

Higher-Order Decision Theory

Jules Hedges¹, Paulo Oliva², Evguenia Shprits³,
Viktor Winschel⁴, and Philipp Zahn⁵

- ¹ Department of Computer Science, University of Oxford, Oxford, UK
julian.hedges@cs.ox.ac.uk
- ² Department of Electronic Engineering and Computer Science, Queen Mary
University of London, London, UK
p.oliva@qmul.ac.uk
- ³ Department of Economics, University of Mannheim, Mannheim, Germany
eugeniaw@rumms.uni-mannheim.de
- ⁴ OICOS GmbH, Mannheim, Germany
viktor.winschel@gmail.com
- ⁵ Department of Economics, University of St. Gallen, St. Gallen, Switzerland
philipp.zahn@unsig.ch

Abstract. This paper investigates a surprising relationship between decision theory and proof theory. Using constructions originating in proof theory based on higher-order functions, so called *quantifiers* and *selection functions*, we show that these functionals model choice behavior of individual agents. Our framework is expressive, it captures classical theories such as utility functions and preference relations but it can also be used to faithfully model abstract goals such as coordination. It is directly implementable in functional programming languages. Lastly, modeling an agent with selection functions and quantifiers is modular and thereby allows to seamlessly combine agents bridging decision theory and game theory.

Keywords: Decision theory · Utility functions · Preferences · Higher-order functions · Quantifiers

1 Introduction

A *higher-order function* (or *functional*) is a function whose domain is itself a set of functions. In this paper we investigate a surprising and deep connection between decision theory on one side and a higher-order construction that originated in proof theory [3, 4] on the other side: We use a particular class of higher-order functions as a way of describing the choice behavior of individual agents.

Assume X is the overall set of alternatives, and R a set of observable outcomes or measures. For instance, X could be the set of all books, and $R = \mathbb{R}^+$ the non-negative reals representing prices. We then see functions of type $X \rightarrow R$ as describing the agent's *decision context*, e.g. $X \rightarrow \mathbb{R}^+$ as the mapping from books to prices.

The core idea we explore here is to model agents’ decision goals as higher-order functions of type $(X \rightarrow R) \rightarrow R$. Functionals of this type have been called *quantifiers* [14], since the standard \exists and \forall logical quantifiers are particular cases of these when $R = \mathbb{B}$ is the type of booleans. Going back to the example of books and prices, an agent who prefers the cheapest book will be modeled by the quantifier

$$\min: (X \rightarrow \mathbb{R}^+) \rightarrow \mathbb{R}^+$$

saying that given any catalogue of prices $p: X \rightarrow \mathbb{R}^+$ the agent will choose to pay the cheapest price on the catalogue $\min(p)$.

Therefore, quantifiers describe the outcome (e.g. price paid for the choice of book) an agent considers to be good in any given decision context. A corresponding notion is that of a *selection function*, i.e. a higher-order function of type $(X \rightarrow R) \rightarrow X$ which calculates a concrete choice that meets the desired goal. In the example above, the corresponding selection function would pick one of the cheapest books according to the prices given in $p: X \rightarrow \mathbb{R}^+$.

Since the $\max: (X \rightarrow \mathbb{R}) \rightarrow \mathbb{R}$ operator is also a quantifier and the corresponding $\arg \max: (X \rightarrow \mathbb{R}) \rightarrow X$ operator is a selection function, we have that the standard approach of modeling preferences via utility functions or preference relations are instances of our modeling framework. But, quantifiers and selection functions can capture alternative decision criteria as well as decision heuristics. We demonstrate this through several simple examples.

As we see it, there are three crucial advantages of adopting a higher-order modeling of choice behavior: *expressivity*, *modularity*, and *computability*:

- *Expressivity*: We show that describing agent’s choices through higher-order functionals captures existing standard approaches such as utility maximization. Moreover, we can directly model at the level of higher-order functionals and thereby faithfully describe the agent’s behavior. As an example, we introduce a fixed-point operator which captures an agent’s goal to coordinate with other agents.
- *Modularity*: By taking into account the decision context $p: X \rightarrow R$ when describing an agent’s goal, we can fully describe the agent “locally”, without having to refer to “global” utility functions or similar constructs. This allows us to seamlessly combine individual agents into strategic games, bridging the gap between decision theory and game theory.
- *Computability*: Finally, such higher-order function, as abstract as they might appear, can be directly expressed and coded into functional programming languages such as Haskell and Ocaml. In fact, most of the theory we have developed here and in [8] has been implemented and tested in Haskell. This Haskell code has also served to guide us in discovering new forms of game equilibria which are available in this setting of higher-order decisions and games.

The plan for the paper is as follows. We give a brief introduction to higher-order functions in the next section. We then show how quantifiers can be used to model behavior directly at the level of these higher-order functions. In Section

3, we also show how to instantiate preference relations and utility maximization as special cases. In Sections 4 and 4.3 we introduce a series of non-optimizing examples and show how they can be represented by higher-order functions. These include coordination goals as fixpoint operators. We conclude in Section 5.

Elsewhere we show that selection functions are very powerful building blocks for game theory and implementations thereof [8]. Moreover, in [3–5, 7, 10] selection functions are a building block for a compositional approach to game theory. These papers rely on simple instances of selection functions. All instances of selection functions and quantifiers introduced in this paper, can be directly integrated into these game theoretic models.

2 Agents as Quantifiers

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 $X \rightarrow Y$ the set of all functions with domain X and codomain Y . A higher-order function is therefore a function $f : (X \rightarrow Y) \rightarrow Z$ where X , Y and Z are sets.

Here are some well known examples of higher-order functions. In case of the maximization of a utility function $u : X \rightarrow \mathbb{R}$

$$\max_{x \in X} u(x)$$

the max operator takes the utility function $u : X \rightarrow \mathbb{R}$ as its input and returns a real number $\max_{x \in X} u(x)$ as the output. Therefore, the max operator has type

$$\text{max} : (X \rightarrow \mathbb{R}) \rightarrow \mathbb{R}$$

In a similar vein, the arg max operator is also a higher-order function of a particular type:

$$\text{arg max} : (X \rightarrow \mathbb{R}) \rightarrow \mathcal{P}(X)$$

where $\mathcal{P}(X)$ is the set of subsets of X . For a given function $u : X \rightarrow \mathbb{R}$ we have that $\text{arg max}(u)$ is the set of points where u attains its maximum value.

2.1 Agent Context

We want to model an agent \mathcal{A} in a choice *situation* or *context* and formulate his motivations and his choices. We shall model such contexts as mappings $X \rightarrow R$ that encode for choices in X their effects on the outcomes in R .

Definition 1 (Agent context). *We call any function $p : X \rightarrow R$ a possible context for the agent \mathcal{A} who is choosing a move from a set X , having in sight a final outcome in a set R ,*

For instance, X could be the set of available flights between two cities, and $R = \mathbb{R}^+$ could be the set of positive real numbers that represent prices. An agent who is interested in choosing a flight having in mind only the cost of the flight

will consider the price list $X \rightarrow \mathbb{R}^+$ as a sufficient context for his decision. If, however, the number of stops (or changes) is an important factor in the decision of the agent, we could take $R = \mathbb{R}^+ \times \mathbb{N}$ and the agent's context would then be $X \rightarrow \mathbb{R}^+ \times \mathbb{N}$.

2.2 Quantifiers

Suppose the agent \mathcal{A} has to make a decision in the context $p: X \rightarrow R$. The agent will consider some of the possible outcomes to be *good* (or *acceptable*), and others to be bad (or *unacceptable*). Such choices define a higher-order function of the following type:

Definition 2 (Quantifier, [3, 4]). *Mappings*

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

from contexts $p: X \rightarrow R$ to sets of outcomes $\varphi(p) \subseteq R$ are called quantifiers.

The terminology comes from the observation that the usual existential \exists and universal \forall quantifiers of logic can be seen as operations of type $(X \rightarrow \mathbb{B}) \rightarrow \mathbb{B}$, where \mathbb{B} is the type of booleans. Mostowski [14] has called arbitrary functionals of type $(X \rightarrow \mathbb{B}) \rightarrow \mathbb{B}$ *generalized quantifiers*. This was generalized further in [3] to the type given here.

We model agents \mathcal{A} as quantifiers $\varphi_{\mathcal{A}}$ and take $\varphi_{\mathcal{A}}(p)$ as the set of outcomes that the agent \mathcal{A} considers preferable in each context $p: X \rightarrow R$. For instance, consider again the example where X is the set of flights, and $R = \mathbb{R}^+ \times \mathbb{N}$ indicates prices together with the number of stops. An agent that wishes to minimize the cost of the flight, but does not wish to make more than two stops will be modeled by the quantifier:

$$\varphi(p) = \{p(x) : x \in X, p_0(x) = \min p_0, p_1(x) \leq 2\}$$

where $p_0: X \rightarrow \mathbb{R}^+$ and $p_1: X \rightarrow \mathbb{N}$ are the two projections of $p: X \rightarrow \mathbb{R}^+ \times \mathbb{N}$.

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.

The classical example of a quantifier is utility maximization. Suppose an agent has a utility function $u': R \rightarrow \mathbb{R}$ mapping outcomes into utilities. Composing the context $p: X \rightarrow R$ and $u': R \rightarrow \mathbb{R}$ we get a new context that maps actions directly into utility $u: X \rightarrow \mathbb{R}$. Given this new context, the good outcomes for the player are precisely those for which his utility function is maximal. This quantifier is given by

$$\max(u) = \{r \in \text{Im}(u) \mid r \geq u(x') \text{ for all } x' \in X\}$$

where $\text{Im}(u)$ denotes the image of the utility function $u: X \rightarrow \mathbb{R}$.

2.3 Context-dependence

In general, we are going to allow the set of outcomes that the agent considers good to be arbitrary. It is reasonable, however, to assume that for each context $p: X \rightarrow R$ we have $\varphi(p) \neq \emptyset$. This is to say that in any given context the agent must have a preferred outcome (even if this would be the least bad one). We will call such quantifiers *total*. Another more interesting class of quantifiers consists of those we call *context-independent*:

Definition 3 (Context-independence). A quantifier $\varphi: (X \rightarrow R) \rightarrow \mathcal{P}(R)$ is said to be context-independent if the value $\varphi(p)$ only depends on $\text{Im}(p)$, i.e.

$$\text{Im}(p) = \text{Im}(p') \implies \varphi(p) = \varphi(p').$$

Dually, a quantifier φ will be called context-dependent if for some contexts p and p' , with $\text{Im}(p) = \text{Im}(p')$, the sets of preferred outcomes $\varphi(p)$ and $\varphi(p')$ are different.

Intuitively, a context-dependent quantifier will select good outcomes not only based on which outcomes are possible, but will also take into account how the outcomes are actually achieved. It is easy to see that the quantifier max is context-independent, since it can be written as a function of $\text{Im}(p)$ only.

Our prototypical example of a context-dependent quantifier is the fixpoint operator

$$\text{fix} : (X \rightarrow X) \rightarrow \mathcal{P}(X)$$

Recall that a fixpoint of a function $f : X \rightarrow X$ is a point $x \in X$ satisfying $f(x) = x$. If the set of moves is equal to the set of outcomes then there is a quantifier whose good outcomes are precisely the fixpoints of the context. If the context has no fixpoints we shall assume that the agent will be equally satisfied with any outcome. Such a quantifier is given by

$$\text{fix}(p) = \begin{cases} \{x \in X \mid p(x) = x\} & \text{if nonempty} \\ X & \text{otherwise.} \end{cases}$$

Clearly $\text{fix}(\cdot)$ is context-dependent, since we could have different contexts $p, p' : X \rightarrow X$ having the same image set $\text{Im}(p) = \text{Im}(p')$ but with p and p' having different sets of fixpoints. For example, if we take $p, p' : \mathbb{R} \rightarrow \mathbb{R}$ to be given by $p(x) = x$ and $p'(x) = -x$ then $\text{Im}(p) = \text{Im}(p') = \mathbb{R}$, but $\text{fix}(p) = \mathbb{R}$ and $\text{fix}(p') = \{0\}$. We will discuss the relevance of this quantifier in Section 4.3.

2.4 Attainability

Another important property of quantifiers that we shall consider is that of *attainability*:

Definition 4 (Attainability). A quantifier $\varphi : (X \rightarrow R) \rightarrow \mathcal{P}(R)$ is called attainable if, for every context $p : X \rightarrow R$, for some $r \in \varphi(p)$ there exists an x such that $p(x) = r$. (In particular, attainable quantifiers are total.)

In other words, an agent modeled by an attainable quantifier will select at least one preferred outcome r that is actually *achievable* by some move x . An equivalent definition is that $\varphi: (X \rightarrow R) \rightarrow \mathcal{P}(R)$ is attainable if and only if

$$\varphi(p) \cap \text{Im}(p) \neq \emptyset.$$

Remark 1. We could also define a *strong attainability* notion whereby all $r \in \varphi(p)$ need to be achievable by some $x \in X$, i.e.

$$\varphi(p) \subseteq \text{Im}(p).$$

For our purposes the weaker notion of Definition 4 has been sufficient and reasonably well-behaved.

Attainable quantifiers bring out the relevance of moves in the decision making process. Sometimes an agent might actually wish to spell out the preferred *moves* instead of the preferred *outcomes*. This leads to the definition of another class of higher-order functions:

Definition 5 (Selection functions). A selection function is any function of type

$$\varepsilon: (X \rightarrow R) \rightarrow \mathcal{P}(X)$$

Remark 2. In the computer science literature, where quantifiers and selection functions have been considered previously, the initial focus was on single-valued ones [3]. Multi-valued quantifiers were first considered in [4]. As the definition above shows, we also make use of multi-valued selection functions, as these are extremely important in our examples and also in game theoretic approaches that take selection functions as a building block [5, 7, 10].

Similarly to quantifiers, the canonical example of a selection function is maximizing \mathbb{R} , defined by

$$\arg \max(p) = \{x \in X \mid p(x) \geq p(x') \text{ for all } x' \in X\}$$

The arg max selection function is naturally multi-valued: a function may attain its maximum value at several different points.

Proposition 1. A quantifier $\varphi: (X \rightarrow R) \rightarrow \mathcal{P}(R)$ is attainable if and only if there exists a total selection function $\varepsilon: (X \rightarrow R) \rightarrow \mathcal{P}(X)$ such that, for all $p: X \rightarrow R$,

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

If such a relationship between a quantifier $\varphi: (X \rightarrow R) \rightarrow \mathcal{P}(R)$ and a selection function $\varepsilon: (X \rightarrow R) \rightarrow \mathcal{P}(X)$ holds then we shall say that ε *attains* φ . The attainability relation holds between the quantifier max and the selection function arg max. The fixpoint quantifier is also a selection function, and it attains itself since

$$x \in \text{fix}(p) \implies p(x) \in \text{fix}(p).$$

3 Preference Relations and Utility Maximization

In this section we relate the concepts of quantifiers and selection functions to standard concepts of classical (economic) decision theory: utility functions and preference relations. In particular, we show that both correspond to context-independent quantifiers that have the same structure. We now want to characterize the relationship between preference relations and context-independent quantifiers.

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

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

$$(X \rightarrow R) \xrightarrow{\text{Im}} \mathcal{P}(R) \xrightarrow{f} \mathcal{P}(R)$$

so that $f \circ \text{Im}: (X \rightarrow R) \rightarrow \mathcal{P}(R)$ are quantifiers.

Proposition 2. *Assume $|X| \geq |R|$, i.e. the number of choices is bigger than the number of possible outcomes. Then a quantifier $\varphi: (X \rightarrow R) \rightarrow \mathcal{P}(R)$ is context-independent if and only if $\varphi = f \circ \text{Im}$, for some choice function $f: \mathcal{P}(R) \rightarrow \mathcal{P}(R)$.*

Proof. If $\varphi = f \circ \text{Im}$ then φ is context-independent. For the other direction, note that since $|X| \geq |R|$ we have for any subset $S \subseteq R$ a map $u_S: X \rightarrow R$ such that $\text{Im}(u_S) = S$. Assume φ is context-independent and let us define $f(S) = \varphi(u_S)$. Clearly,

$$\varphi(p) = \varphi(u_{\text{Im}(p)}) = f(\text{Im}(p))$$

where the first step uses that φ is context-independent and that $\text{Im}(p) = \text{Im}(u_{\text{Im}(p)})$ by the assumption on the family of maps u_S , while the second steps simply uses the definition of f .

Agents who are defined by context-independent quantifiers are choosing the set of good outcomes simply by ranking the set of outcomes that can be achieved in a given context but are ignoring 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 classical agent who cares only about the direct consequences of his decision and is defined in a context-independent way would choose the outcome that gives himself the maximum sum of money, regardless of the nature of action. If however the agent also cares about the actions themselves and their indirect consequences, he might not consider the largest amount of money as preferable.

The following proposition guarantees the attainability of context independent quantifiers arising from preference relations:

Proposition 3. *Whenever f_{\succeq} is a choice function arising from a partial order \succeq , then the context-independent quantifier $\varphi = f_{\succeq} \circ \text{Im}$ is attainable.*

Proof. By the definition of φ we have that if $r \in \varphi(p)$ then r is a maximal element in $\text{Im}(p)$. Hence we must have an $x \in X$ be such that $p(x) = r$.

Another example of a context-independent quantifier is the maximization of a utility function. A utility function can be characterized as the context $p: X \rightarrow \mathbb{R}$ that attaches a real number to each element of the set of choices X with the quantifier defined as

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

Moreover, this quantifier is attained by the selection function

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

Note the types $\varphi: (X \rightarrow \mathbb{R}) \rightarrow \mathcal{P}(\mathbb{R})$ and $\varepsilon: (X \rightarrow \mathbb{R}) \rightarrow \mathcal{P}(X)$ respectively. And indeed we have that

$$x \in \varepsilon(p) \implies p(x) \in \varphi(p).$$

Thus, max and arg max operators, are the prototypical examples of a context-independent quantifier and a selection function attaining it.

4 Alternatives to Optimization

We have seen how the higher-order notion of a context-independent quantifiers is able to model choices based on rational preferences (or equivalently on utility maximization). For simplicity, we consider decisions under certainty but it is straightforward to consider uncertainty and consider expected utility theory. Other decision criteria such as regret minimization [13, 15] or maximin choices [16] are also captured by our framework.

Instead of going in this direction, in the following we consider another direction: Utility functions as well as preference relations are intimately linked to the assumption that the agent fully optimizes. The behavioral economic literature as well as the psychological literature have documented deviations from optimizing behavior [2, 11]. Quantifiers provide a direct way to model such deviations. Here we give a few examples by allowing for a different structure on the set of outcomes R or by allowing for a different mapping $f: \mathcal{P}(R) \rightarrow \mathcal{P}(R)$, or by relaxing both.

4.1 Context-independent agents

Example 1 (Averaging Agent). Consider an agent who prefers the outcome to be as close as possible to the average of all achievable outcomes. Given a decision

context $p: X \rightarrow \mathbb{R}$, the average amongst the possible outcomes can be calculated as

$$A_p = \frac{\sum_{r \in \text{Im}(p)} r}{|\text{Im}(p)|}$$

Therefore, such agent can be directly modeled via the averaging quantifier $\varphi^A: (X \rightarrow \mathbb{R}) \rightarrow \mathcal{P}(\mathbb{R})$ as

$$\varphi^A(p) = \{r \in \text{Im}(p) \mid |r - A_p| \text{ is minimal}\}$$

The next example represents the second best decision problem discussed in [12].

Example 2 (Second-best Agent). Consider a simple heuristic of a person ordering wine in a restaurant whereby he always chooses the second most-expensive wine. In terms of quantifiers, let X be the set of wines available in a restaurant, and $p: X \rightarrow \mathbb{R}$ the price attached to each wine x_i ($i = 1, \dots, N$) on the menu, so that $r_i = p(x_i)$ denotes the price of wine x_i . Given a maximal strict chain $r_n > r_{n-1} > \dots > r_1$ in \mathbb{R} , let us call r_{n-1} a sub-maximal element. The goal of the agent can be described by the quantifier

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

A crucial point of the above examples is the additional degree of freedom of modeling as it is possible to vary the choice operator itself and not being automatically restricted to the max operator.

4.2 Context-dependent agents

So far, we have focused only on context-independent quantifiers. Yet, we can do more. As we have discussed in Section 2, we can allow for quantifiers that do not only take the image of p as input but the complete function. Again, we consider several examples.

Example 3 (Ideal-move Agent [6]). Let $r > 0$ be a fixed real number. For a point $v \in \mathbb{R}^n$ we define the closed ball with centre v and radius r by

$$B(v; r) = \{w \in \mathbb{R}^n \mid d(v, w) \leq r\}$$

where d is the Euclidean distance. Let the set of choices X have a distinguished element $x_0 \in X$. Define the quantifier $\varphi: (X \rightarrow \mathbb{R}^n) \rightarrow \mathcal{P}(\mathbb{R}^n)$ by

$$\varphi(p) = B(p(x_0); r)$$

This quantifier is attained by the constant selection function $\varepsilon(p) = \{x_0\}$.

The last example illustrates Simon's satisficing behavior. The value $r > 0$ can be considered as a satisficing threshold around outcomes that are close to the outcome of an ideal point. Such an agent is equally satisfied with all outcomes which are close enough to the outcome of the ideal choice.

Example 4 (Averaging – revised). Consider again an agent who prefers the outcome to be as close as possible to the average outcome. But this time we assume that he takes into account the number of possible ways an outcome may be attained. Given a decision context $p: X \rightarrow \mathbb{R}$, the weighted average in this case can be calculated as

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

Such agent can be modeled via the weighted averaging quantifier $\varphi: (X \rightarrow \mathbb{R}) \rightarrow \mathcal{P}(\mathbb{R})$ as

$$\varphi(p) = \{r \in \text{Im}(p) \mid |r - A_p| \text{ is minimal}\}$$

It is easy to check that this is a *context-dependent* quantifier.

Now, consider the example where the set of actions allows an agent to earn some money but some actions are illicit and hence not considered to be a permissible behavior. If we care about the actions themselves, we might not necessarily consider the largest sum of money as preferable.

Example 5 (Honest Agents). Consider an agent with a set of possible actions X leading to monetary outcomes $M \subseteq \mathbb{R}$. Assume some of these actions $I \subset X$ are illegal or dishonest. Hence, the set $L = X \setminus I$ consists of the legal, or honest, actions. In the first instance consider an honest agent who maximizes over the outcomes which follows from honest actions. Such a honest agent can be modeled by the quantifier:

$$\varphi^h(p) = \{r \mid r \text{ a maximal element in the set } p(L)\}$$

where $p(L)$ is the image of L under p . Consider, however, a more complicated case where the agent is prepared to consider dishonest or illegal actions when the reward associated with some of these actions is above a threshold T . This subtler preference can be directly modeled as

$$\varphi^d(p) = \begin{cases} \{r \mid r \text{ is maximal in } \text{Im}(p)\} & \text{if } \max_{x \in I} p(x) > T \\ \varphi^h(p) & \text{otherwise} \end{cases}$$

so that the dishonest agent will behave as the honest one if the maximal reward for a dishonest action is low, but he will consider any action to be acceptable if the gain from a dishonest or illegal action is high enough.

In the next example we introduce an extreme case of an agent who decides on preferred outcomes solely based on the set of moves that lead to that outcome.

Example 6 (Safe Agents). Given a decision context $p: X \rightarrow R$ and an outcome $r \in \text{Im}(p)$, we can calculate the number of different ways r can be attained by

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

We say that an outcome r is most unavoidable if n_r^p is maximal over the set of possible outcomes $\text{Im}(p)$. We say that an agent is safe if he prefers most unavoidable outcomes. Such agents are modeled by the quantifier

$$\phi(p) = \{r \in \text{Im}(p) \mid n_r^p \text{ maximal} \}$$

In order to illustrate this quantifier, suppose there are three 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 has to choose which highway to take, and the outcome is the beach that the agent goes to. The safe agent decides to visit the beach which can be reached by the most different routes, which is the third, in order to avoid the risk of being stuck in a traffic jam.

4.3 Fixed points as coordination

We now discuss the specific situations where the set of actions X and outcomes R are the same $X = R$. In this case elements of the type

$$(X \rightarrow X) \rightarrow \mathcal{P}(X)$$

can be either viewed as quantifiers or selection functions. Agents of this type are common in voting contests:

Example 7 (Voting Agent). Consider three judges $J = \{J_1, J_2, J_3\}$ voting for two contestants $X = \{A, B\}$. The winner is determined by the simple majority rule of type $\text{maj} : X \times X \times X \rightarrow X$. The set X denotes both the set of choices and the set of possible outcomes of the contest. We first assume that the judges rank the contestants according to a preference ordering. For example, suppose judges 1 and 2 prefer A and judge 3 prefers B . Consider the decision problem of the first judge. He has an ordering on the set X , namely $A \succeq_1 B$, and his goal is to maximize the outcome with respect to this ordering. Hence, he is modeled via the quantifier:

$$\varphi_1^J(p) = \max_{x \in (X, \succeq_1)} p(x)$$

The set X is equipped with a partial order and the max operator $(X \rightarrow X) \rightarrow \mathcal{P}(X)$ describes the agent.

Another very interesting example of an agent with an important economic interpretation, is the fixpoint operator, that we have already mentioned in Section 2.3.

Example 8 (Keynesian Agent). Consider the last example but now assume that judge 1 has different preferences: he prefers to support the winner of the contest. He is only interested in voting for the winner of the contest and he has no preferences for the contestants per se. The selection function of such a Keynesian agent can be described by a fixpoint operator as

$$\varepsilon_1^K(p) = \text{fix}(p) = \{x \in X \mid p(x) = x\}.$$

Interestingly, such an agent is best described by a selection function, rather than via the corresponding quantifier

$$\varphi_1^K(p) = \{p(x) \mid p(x) = x\}.$$

We note, it is perfectly possible to model such a Keynesian agent via standard utility functions, attaching say utility 1 to good outcomes and 0 to the bad ones, so that the judges maximize over the set of monetary payoffs. In this process of attaching utilities to the decision, however, one has to compute the outcome of the votes, then to check for the second and the third judges whether their vote is in line with the outcome, and finally to attach the utilities. If, instead, we use the fixpoint operator to represent the agent's goals, no such calculation is necessary.

As briefly discussed above, most functions $p: X \rightarrow X$ do not have a fixpoint and the fixpoint operator will often give the empty set. For the purposes of modeling a particular situation we might want to totalize the fixpoint operator in different ways and describe what an agent might do in case that no fixpoint exists. The fixpoint goals are far more interesting when we consider a game with several agents with different concerns, for instance some with usual preferences and some with fixpoint goals. We analyze such a game in [8].⁶

Let us conclude with another example of a reflexive agent.

Example 9 (Coordinating Agent). Consider two players, $\{0, 1\}$, who want to coordinate, for instance, about the restaurant where to meet for lunch. The set of actions $X_0 = X_1 = \{A, B\}$ denotes the different restaurants at choice. The set of outcomes $R = X_0 \times X_1$ denotes the two restaurants where the agents might end up. The fact that these two agents want to meet in the same restaurant can be directly described by another sort of fixpoint operator:

$$\varepsilon_i(p) = \{x \in X_i \mid x = (\pi_{1-i} \circ p)(x)\}$$

where $\pi_i: X_0 \times X_1 \rightarrow X_i$ are the projection functions. The preferred move of agent i is the one which leads him to the same place as the other agent $1 - i$.

These two examples above show that the overall goal of the Keynesian and the coordinating agent are similar, and can be captured by some variants of the same fixpoint operators. Even though it is possible to use utility functions in order to model these concerns in the particular examples, it is not obvious that this commonality can be made explicit when modeling with utility functions. In our more abstract formalization via higher-order functions, it is possible to detect patterns across problems that are hard to find when one only looks at the compiled level of utility maximization.

5 Conclusions

We introduced quantifiers and selection functions to model agents' choices. We illustrated that classical and standard approaches such as utility functions can

⁶ See also the working paper version [9] for more details.

be instantiated in our framework as examples. Alternatives to optimization can be similarly captured. Lastly, one can directly model at the level of these higher-order functions. Overall, higher-order functions provide a possibility to abstraction of lower-level instantiations and by that realize commonality between seemingly different approaches.

In this paper, we limit ourselves to show that quantifiers and selection functions do capture different deterministic approaches. We already noted above that decision-making under uncertainty such as expected utility theory can also be represented in our framework. In fact, analogous to the deterministic case, different theories can be dealt with. What is more, there exist non-deterministic and probabilistic extensions of the quantifiers and selection functions based on monads. In future research we will explore how these constructs can be used to model decisions under uncertainty and under risk more generally.

Also not explored in this paper but on our agenda for the future is the composition of selection functions which can be naturally defined. By that one can consider the aggregation of different individuals, as for instance in the literature of social choice, or model individuals as “multiple selves”, as for instance in [12, 1], where different dimensions of an agent are aggregated and the different dimensions taken together determine his final choices.

Lastly, and probably most importantly, selection functions are a building block for game theoretic approaches built on high-order functions. Thus, the ability to express various goals as discussed here scales to strategic interactions. In [8] we show that goal functions such as the fixed point player introduced above can be fruitfully applied to model voting contests. More generally, we show that the Nash equilibrium extends to games based on quantifiers and selection functions; we show that selection functions and quantifiers yield the same set of equilibria in the case of max and argmax operators; but we also show that for other classes of goal functions, such as the fixed point agent, quantifiers and selection functions yield different equilibria. In fact, we show that the equilibria induced by selection functions are a refinement of the equilibria induced by quantifiers. As a last point, we also explore selection functions as a building block to *open games* in [5, 7, 10]. Open games are a further abstraction based on category theory. Selection functions are an essential component because as in this paper they represent the individual agent’s goal function.

References

1. Attila Ambrus and Kareen Rozen. Rationalising choice with multi-self models. *The Economic Journal*, 125(585):1136–1156, 2015.
2. C.F. Camerer. *Behavioral Game Theory: Experiments in Strategic Interaction*. The Roundtable Series in Behavioral Economics. Princeton University Press, 2011.
3. Martín Escardó and Paulo Oliva. Selection functions, bar recursion and backward induction. *Mathematical structures in computer science*, 20(2):127–168, 2010.
4. Martín Escardó and Paulo Oliva. Sequential games and optimal strategies. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science*, 467(2130):1519–1545, 2011.

5. Neil Ghani and Jules Hedges. A compositional approach to economic game theory. *CoRR*, abs/1603.04641, 2016.
6. Jules Hedges. A generalisation of Nash’s theorem with higher-order functionals. *Proc R Soc A*, 469(2154), 2013.
7. Jules Hedges. String diagrams for game theory. *CoRR*, abs/1503.06072, 2015.
8. Jules Hedges, Paulo Oliva, Evguenia Shprits, Viktor Winschel, and Philipp Zahn. Selection equilibria of higher-order games. In *Practical Aspects of Declarative Languages - 19th International Symposium, PADL 2017, Paris, France, January 16-17, 2017, Proceedings*, pages 136–151, 2017.
9. Jules Hedges, Paulo Oliva, Evguenia Sprits, Viktor Winschel, and Philipp Zahn. Higher-order game theory. *CoRR*, abs/1506.01002, 2015.
10. Jules Hedges, Evguenia Shprits, Viktor Winschel, and Philipp Zahn. Compositionality and string diagrams for game theory. *CoRR*, abs/1604.06061, 2016.
11. D. Kahneman. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, 2011.
12. Gil Kalai, Ariel Rubinstein, and Ran Spiegler. Rationalizing choice functions by multiple rationales. *Econometrica*, 70(6):2481–2488, November 2002.
13. Graham Loomes and Robert Sugden. Regret theory: An alternative theory of rational choice under uncertainty. *The Economic Journal*, 92(368):805–824, 1982.
14. Andrzej Mostowski. On a generalization of quantifiers. *Fundamenta Mathematicae*, 44(1):12–36, 1957.
15. L. J. Savage. The theory of statistical decision. *Journal of the American Statistical Association*, 46(253):55–67, 1951.
16. Abraham Wald. *Statistical decision functions*. Wiley, Oxford, England, 1950.