Lecture Notes on Preferences

Andrew T. Little*

June 2017

Most textbook prefaces are alike. The author taught a class for years, and was always dissatisfied with the existing books. Either they are too technical or too hand-wavy, have applications in the wrong field, or do not present the material in the right order. So the author turns their notes into a book, complains about how much longer it took than expected, and perhaps a few other instructors are spared the effort. (Or the process repeats.)

The origin of these notes is similar but less ambitious. I start all of my game theory classes with a couple hours of lecture on what it means for individuals to have preferences, and when it makes sense to give them a "utility function" to represent these preferences.[1] Most "serious" treatments of this material are quite technical and geared towards readers who are already on board with mathematical models of decision-making. At the other end, the books I have assigned (and all others I have read) generally treat this topic superficially or informally, or maybe stick a bit of technical discussion in an appendix. I think this is a mistake, particularly in the context of political science, a discipline where many are skeptical that formal models where actors optimize utility functions is a worthwhile endeavor.

The challenge of teaching this material is that it is quite abstract, and involves unfamiliar notation. So, the primary motivation of writing these notes up is to have something my students can refer to beyond notes they took in class, which, given the size of classes at a large university and the poor handwriting of their instructor, may be unreliable.

If the notes are useful to other instructors, students, or any interested reader, all the better.

*Department of Political Science, UC Berkeley. andrew.little@berkeley.edu.

[1] When teaching undergraduates, I also like to do a few lectures about when it makes sense to talk about group preferences (i.e., some basic social choice theory). At some point I may get around to writing these up as well.

Scope

Most social science boils down to people making choices. When studying democratic politics, perhaps the core choice is citizens deciding whether to vote in an election and which candidates to choose if they do. Politicians also make many choices, such as setting policy when in power, or deciding what actions to take to try and win the next election when out of power. In less democratic contexts, citizens and opposition parties may not be able to influence policies with these chocies, so choosing whether to join a protest movement or launch a civil war may be the more relevant decisions.

A challenge for individuals making choices is that how much they like the *outcome* of these decisions usually depends on factors outside of their control. These outside factors can be lumped into two groups. First, the wisdom of an action may depend on what others do. For example, an aggrieved citizen deciding whether to join a protest against a repressive regime may only want to take to the streets if others join. She also will likely care about whether the government will order the police to use violence to disperse protests, and if the police will carry out this order. Second, other factors outside of any individual's control affect the desirability of protesting, like whether it rains.

These challenges to individual decision-making also pose a challenge for researchers. If someone's protest decision depends on factors outside of their control, we need not only know what outcomes people like, but their conjectures about how these outside factors will play out. Game theory is a tool developed to address these complications.

To build a simple theory of preferences, here we will ignore these complications. Rather, we will focus on simple choice problems where individuals have full control over the outcome. This is often a fair simplification. For example, when deciding what to eat for lunch, we often don't care what other people are going to order. Since politics is inherently about interactions with other people, clear examples are a bit harder to come by. But as a first approximation, when choosing who to vote for, one could simply want to vote for whoever they like best. (And in a large electorate where individual votes are unlikely to sway the outcome, this may be perfectly reasonably even when thinking about who others are likely to vote for.)

If this seems unrealistic or highly limiting, you are right! But understanding preferences in this simpler context creates a useful scaffolding for analyzing more complex choices.

Choices

To define a choice problem, we first need to know what options are on the table. The mathematical way we represent this is with a *set*. A set is simply a collection of objects. We write sets with curly brackets, and within the brackets list the objects separated by commas. In an an election with two candidates – say, a republican and a democrat – the options for a voter could be represented with the choice set $\{r, d\}$. If we also give the voter the option to abstain, we could write the choice set $\{r, d, a\}$.[2]

Sets can have any number of objects. This number need not be finite: for example, the set "all positive integers," which we usually write $\{1, 2, 3, \ldots\}$, has an infinite number of objects. Choice sets can also be *continuous*, like when deciding how many hours a day to spend thinking about preferences. The common way to represent this set is not with curly brackets – where would we even start listing every possible number in this range? – but as an *interval* $[0, 24]$.[3]

As with pets, it helps to name our sets. In the voting example, a natural name is to let $V = \{r, d, a\}$ be the choice set. We can then describe the choice problem as choosing a $v \in V$, where $\in$ is a symbol which can be read as "in", or, more technically, "an element of".

Relations

At their core, preferences are about comparisons between objects, asking the question *which do you like best*. This can be quite complicated when the choice set is large ("how can I decide between the 45 brands of peanut butter in the supermarket?") An easier problem is to start by assuming that people can compare *pairs* of objects, and see how far that takes us.

A mathematical way to represent comparisons of two objects in general is with a **binary relation**. (We'll usually drop the "binary" and just write "relation"). Generally, think of a relation as a well-posed question with a yes or no answer. If the answer to "do a and b have relation R", is "yes", we write $a\, R\, b$. If not, write $a \not{R} b$. To say a relation R is well-defined on a set S, we mean that for any two objects $a \in S$ and $b \in S$, either $a\, R\, b$ ("yes") or $a \not{R} b$ ("no").[4]

---

[2]The order in which we write the objects doesn't matter: $\{r, d, a\}$ is the same set as $\{a, r, d\}$.

[3]We could get into a debate about whether time is discrete or continuous. Let's not.

[4]A bit more formally, let S^2 be the set of all ordered pairs (a, b), $a, b \in S$. A relation is a subset of S^2 ($R \subseteq S^2$), which identifies the (ordered) pairs that satisfy the relation. By convention we write $a\, R\, b$ if $(a, b) \in R$, and $a \not{R} b$ if $(a, b) \notin R$.

Let's think through some examples. Suppose our set under considerations is some restaurants in the city of Berkeley. One relation we could apply to this set is the "on the same street" relation. Call this the SS relation, and if restaurants a and b are on the same street, write $a\,SS\,b$. If a and b are not on the same street, write $a \not{S}\!\not{S}\, b$. If we know the ages of all of the restaurants (or can look them up), we could also define the "is older than relation", and write $a\,O\,b$ if a is older than b. If a is not older than b, write $a \not{O}\, b$.

Closer to the kinds of questions we study here, consider a student who likes eating out and has been to all the restaurants in question enough to form an opinion. A relation of interest is a relation P, where $a\,P\,b$ means she thinks restaurant a is at least as good as restaurant b.

Preference Relations

We will build our theory of preference by assuming our decision-maker has *weak preference relation* P over her choices S (I will drop the "weak" unless it helps clarify). The fact that this is relation means that whenever we confront her with two choices a and b, she can either tell us "a is at least as good as b", in which case we write $a\,P\,b$, or she can say "a is not at least as good as b", in which case we write $a \not{P}\, b$.

In order for the relation to capture the idea of *preferences*, we will need to add a bit of structure. In particular, we say that P is a weak preference relation if and only if it has the following two properties:

1. (Comparability) For all $a, b \in S$, $a\,P\,b$ or $b\,P\,a$ (or both)[5]

2. (Transitivity) For all $a, b, c \in S$, $a\,P\,b$ and $b\,P\,c$ implies $a\,P\,c$

As the name indicates, the first condition means that all objects are comparable, in the sense that either a is at least as good as b, or b is at least as good as a.[6] Both of these can be true: this will capture a scenario where a and b are equally good (more on this soon). What is ruled out is for neither relation to hold: i.e., you can't say "a is not at least as good as b, **and** b is not at least as good as a".

The second condition is about preferences having a natural order. It states that if a is at least as good as b, and b is at least as good as c, then a is at least as good as c.

[5]Implicit in this definition is that if we pick the same object to compare to itself, it must be the case that $a\,P\,a$ ("a is at least as good as itself"). Sometimes comparability is only defined as "for distinct a and b" and this property ("Reflexive") is explicitly added to the definition of a preference relation. Both formulations end up being the same.

[6]This is sometimes called "completeness."

Counterexamples (skippable) It may help to think through some relations which do not have these properties.

For comparability, take a set of people, and the relation G which represents "grew up in the same town as." Clearly, if we compare two people a and b who grew up in a different town, then $a \not G\, b$ and $b \not G\, a$. A counterexample closer to a relation we will use later is "owns strictly more cats than" (i.e., if a has c_a cats and b has c_b cats, then $a\, MC\, b$ if and only if $c_a > c_b$). A problem arises if we try and compare two people with no cats. Then $a \not{MC}\, b$ and $b \not{MC}\, a$, so the relation is not comparable.

An illustrative example of a kind of relation which is not transitive are ones that ask whether members of the set are "close" to each other. For example, if our set is a group of people sitting on a bench, and the relation captures "is sitting next to." If we number the people on the bench from left to right and call this relation NEXT, it will be true that 1 NEXT 2 and 2 NEXT 3, but $1 \, \cancel{\text{NEXT}} \, 3$.

Commentary on comparability and transitivity (skippable) Would it make sense for someone to have preferences which aren't comparable? One could combatively say no: that whenever faced with two choices, people must be able to say that one is at least as good as the other. Whether this hard response is justified is a matter of taste. Personally, saying everyone should always have comparable preferences over the choices they face strikes me as reasonable, if not obviously true.

Even if we shy away from the scorched earth response to the notion of incomparability, recall our aim is to make claims about how people will behave given their preferences. If there are choices that can't be compared, then we *shouldn't* have anything to say about what choice will be made. So, the fact that we say someone who won't compare choices in this manner does not have proper preferences is on firm ground.

Would it make sense for someone to have non-transitive preferences? That is, should it be possible to think a is better than b, b is better than c, but a is not better than c? Like with comparability, a hard response is defensible: the impossibility of a such a "cycle" flows directly what it means to have preferences. And again, if a decision-maker is choosing between a, b, and c with these cyclical preferences, they we *shouldn't* be able to make predictions about what she should do, so we might as well say her preferences aren't coherent.

If that doesn't convince you, here is a common argument against intransitive preferences. It is called the *money pump*. Suppose Inigo the Intransitive has preferences such that a is strictly better than b, b is strictly better than c, and c is strictly better than a. Also, assume that Inigo likes to have more money (or any easily divisible object: cheese, beer, time to do their next exam, etc.). To make the argument clear, suppose the difference between a and b, b and c, and c and a are worth at least one dollar to Inigo.

Now, imagine I am currently in possession of an a, and Inigo has a b and a c, and ten dollars. I approach Inigo, and make him an offer: "How about you give me your mediocre b and a dollar for this shiny a?" Since Inigo thinks a is strictly better than b (and the difference is at least worth a dollar), he will accept this trade. After trading, I have b and a dollar, and Inigo has a, c, and is down to nine dollars. Now I can say "You may have just given up your b, but you'd rather get it back in exchange for c and a dollar, right?" Since Inigo likes b strictly more than c, he accepts this trade too. Now I have c and two dollars; Inigo has a a, b, and 8 dollars. Finally, since Inigo values c more than a, he will trade me the a and a dollar for the c. Now we are back to where we started in that I have a and Inigo has b and c, though I have just fleeced him of three dollars. And can repeat this process until Inigo is broke.

That we used dollars as an example is immaterial; the point is that someone with transitive preferences can be easily exploited by a series of trades he purportedly always likes.

Rationalizability

What is so special about linking transitivity and comparability to our notion of preferences? First, they seem like things that should be true for our relation to capture our intuitions about what it means for something to be at least as good as something else. Second, and more importantly, by assuming people have a preference relation meeting these assumptions, we will be able to make some clear statements about (1) whether they have a "best choice", and (2) also be able to represent this preferences with a utility function.

First consider the question 1. What does it mean for a choice to be "best"? A natural definitions is the following:

Definition Suppose a decision-maker has a preference relation P over choice set S. Then a choice $s^* \in S$ is **rationalizable** if $s^* \, P \, s'$ for all $s' \in S$.

This definition states that a choice is rationalizable if it is at least as good as any other choice. (We often add stars to choices which are optimal. Similarly, we add a "$'$", pronounced "prime", to variables that indicate alternative choices.) Does having a preference relation over the choices imply there is a rationalizable choice? With a finite number of choices, the answer is yes!

Theorem 1. *Suppose a decision-maker has a preference relation P over a finite choice set S. Then there exists at least one $s^* \in S$ which is rationalizable.*

Here is an intuitive proof. Pick any element of the set, and call that our candidate for being rationalizable. Compare this candidate (call it s_1) to each other element of the set. If our candidate is preferred to every thing else, then we are done. If there is a s_2 such that $s_1 \not P s_2$, then call s_2 our second candidate. By comparability, s_2 is preferred to s_1 ($s_2 \, P \, s_1$). Now check if there is anything which s_2 is not preferred to. If not, we are done. If so, then there is an s_3 such that $s_3 \, P \, s_2$, and by transitivity $s_3 \, P \, s_1$. More generally, our new candidate choice is always preferred to all previous candidates. So, no choice can become a candidate more than once. Finally, since there are a finite number of choices (say n), either we find a rationalizable choice from the first $n-1$ candidates, or the last remain choice which has not yet been a candidate is preferred to s_{n-1} and hence all other choices, and is rationalizable.

So what? The main point of this result is that as long as we can find a reasonable choice when comparing two options at time, we can find a reasonable choice among any finite number of options. We'll make a connection between this idea and assigning "utilities" to each choice soon.

Why is the finite part important? An important part of the proof is that we can eventually run out of potential "better options". Here is a simple example of how things can go wrong if not. Imagine your choice set is "how much money will the author give me", and you can pick any number. As long as your desire for money is insatiable, for any proposed choice x, you would rather pick $x+1$.[7]

Another important aspect of this result is there might be more than one rationalizable choice. At first this may seem unsatisfying, but remember we are allowing for the possibility that two choices are at least as good as each other. So, it is possible that two or more options will be "tied for best".

[7]This is not to say having an infinite number of choices always means there is no rationalizable choice. There are lots of useful results about what assumptions on the choice set and preferences lead to a rationalizable choice, an in applied work we generally set up problems in way that these hold.

Strict Preferences and Indifference

Before getting to utility, it is useful to define some new relations. Suppose we are comparing two objects a and b. Since the order of the comparison matters, for any relation R there are four possibilities for how their relations shake out. Either (1) $a\,R\,b$ and $b\,R\,a$, (2) $a\,R\,b$ and $b \not R\, a$, (3) $a \not R\, b$ and $b\,R\,a$, or (4) $a \not R\, b$ and $b \not R\, a$. The following table illustrates how we might think of these four possibilities with a preference relation:

	$b\,P\,a$	$b \not P\, a$
$a\,P\,b$	equally good	a strictly better
$a \not P\, b$	b strictly better	not comparable

The top left cell is the case where a is at least as good as b and b is at least as good as a. So, they must be equally good. The word we generally use for this is "indifference", and we will write $a\,I\,b$ when this is true.

In the top right cell, a is at least as good as b and b is not at least as good as a. In this case we say that a is strictly preferred to b, and write this $a\,SP\,b$. Conversely, in the bottom left cell b is strictly preferred to a: $b\,SP\,a$.

How about the bottom right cell? For a general relation this would be possible. Returning to the "is on the same street" relation, for two restaurants on different streets it will be true that $a \not{S}\!\not{S}\, b$ and $b \not{S}\!\not{S}\, a$. However, the fact that preference relations are comparable (in the sense of property 1) means that this combination is not possible. To reiterate, if I ask you to compare two objects and you are not allowed to say "a is not as good for me as b, and b is not as good for me as a". Or, more precisely, if you were to tell me this, then you do not have a proper preference relation over your choices.

Put another way, having a preference relation means that if I pick out two options, exactly one of the following three options must be true:

1. $a\,SP\,b$
2. $b\,SP\,a$
3. $a\,I\,b$

It is often more convenient to work with this notion of preferences.[8]

[8] A slightly less elegant way to define a preference relation would be to start at this point, saying when comparing any two objects one of the three options above must be true. For this to make sense as a preference, we also need to assume *SP and I* are transitive, which is again natural.

Some notes on Indifference and Strict Preference (skippable) We derive I and SP to make some arguments easier, but it may be instructive to think about some properties of their relations.

For example, we can ask if they are weak preference relations as defined above. While they are relations which we use to describe preferences, what matters is if they are comparable and transitive.

Transitivity is straightforward to check for both. If $a\,I\,b$ and $b\,I\,c$, then we know $a\,P\,b$ and $b\,P\,c$, and so by the transitivity of P, $a\,P\,c$. Similarly, we know $c\,P\,b$ and $b\,P\,a$, so $c\,P\,a$. Since $a\,P\,c$ and $c\,P\,a$, $a\,I\,c$. This proves I is transitive. The argument to show SP is transitive is just hair tricker; see if you can figure it out before reading the footnote.[9]

However, it is easy to see that both are not necessarily comparable.[10] For indifference, suppose we have two objects such that $a\,SP\,b$. From this we know that $a \not I\, b$ and $b \not I\, a$ (since $b \not P\, a$). Similarly, if there are two choices which the decision-maker is indifferent between, she does not strictly prefer one to the other, so strict preference is not comparable either.

Shelves

We are now close to introducing the concept of utility. As an intermediate step, I find it helpful to first think of a visualization of preference relations. A nice physical analog to what a preference relation does is a make a "shelving" system to display the choices.

Suppose we have a set of objects $\{a, b, c, d\}$, and a preference relation over the set. First, compare a and b. If $a\,SP\,b$, we put a on a higher shelf than b. If $b\,SP\,a$, b goes on a higher shelf. If $a\,I\,b$, they go on the same shelf. (Remember exactly one of these are true!) For illustration, suppose $a\,SP\,b$, so our shelves look like the left half this this figure:

[9] To show SP is transitive, we need to show that $a\,SP\,b$ and $b\,SP\,c$ implies $a\,SP\,c$. Clearly $a\,P\,b$ and $b\,P\,c$, so $a\,P\,c$. However, we also need to show that $c \not P\, a$. We can prove this by contradiction: If $c\,P\,a$, we could combine this with transitivity and $a\,P\,b$ to get $c\,P\,b$, which contradicts $b\,SP\,c$.

[10] For a relation in general to have a property (say, comparability), this property must hold for any set we place it on. For example, if we only place the indifference relation on sets with one choice, it will always be comparable. However, once there are any objects which are strictly preferred to others comparability will fail.

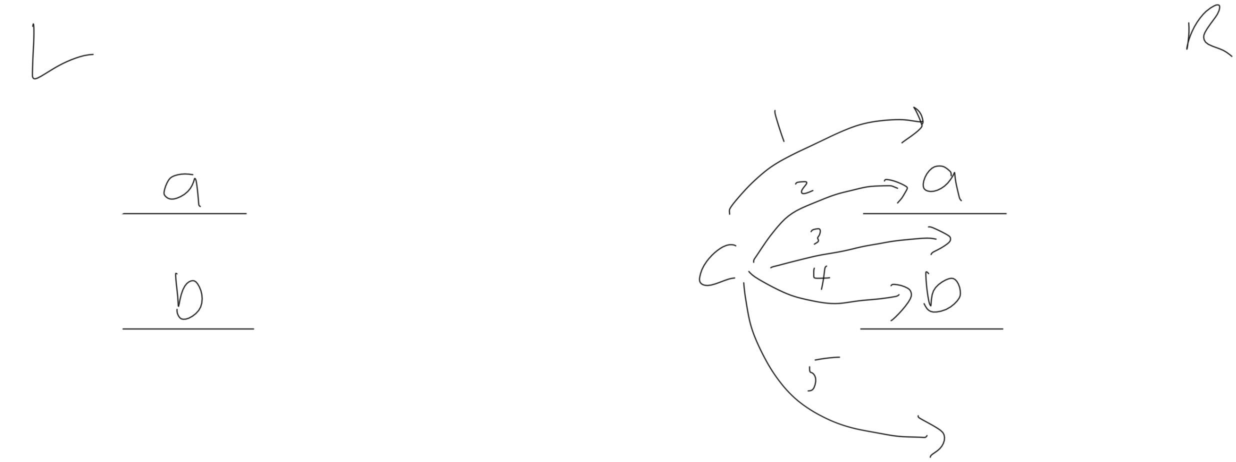

The right half of the figure shows how to think about how c compares to the other choices. There are now five possibilities. Option c could be (1) better than both a and b, (2) equally good as a, (3) better than b but worse than a, (4) equally good as b, or (5) strictly worse than both. Importantly, it can't be the case that, for example, c should go above a but below c; this is just the visual analog of transitivity.

In cases 2 and 4, we put c on the shelf with the choice it is equally good as. In the remaining cases, we "build a new shelf", either above the other two (1), between them (3), or below them (5). If $c\,SP\,a$, things now look like the left panel:

For d, there are again five choices about where to put this option in relation to the others. More generally, the new option to be added to the shelves will either tie a choice already on the shelves, in which case we put it on the shelf with the options it is equally good as.[11] If

[11] In general, if there are n existing shelves, there are $2n+1$ possible places to put the new object: n on an existing shelf, and $n+1$ requiring a new shelf.

it is the best option or lies between two existing shelves, we build a new shelf.

It should seem intuitive that we can "build a shelving unit" in this manner for any possible finite set with a preference relation.[12] With each new object we add, it either goes on an existing shelf, or we build a new shelf which represents the order of the preferences.

Utility

Once we have built the shelving unit for our preference relation, coming up with a "utility function" which assigns higher numbers to more preferred options is trivial. We just want to give higher numbers to choices on higher shelves.

Using our example from before, suppose d was better than all previous choices, so our shelves look like this:

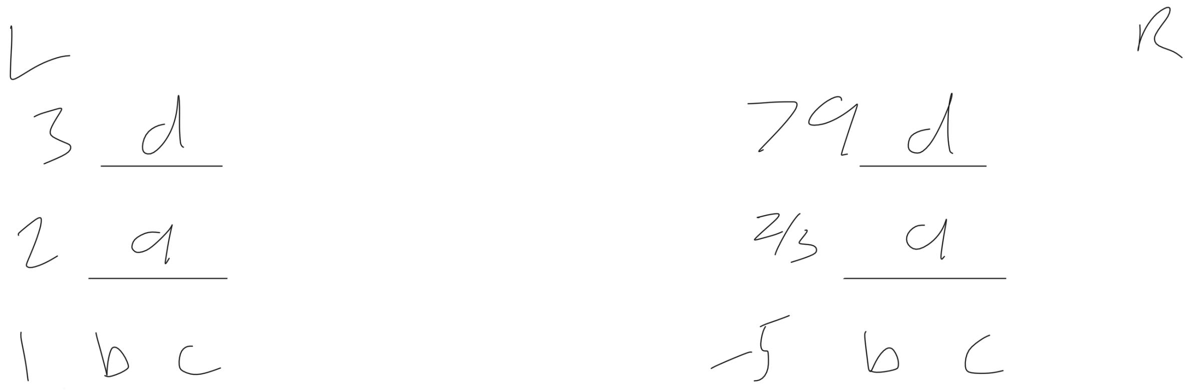

The left panel shows a simple way we can construct a "utility function" from our shelves. Options b and c, on the lowest shelf, get a utility of 1. The middle, a, gets 2. Finally d, at the top, gets 3.

If all we care about is giving higher numbers to better options, the utility function in the right panel works just as well. We can use negative numbers, fractions, or any number we want, as long as the order matches the shelving.

This clearly extends to any shelving system build from a finite number of choices. We can assign 1 to choices on the lowest shelf, 2 to choices on the second lowest, 3 to those above that, etc. until we reach the top shelf. If we want to know whether one choice is preferred to another, we just need to know if it is on a higher(-numbered) shelf. Or we can pick any other numbering system that puts higher numbers on higher shelves.

[12] We could do the same with an infinite set, it would just take an infinite amount of time to finish.

In our simple environment of choices made in isolation, this is all we need to make sense of utility. Formally:

Definition Suppose a decision-maker has a preference relation P on a set S. Then the utility function u **represents** P if for all $a, b \in S$, $u(a) \geq u(b)$ if and only if $a\,P\,b$.

An immediate consequence of this is that if $a\,SP\,b$, $u(a) > u(b)$, and if $a\,I\,b$, $u(a) = u(b)$. (Why?)

As indicated above, there are always many – in fact, an infinite – number of utility functions which can represent the same preference relation. While this may seem indeterminate, it is good in the sense that we need not fret about the fact that it seems weird to assign particular numbers to particular choices ("what does it mean to get utility of 1 from protesting and 3 from staying home?"). There is no inherent scale here, just an ordering. We can pick any utility numbers which preserve that ordering.

Now that we are in the world of utilities, we can stop thinking about individually comparing every pair of choices to see how good they are: each choice now gets a measure of how much we like it. An immediate consequence of this is a more convenient definition of rationalizability:

Definition Suppose a decision-maker has a preference relation P over choice set S represented by a utility function u. Then a choice $s^* \in S$ is **rationalizable** if $u(s^*) \geq u(s')$ for all $s' \in S$

While this is just a restatement of the definition using preference relations, one consequence is that it is now much easier to explain why there is a rationalizable choice. Once we have a utility function in place, all we need to do is find the choice with the highest utility, which always exists if we have a finite number of choices.

Importantly, we need not think that people actually assign utility numbers to all of their decisions before picking. However, if they pick a rationalizable choice with our first definition, that will be observationally equivalent to maximizing their utility. Using utility to represent preference is an extremely useful tool for social scientists even if it is not something consciously used by the people we study.

Commentary

We have arrived at our main goal: showing that representing how much a decision-maker likes choices with numbers called "utility" is just a convenient way to represent her preferences, where we placed pretty minimal demands on what it means to have coherent preferences. All our decision-maker needs to be able to do is compare any pair of choices, and in a way that is transitive.

We also showed why picking a decision which maximizes her utility is a reasonable thing to do, in the sense that any other choice would imply picking something she would like strictly better.

To wrap up, let's consider some more big picture questions.

What does it mean to be rational? A lot of criticism of formal theory is of the form "these models assume people are rational utility maximizers, but they aren't, so the models are useless." But whether people meet the common-sense notion of rationality – even if this clearly defined – has little to do with whether they have coherent preferences, which generate a utility function, which they maximize when picking a rationalizable choice.

Many accusations of irrational behavior boil down to saying "in his shoes, I would have done something differently." But with our simple conception of preferences above, "being in someone else's shoes" is actually quite demanding: it would mean having the exact same preference relation over the choices. While straying a bit from what formal analysis can tell us, a reasonable conclusion from this style of thinking is that if someone makes a choice we find baffling, we shouldn't label them irrational but try and understand why their preferences led them to make a particular choice.

A common misconception is that "rational choice theory" assumes people always try to maximize their wealth or power. I find the terminology "rational choice theory" counter-productive, precisely because it tends to conjure up unproductive debates about whether we should assume people are rational based on particular kinds of preferences. Better to think of what we are doing here as just "choice theory", where one criteria we might use to make predictions about choices is rationalizability.[13] We have said nothing here about what the "rational" preferences are to have; we simply said preferences should be comparable and transitive. Wanting less money or power is a perfectly coherent set of preferences given these requirements.

In practice, formal models generally tend to make assumptions that people want more money, more respect, and policies which more closely match their ideals. Why? First, these assumptions tend to line up with how we believe people think from introspection and our

[13]Without having done a survey, I don't think I am alone in this preference: I have never heard someone who uses formal models in their research and is below the age of 40 describe themselves as a "rational choice theorist."

lived experiences. Second, models with these kinds of assumptions are often tractable and lead to good predictions of how people actually behave. These explanations are probably inter-related: these assumptions seem reasonable precisely because the behavior we observe among other people generally seems to match wanting more money, power, etc. One aim of writing formal models is to sharpen these intuitions and conclusions.

However, models that consider other types of preferences can be powerful as well. It seems pretty clear both from casual observation and rigorous research that people sometimes care about the welfare of others, particularly those close to them. People also may prefer to take actions which they view as morally correct, even if it leads to a material loss. This may pose a problem for particular models which assume people don't hold these preferences, but it poses **no** problem for this general approach to modeling. It just means people hold a different set of preferences.

Where things get dicier. Now, the less good news.

Recall a nice property of assigning utility functions to represent preferences is that it doesn't really matter which numbers assign as long as they preserve the order implied by the preferences. This is because preference relations don't tell us how much better one option is than another. And, fortunately, in the simple decision-making environment we studied, that is fine: all that matters is the order of the options and which is (or are) best.

A reason this doesn't cause any problems is that we worked in a world with no uncertainty: you choose option b, you get option b for sure. In real-world decision-making, we are often unsure about what we are going to get. We won't go into a full discussion here, but just to see why magnitudes are important with uncertainty, consider a decision-maker who thinks option a is best, option b is a neutral, and option c is bad. However, she can't choose a directly, but can only choose option b (maybe not investing in a project), but can only choose a *lottery* over a (investing and the project goes well) and c (investing and the project does not go well). If there is more "downside" to going from b to c than there is "upside" to going from b to c, she may prefer b. If the opposite is true, she might take the gamble.

More importantly, our preference relation does not tell us about these magnitudes, so it can't by itself predict how people behave under uncertainty. For that we need to make stronger assumptions about preferences and attitudes towards risk.

Interacting with other people can generate similar problems. If what is optimal for me to do depends on not just my choice but others' choices, I need to form conjectures about what they will do to know what choice is best. Again, this is essentially what game theory is about, which builds on our theories here.

A final complication is that we often want to simplify our analysis by talking about what groups want. We might ask what the preferences of a political party are, or an entire electorate. Even if all the individual members of the party or country have coherent preferences, it may be tricky to say what they collectively want.