University of Mannheim / Department of Economics

Working Paper Series

Context Dependent Games as Quantifiers and Selection Functions

Jules Hedges Paulo Oliva
Evguenia Winschel Viktor Winschel Philipp Zahn

Working Paper 14-13

April 2014

Context Dependent Games as Quantifiers and 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 22, 2014

Abstract

We use *quantifiers* and *selection functions* to represent simultaneous move games. Quantifiers and selection functions are examples of higher-order functions. A *higher order function* is a function whose domain is itself a set of functions. Thus, quantifiers and selection functions allow players to form goals not only about outcomes but about the whole (or parts) of the game play. They encompass standard preferences and utility functions as special cases, but also extend to non-maximizing behavior and context-dependent motives. We adapt the Nash equilibrium concept to our new representation and also introduce a refinement to capture the essential features of context-dependent motives. As an example, we discuss fixpoint operations as context dependent goals of coordination and differentiation in simultaneous game variants of Keyne's beauty contest and the minority game.

JEL codes: C0, D01, D03, D63, D64
Keywords: context dependent refinement of Nash equilibrium, higher order functions, quantifiers, selection functions, beauty contest, minority game, endogenous economist

Contents

1 Introduction

We propose a generalized representation of finite, simultaneous games. Our framework is based on quantifiers and selection functions which have been introduced in proof theory and computer science [4, 5]. Quantifiers express agents' motivation in a natural and high-level way. They can express motives that cannot be modelled directly by preference relations and utility functions, in particular motives that are dependent on the context in which a decision is made. Since quantifiers and selection functions are higher-order functions, there is a close connection to high-level functional programming languages. As such, the formalism we propose is constructive and equilibria are computable by directly implementable functional programs.

The motives we address in this paper are coordination and differentiation that underly Keynes beauty contest and the minority game. The new feature we add to game theory is that we express in a high-level way the goal of coordination via fixpoints, and differentiation via non-fixpoints. The intuition for agents pursuing such goals is that they are solving the game within the game, using the fact that an equilibrium is itself a fixpoint.

Once agents are modeled to reason about the game itself, via selection functions that represent fixpoints of the game, it becomes necessary to adjust the equilibrium concept of the Nash equilibrium. We propose an equilibrium refinement, which we call a *context dependent Nash equilibrium*, that takes into account that fixpoint goals are context dependent, that is, they depend not only on what outcomes can be attained but how they are attained. We contrast the fixpoint goals to the usual context independent quantifiers and selection functions of max and argmax in classical, i.e. context independent, game theory. In this special case, every Nash equilibrium is a context dependent equilibrium. In our examples, the context dependent variant of the Nash equilibrium perfectly captures the intuitively plausible equilibria in games where agents take into account the actions of other agents in a context dependent way.

2 Context Dependent Game Theory

2.1 Quantifiers

Quantifiers are a very general way of representing the aims or goals of a decision-making agent. Whenever an agent decides, there is an associated *context*, which is a function mapping the agent's moves to final outcomes. Suppose the agent must choose between a set of moves X. A rational agent operating under common knowledge of rationality can associate an outcome $q(x)$ to every move x that may be chosen. An example of a commonly used context is when we fix moves by all but one player in a simultaneous game, and one player is unilaterally changing his strategy. In this situation or context we can unambiguously assign an outcome to every possible move.

A quantifier is a rule for assigning to *every context* a set of 'good' outcomes. A quantifier is therefore a *higher order function* (or *functional*), that is a function whose domain is itself a set of functions. If an agent makes a choice between a set of moves X and the outcomes form a set R, then a quantifier for the agent is any higher order function

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

that takes a function of type $X \to R$ as an argument and returns an element from the power-set $\mathcal{P}(R)$ of R.

The classical example of a quantifier is utility maximisation. In this case the set of outcomes is $R = \mathbb{R}^n$, where the ith element represents the utility of the ith player. Given a context $q : X \to \mathbb{R}^n$, the good outcomes for the ith player are precisely those for which the ith coordinate, i.e. his utility function, is maximal. This quantifier is given by

$$\varphi_i(q) = \{r \in \mathrm{Im}(q) \mid r_i \geq q(x')_i \text{ for all } x' \in X\}$$

where $\mathrm{Im}(q)$ denotes the image of the function $q\colon X \to R$.

There are quantifiers very different from utility maximisation, which we considered in our accompanying paper [9].

An extension of the classical decision theoretical setting that we are pursuing in [9] and extend here into a game theoretical setting is the notion of context independent quantifiers and context dependent ones. Classical decision theory relies on preferences and utility functions which only consider outcomes and cannot directly consider the context in which a decision is made. Yet, there are many economic situations where agents do not only care about final outcomes but also (or even exclusively) care about the process how results come about.[1] Context dependent quantifiers capture exactly goals which are dependent on the context.

One specific example of a context dependent quantifier which we will use is the fixpoint quantifier. Recall that a fixpoint of a function $f : X \to X$ is a point $x \in X$ satisfying $f(x) = x$. When the set of moves is equal to the set of outcomes (for example in an election) there is a quantifier whose good outcomes are precisely the fixpoints of the context. If the context has no fixpoint and the agent will be equally happy (or equally unhappy) with any outcome, then the quantifier is given by

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

$$\varphi(q) = \begin{cases} \mathrm{fix}(q) & \text{if } \mathrm{fix}(q) \neq \varnothing \\ X & \text{otherwise} \end{cases}$$

where

$$\mathrm{fix} : (X \to X) \to \mathcal{P}(X)$$
$$\mathrm{fix}(q) = \{x \in X \mid q(x) = x\}$$

2.2 Selection functions

Just as a quantifier tells us which outcomes an agent considers good, selection functions tell us which moves are good. Thus a selection function is any

[1] In particular in behavioral economics there is wide evidence that subjects take into account the context in which they make a decision, for instance when fairness concerns play a role or when deciding in uncertain contexts. See [3, 10] for overviews.

function of the form

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

To this we add the condition that $\varepsilon(q) \neq \varnothing$, because the agent must always select a move.

In the computer science literature where selection functions have been considered previously the focus was on single-valued ones. However, as multi-valued selection functions are extremely important in our examples we have adapted the definitions accordingly.

As for quantifiers, the canonical example of a selection function is maximising one coordinate in $\mathbb{R}^n$, defined by

$$\varepsilon_i(q) = \{x \in X \mid q(x)_i \geq q(x')_i \text{ for all } x' \in X\}$$

Even in one dimension, the $\arg\max$ selection function is naturally multi-valued: a function may attain its maximum value at several different points.

There is an important relation between quantifiers and selection functions called attainment. Intuitively this means that the outcome of a good move should be a good outcome. Formally, given a quantifier $\varphi : (X \to R) \to \mathcal{P}(R)$ and a selection function $\varepsilon : (X \to R) \to \mathcal{P}(X)$, we say that ε attains φ iff for all contexts $q : X \to R$ it is the case that

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

For example, this relation holds between the quantifier and the selection function which maximise over $\mathbb{R}^n$. The fixpoint quantifier is also itself a selection function, and it attains itself since

$$x \in \mathrm{fix}(q) \implies q(x) \in \mathrm{fix}(q)$$

When modelling a situation, we consider that a quantifier describes the goals of an agent and a selection function describes the strategy of the agent attaining the quantifier. Seen in this way, attainment becomes a general notion of *rationality*.

Given a selection function ε, we can form the 'smallest' quantifier which it attains. This is called $\overline{\varepsilon}$ and is defined by the equation

$$\overline{\varepsilon}(q) = \{q(x) \mid x \in \varepsilon(q)\}$$

Thus the good outcomes according to the quantifier $\overline{\varepsilon}$ are exactly the outcomes resulting from good moves according to the selection function ε.

2.3 Simultaneous games

A simultaneous game can be solely described by a family of selection functions and an outcome function. For instance, consider a n-player game with a set R of outcomes and sets X_i of strategies for the ith player. Such game can be described in two parts:

1. the outcome function
$$q : \prod_{i=1}^{n} X_i \to R$$
i.e., a mapping from the strategy profile to the final outcome,

2. for each player $1 \leq i \leq n$, a selection function $\varepsilon_i : (X_i \to R) \to \mathcal{P}(X_i)$ representing that player's prefered strategy in the game.

Intuitively, we think of the outcome function q as representing the 'situation', or the rules of the game, while we think of the selection functions as describing the agents. Thus we can imagine the same agent in different situations, and different agents in the same situation. This allows us to decompose a modelling problem into a global and a local part: modelling the situation and modelling the players.

In the specific example in which $R = \mathbb{R}^n$ and ε_i is the selection function which maximises the ith coordinate, we obtain ordinary normal-form games.

2.4 Unilateral contexts and equilibria

Consider a strategy profile $\mathbf{x} \in \prod_i X_i$. The outcome of this strategy profile is $q(\mathbf{x})$. We want to define the context in which one player unilaterally changes his strategy. This is given by

$$\mathcal{U}_i^q(\mathbf{x})(x) = q(\mathbf{x}[i \mapsto x]).$$

The functions

$$\mathcal{U}_i^q : \prod_{j=1}^{n} X_j \to (X_i \to R)$$

are called *unilateral maps*. They were introduced in [7] in which it is shown that the proof of Nash's theorem amounts to showing that the unilateral maps have certain topological (continuity and closure) properties. The concept of a context was introduced later in [8], so now we can say that $\mathcal{U}_i^q(\mathbf{x}) : X_i \to R$ is the context in which the ith player has unilaterally changed his strategy, so we call it a unilateral context.

Using these concepts we can generalize the definition of Nash equilibrium to games defined by selection functions.

Definition 2.1 (Nash equilibrium) *A strategy profile* $\mathbf{x}$ *is a Nash equilibrium iff it results in a good outcome in every context in which a player has unilaterally changed strategy. Since the set of good outcomes for the* i*th player in the context* $\mathcal{U}_i^q(\mathbf{x})$ *is given by*

$$\overline{\varepsilon_i}(\mathcal{U}_i^q(\mathbf{x})) \in \mathcal{P}(R)$$

the condition for $\mathbf{x}$ *to be a Nash equilibrium is*

$$q(\mathbf{x}) \in \overline{\varepsilon_i}(\mathcal{U}_i^q(\mathbf{x}))$$

for all players $1 \leq i \leq n$.

This definition of a Nash equilibrium says that the outcome resulting from a strategy profile is a good outcome in the context in which one player unilaterally changes strategy.

Definition 2.2 (Context dependent Nash equilibrium) *Each player's* *move is a good move in the unilateral context, that is,*

$$x_i \in \varepsilon_i(\mathcal{U}_i^q(\mathbf{x}))$$

for all players $1 \leq i \leq n$*. A strategy profile satisfying this condition for each player is called a context dependent Nash equilibrium.*

In classical game theory in which all quantifiers are max and all selection functions are arg max these two definitions are equivalent. However in general, and in the specific examples we use, context dependent equilibria are an equilibrium refinement of Nash equilibria. Moreover they are important in practice: our specific examples contain implausible Nash equilibria which are not context dependent equilibria. Indeed, we have examples in which every strategy is a Nash equilibrium, but there are few context dependent equilibria.

The general concept is qualitative rather than quantitative: in an equilibrium no player can unilaterally improve his situation by changing from a bad outcome to a good outcome.

What is the relation between Nash equilibria and context dependent equilibria? In fact, it is easy to prove that context dependent equilibria are an equilibrium refinement of Nash equilibria. Recall that by definition, for every context p we have

$$x \in \varepsilon_i(q) \implies q(x) \in \overline{\varepsilon_i}(q)$$

Consider in particular the unilateral context $\mathcal{U}_i^q(\mathbf{x})$. Assuming that $\mathbf{x}$ is a context dependent equilibrium we have

$$x_i \in \varepsilon_i(\mathcal{U}_i^q(\mathbf{x}))$$

Therefore

$$\mathcal{U}_i^q(\mathbf{x})(x_i) \in \overline{\varepsilon_i}(\mathcal{U}_i^q(\mathbf{x}))$$

It remains to note that $\mathcal{U}_i^q(\mathbf{x})(x_i) = q(\mathbf{x})$, because $\mathbf{x}[i \mapsto x_i] = \mathbf{x}$.

These definitions are rather abstract and we will illustrate them with specific examples. In the next section we will prove that for classical games both definitions reduce to ordinary Nash equilibrium, and also provide an example that shows that this does not hold when investigating more general games.

2.5 Equivalence for classical games

The reason why the distinction between Nash equilibria and context dependent equilibria is surprising is that both coincide in classical game theory. Fix a strategy profile $\mathbf{x}$ and let $f : X \to R$ be the unilateral context $f = \mathcal{U}_i^q(\mathbf{x})$. In order for Nash equilibria and context dependent equilibria to coincide it must be the case that the condition

$$f(x_i) \in \overline{\varepsilon_i}(f)$$

implies the condition

$$x_i \in \varepsilon_i(f).$$

Consider the particular selection functions

$$\varepsilon_i(q) = \{x \in X \mid q(x)_i \geq q(x')_i \text{ for all } x' \in X\}$$

and the quantifiers

$$\overline{\varepsilon_i}(q) = \{q(x) \mid q(x)_i \geq q(x')_i \text{ for all } x' \in X\}$$

which are used in classical game theory. Now both of the conditions are equivalent to

$$f(x_i) \geq f(x'), \text{ for all } x' \in X$$

To see that this is just the definition of a Nash equilibrium, again we note that $f(x_i) = q(\mathbf{x})$.

If we suppose for simplicity that our quantifiers and selection functions are single-valued, we can see that this is a cancellation argument, namely we go from

$$f(x_i) = \max(f) = f(\arg\max f)$$

to

$$x_i = \arg\max f$$

Typically a cancellation of this form is invalid unless f is an injective function. Thus from an algebraic point of view it is a remarkable property of arg max that this cancellation

$$f(x) = f(\arg\max f) \implies x = \arg\max f$$

is always valid, even when f is not injective. From an alternative point of view, this 'remarkable property' is nothing but the definition of arg max: if the value of f at x_i is maximal then x_i is in arg max of f.[2]

It is also easy to see that the same cancellation argument fails if we replace arg max with another single-valued selection function, and this is the reason that Nash equilibria and context dependent equilibria are distinct. Suppose we replace arg max with a single-valued fixpoint operator (ignoring the fact that some functions have no fixpoint). Then the equivalent cancellation

$$f(x) = f(\text{fix} f) \implies x = \text{fix} f$$

is invalid. Consider in particular the constant function $f : \mathbb{R} \to \mathbb{R}$ given by $f(x) = 1$, and the particular value $x = 0$. The unique fixpoint of f is $\text{fix} f = 1$. Therefore it is the case that $f(x) = 1 = f(\text{fix} f)$, but it is not the case that $x = \text{fix} f$.

Overall we have proved that every classical Nash equilibrium is a context dependent equilibrium, and every context independent equilibrium is a (generalized) Nash equilibrium. This can be represented by the inclusion chain

$$\text{classical Nash} \subsetneq \text{context dependent Nash} \subsetneq \text{generalized Nash equilibria.}$$

[2]This 'deep triviality' is reminiscent of the Curry-Howard isomorphism in logic, which is the connection between proofs and computer programs. A simple proof system (natural deduction for minimal logic) and a simple programming language (simply-typed λ calculus) are equal by definition, but the philosophical implications are huge and the consequences are still being explored.

3 Examples

We focus our examples on a simple situation: there are three players, the judges $J = \{J_1, J_2, J_3\}$. The set of strategies $X = \{A, B\}$ are the same for each judge which are votes for two contestants A and B. The set of outcomes is given by $X = \{A, B\}$ denoting the winner of the contest. The winner is determined by the majority rule of type $\mathrm{maj} : X \times X \times X \to X$. We analyse several instances of this game with different motivations of players in order to illustrate the expressiveness of selection functions.

3.1 Classical game

In a classical game, the judges rank the contestants according to a preference ordering. For example, suppose judges 1 and 2 prefer A and judge 3 prefers B. Thus for each judge we have an order relation on X. The order relation of the first judge is $B <_1 A$, the second judge is $B <_2 A$ and the third is $A <_3 B$.

The judges now attempt to maximise the outcome with respect to their preferred ordering. Thus we obtain 3 different selection functions, which are maximisation with respect to each ordering. The three selection functions we obtain are

$$\begin{aligned}
\varepsilon_1(q) &= \{\arg \max_{x_1 \in (X, \leq_1)} q(\mathbf{x})\} \\
\varepsilon_2(q) &= \{\arg \max_{x_2 \in (X, \leq_2)} q(\mathbf{x})\} \\
\varepsilon_3(q) &= \{\arg \max_{x_3 \in (X, \leq_3)} q(\mathbf{x})\}.
\end{aligned}$$

In this particular example (but not in general) we can fix a 'global' order $B < A$ and notice that $\leq_3$ is the dual order. Thus for short we refer to ε_1 and ε_2 as $\arg\max$ and ε_3 as $\arg\min$.

The game is represented in table 1. Notice that Nash and context dependent equilibria coincide for this game, because it is classical.

Table 1: Agents: max, max, min

Strategy	Winner	NE	Defects	cd NE	Defects
AAA	A	X		X	
AAB	A	X		X	
ABA	A	-	J_3	-	J_3
ABB	B	-	J_2	-	J_2
BAA	A	-	J_3	-	J_3
BAB	B	-	J_1	-	J_1
BBA	B	-	$J_1,\ J_2$	-	$J_1,\ J_2$
BBB	B	X		X	

There is a subtle difference between this setup and the usual approach of classical game theory. In the classical approach, each judge's ordering would be seen as a preference relation. This would be used to derive utilities, which amounts to an order embedding of X into $\mathbb{R}$. Here there are no utilities: we directly maximise over the discrete order X.

We now want to give the calculations of the Nash equilibria of table 1 in the notation of selection functions and unilateral contexts. First we take a look at the Nash equilibrium BBB with outcome $q(BBB) = B$ and give the rationals of player 1. The unilateral context of player 1 is $\mathcal{U}_1^q(BBB)(x) = \mathrm{maj}(xBB) = \{B\}$ meaning that the outcome is (still) B if player 1 unilaterally changes from B to A. The minimal quantifier is $\overline{\varepsilon_1}(\mathcal{U}_1^q(BBB)(x)) = \mathrm{maj}(xBB) = \{B\}$ meaning that B is the outcome resulting from an optimal choice. Hence, we can conclude by $B = q(BBB) \in \overline{\varepsilon_1}(\mathcal{U}_1^q(BBB)(x)) = \{B\}$ that B is a Nash equilibrium strategy for player 1. This condition holds for each player and allows us to conclude that BBB is a Nash equilibrium.

In the same way we see in $B = q(BBA) \notin \overline{\varepsilon_1}(\mathcal{U}_1^q(BBA)(x)) = \overline{\varepsilon_1}(\{A\}) = \{A\}$ that BBA is not a Nash equilibrium.

3.2 Keynes beauty contest

We consider Keynes beauty contest as the paradigmatic example of a context dependent game. The first judge J_1 ranks the candidates according to a preference ordering $B < A$. The second and third judges, however, are 'Keynes agents': they have no preference relations over the candidates per se, but want to vote for the winning candidate.

The context dependent equilibria are precisely those in which J_2 and J_3 are coordinated, and J_1 is not pivotal in any of these. In the next section we will explain in more detail how the fixpoint selection function models coordination. Table 2 contains a summary of the equilibria.

Consider the strategy AAA, which is a context dependent equilibrium of this game. Suppose the moves of J_1 and J_2 are fixed, but J_3 may unilaterally change strategy. The unilateral context is

$$\mathcal{U}_3^{\mathrm{maj}}(AAA)(x) = \mathrm{maj}(AAx) = A$$

Thus the unilateral context is a constant function, and its set of fixpoints is

$$\mathrm{fix}(\mathcal{U}_3^{\mathrm{maj}}(AAA)) = \{A\}$$

This tells us that J_3 has no incentive to unilaterally change to the strategy B, because he will no longer be voting for the winner.

On the other hand, for the strategy ABB the two Keynes agents are indifferent, because if either of them unilaterally changes to A then A will become the majority and they will still be voting for the winner. This is still a context dependent equilibrium (as we would expect) because the unilateral context is the identity function, and in particular B is a fixpoint.

There are two context dependent equilibria, BAA and BBB, which are implausible in the sense that J_1 is not voting for his preferred candidate. We will discuss the issue of equilibrium selection in Section 3.5.

We now calculate the Nash and the context dependent rational for the strategy profile AAB of the Keynes player 3. The outcome of AAB is

Table 2: Agents: max, fix, fix

Strategy	Winner	NE	Defects	cd NE	Defects
AAA	A	X		X	
AAB	A	X		-	J_3
ABA	A	X		-	J_2
ABB	B	X		X	
BAA	A	X		X	
BAB	B	-	J_1	-	$J_1,\ J_2$
BBA	B	-	J_1	-	$J_1,\ J_3$
BBB	B	X		X	

$q(AAB) = A$. The unilateral context of player 3 is

$$\mathcal{U}_3^q(AAB)(x) = \text{maj}(AAx) = A$$

meaning that the outcome is (still) A if player 3 unilaterally changes from B to A. The minimal quantifier is

$$\overline{\varepsilon_3}(\mathcal{U}_3^q(AAB)(x)) = \text{fix}(\text{maj}(AAx)) = \{A\}$$

meaning that A is the outcome resulting from an optimal choice. Hence, we can conclude by

$$A = q(AAB) \in \overline{\varepsilon_3}(\mathcal{U}_3^q(AAB)(x)) = \{A\}$$

that B is a Nash equilibrium strategy for player 3.

The rational for the context dependent Nash equilibrium is as follows: the strategy $B = x_3 \notin \varepsilon_3(\mathcal{U}_3^q(AAB)(x)) = \text{fix}(\text{maj}(AAx)) = \{A\}$ meaning that AAB is not a context dependent Nash equilibrium.

In the classical approach a modeller usually writes down a payoff matrix directly from considerations about the situation as given by a story told in natural language. The modeller solves the game for the fixpoint outcomes and provides the players with the payoffs such that the maximizing behav-

Table 3: Agents: max, fix, fix

Strategy	Winner	NE sim. cd NE	Defects	P_{J_1}	P_{J_2}	P_{J_3}
AAA	A	X		1	1	1
AAB	A	-	J_3	1	1	0
ABA	A	-	J_2	1	0	1
ABB	B	X		0	1	1
BAA	A	X		1	1	1
BAB	B	-	J_1, J_2	0	0	1
BBA	B	-	J_1, J_3	0	1	0
BBB	B	X		0	1	1

ior mimicks the Keynes agent's goal and the classical pure Nash equilibria coincide with our context dependent equilibria.

An example of this classical approach is depicted in table 3 where the last three columns depicts the payoff matrix. In column 'NE sim. cd NE' we denote the Nash equilibria that simulate the context dependent Nash equilibria. As opposed to the classical approach we can equip players with the modeler's ability to reason about the game and to solve for any context dependent goal such as a fixpoint goal in the Keynes beauty contest example. The fixpoint selection function, combined with context dependent equilibria, perfectly model our intuition about this game.

Note that by solving the game the economist changes the payoffs such that the context dependency is internalized into the payoff matrix. Consider the strategies profiles AAA and AAB, here, working with selection functions the outcome does not change, meaning that both of them are Nash equilibria but only one of them is a context dependent Nash equilibrium. Creating the payoff matrix the modeller introduces differences in outcomes for the strategy profiles such that the third player strictly prefers to play A. In the selection function approach the context changes but results in the same (observable) outcome whereas in the classical approach the internalization of the context dependency changes the outcome in terms of utilities. Internalizing the context dependency makes AAA into a Nash equilibrium but not

AAB. Selection functions do not convey implicitly the context dependency but explicitly highlights it.

Compiling games In table 3 we have calculated the payoff matrix such that the Nash equilibrium mimicks the context dependent one. There is a very simple procedure for this: we write 0 whenever the player defects and 1 whenever they do not defect. The reason this is important is that we are unable to define mixed strategies inside the selection functions framework, but context dependent games could still be intuitively expected to have mixed equilibria. A good example of this is the beauty context variant $(\max, \min, \text{fix})$. In this game J_1 prefers A, J_2 prefers B and J_3 would like to vote for the winner. Intuitively we expect this game to have an equilibrium in which J_1 votes for A, J_2 votes for B and J_3 mixes with arbitrary probability. By 'compiling' the game to a classical utility function representation we regain these mixed equilibria.

Underlying this process there is a general pattern of compiling high level languages of high expressivity into languages with a lower expressivity. In general, any mathematical or logical language for knowledge representation [6] faces a trade off between the goals of representation and reasoning [2], page 327: "[...] why do we not attempt to define a formal knowledge representation language that is coextensive with a natural language like English? [...] Although such a highly expressive language would certainly be desirable from a *representation* standpoint, it leads to serious difficulties from a *reasoning* standpoint."

The difference between the classical and the context dependent approach is that both differ in their representational power. The natural language description of an economic situation is rather directly translatable into the high level formal system of selection functions as opposed to the low level classical approach. We have depicted both in figure 1. In our approach, the payoff matrix as in table 3, can be automatically computed, and the modeller needs only to decide which selection functions represent the agents and which outcome function represents the situation. The classical approach is depicted in figure 1 by a translation into the payoff matrices.

Figure 1: Context depended and classical modelling

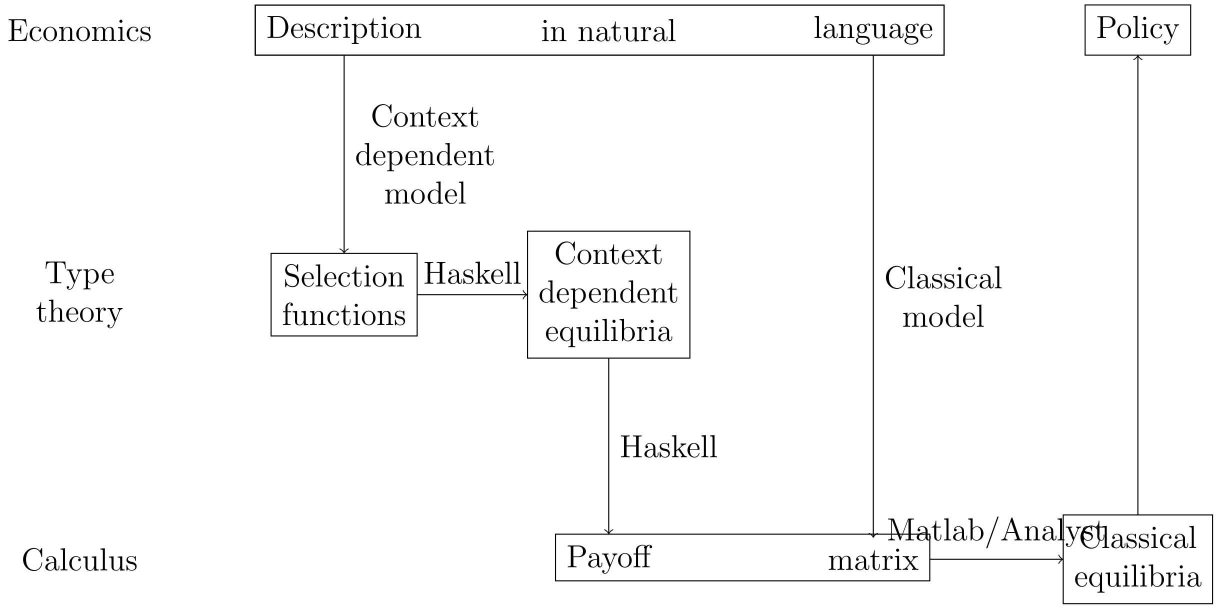

Our approach of context dependent modeling of the left part of figure 1 can thus be seen as a modelling technique to reduce the gap between the high level description of an economic situation in natural language and the formal modelling language. Selection functions are mid-level in between the high level natural language and the low level language of utility maximization represented traditionally in calculus in classical games.

An account for the representational power of languages is an involved research topic and even more raising the expressivity while not sacrificing reasonability. However, we hypothesis that our approach increases the expressivity of the representation language for game theory while not sacrificing reasoning possibilities, but quite the opposite, our framework is ready for automated reasoning, as opposed to calculus [12].

Our approach is directly programmable in modern functional languages like Haskell that has been developed within the high level modern mathematical type theory [13] in order to increase the expressivity of imperative languages such as Fortran or Matlab. The essence of functional languages is the usage of higher order functions that take and output other functions like selection functions, quantifier and outcome functions. Functional programming languages can be understood as languages to design languages. A typical approach to programing via functional languages is to design a domain specific language that allows to express the problems in a most direct and natural way while the program is compiled with its problem declaration into the low level solution steps of an intermediate imperative or a very low level machine language. This is much in the spirit of our approach to compile the selection function model into the payoff matrix of a classical representation of the Keynes beauty contest game in table 3.

We have heavily taken advantage of a prototype Haskell tool that calculates equilibria by brute force (enumeration of all strategies) for the games we have discussed in this paper. In fact the discovery of the notion of context dependent Nash equilibrium has been a direct consequence of using the tool. Before using the software we were misguided by our intuition, and did not recognise the difference between Nash equilibria and context dependent equilibria.

3.3 Coordination game

We consider a game where all agents act according to a fixpoint goal. Judges J_1, J_2 and J_3 want to vote for the winner, so the selection functions are given by the fixpoint operator $(R \to R) \to R$. For this particular case we do not need to consider multi-valued fixpoint operators, since for this game fixed points always exist and are unique.

As can be seen in table 4, the context dependent equilibria are exactly the coordinated strategies. Note that the fixpoint selection function models coordination, and the game in which all selection functions are fixpoints is a coordination game. This gives us a new perspective on the Keynes beauty contest as a one-sided coordination game: the Keynes agent would like to coordinate with the group, whereas the agents of the group are not interested in coordination.

This game is a good example of why ordinary Nash equilibria are not suit-

Table 4: Agents: fix, fix, fix

Strategy	Winner	NE	Defects	cd NE	Defects
AAA	A	X		X	
AAB	A	X		-	J_3
ABA	A	X		-	J_2
ABB	B	X		-	J_1
BAA	A	X		-	J_1
BAB	B	X		-	J_2
BBA	B	X		-	J_3
BBB	B	X		X	

able for modelling context dependent games: it can be seen in the table that every strategy is a Nash equilibrium of this game, but the context dependent equilibrium captures our intuition perfectly that the equilibria should be the strategy profiles that are maximally coordinated, namely AAA and BBB.

3.4 Anti-coordination game

Just as the fixpoint selection function models coordination, so there is a 'non-fixpoint' selection function which models *anti-coordination* (or *differentiation* as in the minority game [1, 11]).

The set of non-fixpoints of a function $p : X \to X$ is

$$\text{non-fix}(q) = \{x \in X \mid x \neq q(x)\}$$

We extend this to a selection function by specifying that the player is indifferent in the event that there are no non-fixpoints (such as if p is the identity function)

$$\varepsilon(q) = \begin{cases} \text{non-fix}(q) & \text{if non-fix}(q) \neq \varnothing \\ X & \text{otherwise} \end{cases}$$

Unlike for fixpoints, this selection function does not attain itself when considered as a quantifier. The corresponding quantifier is instead given by the

Table 5: Agents: non-fix, non-fix, non-fix

Strategy	Winner	NE	Defects	cd NE	Defects
AAA	A	X		-	J_1, J_2, J_3
AAB	A	X		X	
ABA	A	X		X	
ABB	B	X		X	
BAA	A	X		X	
BAB	B	X		X	
BBA	B	X		X	
BBB	B	X		-	J_1, J_2, J_3

set

$$\{q(x) \mid x \neq q(x)\}$$

An agent whose selection function is non-fix is a 'punk' who aims to be in a minority. We consider the game in which all three judges are punks (see Table 5). Of course only one can actually be in a minority, so the context dependent equilibria are precisely the 'maximally anti-coordinated' strategy profiles, namely those in which one judge differs from the other two. This is another example of a game in which every strategy is a Nash equilibrium, but the context dependent equilibrium corresponds perfectly to our intuition.

3.5 Strategic voting models

The classical game considered in Section 2.5, in which J_1 and J_2 maximise and J_3 minimises, contains two implausible equilibria, AAA and BBB. In the first, J_3 is voting against his preferred candidate because he is not pivotal. The second is stranger: both J_2 and J_3 would prefer B, but neither can unilaterally change the outcome.

One way to remove these equilibria is assuming that agents vote for the weakly dominant option. With the conception of a context dependent equilibrium there is a new perspective on this issue. Using a classical Nash-equilibrium implies that an agent only has an incentive to deviate when he is pivotal, that is, when he can change the global outcome. In fact this implicitly represents a form of context dependency but one which is not intended. The context is: vote for your preferred outcome if you are pivotal, choose any of all options if you are not pivotal. Assuming 'on top' that the agent does not vote for the dominated option, is then just a remedy to get around this context dependency while using the classical, context independent, Nash equilibrium concept to solve the game.

In order to account for this context dependency directly, an appropriate selection function for J_1 would be to vote for B iff voting for B is beneficial and voting for A is not beneficial. The selection function is

$$\varepsilon_A(q) = \begin{cases} B & \text{if } q(A) = B \text{ and } q(B) = A \\ A & \text{otherwise} \end{cases}$$

The selection function for B is

$$\varepsilon_B(q) = \begin{cases} A & \text{if } q(B) = A \text{ and } q(A) = B \\ B & \text{otherwise} \end{cases}$$

The selection function ε_A is a *selection refinement* of $\arg\max$, in the sense that it is consistent with $\arg\max$ but reduces the amount of indecision. In this particular example, this selection function leads to the same outcome as a constant selection function where an agent always votes for his preferred party indepedently of the context. The only difference would arise, if voting for the opponent resulted in the own preferred candidate winning. But in the context of voting games with majority and only two alternatives this cannot occur.

Formally, the condition for a selection function ε to be a selection refinement of δ is that $\varepsilon(q) \subseteq \delta(q)$ for every context q. If J_1 and J_2 use ε_A and J_3 uses ε_B then the resulting game has exactly one context dependent equilibrium, which is AAB. Even in such a simple situation as this, the context dependent equilibrium is important: there are two more Nash equilibria which are not context dependent equilibria (see Table 6).

When there are a small number of moves we can actually enumerate every

Table 6: Agents: ε_A, ε_A, ε_B

Strategy	Winner	NE	Defects	cd NE	Defects
AAA	A	X		-	J_3
AAB	A	X		X	
ABA	A	-	J_3	-	J_2, J_3
ABB	B	-	J_2	-	J_2
BAA	A	-	J_3	-	J_1, J_3
BAB	B	-	J_1	-	J_1
BBA	B	-	J_1, J_2	-	J_1, J_2, J_3
BBB	B	X		-	J_1, J_2

Table 7: Enumerating arg max, ε_A and constant

p (A)	p (B)	$\arg\max(p)$	$\varepsilon_A(p)$	always A
A	A	A, B	A	A
A	B	A	A	A
B	A	B	B	A
B	B	A, B	A	A

context to see this, as in table 7.

We can use selection functions to model more elaborate voting preferences. Suppose there are three political parties A, B, C standing in an election. A voter ranks the parties according to a preference relation $A > B > C$. Suppose A is a minority party for which there may be no chance of being elected. Parties B and C are mainstream parties, of which the voter would always prefer B. The voter votes according to the following rule: vote for A if there is any chance of A being elected, otherwise vote for B. In no situation will the voter vote for C. The context q of the decision can tell us whether voting for A will bring about A winning the election, and the selection function we are looking for is

$$\varepsilon(q) = \begin{cases} A & \text{if } q(A) = A \\ B & \text{otherwise} \end{cases}$$

This implicitly assumes that the outcome function is the majority function (or at least does not lead to strange contexts), but it can be argued that this condition is exactly what we mean by the term 'voting'.

4 Conclusions

In our complementary paper on decision theory [9] we have generalized the notions of max, argmax, utility functions and preferences. These classical, context independent concepts are superseded by the context dependent selection functions and quantifiers in order to model for example social preferences, conformity or differentiation goals.

In this paper we have extended the usage of context dependent concepts to finite, simultaneous stage games. We reproduce classical context independent games and go beyond to models of players with context dependent goals. In order to solve context dependent games we have refined the classical Nash equilibrium into the so called context dependent Nash equilibrium. The specific context dependent goals we have discussed are the fixpoint goals of coordination and differentiation within Keynes beauty contest and minority games.

This paper also contributes to the general modelling principle of economics in that we can reduce the gap between the high level description of economic situations in natural languages and their representation in formal models. As a consequence, we can automate a larger part of the modelling process and implement our models directly in modern functional programming languages that have been developed by the mathematical tools of type theory.

In future work we plan to extend the context dependent framework to repeated and sequential games with imperfect and incomplete information and to endogeneize the context itself towards a framework with learning players and for institutional dynamics. When extending the framework to infinite games such as repeated games there are computability issues to consider, and in particular both $\arg\max$ and fixpoint must be replaced with approximation algorithms. We see this as a positive aspect, since the type-theoretic foundation forces us to only discuss objects that are computable.

Moreover, the general idea of building the game theoretic formalism on computable functions implies that "reflexive" phenomena that are omnipresent in economic problems can be directly approached. As Thomas Sargent ([14]) has put it: "[...] the agents in the model should be able to forecast and profit-maximize and utility-maximize as well as the economist - or should we say the econometrician - who constructed the model." As the lambda calculus allows functions to manipulate themselves, it is the tool of choice to approach these questions. In our paper, we are making a first step in this direction as the agents who want to act according to the fixpoints of the game they find themselves in, just as we, the modellers, solve the games for their fixpoints. We will make further steps in this direction in future work.

References

[1] W. B. Arthur. Inductive reasoning and bounded rationality. *The American economic review*, 84(2):406411, 1994.

[2] Ronald Brachman and Hector Levesque. *Knowledge Representation and Reasoning*. Morgan Kaufmann, 1 edition, May 2004.

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

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

[5] 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.

[6] Frank Van Harmelen, Vladimir Lifschitz, and Bruce Porter. *Handbook of Knowledge Representation*. Elsevier New York, 2007.

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

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

[9] Evguenia Winschel Viktor Winschel Philipp Zahn Jules Hedges, Paulo Oliva. A new perspective on classical choice problems using selection functions. Technical report, 2014.

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

[11] Willemien Kets. The minority game: An economics perspective. *arXiv:0706.4432 [physics, q-fin]*, June 2007.

[12] D. Pavlovic and M. H. Escardo. Calculus in coinductive form. In *Logic in Computer Science, 1998. Proceedings. Thirteenth Annual IEEE Symposium on*, page 408417, 2002.

[13] Benjamin C Pierce. *Types and programming languages*. MIT press, 2002.

[14] Thomas J. Sargent. *Bounded Rationality in Macroeconomics*. Clarendon Press, November 1993.