UNIVERSITY OF MANNHEIM

University of Mannheim / Department of Economics

Working Paper Series

# *A New Perspective on Classical Choice Problems Using Selection Functions*

Jules Hedges Paulo Oliva
Evguenia Winschel Viktor Winschel Philipp Zahn

Working Paper 14-14

April 2014

# A New Perspective on Classical Choice Problems Using Selection Functions

Jules Hedges, Paulo Oliva
Department of Computer Science
Queen Mary University London

Evguenia Winschel, Viktor Winschel, Philipp Zahn
Department of Economics
University of Mannheim

April 27, 2014

**Abstract**

We use *quantifiers* and *selection functions* to generalize the classical economic approach to choice. Our framework encompasses preference and utility based approaches as special cases, but also extends to non-maximizing behavior and context-dependent motives such as social concerns. We adapt the method of quantifiers and selection functions which is based on higher-order functions and originate in computer science.

**JEL codes:** C0, D01, D03, D63, D64
**Keywords:** behavioral economics, social preferences, beauty contest, higher order functions, quantifiers, selection functions

# Contents


**1 Introduction** **3**

**2 Quantifiers and Selection Functions** **4**
2.1 Higher Order Functions . . . . . 4
2.2 Quantifiers . . . . . 6
2.3 Selection functions . . . . . 8
2.4 Context-Independent Quantifiers . . . . . 10
2.4.1 Rational Preferences and Utility Functions as Special Cases . . . . . 10
2.4.2 Beyond Rational Preferences . . . . . 12
2.4.3 Beyond Maximization and Standard Rationality . . . . . 13
2.5 Context-Dependent Quantifiers . . . . . 14
2.6 Trading off between selection functions and context . . . . . 17

**3 Combining Quantifiers** **18**
3.1 Example: Social Preferences . . . . . 19
3.1.1 Social Concerns I . . . . . 20
3.1.2 Social Concerns II . . . . . 20

**4 Discussion and Conclusions** **22**

# 1 Introduction

We use *quantifiers* and *selection functions* to generalize the classical economic approach to choice. The framework we use has been developed as a game theoretical approach to proof theory[1] [6, 7] and we adopt it to economics. We introduce *quantifiers* to represent an agent's goals and *selection functions* as the behavior that achieves his goals represented as an algorithm.

As we will show, our framework encompasses preference and utility based approaches as special cases, with quantifiers representing max and selection functions arg max. But more importantly, we can go beyond these special operators. First, it is possible to impose less structure on preferences, as for instance incompleteness. Secondly, it is possible to consider alternatives to maximization. Thirdly, goals can be implemented that take not only the outcomes into consideration but also how outcomes come about from actions. Our framework based on quantifiers and selection functions provides a unifying framework for all these concerns.

Quantifers and selection functions are based on the theory of higher-order functions, which in turn, form the foundation of functional programming languages. Thus, our framework can be directly implemented in languages such as Haskell in order to compute choices. The machinery we propose is readily extendable to interactive situations in a game theoretical setting which we explore in a companion paper [12]. In the current paper we set the scene and take a look at the new formalism within the decision theoretical and non-strategic setting.

The organization of the paper is the following. We first introduce the formal concept of a quantifier and a selection function. We show that a utility maximizing agent and his preferences can be instantiated as a special form of a selection function and quantifier. Then, we show that the selection function approach can be easily extended to choice behavior not covered by utility functions. Moreover, we illustrate that selection functions are a very

[1]Proof theory is a branch of mathematical logic which investigates the structure and meaning of formal mathematical proofs. It has been recently discovered that certain proofs of high logical complexity can be interpreted as computer programs which compute equilibria of suitable generalised games.

natural way of describing choices in general, be it within the relatively narrow frame of an only selfish, rational, and materially interested decision-maker or within alternative settings.

# 2 Quantifiers and Selection Functions

In this section we give a brief overview of the theory of *higher order functions*, i.e. functions that take other (possibly higher order) functions as input. We then use this theory in order to define selection functions and quantifiers. Subsequently, we represent the classical approach to decision theory via preference relations and argmax and max operators within the new formalism.

## 2.1 Higher Order Functions

A *higher order function* (or *functional*) is a function whose domain is itself a set of functions. Given sets $X$ and $Y$ we denote by $X \to Y$ the set of all functions with domain $X$ and codomain $Y$ (this is often denoted $Y^X$ but we avoid this notation). A higher order function is therefore a function $f : (X \to Y) \to Z$ where $X$, $Y$ and $Z$ are sets.

A simple example of a higher order function is the function that evalulates its argument at a constant point. To give a specific example, we take the sets $\mathbb{R}$ (real numbers) and $\mathbb{Z}$ (integers), and pick a constant real number, such as $\pi$. We can then define a function $\Phi : (\mathbb{R} \to \mathbb{Z}) \to \mathbb{Z}$ by the equation $\Phi(f) = f(\pi)$. We can illustrate the behaviour of $\Phi$ by giving it a specific function $f : \mathbb{R} \to \mathbb{Z}$ as an input. For example, let $f$ be the function that takes a real number to its integer lower bound. The integer lower bound to $\pi$ is 3, therefore $\Phi(f) = 3$.

We will sometimes use a particular notation for higher order functions called $\lambda$-*notation*. The example above would be written in $\lambda$-notation as $\Phi = \lambda f.f(\pi)$. The symbol $\lambda$ denotes the *abstraction* or *binding* of the input to the function, and the name of the variable that is bound is not important, i.e. $\lambda f.f(\pi)$ and $\lambda g.g(\pi)$ represent the "same" function. There are familiar examples of this, for example in the maximum operator the name of the variable which ranges over the set $X$ is not relevant, i.e.

$$\max_{x \in X} f(x) = \max_{y \in X} f(y)$$

and similarly with the variable over which one performs integration, i.e.

$$\int f(x)\,\mathrm{d}x = \int f(y)\,\mathrm{d}y$$

(we will see later that it is not a coincidence: both of these examples turn out to be special cases of $\lambda$ expressions). We sometimes write the domain of the function as a superscript, as in $\Phi = \lambda f^{\mathbb{R}\to\mathbb{Z}}.f(\pi)$. $\lambda$-notation is also sometimes useful for writing functions which are not higher-order, for example the function which inputs an integer and squares it can be written $\lambda n^{\mathbb{Z}}.n^2$. It is also possible for the bound variable to not appear in the body of the $\lambda$ expression, giving a constant function such as $\lambda x.5$.

An important result, called *combinatory completeness*, is that the functions that can be represented only using $\lambda$ expressions and nothing else are precisely the computable functions[2]. That is, $\lambda$ expressions are a programming language, equivalent in power to Turing machines and other standard models of computation. However $\lambda$ expressions have an advantage that it is easy to extend them with noncomputable functions when necessary (for example, a $\lambda$ expression can refer to a known noncomputable function), and in practice we will do this sometimes. For example, there is a noncomputable function $E$ which inputs a description of a finite game, and outputs a mixed strategy Nash equilibrium of that game. This function can be freely combined with $\lambda$ expressions, for example $\lambda x, y.E(x)$. However, if we have a function $f$ which is defined entirely in terms of $\lambda$ expressions and other functions known to be computable, then we know that $f$ is computable. This has two important advantages: firstly, we can easily identify which objects we are

[2]One must distinguish at this point the *typed* from the *untyped* $\lambda$-calculus. For this paper we will be mostly working with the typed version, where every term has a precise type, and an application $f(x)$ is only allowed when $f$ has type $X \to Y$ and $x$ has type $X$, so that an application such as $f(f)$ is not allowed. However, in order to obtain combinatorial completeness one needs to work with the untyped $\lambda$-calculus, so as to obtain fixed point operators such as $\Phi(f) = (\lambda x.f(x(x)))(\lambda x.f(x(x)))$.

discussing are computable, and secondly, for those which are computable, we can see the $\lambda$ expression as an implementation in a proto-programming language which can easily be converted into code of a suitable real programming language such as Haskell.

## 2.2 Quantifiers

In this section we will define quantifiers and selection functions as particular classes of higher order functions. In subsequent sections we will obtain quantifiers as a series of generalisations from utility functions.

Suppose we have an agent $\mathcal{A}$. We know nothing about $\mathcal{A}$'s motivations, but we know that it is *deterministic* (or *predictable*) in the sense that its moves are not dependent on chance (this is without loss of generality, because we can always allow the set of outcomes to be a set of probability distributions). Consider $\mathcal{A}$ as a black box: we can insert it into any *situation* and observe what it does.

We need to define what we mean by 'situation' here. A situation, which we will from now on call a *context*, should be an object that encodes all of the relevant information that the agent could consider when choosing a move. Assume our agent is choosing a move in the set $X$, and the set of possible final outcomes is $R$. The context will normally include other agents and all the other choices that together with the choice of our agent $\mathcal{A}$ will determine a final outcome. If all we care about is the final outcome, then our context can be modelled simply by a function $p\colon X \to R$ that maps each of the agent's move to a specific outcome. In other words, to give the context of an agent is the same as to define precisely what final outcomes will result after each of the agent's choice. That is all that our agent needs to know about this "context" in order to make the good choice.

Therefore, for an agent choosing a move from a set $X$, having in sight a final outcome in a set $R$, we call any function $p : X \to R$ a possible context for that agent.

Suppose that $\mathcal{A}$ makes a decision in the context $p$. Then the agent will consider some outcomes to be *good* (or *acceptable*), and other outcomes to be bad. We are going to allow the set of outcomes that the agent considers good to be totally arbitrary. Thus, to each context $p : X \to R$, we associate a set of outcomes $\varphi(p) \subseteq R$. This defines a higher order function

$$\varphi : (X \to R) \to \mathcal{P}(R)$$

where $\mathcal{P}(R)$ is the set of all subsets of $R$. The function $\varphi$ is precisely what is known as a *quantifier*. Considering our discussion above about contexts, a quantifier can be seen as a description of what the good outcomes are for each possible context. Our main objective in this paper is to convince the reader that this is a general, modular, and highly flexible way of describing an agent's goal or objective.

For some quantifiers, there will always be exactly one good outcome. This is the case for our motivating example, which is maximization. Suppose the set of moves $X$ is finite, and $R = \mathbb{R}$ is the real numbers, representing profit. If a decision is made in the context $p : X \to \mathbb{R}$ then the good outcome is precisely the maximal one, i.e. given a context $p \colon X \to \mathbb{R}$, the outcome in that context which our agent would consider good is the maximum value of $p$. Thus our quantifier is defined by the equation

$$\varphi(p) = \{\max_{x \in X} p(x)\}$$

Often we will omit the set brackets when writing a single-valued quantifier, so $\varphi$ will be written

$$\varphi(p) = \max_{x \in X} p(x)$$

Another common situation is that there is at most one good outcome. The motivating example for this case is maximization over an infinite set, such as the unit interval $[0, 1]$. Then the maximum of $p : [0, 1] \to \mathbb{R}$ in general exists only if $p$ is a continuous function. Thus we can consider a partially defined quantifier that only chooses good outcomes when given a continuous context, i.e.

$$\varphi(p) = \begin{cases} \{\max_{x \in X} p(x)\} & \text{if } p \text{ is continuous} \\ \varnothing & \text{otherwise} \end{cases}$$

We define the *domain* of a quantifier $\varphi$ to be the set of all contexts $p$ such that $\varphi(p)$ is nonempty, that is, $p$ has a good outcome. In symbols the definition is

$$\mathrm{dom}(\varphi) = \{p : X \to R \mid \varphi(p) \neq \varnothing\}$$

## 2.3 Selection functions

While considering 'good' outcomes may be considered problematic (since they are subjective to the agent and cannot be observed) there is no problem in observing the moves made by the agent. If the agent makes a decision in the context $p : X \to R$ then it makes a move $x \in X$. This defines a higher order function

$$\varepsilon : (X \to R) \to X$$

which we call a *selection function*. In general, we consider a quantifier to describe the *goals* or *motivations* of the agent (what the agent likes) and a selection function to describe the *behavior* of the agent (what she would do to get what she likes). We can give a very general definition of what it means for an agent to be rational: their behavior is consistent with their motivation. For any context, the move they make according to the selection function should result in a good outcome according to their quantifier. Formally, for every context $p : X \to R$ it should be the case that

$$p(\varepsilon(p)) \in \varphi(p)$$

However if $p$ is a domain with no good outcomes then we cannot expect there to be a good move. Thus we only require this condition to hold for contexts $p \in \mathrm{dom}(\varphi)$. When this is the case we say that the selection function $\varepsilon$ *attains* the quantifier $\varphi$. The motivating example of a selection function is arg max, which defines the selection function

$$\varepsilon(p) = \arg\max_{x \in X} p(x)$$

Then $\varepsilon$ attains the max quantifier.

Suppose we have a quantifier $\varphi$ which describes the outcomes that an agent considers to be good. The quantifier might be *unrealistic* in the sense that it has no attainable good outcome. For example, in my current context I would consider it a good outcome if I received a million dollars, but I have no move at the moment which will lead to this outcome. Given a context $p$, the set of attainable outcomes is precisely the image of $p$. A *realistic* quantifier is simply a quantifier in which every context with a good outcome has an attainable good outcome. We can write it in symbols as

$$\varphi(p) \neq \varnothing \implies \varphi(p) \cap \mathrm{Im}(p) \neq \varnothing.$$

In fact the following are equivalent:

- $\varphi$ is realistic
- there is a selection function which attains $\varphi$.

Thus selection functions are a way to describe realistic quantifiers.

The theory of quantifiers and selection functions has been developed in stages. Selection functions and single-valued quantifiers first appeared in [6], unifying earlier definitions in proof theory and type theory. The connection between selection functions and game theory also first appeared there. General quantifiers appeared in [7], which allows us to capture more important examples in a more natural way. The connections between selection functions and game theory were explored in more depth in [8] and [10], and the latter contains the definition of attainment given here. Finally [11] contains the terminology *context* and the definition of a realistic quantifier.

## 2.4 Context-Independent Quantifiers

Suppose $R$ is the set of possible final outcomes, and each agent $i$ has a partial order relation $\succeq_i$ on $R$, so that $x \succeq_i y$ means that agent $i$ prefers the outcome $x$ to $y$. These partial orders lead to choice functions $f_i : \mathcal{P}(R) \to \mathcal{P}(R)$ where $f_i(S)$ are the maximal elements in the set of possible outcomes $S$ with respect to the order $\succeq_i$. Note that these $f_i$ satisfy $f_i(S) \subseteq S$, and $f_i(S) \neq \varnothing$ for non-empty $S$.

Every such $f_i$ can be turned into a quantifier $\varphi_i$ in a generic way, using the fact that the image operator is a higher-order function $\mathrm{Im} : (X \to R) \to \mathcal{P}(R)$:

$$(X \to R) \xrightarrow{\mathrm{Im}} \mathcal{P}(R) \xrightarrow{f_i} \mathcal{P}(R)$$

so that $f_i \circ \mathrm{Im} \colon (X \to R) \to \mathcal{P}(R)$ are quantifiers. We call quantifiers factoring as $f \circ \mathrm{Im}$ as *context-independent quantifiers*. Player's defined by context-independent quantifiers are choosing the set of good outcome simply by ranking the set of outcomes that can be achieved in a given game context. But are forgetting all the information about how each of the outcomes arise from particular choices of moves. For instance, we might have a set of actions that will lead us to earn some large sums of money. Some of these, however, might be illicit. A maximising agent defined in a context-independent way would choose the outcome that gives himself the maximum return. If we have control over which actions lead to which outcomes, we might consider other choices as preferable.

It is easy to show that whenever $f_i$ is a choice function arising from partial order $\succeq_i$, then the a context-independent quantifier $f_i \circ \mathrm{Im}$ is realistic, in the sense of the previous section.

### 2.4.1 Rational Preferences and Utility Functions as Special Cases

The usual approach to model behavior in economics is to either postulate a preference relation on the set of alternatives or to directly assume a utility function [14]. Typically, structure is imposed on preference relations. These assumptions are made due to two reasons: either because additional structure deems to be a characteristic of an agent's rationality[3] or because one wishes to work with utility functions. It is a classical result that for utility functions to exist, preferences relations have to be *rational* [14].

Now, rational preferences and utility functions are special cases of the generic construction of a context-independent quantifier we outlined in the last section. They are special because (i) we impose additional structure on $R$, that is, $\succeq_i$ is a total preorder and (ii) we focus on one particular $f_i$, that is, $f_i : \mathcal{P}(R) \to \mathcal{P}(R)$ defined by

$$f_i(S) = \{\succeq_i \text{-maximal elements of } S\}$$

A rational preference relation can always be represented by a utility function. Translated into the selection function approach, the utility function can be characterized as the environment which is a mapping $p\colon X \to \mathbb{R}$, attaching a real number to each element of the set of choices $X$. So, we can define the quantifier

$$\phi(p) = \max p$$

which is attained by the selection function

$$\varepsilon(p) = \arg\max p$$

Note the types $\phi\colon (X \to \mathbb{R}) \to \mathcal{P}(\mathbb{R})$ and $\varepsilon\colon (X \to \mathbb{R}) \to X$ respectively and that $p(\varepsilon(p)) \in \phi(p)$. Thus, max and arg max operators, which are universally used in economic literature, become the prototypical examples of an context-independent quantifier and a selection function attaining it.

**Minimizing Deviations from a Target or Ideal Point Preferences** Another example of a concern which is relevant for economics - in particular in political economy - is a minimisation of a distance from a target. Such a concern can be easily represented by a utility function but can be directly represented by a quantifier as well.

Suppose the outcomes of a decision problem form a set $R$ with a distance

[3] This issue has been intensely debated, see [17, 15, 14].

measure $d$, a metric. For example $\mathbb{R}^n$ has a metric given by Euclidean distance, but there are many other examples including metrics on finite sets. Suppose the decision maker identifies a set of outcomes $S \subseteq R$ which are considered to be optimal outcomes. Then the distance of a particular outcome $r$ to the set $S$ is defined by the equation

$$d(r, S) = \inf_{s \in S} d(r, s)$$

Now the decision maker can choose a move minimising the distance to the target set:

$$\varphi(p) = \min_{x \in X} d(p(x), S)$$

Although this quantifier can be described by utility functions, it can potentially simplify the modelling of several important situations. For example, a decision maker focussing only on fairness, at the expense of profit, is equivalent to minimising distance to the diagonal

$$S = \{(x, \ldots, x) \in \mathbb{R}^n \mid x \in \mathbb{R}\}$$

### 2.4.2 Beyond Rational Preferences

The generic construction of context-independent quantifiers instantiates choices based on rational preferences (or equivalently on utility maximization) as special cases. In this section we show that we can go beyond these cases by allowing for a different structure on $R$ or by allowing for a different $f_i$ (or by relaxing both).

Utility functions are considered as a very convenient tool to represent and analyze choice behavior. Still, the assumption that the preorder is total, which guarantees the existence of a utility function, is demanding and in fact more demanding than is necessary to rationalize choice behavior [17]. Secondly, when taking the perspective of preferences, from a positive as well as a normative viewpoint, there are good reason why a rational decision-maker may exhibit "indecisiveness", meaning that his preference for a pair of outcomes is not defined [1]. Thirdly, consider a situation where the economist or some other agents/principal has only partial information about the preferences of an agent and considers him "as if" he has incomplete preferences [5]. Lastly, $R$ may be a set of alternatives to be chosen by a group of agents. Even if each individual's preferences are complete, the aggregate social welfare ordering does not have to be [16].

There have been various attempts to change standard formalisms to allow for a utility theory without the need to fulfill the completeness assumption.[4]

When working with quantifiers and selection functions, the set of outcomes $R$ can have *any* order. In particular, the preference relation does not have to be complete. That is, given any preference relation $\succeq \subseteq R \times R$, an agent chooses the best alternatives as outlined in Section 2.4 above. So, one can very easily consider choices not in the scope of utility functions without the need to change the framework.

### 2.4.3 Beyond Maximization and Standard Rationality

The utility approach is intimately linked to the assumption that the agent fully optimizes. The behavioral economic literature as well as the psychological literature have documentated deviations from optimizing behavior, and have collected various decision "heuristics" [3, 13]. Quantifiers provide a nice way to model such deviations. Morover, even situations that can be modelled with utility functions may have (more) natural representations in the quantifier framework.

**Decisions Heuristic** Consider a simple heuristic of a person ordering wine in a restaurant. Suppose he always chooses the second most-expensive wine. In terms of selection functions, let $X$ be the set of wines available in a restaurant, and $p : X \to \mathbb{R}$ the price attached to each wine $x_i$ $(i = 1, ..., N)$ on the menu. Denote with $r_i$ the price of wine $x_i$. Given a maximal chain $r_n > r_{n-1} > \ldots > r_1$ in $\mathbb{R}$, let us call $r_{n-1}$ a sub-maximal element. The

[4]For an important early contribution see [1]. More recent contributions include [16] for utility representations in certain environments and [5] for uncertain environments. See also references in [16].

"goal" of the agent can be described by the quantifier

$$\phi_{>}(p^{X \to \mathbb{R}}) = \{\text{sub-maximal elements with respect to} > \text{within img}(p)\}.$$

Such quantifiers are attainable with selection function

$$\varepsilon_{>}(p^{X \to \mathbb{R}}) = \text{any } x \text{ such that } p(x) \text{ is a sub-maximal element of img}(p)$$

since clearly $p(\varepsilon_{>}(p)) \in \phi_{>}(p)$.

## 2.5 Context-Dependent Quantifiers

So far, we have focused only on the generic context-independent quantifier. As the last examples illustrate with this construction we can already go beyond choices motivated by rational preferences. Yet, we can do more. We can allow for quantifiers that do not only take the image of $p$ as input but the complete function. Next, we provide several examples to illustrate that this opens up a complete new dimension. Indeed, with context-dependent quantifiers it is possible to go far beyond what can be modelled using utility functions.

**Majority and minority** The majority quantifier selects outcomes which are "most unavoidable". Given a function $p : X \to R$, where $X$ and $R$ are finite sets, we can rank the elements $r \in R$ by number of different way it can be achieved:

$$n_r = |\{x \in X \mid p(x) = r\}|$$

If some $r$ has $n_r > n_{r'}$ for all $r' \neq r$ then we can say that $r$ is the most unavoidable outcome. If there is a tie we can simply select the set of all outcomes which are tied for most unavoidable. This defines a quantifier $\varphi : (X \to R) \to \mathcal{P}(R)$, which is attained by a selection function because we only consider outcomes in the image of $p$. Technically, this quantifier depends on the *multiset* image of $p$: not just which outcomes are attainable, but how many ways they can be attained. Later we will consider quantifiers even more general than depending on the multiset image of $p$.

To illustrate this quantifier, suppose there are 3 identical beaches, and the agent is indifferent between them. The first can be reached by one highway, the second by two highways and the third by three highways. The agent must choose which highway to take, and the outcome is the beach that the agent goes to. The agent decides to visit the beach which can be reached by the most different routes, which is the third, for example to avoid the risk of being stuck in a traffic jam.

Just as easily we can define the minority quantifier, which selects the set of all outcomes that are easiest to avoid, as they have the least number of moves which lead to them (ignoring outcomes which are never chosen).

**Average** Given a finite set $X$ and a function $p : X \to \mathbb{R}$, the average value of $p$ is given by the quantifier

$$\varphi(p) = \frac{1}{|X|} \sum_{x \in X} p(x)$$

Unfortunately this quantifier is not realistic. However if we replace $X$ with the unit interval $[0, 1]$ then the average value of a *continuous* function $p : [0, 1] \to \mathbb{R}$ is given by the integral

$$\varphi(p) = \int_0^1 p(x) \, \mathrm{d}x$$

Moreover there is a selection function which attains this quantifier, since by the mean value theorem for every continuous function $p : [0, 1] \to \mathbb{R}$ there always exists a point $a$ such that

$$p(a) = \int_0^1 p(x) \, \mathrm{d}x$$

**Fixed points** Consider an agent who is part of a group that has to choose an element of out of a set $X$. Assume that he observes the choices of other people and wants to make the same choice that the majority makes. One example for such goals is Keynes's Beauty contest where the agent is only interested in voting for the winner of the contest and he has no preferences for the contestants per se. The set of all possible choices that can be made by the group is $X^n$, and the outcome of the contest is the majority function $\text{maj} : X^n \to X$ which selects the most common element of the tuple (this *outcome* function should not be confused with the majority quantifier considered above, which is different). Thus the outcome is the most popular choice.

The agent in the Keynes beauty contest can be described by a quantifier which is a fixed point operator. Let $X$ be any set. A fixed point of a function $f : X \to X$ is a point $x \in X$ satisfying the equation $f(x) = x$. We can define a multi-valued quantifier $\varphi : (X \to X) \to \mathcal{P}(X)$ which selects the set of all fixed points of a function:

$$\varphi(p) = \{x \in X \mid p(x) = x\}$$

In practice, most functions do not have a fixed point and so this quantifier will often give the empty set. For the purposes of modelling a particular situation we might want to 'complete' $\varphi$ in different ways, describing what an agent might do in the event that no fixed point exists. The simplest situation is that we fix a 'default option' $x_0 \in X$, and define a multi-valued quantifier by

$$\varphi(p) = \begin{cases} \{x \in X \mid p(x) = x\} & \text{if } p \text{ has a fixed point} \\ \{x_0\} & \text{otherwise} \end{cases}$$

There are other possibilities, such as the agent having a preference ordering on $X$ and choosing the maximum in the event of there being no fixed point.

Suppose the other votes $(x_2, \ldots, x_n)$ are fixed, so the context of the agent's move is $\lambda x_1.\, \text{maj}(x_1, x_2, \ldots, x_n)$. Thus the agent will aim to find an $x_1$ that satisfies the equation

$$x_1 = \text{maj}(x_1, x_2, \ldots, x_n)$$

where the decisions $x_2, ..., x_n$ are exogenously given. The situation becomes far more interesting when considered as a game in which several agents are voting. We analyse this in detail in the follow-up paper [12].

Another possibility is to work in an alternative foundation to set theory, such as domain theory or a specific programming language, in which every function has a canonical fixed point. Such a foundation relies on every set having an element $\bot$, which represents a computer program which does not terminate. Such an element can be interpreted as the 'move' of an agent who is indecisive. We leave this discussion for future work.

## 2.6 Trading off between selection functions and context

Modelling a decision problem using a selection function introduces an additional degree of freedom over utility functions, namely we can trade off between a *selection function* and a *context*. Suppose we have a set of moves $X$, a set of outcomes $R$ and a context $p : X \to R$. Consider an agent who acts according to the selection function $\varepsilon : (X \to R) \to X$, so the move the agent will make in this particular situation is $\varepsilon(p)$.

However, we can model the same situation in a different way. Suppose the set of outcomes is also $X$, and the context is the identity function $X \to X$, so the outcome of a move is the move itself. Now consider a selection function $\delta : (X \to X) \to X$ defined by

$$\delta(id) = \varepsilon(p \circ id)$$

The move made by an agent playing this selection function is

$$\delta(id) = \varepsilon(p \circ id) = \varepsilon(p)$$

Thus we have moved all the information encoded by the context 'inside' the selection function $\delta$.

In the other direction, when we have a selection function that can be modelled using utility functions then we can move some of the information contained in the selection function into the context, which is the utility function, reducing the selection function to $\arg\max$ at the expense of a more complicated context. This can be used to characterise the quantifiers that are equivalent to utility functions.

Usually there will be an infinite variety of ways to make this trade off. To some extent this also occurs when considering interactions between several agents.

In practice we resolve this by using selection functions and contexts to model different aspects. Roughly speaking, quantifiers and selection functions should be used to model what is intrinsic to the *agent*, while contexts should be used to model what is intrinsic to the *situation*. So long as the move sets and outcome sets are equal we can then consider different agents making the same decision, and the same agent making different decisions.

A simple example of this is modelling fairness concerns, which we consider in the next section. What is intrinsic to the situation is the amount of profit that the agent receives. What is intrinsic to the agent is whether the agent cares only about profit or will take fairness into consideration. A suitable context is $p : X \to \mathbb{R}^n$, where the tuple $p(x)$ represents the profits received by $n$ recipients after the move $x$. An agent who cares only about his own profit $p(x)_i$ will use the quantifier

$$\varphi_i(p) = \max_{x \in X} p(x)_i$$

An agent who cares only about fairness will use the quantifier

$$\varphi_j(p) = \min_{x \in X} f(p(x))$$

where $f : \mathbb{R}^n \to \mathbb{R}$ is a function that the agent uses to measure the unfairness of an outcome. We can then consider different agents who measure unfairness by different criteria. In the next section we will consider ways of combining these quantifiers to take both concerns into consideration.

# 3 Combining Quantifiers

Quantifiers allow for a flexible modelling of choice behavior. But there is the additional important aspect of compositionality. We can combine different agents/goals into new agents/goals by building the product of quantifiers. Thereby we can make explicit the way in which differing motives interact. In this section, we show the general approach to composition and we apply it to an example where an agent cares for his and another agents' payoffs.

There is one important way of combining single-valued quantifiers, called the *product of quantifiers*, and a corresponding operation on selection functions called the *product of selection functions*. These are central to the applications of selection functions to proof theory [6], but as they are closely connected to extensive-form games of complete information we will not consider them in this paper. Another method of combining selection functions appears in [10], called the *binary Berardi-Bezem-Coquand functional*, is also directly based on the needs of proof theory and is connected to normal-form games, so they will also not be considered.

## 3.1 Example: Social Preferences

A standard assumption in economics is that people are selfish. Accordingly, they only care for their own private benefit. However, experimental evidence as well as evidence from the field indicates that individuals regularly deviate from selfish decisions. The evidence has triggered a theoretical literature which tries to incorporate social concerns into agents' motivation. We will focus on this example here. Our approach is not limited to this particular application but we consider the issue of social preferences to by prototypical.

A prominent class of models assumes that individuals take the allocative consequences of their decisions for others into account [2, 4, 9]. Consider a set of agents $\{1, ..., N\}$. One agent, 1, has to make a decision $x \in X$ which potentially influences the outcomes $R$ for all agents. The type of outcomes is then $R = R_1 \times \cdots \times R_N$.

The models of [2, 4, 9] are restrictive in two dimensions. First, they focus only on outcomes which involve monetary transfers, that is $R_i = \mathbb{R}$. Secondly, they only consider standard utility functions. Hence, they are restricted to depict social behavior by agents as a maximization of some utility function, usually of the following form: $U(r_1, ..., r_N) = f(r_1) + g(r_1, ..., r_N)$, where

$r_i \in R_i$ for all $i$ and the functional form of $f$ and $g$ depends on the motives of agent $i$.

### 3.1.1 Social Concerns I

Suppose we have $n$ agents, together with one benevolent agent who must choose a move in a set $X$. We can describe the goals of this agent by a quantifier $\varphi : (X \to \mathbb{R}^n) \to \mathbb{R}^n$. A simple situation is that the agent moves to maximise the sum of all payoffs (see, for instance [4]):

$$\varphi(p) = \max_{x \in X} \sum_{i=1}^{n} p(x)_i$$

However we can just as easily describe more complex behaviors for the agent, for example minimising inequality.

### 3.1.2 Social Concerns II

Consider an agent be interested in both a fair allocation and in maximizing his payoffs. Suppose he must make a move which determines a split of a fixed pie $\pi$ between himself and another recipient[5], so the context is $p : X \to \mathbb{R}^2$. Hence, if the agents decides to allocate $x$ to himself, then the recipient receives $\pi - x$. So, it holds that $p(x) = (x, \pi - x)$.

Suppose we consider the agent to have two aspects to his personality, one which cares only about profit and one which cares only about fairness. In general, such an agent may have a pair of quantifiers $\varphi, \psi : (X \to R) \to \mathcal{P}(R)$. After making a move in the context $p$, the 'first personality' is satisfied with those outcomes in $\varphi(p)$, and the 'second personality' is satisfied with those outcomes in $\psi(p)$. Therefore the moves which satisfy both personality aspects are precisely the elements of $\varphi(p) \cap \psi(p)$. This is captured by a new quantifier $\varphi \cap \psi$ by

$$(\varphi \cap \psi)(p) = \varphi(p) \cap \psi(p)$$

[5]In the literature, this type of situation is often referred to as a "Dictator Game". See [3] for an overview of the empirical evidence.

However typically this new quantifier will not be realistic, even if $\varphi$ and $\psi$ are both realistic. Therefore the agent will often face a conflict between the two aspects of his personality. Thus, we have to resort to resolving the conflict in different ways for each agent and for each situation.

In our specific example, the first personality cares only about own payoff, and so uses the single-valued quantifier

$$\varphi(p) = \max_{x \in X} p(x)_1$$

Suppose the second personality has identified a set of outcomes $F \subseteq \mathbb{R}^2$ which are deemed to be fair. The second personality is then satisfied with any fair outcome:

$$\psi(p) = F$$

One way in which the agent may resolve the conflict is to choose a maximal fair outcome:

$$\chi(p) = \max(F \cap \operatorname{Im}(p))_1$$

For a pair of selection functions (or single-valued quantifiers) there is a very general way of resolving the conflict by introducing an 'arbiter'. Suppose the agent consists of a committee of two decision-makers, each of which suggests a course of action. The committee members are described by the selection functions $\varepsilon, \delta : (X \to R) \to X$, so given the context $p : X \to R$ the first member will suggest the move $\varepsilon(p)$ and the second will suggest the move $\delta(p)$. The arbiter will decide based on the outcomes which would result from these moves.

In general, the arbiter will be described by a binary relation $A$ on $R$. If the moves suggested are $x_1$ and $x_2$ respectively, the arbiter will check whether the relation $(x_1, \pi - x_1)A(x_2, \pi - x_2)$ holds. If it does, the move $x_1$ will be selected, otherwise the move $x_2$ will be selected.

In our example, there are several criteria which the arbiter could use to select moves. For example, there could be a maximum amount of profit $c$ which is considered acceptable to lose in the name of fairness. Then given outcomes $(x_1, \pi - x_1) \notin F$ and $(x_2, \pi - x_2) \in F$, the relation $(x_1, \pi - x_1)A(x_2, \pi - x_2)$

will hold iff the loss incurred by the fair outcome is too great, that is,

$$(x_1, \pi - x_1)A(x_2, \pi - x_2) \iff x_1 - x_2 > c$$

Another possibility is to use lexicographic preferences, in which case $(x_1, \pi - x_1)A(x_2, \pi - x_2)$ holds iff one of the following holds:

- $(x_1, \pi - x_1) \in F$ and $(x_2, \pi - x_2) \notin F$ (prefer a fair outcome to an unfair one)
- $(x_1, \pi - x_1) \in F \iff (x_2, \pi - x_2) \in F$ and $x_1 \geq x_2$ (all other things equal, prefer greater profit)

We can consider an arbiter to be unbiased if the result does not depend on the order of $\varepsilon$ and $\delta$, which is equivalent to the formal condition of antisymmetry: $x_1 A x_2$ holds iff $x_2 A x_1$ does not hold. This method thus allows us to consider biased arbiters, by considering relations which are not antisymmetric. There is also no formal need to require that $A$ is transitive. For any arbiter $A$, the resulting behavior of the agent as a whole is described by the selection function

$$\varepsilon'(p) = \begin{cases} \varepsilon(p) & \text{if } p(\varepsilon(p))Ap(\delta(p)) \\ \delta(p) & \text{otherwise} \end{cases}$$

# 4 Discussion and Conclusions

In this paper we propose the use of quantifiers and selection functions to model individual choices. Our framework instantiates the standard approach in economics, preferences and utility functions, as one special case. In our framework we can also model agents whose preferences are incomplete, who deviate from maximization, whose motives are not only influenced by the outcome but also by the way outcomes realize, or any combination thereof.

Typically, economists restrict themselves to choice settings where an equivalence between the preference and the utility maximization approach holds

[14]. In practice, this results in the tendency to use utility functions which excludes interesting choice phenomena from analysis. Quantifiers and selection functions provide means to address such phenomena.

Even when considering questions in the realm of the standard framework, where formal equivalence between different approaches exists, the naturalness of one form versus the other may be different. While representing choice problems as a maximization problem of a utility function is typically considered to be advantageous as the mathematical toolbox economists are used to can be applied, representing choices in a different form may be more natural and more insightful nevertheless. With the developments in modern theoretical computer science many of the supposedly benefits of sticking to optimization methods deem less advantageous to us.

Our approach is based on higher-order functions which also form the foundation for high-level functional programming languages. The idea behind these languages is that they represent mathematical objects themselves. Having common ground has the consequences that the models within our framework are implementable in these languages.

Lastly, the roots of the quantifier and selection functions approach with respect to economics are in game theory, that is, interaction between different agents. What we have presented in this paper is readily extended to interactions. It is here, where we think the expressiveness of this approach as well as its implementability in programming languages will come to particular fruition.

# References


[1] Robert J. Aumann. Utility theory without the completeness axiom. *Econometrica*, 30(3):pp. 445–462, 1962.

[2] Gary E Bolton and Axel Ockenfels. Erc: A theory of equity, reciprocity, and competition. *The American Economic Review*, 90(1):pp. 166–193, 2000.

[3] C.F. Camerer. *Behavioral Game Theory: Experiments in Strategic Interaction*. The Roundtable Series in Behavioral Economics. Princeton University Press, 2011.

[4] Gary Charness and Matthew Rabin. Understanding Social Preferences with Simple Tests. *Quarterly Journal of Economics*, 117(3):817–869, 2002.

[5] Juan Dubra, Fabio Maccheroni, and Efe A. Ok. Expected utility theory without the completeness axiom. *Journal of Economic Theory*, 115(1):118–133, March 2004.

[6] Martin Escardó and Paulo Oliva. Selection functions, bar recursion and backward induction. *Mathematical structures in computer science*, 20(2):127–168, 2010.

[7] Martín Escardo and Paulo Oliva. Sequential games and optimal strategies. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science*, 467(2130):1519–1545, 2011.

[8] Martin Escardó and Paulo Oliva. Computing Nash equilibria of unbounded games. *Proceedings of the Turing centenary conference*, 2012.

[9] Ernst Fehr and Klaus M Schmidt. A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics*, 114(3):817–868, 1999.

[10] Jules Hedges. A generalisation of Nash's theorem with higher-order functionals. *Proc R Soc A*, 469(2154), 2013.

[11] Jules Hedges. Monad transformers for backtracking search. 2014.

[12] Evguenia Winschel Viktor Winschel Philipp Zahn Jules Hedges, Paulo Oliva. Context dependent games as quantifiers and selection functions. Technical report, 2014.

[13] D. Kahneman. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, 2011.

[14] D.M. Kreps. *Microeconomic Foundations I: Choice and Competitive Markets*. Number Bd. 1 in Microeconomic Foundations. Princeton University Press, 2012.

[15] Andreu Mas-Colell, Michael D. Whinston, and Jerry Green. *Microeconomic Theory*. 1995.

[16] Efe A. Ok. Utility representation of an incomplete preference relation. *Journal of Economic Theory*, 104(2):429–449, June 2002.

[17] M. Richter. Rational choice. In John Chipman, L. Hurwicz, and H. Sonnenschein, editors, *Preferences, utility, and demand*. Harcourt Brace Jovanovich, New York, 1971.