it is maximized when all individuals have the same level of income. At the two person level this is easy to show. Imagine we have $2x$ to split between the individuals. If we split it equally and each receives x then the Nash SWF equals x^2. However, if we split it unequally, giving ϵ more to the first person and taking that away from the second we end up with the

Nash SWF value of $(x + \epsilon)(x - \epsilon) = x^2 - \epsilon^2$. Obviously, the generalized model could alter this somewhat but will push the SWF away from equal distribution of incomes (ironically then, as opposed to the generalized utilitarian model, the generalized Nash model always *reduces* equality).

A different though related approach is the Rawlsian utility function, which follows John Rawls' maximin principle that any change in the distribution of income must increase the income of the poorest citizen first. Mathematically this is easily expressed as $W_R(.) = min\{u_i, ...u_N\}$. It is easy to see here that if we increase the income / utility of any citizen but the least well-off, the Rawlsian social welfare function will not change - it will remain constant unless the utility of the poorest individual is increased. Thus the social welfare function is focused on 'leveling up'.

A brief note on the *shape* of the utilitarian, Nash, and Rawlsian social welfare functions is now worth making. Imagine a standard Cartesian plot (i.e. x, y) with the x-axis being the utility of one citizen and the y-axis being the utility of another citizen (and ignore all the other citizens for the moment). Indifference curves for the social welfare function are pairs of utilities between the two citizens (u_1, u_2) between which the social welfare function is indifferent. Put simply with the utilitarian social welfare functions, indiffference curves are negatively sloping straight lines, with a slope reflecting the ratio of the a_i's attached to each citizen by the social welfare function. Importantly, the utilitarian social welfare functions will be indifferent between giving all resources to one citizen or all to the other if $a_1 = a_2$. The Nash social welfare functions appear differently - as downwardly sloping curves that bend inward towards the origin and asymptote toward the axes. These imply that higher utilities always emerge from a more equal distribution of resources. Finally, the Rawlsian utility functions look like a capital L, they are flat for a given utility level of one of the citizens for any increase in the utility of the other. Rawlsian utility can only increase if both citizens receive the same amount of resources.