Snapshots of modern mathematics
from Oberwolfach

№ 9/2019

On Logic, Choices and Games

Paulo Oliva

Can we always mathematically formalise our taste and preferences? We discuss how this has been done historically in the field of game theory, and how recent ideas from logic and computer science have brought an interesting twist to this beautiful theory.

1 What do I want? I want the best!

Suppose that you need to ask a friend to go to a shop and buy something for you, say an umbrella. Your friend might ask, "Sure, no problem, what kind of umbrella do you want"? You might really like to say, "Just bring me the best you can find". But what might be "best" for you, might not be "best" for your friend. So you need to be more specific, "A large one, with light colours, and costing no more than 20 pounds". This should probably give your friend enough information for her to make this purchase on your behalf.

Figure 1: A Choice of Umbrellas. It is clear that, depending on their price or other factors, there could be more than one "best" umbrella, so putting them in a total order might not be straightforward.

What you have done here, in mathematical terms, is to convey to your friend your *preference* over any collection of umbrellas. Given two umbrellas that cost 15 and 30 pounds, respectively, you prefer the one that costs 15 pounds. Given an umbrella that is large and one that isn't, you prefer the large one, and so on. This is a way of *ordering* the choice of umbrellas. It does not define a *total order* on the set of umbrellas, which would mean that every umbrella can be rated as better or worse than each the others, because using your stated rules of preference, there might be many equally good (or equally bad) umbrellas.

One trivial but extremely interesting observation here is that you do not know which umbrellas are actually available at the time you express your preference. In order to make this point clear, let us simplify our preference and say that we just want the cheapest available umbrella. Suppose our friend walks into a shop and sees the following options:

Umbrellas	Price
Umbrella 1	9.99
Umbrella 2	12.00
Umbrella 3	5.25
Umbrella 4	11.10
Umbrella 5	9.99

In this case, she will know exactly which umbrella to choose, namely, Umbrella 3, the cheapest one. If, on the other hand, Umbrella 3 was no longer available, then the cheapest umbrellas would be 1 and 5 instead, and there would be no way to choose between them according to our preference.

Let us carry on making this process a bit more formal and mathematically precise. Say U is the set of all umbrellas and let us take the positive real numbers $\mathbb{R}^+$ for the set of possible prices. Each shop contains a collection $X \subseteq U$ of umbrellas together with a price tag for each of these. We can model the price tags as a function $P\colon X \to \mathbb{R}^+$, since in each shop each umbrella has exactly one price.

Now, when we say that we want the "cheapest umbrella" what we are saying is that given any collection of umbrellas $X \subseteq U$ with an associated price function $P\colon X \to \mathbb{R}^+$, we prefer the cheapest of these, that is, our set of preferred umbrellas is

$$\left\{ u \in X \,:\, P(u) = \min_{u' \in X} P(u') \right\}.$$

This type of mapping, from the price function to the set of cheapest umbrellas, is one that shows up so much in mathematics that it has been given a name: *argmin*. Given any function $f\colon X \to \mathbb{R}^+$ (not just the price function), the map argmin is defined to be

$$\mathrm{argmin}(f) = \left\{ x \, : \; x \in X, f(x) = \min_{x' \in X} f(x') \right\}.$$

Alternatively, we could have said that we would like the biggest available umbrella, which in this case, if we define the function $f : X \to \mathbb{R}^+$ to be the size of each umbrella, would be making use of the similar *argmax* function:

$$\mathrm{argmax}(f) = \left\{ x \, : \; x \in X, f(x) = \max_{x' \in X} f(x') \right\}.$$

There are many interesting aspects of functions such as argmin and argmax. One is that these are *higher-order functions*, which means that they are functions that take other functions as input. Remember that $f \colon X \to \mathbb{R}^+$ itself is a function.

Another interesting fact is that one can always replace an argmin by an argmax by simply inverting the preference order: the cheapest umbrella is also the one that leaves me with most cash after the purchase. In fact, one can almost always view a person making a decision, or, as we will see later, a player of a game, as a *maximising* agent, since we can, in most cases, describe a preference by maximising over a suitable ordering on X.

2 Dining with Keynes

The notion of a maximising agent is important, so let us a give another concrete example. Suppose three people (John, Mary, and Peter) need to decide at which restaurant they are all going to eat tonight. Let us further suppose that they have two options, either a pizza place or a steak house, so $X = \{$Pizza, Steakhouse$\}$. They decide to hold a democratic vote, and they will go wherever the majority chooses. They all write their choices on a piece of paper, and reveal it at the same time.

It could be that John prefers Pizza to Steakhouse, in which case John is a maximising agent over the relation Pizza $>$ Steakhouse. On the other hand Mary could prefer Steakhouse to Pizza, so Mary is also a maximising agent but over the inverse ordering Steakhouse $>$ Pizza. And Peter might not actually care. He is also a maximising agent over the ordering in which Steakhouse and Pizza are both equally good, that is, are both maximal.

So what we see here is that by choosing different orderings on the set of alternatives X we can actually view all three decision makers as maximising agents. Except, unfortunately, this does not always work. In 1936 an economist called John Maynard Keynes (1883–1946) wrote a very influential book called

"The General Theory of Employment, Interest and Money". In Chapter 12 of that book Keynes described an interesting puzzle which is related to the example above. We know that John prefers Pizza and Mary prefers Steakhouse. But suppose now we replace Peter with Keynes, and Keynes does not mind which restaurant is chosen, as long as it is the one he voted for! It might sound like a strange preference at first, but being on the winning side has its appeal.

Then how do we describe Keynes preference in this case? What is he maximising? Is it possible to choose an ordering on the set $X = \{$Pizza, Steakhouse$\}$ so that we can describe Keynes as a maximising agent? The answer is no, we can't. Keynes does not really care whether Pizza or Steakhouse is chosen, all he cares is that the final outcome is the same as his own choice.

One possible solution, and it is the one economists normally adopt in such cases, is to define Keynes preferences over the set of all combinations of votes, so we look at the triples $X \times X \times X$, and define a preference over these. The elements of these triples are the three votes of the three players, say in the order John, Mary and Keynes. So the triple (Pizza, Steakhouse, Steakhouse) means that John voted for Pizza, Mary voted for Steakhouse, and Keynes also voted for Steakhouse. In this case Steakhouse won. Mary is happy because she indeed prefers Steakhouse. Keynes is also happy because the final outcome was Steakhouse, and that is what he chose.

Although this does work, and we were able to describe Keynes as a maximising agent, it is clear that this is not completely satisfactory. Keynes preference is only over his own choice and the final outcome, so why should we need to consider all the other players' votes as well when defining Keynes preference? What if there are 100 diners involved, or one million. We would need to define Keynes preference over the product space $X \times X \times \ldots$ one million times!

Luckily we have recently come up with a much neater solution to this problem, which we describe in the next section.

3 The power of generalisation

The main problem above is that we were too attached to the argmax function. If we want to see each player as a maximising agent, we then need to artificially change the set of alternatives, in order to be able to describe a suitable order over which the agent will maximise. But what if we forget about argmax, and accept that sometimes people are not really maximising over the set of alternatives? It is a scary thought since maximisation is normally associated with *rationality*, and it might look like we cannot say much about decisions and games if we do not assume that the players are rational. Luckily, since in real-world economics it is not clear at all that players are rational, it turns out that we can say quite a lot.

Let us look again at the *type* of the argmax function, where type is being used here in the sense of computer science, or programming, which is just the description of its input and output sets. When choosing over the set of all umbrellas, argmax has type:

$$\text{argmax}\colon (U \to \mathbb{R}^+) \twoheadrightarrow U,$$

where I am using $\twoheadrightarrow$ to denote a function that returns a subset of values from U. That is, given a price function $P\colon U \to \mathbb{R}^+$, there might be many umbrellas that have the same cheapest price, so argmax would return all of these.

Notice that U is the set of *alternatives*, and $\mathbb{R}^+$ can be seen as the *outcome set*, that is, the set of possible values that we will need to pay if we decide to choose a particular umbrella.

Now, in the dining example, the set of alternatives and the outcome set are both $X = \{\text{Pizza, Steakhouse}\}$. In this case each player chooses one of the possible values of X and, based on the counting of votes, one of the values of X is chosen. So, in analogy with the umbrella example, what corresponds to the price table are functions from X to X, which are called *self-mappings on* X. Just like a price function maps an umbrella to a price, a function $p\colon X \to X$ maps a vote to some outcome of the election.

We can then describe our diners as higher-order functions as well. They take as input self-mappings $X \to X$, and choose some values of X which lead to a "good" outcome. If we use the index J for John's ordering, namely, Pizza $>_J$ Steakhouse, and M for Mary's ordering, that is, Steakhouse $>_M$ Pizza, we can write:

$$\text{argmax}_{>_J}\colon (X \to X) \twoheadrightarrow X$$

and

$$\text{argmax}_{>_M}\colon (X \to X) \twoheadrightarrow X.$$

John would vote for any place as long as the outcome of the vote is Pizza, and Mary would vote for any place as long as the outcome of the vote is Steakhouse. In this case we know they are using a majority vote, so it makes sense for John to vote for Pizza, and Mary to vote for Steakhouse. But describing their preference in this more general way allows us to free ourselves from the actual mechanism by which the final decision is made. If for some reason they decide to change the mechanism of choice, and say that if no one votes for a particular place then that's where they are going, this does not affect the way we have described John and Mary's preferences; they will still choose in such a way that their preferred choices is the final outcome.

More importantly, we can now easily describe Keynes' preference. Keynes is also described by a function of the same type

$$\varepsilon_K\colon (X \to X) \twoheadrightarrow X$$

but now the function ε_K does not consider any ordering on X, it simply declares that he is happy to vote for any restaurant x as long as it is the same as the final outcome $f(x)$, that is,

$$\varepsilon_K(p) = \{x \in X \,:\, f(x) = x\}$$

In mathematical terms ε_K is a *fixed point operator*. The fixed points of a self-map $f\colon X \to X$ are precisely the points x such that $f(x) = x$.

In general, in a situation where a player needs to choose an alternative in the set X, having in mind an outcome in another set R, such a player's preference can be described by a higher-order function of type $(X \to R) \twoheadrightarrow X$, and we have called such functions *selection functions*.

Describing players and preferences via selection functions leads to a beautiful theory of strategic games where all definitions and notions from classical game theory naturally generalise. This idea for modelling players using selection functions was originally proposed in [1] and recently further developed in [2, 3].

4 Where it all started

It might look from what you have read above that this is all about game theory, economics and maybe even finance. And you would be quite right. But for us, it all started by looking at the logic behind mathematical proofs, or, as people call it, *mathematical logic*.[1]

Some of you might be familiar with some aspects of mathematical logic, such as the famous incompleteness theorem due to Kurt Gödel (1906–1978) or the Zermelo-Fraenkel formalisation of set theory. But what has always fascinated me is the connection between *mathematical proofs* and *computer programs*. When a mathematician proves that under certain conditions, for certain classes of spaces, or for particular classes of functions some particular "object" must exist, their proof often comes with a *construction* showing how to build the said "object". For instance, think of Euclid's geometrical proofs: they are really a collection of algorithms or constructions.

There are, however, some proofs that are completely "non-computational". For instance, it is easy to show that there are two irrational numbers a and b such that a^b is rational, without ever constructing a and b. We just need to consider two possible situations: either $\sqrt{2}^{\sqrt{2}}$ is rational, or it is irrational. In the first case, we simply let $a = b = \sqrt{2}$, and so a^b is rational by assumption. In the second case, we can take $a = \sqrt{2}^{\sqrt{2}}$ and $b = \sqrt{2}$, so $a^b = \sqrt{2}^2 = 2$. But

[1] My colleague Martín Escardó initiated the work on selection functions back in 2007, originally motivated by the connection between selection functions and the topological notion of compactness.

we can only know which of the two cases is the one we need if we know whether $\sqrt{2}^{\sqrt{2}}$ is rational or not.[2]

What the above example illustrates is that in some cases a mathematical proof carries with it some extra information, quite often of a computational nature, showing how an object which is claimed to exist can actually be constructed. We call a proof *computational* if it is essentially an algorithm for constructing a solution. For instance, if you have a computational proof that a solution to the equation $f(x) = 0$ exists, that proof will tell you how to construct the solution. Otherwise, the proof is *non-computational*, and no such algorithm seems to be present in the proof.

The work on selection functions started not from a game-theoretic point of view, but rather as a tool to convert non-computational theorems into computational ones. The importance of extracting algorithms from mathematical proofs is that these algorithms can then be used in practice to calculate solutions of concrete problems.

5 The drinker and the central bank

Let us conclude by illustrating this last point with one final example, the *drinker paradox*: In any (non-empty) bar there always exists one person such that if that person is drinking, then everybody is drinking. This might sound strange at first, because it seems to be saying that there is a drinker in the bar who is causing everyone else to drink. We are used to inferring causal relationships in such statements, but here we are only claiming some sort of correlation. Also, the person in the statement is not necessarily always the same. Let us look more closely to see why the statement of the drinker paradox is true.

First, we must agree that either everybody is drinking, or at least one person is not drinking. In the first case, our “one person” can be anyone. In the second case, we choose our person to be (one of) the non-drinkers. Then, since they are not drinking, it will falsify the premise of our theorem, which, in turn, makes it (trivially) true. Strictly speaking, this is not really a paradox, it is a consequence of what in formal logic is called *material implication*. This means we have a statement of the form “If P, then Q”, and this implication can only be false if P is true and Q is false. In other words, “If P, then Q” is defined to be true not only when P and Q are both true, but also whenever P is false. Bear in mind that we are considering here the truth value of the whole statement, not of Q. This is admittedly counterintuitive, but makes more sense when we keep in mind that we are not claiming that P causes Q.

[2] It turns out that $\sqrt{2}^{\sqrt{2}}$ is indeed irrational, but this is not an easy result to prove.

Using the notation of formal logic, if $D(x)$ stands for "x is drinking", we can rewrite the drinker paradox as follows:

$$\exists x(D(x) \to \forall y D(y)). \tag{1}$$

Notice that this is a non-computational theorem, assuming we have infinitely many people in the bar.[3] How can we ever decide if everybody is drinking or if there is someone who is not drinking? At any finite time we can only check a finite number of people, and if they are all drinking we have to keep searching.

But luckily there is a computational version of this theorem, which relies on the *axiom of choice*: If for each x there exists a y, then there must exist a choice function p that produces y from x. The axiom of choice stated like this seems to be entirely innocuous, but it has many surprising consequences, such as the one discovered by Stefan Banach and Alfred Tarski in 1924, which states that a solid ball can be chopped into finitely-many subsets (as few as five) which can then be rearranged in space using only rotations and translations to form two solid balls each having the same volume as the first.[4] Returning to our example, we can use the axiom of choice to obtain an equivalent version of the drinker paradox:

$$\forall p \, \exists x \, (D(x) \to D(p(x))) \tag{2}$$

In (2), we are quantifying over all possible choice functions. If X is the set of all "persons", then each choice function $p \colon X \to X$ maps a person c to another person $p(c)$. Let us see why (1) and (2) are equivalent. First, assume that (1) holds, that is, assume that there exists a person x' such that if x' is drinking then everyone is drinking, and let p be an arbitrary choice function. Now, since we have that $D(x') \to \forall y D(y)$, in particular we have that $D(x') \to D(p(x'))$ for our arbitrary choice function p. Now since this p can be any choice function, it follows that (2) holds.

Now for the other direction, we will use the *c*ontrapositive, which means that instead of proving directly that (2) implies (1), we will prove instead that if (1) does not hold, then neither does (2). So let us suppose that for every person x there exists at least one person y such that $D(x)$ and $\neg D(y)$ (that is, y is not drinking). Now we use the axiom of choice to create a choice function p, by choosing amongst the possible values of y for each x, such that for all x we have $D(x)$ and $\neg D(p(x))$. In other words, we have shown that (2) does not hold.

Now what does this have to do with describing a player's preference? In order to see this, let us give a different game-theoretic interpretation to theorem (2) in terms of a strategy for a *central bank*. We let $X = \mathbb{R}$ be the possible

[3] Perhaps this is the bar of Hilbert's famous infinite hotel? See https://en.wikipedia.org/wiki/Hilbert's_paradox_of_the_Grand_Hotel.

[4] We refer to https://en.wikipedia.org/wiki/Banach-Tarski_paradox for more details.

values for inflation, and see $D(x)$ as saying whether inflation x is "good" or not. Let us suppose further that $D(x)$ is outside of the control of the bank, so the bank does not decide on what inflation rate is good or not, that is decided by the market. But the bank can, for any value of x, check whether x is good or not. We claim that the theorem (2) describes a strategy for a central bank who needs to predict next year's inflation. As we know, the central bank's inflation prediction might influence the actual inflation value. This is captured by the self-mapping $p\colon X \to X$, which says for each predicted inflation value $x \in X$, what the actual inflation $p(x) \in X$ will be. Now, what prediction x should the central bank make, so that if its predicted value is good, then the actual inflation is also good? That is, what prediction should the bank make such that $D(x) \to D(p(x))$? The optimal strategy for the central bank is very similar to the one we used to solve the higher-order drinker paradox. The central bank should consider predicting a certain value, say 0.5, and see what the actual inflation would be in that case, that is, calculate $p(0.5)$. If $p(0.5)$ is good, then predicting $x = 0.5$ satisfies condition $D(x) \to D(p(x))$. If on the other hand we find that $p(0.5)$ is not good, then we could simply take that value as our prediction $x = p(0.5)$, since, to be safe, giving a gloomy outlook on inflation is a good way to avoid future disappointment.

So, what we have seen is that higher-order functions, and in particular selection functions, provide a bridge between economics (where it provides a novel and powerful way to describe preferences) and mathematical logic (where it enables the computational interpretation of non-computational theorems). Further computational interpretations of originally non-computational mathematical theorems can be found in [4, 5], where game-theoretic interpretations of these higher-order constructions are also discussed.

Image credits

Fig. 1 Johanna84, CC0 Creative Commons, https://cdn.pixabay.com/photo/2013/09/12/15/46/umbrella-181682_960_720.png, visited on 11 January 2018.

References

[1] M. H. Escardó and P. Oliva, *Selection functions, bar recursion, and backward induction*, Mathematical Structures in Computer Science **20** (2010), no. 2, 127–168.

[2] ———, *Sequential games and optimal strategies*, Royal Society Proceedings A **467** (2011), 1519–1545.

[3] J. Hedges, P. Oliva, E. Winschel, V. Winschel, and P. Zahn, *Higher-order decision theory*, Proceedings of the 5th International Conference on Algorithmic Decision Theory, 2017.

[4] P. Oliva and T. Powell, *A constructive interpretation of Ramsey's theorem via the product of selection functions*, to appear: Mathematical Structures in Computer Science, 2013.

[5] ———, *A game-theoretic computational interpretation of proofs in classical analysis*, Gentzen's Centenary: The Quest for Consistency, 501-531, Springer, 2015.

Paulo Oliva *is a reader of mathematical logic at Queen Mary University of London.*

Mathematical subjects
Algebra and Number Theory, Analysis, Discrete Mathematics and Foundations

Connections to other fields
Computer Science, Finance, Humanities and Social Sciences, Reflections on Mathematics

License
Creative Commons BY-SA 4.0

DOI
10.14760/SNAP-2019-009-EN

Snapshots of modern mathematics from Oberwolfach provide exciting insights into current mathematical research. They are written by participants in the scientific program of the Mathematisches Forschungsinstitut Oberwolfach (MFO). The snapshot project is designed to promote the understanding and appreciation of modern mathematics and mathematical research in the interested public worldwide. All snapshots are published in cooperation with the IMAGINARY platform and can be found on www.imaginary.org/snapshots and on www.mfo.de/snapshots.

ISSN 2626-1995

Junior Editor
Sara Munday
junior-editors@mfo.de

Senior Editor
Sophia Jahns (for Carla Cederbaum)
senior-editor@mfo.de

Mathematisches Forschungsinstitut
Oberwolfach gGmbH
Schwarzwaldstr. 9–11
77709 Oberwolfach
Germany

Director
Gerhard Huisken

Mathematisches Forschungsinstitut Oberwolfach

Member of Leibniz Association

IMAGINARY open mathematics