

Towards
compositional game theory

Submitted in partial fulfilment of
the requirements of the degree of
Doctor of Philosophy,
Queen Mary University of London by

Julian Hedges

Jeni says hi

Abstract

We introduce a new foundation for game theory based on so-called open games. Unlike existing approaches open games are fully compositional: games are built using algebraic operations from standard components, such as players and outcome functions, with no fundamental distinction being made between the parts and the whole. Open games are intended to be applied at large scales where classical game theory becomes impractical to use, and this thesis therefore covers part of the theoretical foundation of a powerful new tool for economics and other subjects using game theory.

Formally we define a symmetric monoidal category whose morphisms are open games, which can therefore be combined either sequentially using categorical composition, or simultaneously using the monoidal product. Using this structure we can also graphically represent open games using string diagrams. We prove that the new definitions give the same results (both equilibria and off-equilibrium best responses) as classical game theory in several important special cases: normal form games with pure and mixed strategy Nash equilibria, and perfect information games with subgame perfect equilibria.

This thesis also includes work on higher order game theory, a related but simpler approach to game theory that uses higher order functions to model players. This has been extensively developed by Martin Escardó and Paulo Oliva for games of perfect information, and we extend it to normal form games. We show that this approach can be used to elegantly model coordination and differentiation goals of players. We also argue that a modification of the solution concept used by Escardó and Oliva is more appropriate for such applications.

Contents

0	Introduction	7
0.1	Acknowledgements	7
0.2	A note on style	7
0.3	Intended audience	8
0.4	Overview of the thesis	9
0.5	Publications	10
0.6	Background: game theory	10
0.7	Background: functional programming	13
0.8	Background: logic for social behaviour	15
0.9	On compositionality	16
0.10	The future of compositional game theory	18
0.11	Notation and conventions	21
1	Higher order game theory	23
1.1	Decision theory	23
1.1.1	Discussion	23
1.1.2	Quantifiers	24
1.1.3	The continuation monad	25
1.1.4	Selection functions	26
1.1.5	Attainment	26
1.1.6	Multi-valued variants	27
1.1.7	Multi-valued attainment	29
1.1.8	Modifying the outcome type	30
1.2	Normal form games	30
1.2.1	Discussion	30
1.2.2	Games, strategies and unilateral continuations	31
1.2.3	Nash equilibria of normal form games	32
1.2.4	Best responses	33
1.2.5	Classical games	34
1.2.6	Mixed strategies	36
1.2.7	Voting games	37
1.2.8	Modelling with selection functions	39
1.2.9	Coordination and differentiation	40
1.2.10	Illustrating the solution concepts	41
1.3	Sequential games	42
1.3.1	Discussion	42
1.3.2	The category of selection functions	43
1.3.3	The product of selection functions	44

1.3.4	Sequential games	45
1.3.5	Subgame perfection	45
1.3.6	Backward induction	46
1.3.7	The inductive step	48
2	The algebra and geometry of games	50
2.1	Open games	50
2.1.1	Discussion	50
2.1.2	The underlying model of computation	51
2.1.3	The category of stochastic relations	52
2.1.4	Open games	53
2.1.5	The best response function	55
2.1.6	Closed games	57
2.1.7	Decisions	57
2.1.8	Preliminary examples of decisions	58
2.1.9	Computations and counit	60
2.2	The category of games	61
2.2.1	Discussion	61
2.2.2	Categorical composition of games	61
2.2.3	Best response for sequential compositions	63
2.2.4	The identity laws	64
2.2.5	Associativity	66
2.2.6	Tensor product of games	68
2.2.7	The monoidal category axioms	69
2.2.8	Strategic triviality	71
2.2.9	Computations as a monoidal embedding	72
2.2.10	The counit law	74
2.2.11	Information flow in games	75
2.3	String diagrams	77
2.3.1	Discussion	77
2.3.2	String diagrams for monoidal categories	77
2.3.3	Compact closed categories	79
2.3.4	Boxing and compositionality	80
2.3.5	The geometry of games	81
2.3.6	Partial duality	82
2.3.7	Covariance, contravariance and symmetries	84
2.3.8	Copying and deleting information	85
2.3.9	A bimatrix game	86
2.3.10	A sequential game	87
2.3.11	Coordination and differentiation games	88
2.3.12	Designing for compositionality	89
3	Game theory in string diagrams	92
3.1	Normal form games in string diagrams	92
3.1.1	Discussion	92
3.1.2	Tensor products of decisions	93
3.1.3	Best response for a tensor of decisions	94
3.1.4	Best response for a tensor of decisions, continued	95
3.1.5	The payoff functions	96
3.1.6	Stochastic decisions	98

CONTENTS

3.1.7	Best response for a tensor of stochastic decisions	99
3.2	Extensive form	101
3.2.1	Discussion	101
3.2.2	Composition with perfect information	101
3.2.3	Building the composition	103
3.2.4	Best response for the composition	105
3.2.5	Best response for the composition, continued	106
3.2.6	Information sets	108
3.2.7	Imperfect information in string diagrams	110
3.3	Solvable games	111
3.3.1	Discussion	111
3.3.2	The definition of solvability	112
3.3.3	Solvable decisions	113
3.3.4	Backward induction for open games	114
3.3.5	Selection functions and open games	115
3.3.6	Tensor does not preserve solvability	116
3.3.7	Failure of compositional Nash's theorem	117
	Bibliography	120

Chapter 0

Introduction

0.1 Acknowledgements

Top billing goes to my family (including Jeni, who says hi); my supervisor Paulo Oliva; my second supervisor Edmund Robinson; my closest collaborators Viktor Winschel and Philipp Zahn; my other coauthors Neil Ghani, Mehrnoosh Sadrzadeh and Evguenia Spritz; and my other colleagues in Theory Group at QMUL, especially my fellow PhD students in CS420 that was.

Over the past three and a half years I have had gained from talking with many people. The following list is probably incomplete: Alessandro Abate, Samson Abramsky, Valeria de Paiva, Martin Escardó, Jeremy Gibbons, Helle Hvid Hansen, Peter Hines, Aleks Kissinger, Clemens Kupke, Alexander Kurz, Pierre Lescanne, Adrian Mathias, Arno Pauly, Dusko Pavlović, Alex Smith.

I gratefully acknowledge that my PhD studies were funded by EPSRC doctoral training grant EP/K50290X/1.

0.2 A note on style

To begin with, in this introduction I will refer to myself in the first person. Once the thesis proper begins in [chapter 1](#) I will return to using the third person.

I feel very fortunate to be writing a thesis in my native language, and I plan to take advantage of it. The writing throughout is intentionally slightly less formal¹ than would be reasonable in a publication, and in this introduction much less so.

When I read ([Atiyah 2008](#)) I was overly influenced by point 5 under ‘style’, namely “Identify papers you have enjoyed reading and imitate their style”, even to the expense of the previous point, “Be as clear and succinct as possible while being clear and easy to understand”. I immediately thought of ([Girard 2001](#)) and the book ([Hoyte 2008](#)), which have both strongly influenced the way I think about logic and programming respectively; both contain a mixture of formal mathematics or computer science with vivid intuitions and downright aggressive personal opinion that borders on philosophy. I have not carried my style to nearly that extreme, but I hope that some of it is visible.

¹For example, there are many footnotes.

My opinion is that the definition-theorem-proof style of mathematics inherited from Bourbaki will soon (but not quite yet) belong in a past age when theorem provers were not practical to use, and that in the future the style of a typical publication in mathematics will need to change to account for the fact that human-checked proofs are unacceptably unreliable² compared with machine-checked proofs. I hope that the outcome is that style will become more informal and focus more on intuition, in contrast to what seems to be happening now (for example, in the homotopy type theory community) with papers written in the ugly syntax of literate proof assistant scripts.

I have intentionally written in continuous prose, rather than dividing into sections labelled ‘definition’, ‘theorem’ and ‘proof’, to reflect the way that mathematics is actually done: plausible definitions, theorems and proof ideas are used to adjust each other until a fixpoint is reached. For example, the proof idea may come first, followed by a definition encapsulating the hypotheses found to be needed to make the proof go through, with the theorem coming last. The proofs in this thesis are indeed checked by hand, although I trust the experimental evidence discussed in §0.7 more than I trust my ability to write correct proofs. Since in some cases several variant or false definitions are given, the ‘official’ one will always be distinguished by a bold font.

0.3 Intended audience

It is very strange to write about ‘intended audience’ in a thesis, when the rule of thumb is that it will be read by (at most) my supervisor and examiners. However I intend to continue using my thesis as a reference on compositional game theory even after papers are published, in particular because chapter 2 contains far more informal text than would be reasonable in a publication, and I think the informal text is very important. Therefore, I will write here about what background knowledge is assumed.

This thesis contains nontrivial amounts of both pure mathematics (by which I mean the study of mathematical objects for no other reason than their inherent beauty) and applied mathematics³ (by which I mean mathematics motivated and influenced by modelling problems). An idealised reader has some background knowledge in both game theory and category theory, but I expect that most readers will be familiar with one, but not the other.

The category theory required to read this thesis is mostly monoidal categories, for which the usual reference is (Mac Lane 1978). Alternatively, a self-contained introduction to monoidal category theory that emphasises the process-oriented view⁴ used in this thesis can be found in (Coecke 2006). Cat-

²The use of uncertified mathematics in economics, for example, should soon come to be seen as reckless behaviour. See the editor’s report on (Lescanne and Perrinel 2010), and Lescanne’s response to it, both currently available from Lescanne’s website, for an interesting anecdote on the response to certified proofs in academic economics.

³My personal view is that beauty can be *unreasonably effective* in applications. For a good example of this, compare purely functional programming to some false start such as object oriented programming. More vividly, there is no reason why the ‘rising sea’ approach might not be applicable to very difficult problems in applied mathematics, for example with better models of human behaviour arising most easily as a special case of some very abstract theory. Monoidal bicategories, mentioned in §2.2.1, certainly fit that bill in terms of abstraction.

⁴The view of arrows of a monoidal category as processes is almost synonymous with the

egory theory in this thesis is largely treated as a means to an end, as an axiomatic approach to compositionality and a way to easily prove the soundness of the string diagram language in §2.3. Readers who are category theorists will be able to tell that I am not a category theorist: in particular, no attempt is made to abstractly study the properties of the category of games, the reasoning being very concrete and often by example.

The game theory that is needed as background is also very small, and is discussed in §0.6. A list of topics reads like half of an undergraduate course in game theory: normal form, extensive form, pure and mixed strategies, Nash equilibria, subgame perfect equilibria. It is more important to have an intuition for game theory than any specific piece of mathematical theory, for which a good introduction is (Kreps 1991).

0.4 Overview of the thesis

The thesis is divided into three large chapters, each of which is divided into three sections, each of which is divided into many subsections. Each section begins with a ‘discussion’ subsection that gives motivation and background. The serious part of the thesis consists of chapter 2 and chapter 3, with chapter 1 as a sort of extended introduction.

The subject of this thesis is two new approaches to game theory, which can be called the ‘higher order’ approach and the ‘compositional’ approach. Higher order game theory is chronologically prior and much simpler, and can serve as an introduction to the modes of thinking needed for compositional game theory, which is much more complicated and unfamiliar. In principle it should be possible to begin reading at chapter 2 and locally follow hyperlinks back into chapter 1 when necessary to refer to definitions and notations. I do not recommend this, however.

The main objects of study in higher order game theory are so-called quantifiers and selection functions, which are introduced and studied in isolation in §1.1. Simultaneous or normal form higher order games are studied in §1.2, and sequential or perfect information games in §1.3. The contents of §1.2 is closely based on (Hedges et al. 2015b), and §1.3 is closed based on (Escardó and Oliva 2011) (and therefore is not my own work), with §1.1 being a mixture of the two.

Open games, the objects of study in compositional game theory, are introduced in §2.1. This is a section heavy on definitions, introducing the definitions of open games, decisions, computations and counits. Sequential and parallel composition of open games are studied in §2.2, including some important theorems about how composition behaves. This is applied in §2.3 to give a string diagram language for specifying games. The whole of chapter 2 is based on the preprint (Ghani and Hedges 2016).

The purpose of chapter 3 is to formally connect open games to standard game theory, which previously is done only informally. In §3.1 normal form games, both with pure and mixed strategies, are shown to be a special case of open games. The same is done for perfect information games with pure strategies in §3.2, together with some sketched ideas for extending to imperfect information.

quantum computation group in Oxford, and I learned it from a steady stream of seminar speakers coming from Oxford to Queen Mary. It is one of two ways that I think about category theory, the other being functional programming (see §0.7).

Finally, §3.3 returns to more theoretical considerations by exploring a possible solution concept for arbitrary open games, which gives a connection between compositional game theory and higher order game theory.

0.5 Publications

At the time of submission I have three publications: (Hedges 2013), (Hedges 2014) and (Hedges 2015a). Of these, the last is on a different topic, and this thesis contains little material from the first two.⁵

On the other hand, this thesis does contain large amounts of material from the preprints (Hedges et al. 2015b) and (Ghani and Hedges 2016). A third preprint (Hedges and Sadrzadeh 2016) is not related to this topic. Two further preprints, (Hedges 2015b) and (Hedges et al. 2015a), are not currently under review and have not been kept up to date.

0.6 Background: game theory

There is a tendency for theoretical computer scientists, when writing about game theory, to refer mostly to the oldest references on the subject, such as (von Neumann and Morgenstern 1944).⁶ To a computer scientist, the term ‘game’ may mean⁷ ‘normal form game’, or it may mean ‘extensive form game’, in the latter case often with information sets quietly ignored. In particular, though, the computer scientist will ignore the fact that game theory is a large research areas within economics, which itself is a subject of comparable size to all of computer science. Another mistake that a computer scientist can make, perhaps even simultaneously, is to equate (academic) economics with game theory. These are both errors that I am still in the process of trying to overcome myself.

With that being said, essentially all of the game theory needed to follow this thesis can be found in (von Neumann and Morgenstern 1944). For a more concise introduction written by (and therefore readable by) computer scientists⁸ I recommend (Leyton-Brown and Shoham 2008). To computer scientists I would also recommend (Kreps 1991) which, being short on mathematics and long on economic intuition, is likely to be very different to the way they think about the subject. Of the various weighty reference books on game theory, the one I use is (Fudenberg and Tirole 1991). Failing that, of course there are endless lecture notes and slides online written for undergraduates in economics, computer science, mathematics, engineering, ...

For the closely related two types of game theory covered in this thesis, namely higher order and compositional, I would like to be clear about how they relate to what I will call classical game theory, by which I mean game theory as covered by these references. The questions are: which features are common? What new

⁵Both of these papers predate when I understood monad algebras. Reworking them to follow (Escardó and Oliva 2015) is work in progress.

⁶To be fair, when economists write about computability, they also tend to write as though the subject began and ended with Turing. Of course these are sweeping generalisations, and I don’t mean to insult members of either subject: on the contrary, I believe that each has a lot to offer and learn from the other.

⁷I am ignoring game semantics here (see §0.8).

⁸Of course, von Neumann also worked in computer science, among many subjects.

features are gained? What features are lost? Which problems are solved, and which are not?

I will begin with the features common to all approaches. A game consists of players or agents, who act in a way that is constrained by some rules or protocol. By this I mean that when each player moves, they have a collection of possible moves, and a collection of possible things they could observe about the past. A strategy for each player is a mapping from possible observations to possible moves, possibly allowing certain side effects such as probabilistic choices or belief updating. Then a ‘solution’ of the game is an equilibrium, which is a choice of strategy for each player that is stable or non-self-refuting, or equivalently is a fixpoint of a best response function. In general a game may have zero, one or many equilibria. In justifying the solution as a prediction of real world behaviour we assume that the rules of the game and the perfect rationality of the players are common knowledge⁹.

A built-in feature of classical game theory is that a choice of strategy for each player will determine a real number for each player called a utility, and the perfect rationality of the players is defined to mean that they act such as to maximise their utility. This is discussed further in §1.2.5. Both higher order and compositional game theory generalise away from this, replacing real numbers with arbitrary objects and allowing rationality to be defined in far more general ways which become part of the specification of a game. I will offer three arguments in favour of doing this. The first is that by abstracting away from a nontrivial but inessential feature, namely real analysis and optimisation, the theory is genuinely simplified, and the significant issues become clearer. The second is that new modelling techniques become available, such as the coordination and differentiation games in §1.2.9 and §2.3.11. The third is that this is a necessary step for compositionality: taking two players who perfectly maximise and composing them together will typically produce a system that does not perfectly maximise.¹⁰

Since this is in some sense a strict generalisation, we can still revert back to the special case by taking outcomes to be real numbers and considering only agents who maximise real numbers. Indeed, §3.1 and §3.2 of this thesis do exactly that. But what is lost is any piece of theory that begins by assuming that outcomes are utilities and that players maximise. It may even be that we lose the vast majority of all of the literature on game theory this way. To give perhaps the most serious but elementary example, it is impossible to define strategic dominance in general for higher order or compositional games.

Higher order game theory, as a subject which naturally grew out of applications in proof theory, is not really intended to solve any problem in game theory. One feature that stands out, however, is the ability to write very short and elegant functional programs that compute subgame perfect equilibria of perfect information games, including certain infinite games (Escardó and Oliva 2010b; 2012). The coordination and differentiation games of §1.2.9 and (Hedges et al. 2015b), in addition, constitute an ‘application’ of higher order game theory.

Compositional game theory, on the other hand, has been consciously de-

⁹The phrase “it is common knowledge that X ” means that “all players know that X ” together with “all players know that all players know that X ”, and so on.

¹⁰It is tempting to link this issue with macroeconomics, for example with recessions in the real economy, but I don’t think that is reasonable.

signed as an attack on a specific problem: compositionality. This is the principle that a system should be built by composing together smaller subsystems, and that the behaviour of a system is entirely determined by the behaviour of the subsystems and the ways in which they are composed, and therefore it is possible to reason about the system by structural induction on its decomposition. To a computer scientist compositionality is such a fundamental idea that it is most often not even mentioned. For example, every serious programming language is compositional: a program is built from code blocks composed using sequencing, loops, functions and so on, and the program's behaviour is determined by the behaviour of the code blocks together with the constructs used to compose them. I will discuss the principle of compositionality in considerable detail in §0.9.

Compositionality, however, is an alien concept in game theory, because there is no meaningful formal sense in which a game is built from composing together smaller components. Put bluntly, this is why it is feasible to create a reasonably robust software system containing millions of lines of code, but it is not feasible to work with game theoretic models of comparable scale and complexity. I am not aware of any literature that has come close to identifying the lack of compositionality as a problem in game theory, but nevertheless it is a very serious problem, and it is the problem that is solved in this thesis.

In order to talk systematically about which problems are not solved, I will refer to chapter 5 of (Kreps 1991) for a discussion of the problems of game theory. With the exception of the section titled 'What are a player's payoffs?', which relates to the generalised rationality described above, neither higher order nor compositional game theory contributes anything to these well-known problems. I will divide these into two classes: the problems relating to 'the rules of the game', and the problems relating to equilibrium analysis.

Problems with the 'rules of the game' include:

- A game-theoretic model must have fixed rules, and game theory is unable to model unrestricted negotiation, for example;
- Similarly, it is difficult to model the ability of players to dynamically modify the rules;
- The predictions of the model can be extremely sensitive to apparently small changes in the rules;
- The structure of the game, by default, is common knowledge.

Regarding the last point, in higher order game theory each player's quantifier or selection function is assumed to be common knowledge, and in compositional game theory the entire structure of a string diagram is assumed to be common knowledge. I hope that the usual technique of using Bayesian games and universal type spaces can be generalised to compositional game theory, but that is entirely work for the future.

The central problem of equilibrium analysis is that there is no generally accepted mechanism by which a particular equilibrium can be selected by the players, in some cases even when there is exactly one equilibrium. Both higher order and compositional game theory are fundamentally equilibrium-based, and suffer from the same, familiar problems with equilibrium analysis, and I will say nothing more about it.

0.7 Background: functional programming

Compositional game theory was almost entirely developed by me during an intense two weeks in February 2015 (when I was living in Mannheim), in which time I felt like the Haskell interpreter `ghci` became a sort of extension of my brain. I already had working implementations¹¹ of the definitions in §2.1 and §2.2, and experimental verification of the results in §3.1 and §3.2, and even more that is not covered in this thesis, before I had even written down the definitions in mathematical language, let alone proved any theorems. The ability to rapidly typecheck and experimentally test many different definitions was crucial in eventually arriving at the definitions that work, and I still genuinely struggle to understand the resulting definitions (particularly those in §2.1) intuitively.

Although the functional programming point of view was built into compositional game theory from the start, I have tried my best to minimise it in the presentation in this thesis, because the intersection of the intended audience (see §0.3) with functional programmers might vanish. In particular, monads have mostly been replaced with their Kleisli categories. However, here I will describe my intuitions for the use of functional programming in game theory.

The fundamental intuition I use is **Moggi’s thesis** (Moggi 1991), which can be paraphrased as follows: there is a correspondence between

1. Computations that input a value of type X , possibly carry out side effects, and output a value of type Y
2. Functions of type $X \rightarrow TY$ for a suitable monad T
3. Morphisms in $\text{hom}_{\mathcal{C}}(X, Y)$ for a suitable category \mathcal{C}

The passage from 2 to 3 is to take the Kleisli category $\mathcal{C} = \text{Kl}_T$. The passage from 1 to 2, which involves choosing a suitable monad, is part of the art of functional programming. See also (Plotkin and Power 2002).

I will give one prototypical example, which is nondeterminism. Nondeterministic choice is the ability of a program to return a result that is not uniquely determined by its input. Instead, for each input the program has a set of outputs that might possibly occur. If this set is empty then the program can never return a result, which can be interpreted either as nontermination or as exceptional termination. The corresponding monad is the powerset monad \mathcal{P} , and functions of type $X \rightarrow \mathcal{P}Y$ form a useful model of nondeterministic programs. The category that corresponds to this is $\text{Kl}_{\mathcal{P}} = \mathbf{Rel}$, the category of sets and relations, via forward images of relations.

An important idea which is foundational to my research is that side effects, or equivalently monads, are ubiquitous in game theory, and that identifying and

¹¹I will make the Haskell code available in the future, after it has been tidied up and rewritten in the light of more recent ideas. By far the biggest difficulty with my prototype is that Haskell’s type system does not unify naturally isomorphic types, and so the type X might appear in an isomorphic form such as

$$((1 \times 1) \times (1 \rightarrow (1 \times 1) \times 1)) \times (1 \times 1 \rightarrow X \times 1)$$

with the element $x : X$ written isomorphically as

$$(((*, *), \lambda *. ((*, *), *)), \lambda (*, *) . (x, *))$$

and the isomorphisms must be written and carried around manually.

classifying them is a useful thing to do. The most important examples are the selection and continuation monads, but they are not intuitive and will be left for [chapter 1](#). More intuitively, consider a player who makes an observation¹² from a type X and then makes a choice from a type Y .

A pure strategy for the player is a function $X \rightarrow \text{Id } Y$, where Id is the identity monad $\text{Id } Y = Y$, whose Kleisli category is **Set**. A mixed strategy is a function $X \rightarrow \mathcal{D}Y$, where $\mathcal{D}Y$ is the set of probability distributions on Y . The monad \mathcal{D} is called the probability distribution or (finitary) Giry monad ([Giry 1982](#)), and it is introduced in [§2.1.3](#), along with its Kleisli category **SRel**. The Haskell implementation of probability I use is described in ([Erwig and Kollmansberger 2006](#)).

From this viewpoint it is natural to also consider players who make truly nondeterministic choices, without a probability distribution. Such a player has a set of possible choices for each possible observation, and her¹³ strategies have type $X \rightarrow \mathcal{P}Y$, and are morphisms of **Rel**. Several computer scientists writing about game theory have independently had the idea of nondeterministic players ([LaValle 2006](#), [Pavlovic 2009](#), [Hedges 2014](#)), although there is little significant theory.¹⁴ In Haskell the most common implementation of nondeterminism uses lists¹⁵, although there are many alternatives.

If the player makes use of a prior of type A then her strategy has type $X \rightarrow \text{Rd}_A Y$, where Rd is the reader monad, which acts on sets by $\text{Rd}_A Y = A \rightarrow Y$. If her strategy moreover has the ability to update the prior with a posterior after making the observation then it has type $X \rightarrow \text{St}_A Y$, where the state monad St is $\text{St}_A Y = A \rightarrow Y \times A$.

This Bayesian updating or learning is more complicated than the other examples, because St_A is the only example listed that is a noncommutative monad. In this thesis only commutative monads will be considered, for simplicity. By corollary 4.3 of ([Power and Robinson 1993](#)), a monad is commutative iff its Kleisli category is symmetric monoidal, which justifies the choice in [§2.1.2](#) to parameterise the definition of open games by an arbitrary symmetric monoidal category. More general premonoidal categories also destroy the connection with string diagrams, although see ([Jeffrey 1997](#)) and discussion in [§0.10](#).

This point of view is shared with ([Pavlovic 2009](#)), which moreover refers to Freyd categories ([Power and Thielecke 1999](#)). The heavier machinery of Freyd categories is avoided in this thesis by using weak products (see [§2.1.2](#)), which is stronger than necessary but is satisfied by the most important examples.

¹²For readers who are game theorists, X is the set of information sets owned by the player, and the strategies we consider are behavioural strategies. The functional programming viewpoint makes behavioural strategies $X \rightarrow \mathcal{D}Y$ far more natural to consider than mixed strategies $\mathcal{D}(X \rightarrow Y)$, which means that the assumption of perfect recall is essential. Throughout this thesis, the term ‘mixed strategy’ really means ‘behavioural strategy’ in the context of dynamic games.

¹³See [§0.11](#).

¹⁴Nondeterminism is related to extreme-case optimisation, in contrast to average-case optimisation, in ([LaValle 2006](#)) and ([Hedges 2014](#)), and fixpoints of monotone functions are used in ([Pavlovic 2009](#)).

¹⁵The covariant powerset functor $\mathcal{P} = (\rightarrow \mathbb{B})$ is not a monad (or even a functor) constructively.

0.8 Background: logic for social behaviour

In this section I will give a brief literature review of applications of logic and theoretical computer science to game theory. For lack of a better name, I will call this research topic ‘logic for social behaviour’ after workshops held in Leiden in 2014, Delft in 2015 and Zürich in 2016. Another event that should be mentioned in this context is the 2015 Dagstuhl workshop ‘Coalgebraic semantics of reflexive economics’ (Abramsky et al. 2015). Many researchers in this area also consider applications to social choice theory, especially preference aggregation, but I will mention only (Abramsky 2015), which links Arrow’s famous impossibility theorem with category theory.

A starting point is (Pavlovic 2009), which proposes to study game theory using ideas from program semantics, in particular viewing games as processes which can have side effects such as state and probabilistic choices. That paper suggests a larger research programme called ‘abstract game theory’ in which this thesis can be located, although see §2.1.1.

An approach to infinitely repeated games using coinduction was introduced in (Lescanne and Perrinel 2012) and continued in (Abramsky and Winschel 2015) and the working paper¹⁶ (Blumensath and Winschel 2013). Infinite and coalgebraic games are mentioned only briefly in this thesis, in §2.2.1. However, given that coinduction and bisimulation are the *correct* techniques for reasoning about infinite processes, it is likely that they will continue to be important in game theory. In particular, if trying naïvely to verify that some strategy of an infinite game is an equilibrium, then infinitely many properties must be checked; however a finitary proof technique based on bisimulation should be expected to work. As yet, coalgebraic game theory has not been connected with the classical approach to repeated games using real analysis, as covered for example in (Fudenberg and Tirole 1991), and the very extensive literature on repeated games. An unrelated application of coalgebra to game theory is (Moss and Iglizzo 2004), which shows that universal Harsanyi type spaces are also final coalgebras.

Another line of work begins with (le Roux 2014), connecting two-valued games and determinacy theorems with real-valued games and existence theorems.¹⁷ In (le Roux and Pauly 2014) this is moreover connected with synthetic topology (Escardó 2004), which is related to ideas I am working on involving computably compact sets of probability distributions in game theory, which are not in the scope of this thesis.

Practical experience of applying functional programming techniques to economic modelling is described in (Botta et al. 2011), (Ionescu and Jansson 2013) and (Boudes et al. 2015). More generally, (Erwig and Kollmansberger 2006) describes the application of functional programming to mathematical modelling in biology. I directly quote the last sentence of that paper: “In particular, the high-level abstractions allowed us to quickly change model aspects, in many

¹⁶In its current form that paper ends with a 10-page essay which amounts to a manifesto for new approaches to game theory.

¹⁷This connects the two cultures that began with the founding work of Zermelo on winning strategies (Schwalbe and Walker 2001) and immediately diverged. One contains Borel games and determinacy in set theory, dialogical semantics in logic and game semantics in computer science. The other is game theory, which is otherwise related to these topics only in trivial ways.

cases immediately during discussions with biologists about the model.” Due to the close connections between game theory and computational effects described in §0.7, I expect the gains of using functional programming to increase in the future.

Finally, algorithmic game theory (Nisan et al. 2007) is a large topic that studies the computational complexity of Nash equilibria and other constructions in game theory, which began with the result in (Daskalakis et al. 2006) that computing approximate Nash equilibria is infeasible. Algorithmic game theory can be contrasted with the semantic approach to game theory that this thesis represents, but will likely be an essential ingredient in the research project outlined in §0.10.

0.9 On compositionality

This section is essentially an essay, loosely based on a talk I gave at Logic for Social Behaviour 2016 in Zürich, after the vast majority of the thesis was written.

The term *compositionality* is commonplace in computer science, but is not well-known in other subjects. Compositionality was defined in §0.6 as the principle that a system should be designed by composing together smaller subsystems, and reasoning about the system should be done recursively on its structure. When I thought more deeply, however, I realised that there is more to this principle than first meets the eye, and even a computer scientist may not be aware of its nuances.

It is worthwhile to spend some time thinking about various natural and artificial systems, and the extent to which they are compositional. To begin with, it is well-known that most programming languages are compositional. The behaviour of atomic¹⁸ statements in an imperative language, such as variable assignments and IO actions, is understood. Functions are written by combining atomic statements using constructs such as sequencing (the ‘semicolon’ in C-like syntax), conditionals and loops, and the behaviour of the whole is understood in terms of the behaviour of the parts together with the ways in which they are combined. This scales sufficiently well that a team of programmers can broadly understand the behaviour of a program consisting of hundreds of millions of individual atomic statements.

When the software industry began software was unstructured, with no intermediate concepts between atomic statements and the entire program, and much of its history has been the creation of finer intermediate concepts: code blocks, functions, classes, modules. Compositionality is not all-or-nothing, but is slowly increased over time; nor is it entirely well-defined, with many tradeoffs and heated debates in the design and use of different language features. Even with a modern well-designed language it is possible to write bad code which cannot be easily decomposed; and even though there are many design patterns and best practice guidelines, good software design is ultimately an art.

Going beyond software, consider a physical system designed by human engineers, such as an oil refinery. An individual component, such as a pump or a section of pipe, may have a huge amount of engineering built into it, with

¹⁸The term ‘atomic’ is used naively here, and does not refer to concurrency.

detailed knowledge of its behaviour in a wide variety of physical situations. It is then possible to connect these components together and reuse knowledge about the components to reason about the whole system. As in software, each component has an ‘interface’, which is a high level understanding of its behaviour, with unnecessary details being intentionally forgotten.

As a third example, an organisation made of human beings, such as a company or university, is also built in a compositional way, demonstrating that engineering is not a requirement. It is possible to understand the behaviour of a department without knowing the details of how the behaviour is implemented internally. For example, a software engineer can use a computer without knowing the exact process through which the electricity bill is paid, and will probably not even be aware if the electricity provider changes. This is another example of reasoning via an interface.

Clearly interfaces are a crucial aspect of compositionality, and I suspect that interfaces are in fact synonymous with compositionality. That is, compositionality is not just the ability to compose objects, but the ability to work with an object *after intentionally forgetting how it was built*. The part that is remembered is the ‘interface’, which may be a type, or a contract, or some other high-level description. The crucial property of interfaces is that their complexity stays roughly constant as systems get larger. In software, for example, an interface can be used without knowing whether it represents an atomic object, or a module containing millions of lines of code whose implementation is distributed over a large physical network.

For examples of non-compositional systems, we look to nature. Generally speaking, the reductionist methodology of science has difficulty with biology, where an understanding of one scale often does not translate to an understanding on a larger scale. For example, the behaviour of neurons is well-understood, but groups of neurons are not. Similarly in genetics, individual genes can interact in complex ways that block understanding of genomes at a larger scale.

Such behaviour is not confined to biology, though. It is also present in economics: two well-understood markets can interact in complex and unexpected ways. Consider a simple but already important example from game theory. The behaviour of an individual player is fully understood: they choose in a way that maximises their utility. Put two such players together, however, and there are already problems with equilibrium selection, where the actual physical behaviour of the system is very hard to predict.

More generally, I claim that *the opposite of compositionality is emergent effects*. The common definition of emergence is a system being ‘more than the sum of its parts’, and so it is easy to see that such a system cannot be understood only in terms of its parts, i.e. it is not compositional. Moreover I claim that *non-compositionality is a barrier to scientific understanding*, because it breaks the reductionist methodology of always dividing a system into smaller components and translating explanations into lower levels.

More specifically, I claim that compositionality is strictly necessary for working at scale. In a non-compositional setting, a technique for solving a problem may be of *no use whatsoever* for solving the problem one order of magnitude larger. To demonstrate that this worst case scenario can actually happen, consider the theory of differential equations: a technique that is known to be effective for some class of equations will usually be of no use for equations removed from that class by even a small modification. In some sense, differential equa-

tions is the ultimate non-compositional theory.

Of course emergent phenomena do exist, and so the challenge is not to avoid them but to *control* them. In some cases, such as differential equations, this is simply impossible due to the nature of what is being studied. The purpose of this thesis is to demonstrate that it is possible to control emergent effects in game theory, although it is far from obvious how to do it. A powerful strategy that is used in this thesis is *continuation passing style*, in which we expand our model of an object to include not only its behaviour in isolation, but also its behaviour in the presence of arbitrary environments. Thus an emergent behaviour of a compound system was already present in the behaviour of each individual component, when specialised to an environment that contains the other components.

As a final thought, I claim that *compositionality is extremely delicate*, and that it is so powerful that it is worth going to extreme lengths to achieve it. In programming languages, compositionality is reduced by such plausible-looking language features as goto statements, mutable global state, inheritance in object-oriented programming, and type classes in Haskell. The demands placed on game theory are extremely strong¹⁹: seeing a game as something fundamentally different to a component of a game such as a player or outcome function breaks compositionality; so does seeing a player as something fundamentally different to an aggregate of players; so does seeing a player as something fundamentally different to an outcome function. This thesis introduces *open games*, which include all of these as special cases.

0.10 The future of compositional game theory

This section is partly based on a funding application I wrote, and will sketch a research programme about the theory and applications of composing games. I will begin with an unreasonably optimistic best-case vision of the future, and work backwards. This work of speculative fiction is based on long discussions with Viktor Winschel and Philipp Zahn, mostly in the pub.

Suppose an economist wants to create a mathematical model of some economic system, say, a new market. Using her experience as a working economist, she analyses the market and divides it into a number of interacting components. Most of these component markets are well known and have been extensively studied, and she recognises the last component as behaving similarly to part of another market that was recently studied by a colleague.

She opens her software suite and loads these known models, most of which come pre-installed, and the last downloaded from a source repository accompanying her colleague's research paper. Each of these specifies an open game, whose ports specify how it communicates with an arbitrary environment, which is graphically represented on the screen by a bead with open strings. Although she has never heard the phrase 'category theory' except in a cryptic footnote in the software's documentation, she draws strings connecting the beads, in a way corresponding to her intuition about how the component markets are communicating and influencing each other.

¹⁹Of course, these demands have been written retrospectively.

Behind the scenes, the software compiles the string diagram to an intermediate representation, and then to Haskell, which is compiled and optimised by `ghc`.

The economist begins by specifying a plausible strategy profile, for example by drawing a finite state automaton, and then plays the game with the strategy profile to obtain a plot showing a probability distribution on outcomes. She tests this against real econometric data and finds that it does not fit the data. She then replaces some parameter in the strategy profile by a variable, and runs an automated optimisation procedure. This takes 30 minutes to complete, and still does not fit the data well.

By tracing the results through the string diagram, viewing plots of the data flowing along the intermediate strings, she determines that her colleague's model is at fault. She 'zooms in' to that model in the editor, seeing visually the information flow inside that model. It consists of many standard components interacting, each of which itself can be opened as a string diagram for editing. She manually tunes several quantitative parameters, and qualitatively changes the logic of the information flow in a few places, in an iterative process, at each stage comparing the simulation results to the econometric data.

Finally, she obtains a model and a strategy profile that is a good fit for the data. The next day, a three-hour computation confirms that the strategy profile is an equilibrium, suggesting that the model is stable. From here, she can answer questions of economic interest. For example, she can simulate economic events such as changing the structure of the market in some way, or changing the response of a neighbouring market. She can test the existing strategy profile with the modified game, and the software will determine which agents now have incentive to deviate. In some cases, she is able, through a mixture of intuition and automatic optimisation procedures, to find a new equilibrium to which the old strategy profile will plausibly settle given the changes made to the game. This constitutes an economic prediction, about how the market will respond to a given event once it settles back into equilibrium.

Realistically, of course, this is very optimistic. However I believe the description of the technology is entirely plausible; the main question is whether nontrivial and reasonably accurate predictions about macroeconomic systems can be made using such a compositional game model, and whether compositionality of games is a good model for compositionality of macroeconomic systems. I conjecture that computations made from a user-supplied strategy profile can be done efficiently, in time linear in the size of the game; although to compute an approximate equilibrium from nothing we run into the fundamental problem that equilibria of both perfect information and normal form games are expensive to compute.

I will now discuss what needs to be done for this vision of the future to come true, with the exception of the economic questions, which are far outside of my expertise. These divide into theoretical problems and implementation.

A direct translation of the definitions in §2.1 and §2.2 into a programming language will consist of combinators for composing games, and each game will implement the play, coplay and best response functions, each taking as one input a strategy profile implemented as a tuple of functions. In particular, the play function will convert a strategy profile into a distribution of strategic plays (in the case of probabilistic choice), and the best response function can be used

to decide whether a given strategy is a Nash equilibrium.²⁰ The prototype described in §0.7 already has these features, and it is not difficult to imagine other features such as search and optimisation algorithms sitting above these primitive features, although it might be more difficult to make these run in a feasible amount of time.

The software *Quantomatic* (Kissinger et al. 2014) is currently able to reason about string diagrams using graph rewrite rules, but it does not convert the string diagram into the logical language of composition and tensor, which is what is needed in game theory.²¹ I propose an extensible, language-independent API implementing the logical language of morphisms of monoidal categories, which can be used as an intermediate language between *Quantomatic* (or similar software) and several backends. For example, different code generators could target **MATLAB** for **FVect**, **Maple** for **Rel** and **Haskell** for **Game(C)**.

Far more work is needed on the theoretical side. A typical example of a game in practical economics may have several awkward features simultaneously: dependent types, infinite repetition, incomplete information, learning, irrationality, ...

For such practical examples, each of these needs to be represented in compositional game theory, in compatible ways. For several of these problems, my preferred approach is to apply intuition from functional programming and change the underlying category to successively add new features. Informally, I think of this in terms of ‘effects stacks’. A simple example may have probabilistic choice at the bottom of the stack, followed by learning, followed by nondeterminism, followed by selection. As the first step, probabilistic choice is the monad

$$T_1X = \mathcal{D}X$$

introduced in §2.1.3. The second step is to apply the state monad transformer

$$T_2X = \text{St}_A^{T_1} X = A \rightarrow T_1(X \times A) = A \rightarrow \mathcal{D}(X \times A)$$

Alternatively, this point can be reached using Lawvere theories rather than monad transformers²². Adding nondeterminism to a stack involving probability is not straightforward, but there are two approaches discussed in (Varacca and Winskel 2006). An alternative possibility is to use synthetically compact sets (Escardó 2004) to represent nondeterminism, which leads to

$$T_3X = \mathcal{J}_{\mathbb{B}}(T_2X) = ((A \rightarrow \mathcal{D}(X \times A)) \rightarrow \mathbb{B}) \rightarrow A \rightarrow \mathcal{D}(X \times A)$$

which is a monad due to the unpublished fact that for every monad T there is a distributive law

$$\lambda : T \circ \mathcal{J}_{\mathbb{B}} \rightarrow \mathcal{J}_{\mathbb{B}} \circ T$$

²⁰The prototype decides Nash equilibria rather than subgame perfect equilibria, by implementing the first variant definition in §2.2.3. Deciding subgame perfection naïvely is impossible in general, because there may be too many subgames to search.

²¹From discussions with Mehrnoosh Sadrzadeh I understand that the lack of this feature is also a bottleneck in quantum linguistics.

²²Monad transformers are deprecated by theoretical computer scientists (with good reason), but in my opinion they are still useful as one tool in a larger toolbox for building effects stacks, especially when the stack involves infinitary effects such as continuations and selections.

Finally, selections can be added using the generalised selection monad in (Escardó and Oliva 2015)²³, resulting in

$$\begin{aligned} T_4X &= \mathcal{J}_R^{T_3} X \\ &= (X \rightarrow R) \rightarrow T_4X \\ &= (X \rightarrow R) \rightarrow ((A \rightarrow \mathcal{D}(X \times A)) \rightarrow \mathbb{B}) \rightarrow A \rightarrow \mathcal{D}(X \times A) \end{aligned}$$

where R is a suitable algebra of T_4 , ultimately derived from the expectation operator $\mathbb{E} : \mathcal{D}\mathbb{R} \rightarrow \mathbb{R}$.

Although this looks complicated, machinery exists in functional programming to work with compound monads like this. I think it will be necessary to make a systematic study of this before serious applications of compositional game theory in economics become possible. In general, the Kleisli categories of noncommutative monads are symmetric premonoidal (and moreover closed Freyd). It may be that for premonoidal categories, string diagrams in the sense of (Jeffrey 1997) are impractical and should be replaced with a form of arrow notation (Paterson 2001). However, in an effects stack in which each player has their own state variable that is private to the other players (such as Bayesian beliefs), any computations done by different players will commute past each other. This results in far more allowed topological moves on string diagrams than for general premonoidal categories, and so I suggest that the software will also need to keep track of allowed moves on a string diagram, based on which morphisms of the category commute.

0.11 Notation and conventions

Function application is written without brackets whenever it is unambiguous, and the function set arrow \rightarrow associates to the right, so if $f : X \rightarrow Y \rightarrow Z$ then $fxy = (f(x))(y)$. λ -abstractions are denoted $\lambda(x : X).t$, where x is the abstracted variable and X its type. The scope of the abstraction extends as far as possible to the right, and binds tighter than everything except parentheses and the equals sign. The application of a selection function to a λ -abstraction, for example, is written $\varepsilon\lambda(x : X).t$.

Binary and arbitrary products are written \times and \prod , binary and arbitrary coproducts are written $+$ and \sum , and both bind tighter than \rightarrow , so for example $A \times B \rightarrow C + D$ means $(A \times B) \rightarrow (C + D)$. Projections from a product are written π , and injections into a coproduct are ι . If $x : \prod_{i:I} X_i$ then I write x_i for $\pi_i x : X_i$. Similarly, if $f : X \rightarrow \prod_{i:I} Y_i$ then I write f_i for $\pi_i \circ f : X \rightarrow Y_i$. The type with one element is 1 , and its element is $*$: 1 . The subscript $-i$, as in x_{-i} , f_{-i} , π_{-i} is used for projection onto $\prod_{j \neq i} X_j$. The notation (x_i, x_{-i}) , common in game theory, additionally uses the natural isomorphism $X_i \times \prod_{j \neq i} X_j \cong \prod_{j:I} X_j$ implicitly.

I sometimes pretend that categories are stricter than they really are, especially in chapter 3, in order to avoid the problems with isomorphisms mentioned in §0.7. For example, I will pretend that $(X \times Y) \times Z = X \times (Y \times Z)$ and write an element of either in the form (x, y, z) . Similarly, if $f : X \rightarrow 1 \rightarrow Y$ I will

²³The selection monad transformer in (Hedges 2014) is the special case when R is a free algebra. Note that $[0, 1]$ is the free \mathcal{D} -algebra on two generators, but \mathbb{R} is a non-free \mathcal{D} -algebra. Algebras of \mathcal{D} are called convex sets, see (Fritz 2009, Jacobs et al. 2015).

sometimes write fx for fx^* , and if $g : X \times 1 \rightarrow Y$ I will sometimes write gx for $g(x, *)$. This is a tradeoff between correctness and manageable notation, and I am still unsure what the best approach is, short of using a proof assistant.

I distinguish carefully between ‘types’ X , which are sets, and ‘sets’ $A \subseteq X$, which are functions $A : X \rightarrow \mathbb{B}$ where X is a type and \mathbb{B} is the type of booleans containing \perp and \top . To confuse matters, however, I often also use the term ‘set’ to refer to a type when the difference is harmless, because there is so much social inertia behind the ‘set’ terminology. If $A \subseteq X$ and $x : X$, then the notation $x \in A$ is shorthand for $Ax = \top$, and $\{x : X \mid \dots\}$ is shorthand for its characteristic function. $\mathcal{P}X = X \rightarrow \mathbb{B}$ is the set of all subsets of the type X .

$\mathbb{P}[\alpha = x] : [0, 1]$ is the probability that a random variable $\alpha : \mathcal{D}X$ is equal to $x : X$. Here \mathcal{D} is the probabilistic analogue of \mathcal{P} , defined in §2.1.3. If $\alpha : \mathcal{D}\mathbb{R}$ then $\mathbb{E}[\alpha] : \mathbb{R}$ is the expected value of α .

Some symbols are reserved for particular uses. Types of plays or moves are denoted X, Y, Z and types of outcomes are denoted R, S, T . Σ is the set of strategy profiles of a game, and σ is a strategy profile. \mathbf{P} is a play function, which converts a strategy profile into a play, and \mathbf{B} is a best response function, which takes a strategy profile to its set of best responses. \mathbf{R} is a response function (§2.1.7) and \mathbf{C} is a coplay function (§2.1.4), which appear only in the context of open games. ε and δ are single-valued selection functions, E is a multi-valued selection function, φ and ψ are single-valued quantifiers, and Φ is a multi-valued quantifier. q is an outcome function, which takes a play to an outcome. \mathcal{U} is the unilateral deviation operator. \mathcal{G} and \mathcal{H} are open games, and \mathcal{D} is a decision.

Chapter 1

Higher order game theory

1.1 Decision theory

1.1.1 Discussion

The purpose of [chapter 1](#) of this thesis, which is largely based on ([Hedges et al. 2015b](#)) and ([Escardó and Oliva 2011](#)), is to introduce the reader to a particular way of thinking about game theory. Although it is logically self-contained, readers unfamiliar with game theory should read it together with another source, such as ([Leyton-Brown and Shoham 2008](#)) or ([Kreps 1991](#)), that gives a more classical introduction with motivation and standard examples. Readers already familiar with game theory could begin reading at [chapter 2](#) and follow hyperlinked references into this part when necessary.

Suppose we have some situation in which an agent is choosing a move of type X . After the choice is made, she receives some outcome, say of type R . For example, if we are representing preferences by utilities, we would take $R = \mathbb{R}$. The outcome depends not just on the agent's choice, but also on the choices of other agents, the 'rules' of the situation, and the agent's own preferences¹. The first key concept of this thesis is that these additional dependencies will be abstracted away into a single function $k : X \rightarrow R$ mapping choices to outcomes. This function will be called a (strategic) **context**².

If we view an agent as computing a move, the context represents the computation done afterwards by the environment. This leads us to the principle that *strategic contexts are delimited continuations*. Furthermore the structure of game theory, and especially the definition of Nash equilibrium ([§1.2.3](#)), is such that many of our definitions have explicit access to the context. A computation which has access to its calling environment by means of a continuation is precisely a continuation passing style computation. Throughout this thesis, and especially in [§2.2](#) we will see that the mathematical structure of game theory can be usefully improved by allowing more things to depend on an arbitrary continuation. This is the principle, perhaps the most important single idea in this thesis, that *game theory wants to be in continuation passing style*.

¹The dependence of the outcome on the agent's preferences comes from the fact that the outcome may represent the agent's subjective 'rating' of the *physical* outcome (see [§1.2.5](#)).

²However, in [§2.1.5](#) we will generalise the term 'context'.

We will describe agents by their behaviour on each context. There are two options, which lead respectively to ‘quantifiers’ and ‘selection functions’: we can map the context to either the good outcomes, or the good moves. We will see in §1.2 that the latter is preferred for technical reasons, and so we will adopt the following slogan: *to know an agent is to know her preferred moves in every context.*

1.1.2 Quantifiers

Consider a sentence of predicate logic of the form

$$\exists(x : X).kx$$

Here k is a predicate, which either holds or does not hold for each element of the domain of quantification X . Thus, we can view k as a function $k : X \rightarrow \mathbb{B}$, where $\mathbb{B} = \{\perp, \top\}$ is the type of booleans. Since the meaning of our sentence is invariant under α -renaming, it depends only on the value of the function k , and thus could be unambiguously³ written with the point free syntax $\exists k$. For comparison, there is a familiar example of point free syntax in measure theory⁴:

$$\int_X k \, d\mu = \int_X k(x) \, d\mu(x)$$

Since $\exists k$ has a value of type \mathbb{B} and depends only on $k : X \rightarrow \mathbb{B}$, we can say that \exists is a particular function

$$\exists : (X \rightarrow \mathbb{B}) \rightarrow \mathbb{B}$$

To be precise, \exists is the function of this type defined by

$$\exists k = \begin{cases} \top & \text{if } kx = \top \text{ for some } x : X \\ \perp & \text{otherwise} \end{cases}$$

We can do the same thing with the universal quantifier: it is a function

$$\forall : (X \rightarrow \mathbb{B}) \rightarrow \mathbb{B}$$

defined by

$$\forall k = \begin{cases} \top & \text{if } kx = \top \text{ for all } x : X \\ \perp & \text{otherwise} \end{cases}$$

Abstracting from these two cases leads to the definition of a **generalised quantifier** in (Mostowski 1957) as an arbitrary function of type $(X \rightarrow \mathbb{B}) \rightarrow \mathbb{B}$.

This is generalised one step further in (Escardó and Oliva 2010a), by allowing the type of booleans \mathbb{B} to be replaced by an arbitrary type R . The most important new example that this gains us is maximisation of real-valued functionals $k : X \rightarrow \mathbb{R}$: by the same reasoning as before, the expression

$$\max_{x:X} kx$$

³However, notice that the type X has become implicit in the syntax; we could more clearly write $\exists_X k$.

⁴See the integration quantifier in section §1.1.5.

can be written as $\max k$, and we can view \max as a function

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

In summary, we define a **quantifier** on a set X as an arbitrary function

$$\varphi : \mathcal{K}_R X$$

where $\mathcal{K}_R X = (X \rightarrow R) \rightarrow R$. We will view φ as a function that takes each context $k : X \rightarrow R$ to an agent's preferred outcome given that context. We should always⁵ think of this as an outcome that is in the image of the context, that is, it is an outcome that can actually occur given that context.

For example, the quantifier $\max : \mathcal{K}_{\mathbb{R}} X$ models a classical economic agent who maximises utility, in the sense that the preferred outcome is the one that is maximal among those that can be attained. In this sense, quantifiers can be seen as a generalisation of utility maximisation, which abstracts away the irrelevant fact that we are working with the ordered real numbers, and allows us to focus on the important structure.

The existential quantifier $\exists : \mathcal{K}_{\mathbb{B}} X$ is also an instance of maximisation, this time over the discrete order $\perp < \top$. We also observe that $\min : \mathcal{K}_{\mathbb{R}} X$ is a quantifier, and that $\forall : \mathcal{K}_{\mathbb{B}} X$ similarly minimises⁶ over the order $\perp < \top$.

1.1.3 The continuation monad

In §1.1.2 we introduced the type $\mathcal{K}_R X = (X \rightarrow R) \rightarrow R$. The operator \mathcal{K}_R is well known in programming language theory where it is called the **continuation monad** (Moggi 1991). This means that we have unit maps $\eta : X \rightarrow \mathcal{K}_R X$, and for each function $f : X \rightarrow \mathcal{K}_R Y$ a Kleisli extension $f^* : \mathcal{K}_R X \rightarrow \mathcal{K}_R Y$. Explicitly, these are given by

$$\eta x = \lambda(k : X \rightarrow R).kx$$

and

$$f^* \varphi = \lambda(k : Y \rightarrow R).\varphi \lambda(x : X).fxk$$

A Kleisli arrow $\varphi : X \rightarrow \mathcal{K}_R Y$ is viewed as a computation of type $X \rightarrow Y$ in **continuation passing style**. This means that after the function $X \rightarrow Y$ has terminated the result is passed to a continuation $Y \rightarrow R$, and the computation is allowed to have first class access to its continuation. This point of view is detailed in (Hedges 2015b).

The structure of the continuation monad is related to game theory in (Escardó and Oliva 2011, Hedges 2014), but we will not do so in this thesis, because of the argument in (Hedges et al. 2015b) and §1.2 for preferring selection functions to quantifiers.

⁵This is always the case for attainable quantifiers (§1.1.5), and all reasonable quantifiers are attainable.

⁶The fact that \forall is an instance of \min is often surprising when seen for the first time. If k only takes the value \top then $\min k = \top$, but if k takes the value \perp somewhere then $\min k = \perp$. This is precisely the definition of $\forall k$.

1.1.4 Selection functions

Just as a quantifier gives the best outcome in a context, so a selection function gives the best move. Thus, a selection function is a function $\varepsilon : \mathcal{J}_R X$ where

$$\mathcal{J}_R X = (X \rightarrow R) \rightarrow X$$

The selection function corresponding⁷ to the existential quantifier is the Hilbert epsilon operator⁸. In the Hilbert calculus $\varepsilon(x : X).kx$ is a term that, by definition, satisfies k if possible. We can informally define $\varepsilon : \mathcal{J}_R X$ by

$$\varepsilon k = \begin{cases} \text{some } x : X \text{ satisfying } kx & \text{if such } x \text{ exists} \\ \text{arbitrary} & \text{otherwise} \end{cases}$$

By the axiom of choice, we can obtain such a function ε satisfying this specification.

Similarly, the selection function corresponding to max is $\arg \max : \mathcal{J}_R X$, which chooses some point at which a function is maximised. This is ordinarily written as

$$\arg \max k = \arg \max_{x:X} kx$$

Again, because a function may attain its maximum at many⁹ points, we generally need the axiom of choice to actually obtain $\arg \max$ as a function.¹⁰

Although mathematically speaking selection functions are often about *optimality*, it will sometimes be useful to think in anthropomorphic terms of *satisfaction*. To talk about the value φk , we say that if an agent is *choosing in the context* $k : X \rightarrow R$, then she is **satisfied** with the outcome φk . Similarly, she is satisfied with making the move εk . If the agent has a quantifier φ that chooses an outcome $\varphi k \notin \text{Im}(k)$ that is not attainable by any move then the agent's preferences in the context k are **unrealistic**, because she will never be satisfied with any outcome that can actually occur.¹¹

1.1.5 Attainment

We have given two types of functions that can be used as models of agents, and now we will study the relationship between them. We begin by noticing that a selection function ε defines a quantifier $\bar{\varepsilon}$ by the equation

$$\bar{\varepsilon} k = k(\varepsilon k)$$

⁷Quantifiers and selection functions generally come in pairs, as explained in §1.1.5.

⁸The Hilbert epsilon symbol was used as a primitive in Hilbert's ε -calculus, and again in Bourbaki's set theory. See for example (Mathias 2002).

⁹A discontinuous function may not have an attained maximum, but because discontinuous functions do not exist in nature, we do not need to worry about this.

¹⁰The function $\arg \max$ is itself discontinuous, and as per the previous footnote it cannot exist in nature. This is not yet an issue in practice, because we can replace the reals with some approximation such as the rationals or floating point numbers. True reals become important with repeated games, where completeness is needed to ensure convergence of the utility functions.

¹¹For a vivid example of unrealistic preferences, suppose I am satisfied only with the outcome in which I become a millionaire, but there is no choice I can make that leads to that outcome.

This defines a map $\bar{\cdot} : \mathcal{J}_R X \rightarrow \mathcal{K}_R X$, which can be proved to be a morphism of monads, see lemma 6.3.6 of (Escardó and Oliva 2010a), where the selection monad \mathcal{J}_R is defined in §1.3.2.

We will say that the selection function ε **attains** the quantifier φ just if $\varphi = \bar{\varepsilon}$, and we will call the quantifier φ **attainable** if it is attained by some selection function.¹²

As a first example, the Hilbert epsilon operator ε attains the existential quantifier. This generalises Hilbert’s definition of the existential quantifier as $\exists k = k(\varepsilon k)$. For if $\exists k = \top$ then we have at least one $x : X$ satisfying $kx = \top$, and so $x = \varepsilon k$ has this property. Conversely, if $\exists k = \perp$ then there is no such x , and so $x = \varepsilon k$ is some arbitrary point, which has $kx = \perp$. As a second example, it is easy to see that arg max attains max , essentially by definition.

For a less trivial example, consider the definite integration operator¹³.

$$\int : \mathcal{K}_{\mathbb{R}}[0, 1]$$

defined by

$$\int k = \int_0^1 kx \, dx$$

The mean value theorem tells us that if $p : [0, 1] \rightarrow \mathbb{R}$ is integrable then there is some point $x : [0, 1]$ with the property that

$$kx = \int k$$

We can apply the axiom of choice to form a selection function $\varepsilon : \mathcal{J}_{\mathbb{R}}[0, 1]$ that takes each k to such an x . Thus, \int is an attainable quantifier. This is given as an example in (Escardó and Oliva 2010a).

For a second interesting example, also from (Escardó and Oliva 2010a) suppose we work in a setting (such as a cartesian closed category of domains) in which every endomorphism $k : X \rightarrow X$ has a canonical fixpoint, computed by a function

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

In this case we have $R = X$, and $\mathcal{J}_X X = \mathcal{K}_X X$, and so fix can be seen as both¹⁴ a quantifier and a selection function. Moreover, because $\text{fix } k$ is guaranteed to be a fixpoint of k we have $k(\text{fix } k) = \text{fix } k$, and hence $\text{fix} : \mathcal{J}_X X$ attains $\text{fix} : \mathcal{K}_X X$.

1.1.6 Multi-valued variants

Quantifiers were generalised yet another step in (Escardó and Oliva 2011), by allowing the quantifier to return a set of results,

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

¹²Typically, a single quantifier is attained by many selection functions, as is the case for \exists and max .

¹³To deal with \int being undefined on arguments that are not integrable, we can restrict its domain (§1.1.6).

¹⁴This is a situation in which we need to be careful about types. Although $\text{fix} : \mathcal{J}_X X$ and $\text{fix} : \mathcal{K}_X X$ might at first appear to be the same function, it is not even a well formed claim to say that they are equal, because the types have different monad structures. Quantifiers and selection functions support different game-theoretic solution concepts, and the games defined by these two (apparently equal) operators have different sets of equilibria, as we will see in §1.2.10.

We will call a function with this type a **multi-valued quantifier**. Similarly, in (Hedges et al. 2015b) and (Hedges et al. 2015a), selection functions were generalised to **multi-valued selection functions**

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

Multi-valued quantifiers were introduced in (Escardó and Oliva 2011) for applications in game theory and the proof of Bekić’s lemma. For example, the max quantifier is single-valued for a total order such as \mathbb{R} , but is multi-valued for a preorder such as \mathbb{R}^n with $x \leq_i y \iff x_i \leq y_i$. The step from single-valued to multi-valued selection functions makes it harder to work with sequential games (because we lose the monad structure, see §1.3), but easier to work with simultaneous games. Thus, §1.2 will focus on multi-valued selection functions, but §1.3 on single-valued selection functions.

We will give two important examples of multi-valued selection functions. The first is the multi-valued variant of arg max,

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

which chooses *all* points at which its argument is maximised:

$$\text{arg max } k = \{x : X \mid kx \geq kx' \text{ for all } x' : X\}$$

Notice that under reasonable hypotheses¹⁵ arg max k is nonempty.

The second example is the multi-valued fixpoint operator

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

defined by

$$\text{fix } k = \{x : X \mid x = kx\}$$

The fixpoint operator on sets (as opposed to posets) is naturally multivalued, because a function may have zero, one or many fixpoints, and no preferred fixpoint.

It is a subtle question whether we should allow multi-valued quantifiers and selection functions to return the empty set on any input. For some applications in game theory it is useful to suppose that the sets are always non-empty. We will call such a multi-valued quantifier or selection function **total**, after (Hedges 2013). For example arg max on a finite set is total, and in §1.2.9 we will extend fix to a total selection function by defining its behaviour differently on a context with no fixpoints. This is the approach taken in (Hedges et al. 2015b).

On the other hand, there are situations where it is unreasonable to require totality, such as when working with games of imperfect information with pure strategies, so that equilibria may not exist. We will define the **domain** of a multi-valued quantifier or selection function to be the subset of $X \rightarrow R$ on which it returns a nonempty set. This will be denoted by $\text{dom}(E)$ or $\text{dom}(\Phi)$. For example, the domain of arg max on a compact space X contains all continuous functions. Of course, a multi-valued quantifier or selection function can be restricted to a total one on its domain. This is the approach taken in (Hedges 2013).

¹⁵If X is finite then arg max k is nonempty for all k . If X is compact then arg max k is nonempty if k is continuous.

1.1.7 Multi-valued attainment

Multi-valued quantifiers and selection functions also support a concept of attainment, namely that $E : (X \rightarrow R) \rightarrow \mathcal{P}X$ attains $\Phi : (X \rightarrow R) \rightarrow \mathcal{P}R$ if for all $k : X \rightarrow R$ we have

$$\{kx \mid x \in Ek\} \subseteq \Phi k$$

This definition gives the expected attainments: for example, the multi-valued $\arg \max$ attains \max (where \max is a single-valued quantifier viewed as a multi-valued quantifier by returning singletons), and the multi-valued fixpoint operator attains itself.

By analogy to the overline operator $\overline{\cdot} : \mathcal{J}_R \rightarrow \mathcal{K}_R$, given a multi-valued selection function $E : (X \rightarrow R) \rightarrow \mathcal{P}X$ we can define the ‘smallest’ multi-valued quantifier $\overline{E} : (X \rightarrow R) \rightarrow \mathcal{P}R$ attained by E , namely

$$\overline{E}k = \{kx \mid x \in Ek\}$$

Equivalently, $\overline{E}k$ is the forward image of Ek under k . However, in the multi-valued case we can also do this in reverse, converting a quantifier Φ into the ‘largest’ selection function $\overline{\Phi}$ attaining it, namely

$$\overline{\Phi}k = \{x : X \mid kx \in \Phi k\}$$

As suggested, the types $(X \rightarrow R) \rightarrow \mathcal{P}X$ and $(X \rightarrow R) \rightarrow \mathcal{P}R$ both carry a partial order structure inherited from the powerset operator. Given multi-valued selection functions $E_1, E_2 : (X \rightarrow R) \rightarrow \mathcal{P}X$ we will say that E_1 **refines** E_2 , and write $E_1 \sqsubseteq E_2$, if for every $k : X \rightarrow R$ we have $E_1k \subseteq E_2k$. Similarly, for quantifiers we will say that Φ_1 refines Φ_2 and write $\Phi_1 \sqsubseteq \Phi_2$. With this notation, we can say that E attains Φ iff $\overline{E} \sqsubseteq \Phi$.

We can view a single-valued selection function $\varepsilon : \mathcal{J}_R X$ as multi-valued by setting $Ek = \{\varepsilon k\}$, and so we can talk about refinement between single-valued and multi-valued quantifiers. Specifically, we say that ε refines E iff for all k we have $\varepsilon k \in Ek$. By the axiom of choice, a multi-valued selection function has a single-valued refinement iff it is total. This also applies to quantifiers.

The overline operators define a Galois connection between the refinement orders. Given a selection function $E : (X \rightarrow R) \rightarrow \mathcal{P}X$ and a quantifier $\Phi : (X \rightarrow R) \rightarrow \mathcal{P}R$,

$$\overline{E} \sqsubseteq \Phi \iff E \sqsubseteq \overline{\Phi}$$

The proof is straightforward, by showing that both sides are equivalent to the claim that for every $k : X \rightarrow R$ and $x : X$, if $x \in Ek$ then $kx \in \Phi k$.

The double overline operator on total quantifiers is the identity, because

$$\overline{\overline{\Phi}}k = \{kx \mid x \in \overline{\Phi}k\} = \{kx \mid kx \in \Phi k\} = \Phi k \cap \text{Im } k$$

On selection functions, following order-theoretic terminology we will think of $\overline{\overline{\cdot}}$ as a *closure operator*, and a selection function E satisfying $E = \overline{\overline{E}}$ will be called **closed**. We will see examples of closed selection functions in §1.2.5, and non-examples in §1.2.9.

1.1.8 Modifying the outcome type

The operator \mathcal{J}_R defined in §1.1.4 is contravariant in the outcome type R , in contrast to the continuation monad \mathcal{K}_R , which is not functorial in R . Moreover, the action of \mathcal{J}_R on morphisms of outcomes has a clear game-theoretic reading.

Given a selection function $\varepsilon : \mathcal{J}_R X$ and a function $f : S \rightarrow R$, let $f \cdot \varepsilon : \mathcal{J}_S X$ be the selection function

$$f \cdot \varepsilon = \lambda(k : X \rightarrow S). \varepsilon(f \circ k)$$

This can easily be shown to commute with the other operations on \mathcal{J}_R , including the product of selection functions (§1.3.3). The selection category operator $\mathcal{J}_R \mathcal{C}$ defined in §1.3.2 can therefore be considered as a presheaf of categories $\mathcal{C}^{\text{op}} \rightarrow \mathbf{Cat}$.

This also applies to multi-valued selection functions, where the selection function $E : (X \rightarrow R) \rightarrow \mathcal{P}X$ changes to

$$f \cdot E = \lambda(k : X \rightarrow S). E(f \circ k)$$

If we take f to be the projection $\pi_i : \mathbb{R}^N \rightarrow \mathbb{R}$ then the multi-valued selection function $\arg \max : (X \rightarrow \mathbb{R}) \rightarrow \mathcal{P}X$ changes to

$$\pi_i \cdot \arg \max = \lambda(k : X \rightarrow \mathbb{R}^N). \{x : X \mid k_i x \geq k_i x' \text{ for all } x' : X\}$$

which is the selection function modelling an agent who maximises the i th coordinate while ignoring the others. This selection function is used in (Hedges et al. 2015b).

We will return to this idea in §3.3.

1.2 Normal form games

1.2.1 Discussion

A definition of game was given in (von Neumann and Morgenstern 1944) that is general enough for most purposes, the so-called extensive form games. At this point we will not introduce extensive form games formally, but we will discuss some of the important concepts.

A game consists of players, who make choices. The choices made by all players, together with the rules of the game, determine an outcome. The choices made by the players are constrained by the fact that each player has (usually different) preferences over the outcomes, and each player acts in such a way as to bring about their preferred outcome. We will model preferences of players by quantifiers or selection functions, which abstracts away more specific definitions such as preference relations or utilities used in standard game theory.

In order to make an informed choice the player needs to know which outcomes will occur for a given choice, but to know this, she needs to know what the other players will choose. However, the other players are reasoning in the same way, and need to know what she will choose, and so we have a circularity. The circularity is resolved by a solution concept, each of which is a proposed definition for what it means for a player to choose rationally, under various assumptions about the other players.

The game may have some dynamic structure in which some players can observe (possibly partial information about) the choices of some other players before making their own choice. A **strategy** for a player is a function that chooses a move, contingent on observed information. Strategies can be used to abstract away the details of how players interact with each other.

Given an arbitrary game, we can define a new game called its **normalisation**. In this new game, the choices are precisely the strategies of the previous game. Given a strategy for each player, we can ‘play out’ the strategies to determine a choice for each player in the old game, which in turn determines an outcome. The new game is played simultaneously, with no player able to make any direct observations. Thus, games with dynamic structure can be disregarded, and we can focus on simultaneous games only.

For example, suppose the first player chooses $x : X$, and then the second player perfectly observes it and chooses $y : Y$, with the outcome being $q(x, y)$. In the normalisation, simultaneously the first player chooses a strategy¹⁶ $\sigma_1 : X$ and the second player chooses a strategy $\sigma_2 : X \rightarrow Y$, with the outcome being $q'(\sigma_1, \sigma_2) = q(\sigma_1, \sigma_2 \sigma_1)$.

The most standard solution concept for simultaneous games is called the Nash equilibrium. This is unable to distinguish between an extensive form game and its normalisation, in the sense that it gives the same ‘solutions’ (rational choices) for both. However there are more refined solution concepts, such as subgame perfect equilibrium, which can distinguish an extensive form game from its normal form. Thus, it is still important to study extensive form games.

This section is based almost entirely on (Hedges et al. 2015b).

1.2.2 Games, strategies and unilateral continuations

We begin with a collection I of players. Each player $i : I$ has a nonempty set X_i of **moves**. We also have a set R of outcomes, and an **outcome function**

$$q : \prod_{i:I} X_i \rightarrow R$$

Since the game is played simultaneously, a **strategy** for each player is simply a move $\sigma_i : X_i$. The function q (together with its type) completely specifies the *rules* of a normal form game. We will leave the specification of the *players* until §1.2.3.

A **strategy profile** is a tuple of strategies for each player,

$$\sigma : \prod_{i:I} X_i$$

A **play** is a tuple of moves (which has the same type) and playing the strategy profile σ results in the play $\mathbf{P}\sigma = \sigma$, so strategies are ‘played out’ by the **play function**¹⁷

$$\mathbf{P} = \text{id} : \prod_{i:I} X_i \rightarrow \prod_{i:I} X_i$$

¹⁶The notation is chosen carefully here. $x : X$ should be thought of as a choice, and $\sigma_1 : X$ as a strategy. For the first player in particular (who makes no observation, and cannot use side effects such as randomness), choices and strategies happen to coincide, but they should still be distinguished carefully.

¹⁷This is our first encounter with play functions, which appear in full generality in §2.1.4.

The play $\mathbf{P}\sigma = \sigma$ is called the **strategic play** of the strategy σ , and is the play that results from all players playing according to σ . Each strategy σ additionally determines an outcome $q(\mathbf{P}\sigma) = q\sigma$.

We will next define **unilateral continuations**, which have proved to be a useful tool for reasoning about higher order games. They were introduced in (Hedges 2013)¹⁸, in which it is shown that the majority of the proof of Nash’s theorem amounts to showing that the unilateral continuations satisfy certain topological properties. They are also heavily used in (Hedges et al. 2015b).

Suppose we have a fixed strategy profile σ . We can now define unilateral continuations in which all but one player use σ , and the remaining player **unilaterally deviates** to some other move. The i th player’s unilateral continuation from σ is the function

$$\mathcal{U}_i^q \sigma : X_i \rightarrow R$$

given by

$$\mathcal{U}_i^q \sigma x_i = q(x_i, \sigma_{-i})$$

The notation (x_i, σ_{-i}) , which we will now introduce here, is standard in game theory and useful for reasoning about unilateral deviation. The subscript in σ_{-i} means that we project σ onto the subspace $\prod_{j \neq i} X_j$. We will sometimes write the projection operator as

$$\pi_{-i} : \prod_{j=1}^N X_j \rightarrow \prod_{\substack{1 \leq j \leq N \\ j \neq i}} X_j$$

The notation (x_i, σ_{-i}) fills the ‘missing’ i th entry with x_i , and is defined by the equation

$$(x_i, \sigma_{-i})_j = \begin{cases} x_i & \text{if } i = j \\ \sigma_j & \text{otherwise} \end{cases}$$

Although this could be ambiguous and is often disliked by those who strive for type safety¹⁹, this notation will come into its own in §3.1 and §3.2.

The purpose of a unilateral continuation is that the behaviour of all other players has been abstracted into a single function, allowing us to reduce a game-theoretic problem to a decision-theoretic one. A different definition of unilateral context that is suitable for sequential games will be given in §1.3.5. The continuations introduced in §2.1.5 are more general, because they can model any number of deviating players.

1.2.3 Nash equilibria of normal form games

The preferences of the player $i : I$ can be modelled either using a multi-valued quantifier

$$\Phi_i : (X_i \rightarrow R) \rightarrow \mathcal{P}R$$

or a multi-valued selection function

$$E_i : (X_i \rightarrow R) \rightarrow \mathcal{P}X_i$$

¹⁸In that paper they are known as ‘unilateral contexts’.

¹⁹See §3.2.1 for remarks on reasoning about this notation with a proof assistant.

In any context $k : X_i \rightarrow R$, the value of the quantifier $\Phi_i k$ is the set of outcomes that player i considers to be good in the context k . Similarly, if we use selection functions then $E_i k$ is the set of moves that player i considers to be good in the context.

The unilateral continuation $\mathcal{U}_i^q \sigma : X_i \rightarrow R$ is the context in which player i is unilaterally deviating from the strategy profile σ . If the player is implemented by a multi-valued quantifier then the set of outcomes that the player considers good, and can be attained by unilaterally deviating, is $\Phi_i(\mathcal{U}_i^q \sigma)$. The outcome that actually occurs if the player does not deviate is $q(\mathbf{P}\sigma) = q\sigma$. Therefore, if

$$q\sigma \in \Phi_i(\mathcal{U}_i^q \sigma)$$

then the player is already satisfied with the outcome and has no incentive to unilaterally deviate. If this condition holds for each player $i : I$ then we will call σ a **quantifier equilibrium**.

On the other hand, if we implement players using multi-valued selection functions, then we have a set of good moves $E_i(\mathcal{U}_i^q \sigma)$. This is the set of moves which it would be rational for player i to choose, given that all other players use σ . The actual move chosen by player i is σ_i , so player i has no incentive to unilaterally deviate if

$$\sigma_i \in E_i(\mathcal{U}_i^q \sigma)$$

If this condition holds for all $i : I$ then we call σ a **selection equilibrium**.

If we could model the players either by quantifiers Φ_i , or by selection functions E_i that attain Φ_i (in the sense of §1.1.7), then every selection equilibrium is a quantifier equilibrium. To see this, if $\sigma_i \in E_i(\mathcal{U}_i^q \sigma)$ and E_i attains Φ_i , then

$$q\sigma = \mathcal{U}_i^q \sigma \sigma_i \in \Phi_i(\mathcal{U}_i^q \sigma)$$

We can therefore say that the selection equilibrium is an **equilibrium refinement** of the quantifier equilibrium. In particular, if we have players modelled by selection functions E_i then we can define quantifier equilibria using the quantifiers \overline{E}_i .

The converse of this holds, and hence quantifier and selection equilibria coincide, if our selection functions are of the form $E_i = \overline{\Phi}_i$, where the overline operator is the one defined in §1.1.7. To see this, suppose σ is a quantifier equilibrium, so we have

$$q\sigma = \mathcal{U}_i^q \sigma \sigma_i \in \Phi_i(\mathcal{U}_i^q \sigma)$$

Then $\sigma \in \overline{\Phi}_i(\mathcal{U}_i^q \sigma)$ and hence we have a selection equilibrium for the selection functions $\overline{\Phi}_i$.

Given a strategy profile σ , we will say that player i is **non-pivotal** if the unilateral continuation $\mathcal{U}_i^q \sigma : X_i \rightarrow R$ is constant. If σ is a quantifier equilibrium then any deviation (x_i, σ_{-i}) by a non-pivotal player results in another quantifier equilibrium. Selection equilibria in general do not share this property. We will see this in practice in §1.2.10.

1.2.4 Best responses

An important concept in the foundations of game theory is that of a **best response function**. This is a function $\mathbf{B} : \Sigma \rightarrow \mathcal{P}\Sigma$, where Σ is the type of

strategy profiles of a game, such that the fixpoints of \mathbf{B} pick out some solution concept. The informal specification of a game's best response function is that $\mathbf{B}\sigma$ should be the set of strategy profiles σ' such that for each player i , it is rational for player i to use the strategy σ'_i , in the sense of having no incentive to unilaterally deviate, under the assumption that the other players are using σ_{-i} .

For a normal form game as defined in §1.2.2, we can define the best response function

$$\mathbf{B} : \prod_{i:I} X_i \rightarrow \mathcal{P} \prod_{i:I} X_i$$

in two different ways, which lead to quantifier and selection equilibria. If our players are defined by multi-valued quantifiers $\Phi_i : (X_i \rightarrow R) \rightarrow \mathcal{P}R$ then we can use the definition

$$\mathbf{B}\sigma = \left\{ \sigma' : \prod_{i:I} X_i \mid \mathcal{U}_i^q \sigma \sigma'_i \in \Phi_i(\mathcal{U}_i^q \sigma) \text{ for all } i : I \right\}$$

A fixpoint of this \mathbf{B} is a strategy profile satisfying

$$\mathcal{U}_i^q \sigma \sigma_i \in \Phi_i(\mathcal{U}_i^q \sigma)$$

for each $i : I$. Since $\mathcal{U}_i^q \sigma \sigma_i = q\sigma$, this is precisely the definition of a quantifier equilibrium.

Alternatively, if our players are specified by multi-valued selection functions $E_i : (X_i \rightarrow R) \rightarrow \mathcal{P}X_i$, we define the best response function by

$$\mathbf{B}\sigma = \prod_{i:I} E_i(\mathcal{U}_i^q \sigma)$$

A fixpoint $\sigma \in \mathbf{B}\sigma$ of this satisfies

$$\sigma_i \in E_i(\mathcal{U}_i^q \sigma)$$

for each $i : I$, and so is a selection equilibrium.

There are two reasons for focussing on best response functions. A technical reason is that it is sometimes possible to prove existence theorems for a solution concept by applying a fixpoint theorem to a suitable best response function. Nash's theorem has this form, and is discussed in §1.2.6, but this is not something that will be emphasised in this thesis.

A second, more philosophical reason is that having the same best response function seems to be a necessary condition to consider two games to be *equivalent*. Although there is no formal theory of this (although see §2.2.1), in §3.1 and §3.2 we will use equality of best response functions as an informal criterion for two games to be considered the same.

Best response functions are introduced in generality in §2.1.5.

1.2.5 Classical games

In classical normal form games, each play determines a real number for each player called that player's **utility**. The utility is a numerical rating of the player's preference for that play, and players with rational preferences, in the

sense of the von Neumann-Morgenstern utility theorem (von Neumann and Morgenstern 1944), act *as if* they maximise some real utility. It is important to realise that utilities need not *physically* exist, and in particular utility should not be identified with monetary profit (which is usually called ‘payoff’). This point is made strongly in section 1.1 of (Kreps 1990).

In this setting, the outcome of a play consists of a real number for each player, and so we take $R = \mathbb{R}^I$, with the outcome function having type

$$q : \prod_{i:I} X_i \rightarrow \mathbb{R}^I$$

We can think of each player $i : I$ as having her own personal outcome function

$$q_i : \prod_{j=1} X_j \rightarrow \mathbb{R}$$

defined by $q_i = \pi_i \circ q$ where π_i is the i th projection $\mathbb{R}^I \rightarrow \mathbb{R}$.

Player i acts to maximise the i th coordinate, and so we could model her with the maximising quantifier

$$\Phi_i : (X_i \rightarrow \mathbb{R}^I) \rightarrow \mathcal{P}\mathbb{R}^I$$

given by

$$\Phi_i k = \{kx \mid k_i x \geq k_i x' \text{ for all } x' : X_i\}$$

or alternatively with the maximising selection function

$$E_i : (X_i \rightarrow \mathbb{R}^I) \rightarrow \mathcal{P}X_i$$

given by

$$E_i k = \{x : X_i \mid k_i x \geq k_i x' \text{ for all } x' : X_i\}$$

(see §1.1.8).

Notice that $E_i = \overline{\Phi_i}$, and hence quantifier and selection equilibria coincide for a classical game. An equilibrium σ (of either kind) satisfies the conditions, for each $i : I$, that

$$q_i \sigma \geq (\mathcal{U}_i^{q_i} \sigma x')_i$$

for each possible unilateral deviation $x' : X_i$. Since $\mathcal{U}_i^{q_i} \sigma x' = q_i(x', \sigma_{-i})$, we see that this is precisely the ordinary definition of a pure strategy Nash equilibrium (Leyton-Brown and Shoham 2008). Thus, our two solution concepts both coincide with the ordinary one in the case of classical games.

The best response function for a classical normal form game is²⁰

$$\mathbf{B}\sigma = \left\{ \sigma' : \prod_{i:I} X_i \mid q_i(\sigma'_i, \sigma_{-i}) \geq q_i(x'_i, \sigma_{-i}) \text{ for all } i : I \text{ and } x' : X_i \right\}$$

We get this irrespective of whether we use the quantifier or selection function forms, and fixpoints of this \mathbf{B} give precisely the pure strategy Nash equilibria.

The same reasoning applies if we replace \mathbb{R}^I with an arbitrary set R with a **rational preference relation** \succeq . A rational preference relation is another

²⁰To be clear, the condition for a best response is *not* $q_i(\sigma'_i, \sigma_{-i}) \geq q_i \sigma$. The function returns ‘best’ responses, not ‘better’ responses.

name for a total preorder, that is, a relation that is transitive and total but not necessarily antisymmetric. We will write $a \approx b$ if we have $a \preceq b$ and $b \preceq a$ for some $a \neq b : R$. For example, if we take \mathbb{R}^I with the order $a \preceq_i b$ iff $a_i \leq b_i$, then \preceq_i is a rational preference relation which has $a \approx b$ iff $a_i = b_i$.²¹

A rational preference relation on R is the same as a total order on \approx -equivalence classes. Thus, given a continuation $k : X \rightarrow R$, we have a set of optimal outcomes, namely the maximal equivalence class

$$\Phi k = \{kx \mid kx \succeq kx' \text{ for all } x' : X\}$$

This defines a multi-valued quantifier, and similarly we have a multi-valued selection function attaining it.

1.2.6 Mixed strategies

In this section we will turn aside from the development of higher order game theory in order to introduce mixed strategies, which will be used in §3.1. It is still an open problem how mixed strategies should be modelled in the higher order framework.

Consider a classical normal-form game with a finite number of players, labelled $1 \leq i \leq N$, so the outcome function has type

$$q : \prod_{i=1}^N X_i \rightarrow \mathbb{R}^N$$

We will also assume that the sets X_i are finite²². A **mixed strategy** for player i is a probability distribution on X_i . We will write this as $\sigma_i : \mathcal{D}X_i$, where the probability distribution operator \mathcal{D} is properly introduced in §2.1.3. A **mixed strategy profile** is a tuple

$$\sigma : \prod_{i=1}^N \mathcal{D}X_i$$

Given a mixed strategy profile σ , for each player i we obtain a probability distribution on utilities. This can be used to obtain an **expected utility** for the i th player, which is given by

$$\mathbb{E}[q_i \sigma] = \sum_{x : \prod_{i=1}^N X_i} \left((qx)_i \cdot \prod_{i=1}^N \mathbb{P}[\sigma_i = x_i] \right) : \mathbb{R}$$

where we use the fact that the X_i are finite to ensure convergence. The notation $q_i \sigma$ means the application of the non-stochastic function q_i to the random variable σ . Formally, the functor \mathcal{D} is acting on q_i (see §2.1.3).

A **mixed strategy Nash equilibrium** is a mixed strategy profile in which no player can increase her expected utility by unilaterally deviating to some

²¹Put another way, a player is indifferent between outcomes that differ only in the utility of other players. This formally encodes the assumption that players are *selfish*.

²²Recall that we have an overriding assumption that sets of moves are nonempty (§1.2.2).

other move²³. Equivalently, a mixed strategy Nash equilibrium is a fixpoint of the best response function

$$\mathbf{B} : \prod_{i=1}^N \mathcal{D}X_i \rightarrow \mathcal{P} \prod_{i=1}^N \mathcal{D}X_i$$

given by

$$\mathbf{B}\sigma = \left\{ \sigma' : \prod_{i=1}^N \mathcal{D}X_i \mid \mathbb{E}[q_i(\sigma'_i, \sigma_{-i})] \geq \mathbb{E}[q_i(x'_i, \sigma_{-i})] \right. \\ \left. \text{for all } 1 \leq i \leq N \text{ and } x'_i : X_i \right\}$$

Nash's famous existence theorem, namely that mixed strategy equilibria always exist under only the finiteness condition we have given, is proved in (Nash 1951) by applying the Kakutani fixpoint theorem (Kakutani 1941) to this function \mathbf{B} . A generalisation of this method for certain higher order games was given in (Hedges 2013). However, there is no known sense in which we can consider mixed equilibria of an arbitrary higher order game. This will be called the **problem of mixed extensions**. For example, the passage from the deterministic to the probabilistic arg max operator appears not to be categorically natural, and for other selection functions it is unclear what the probabilistic equivalent should even be.

1.2.7 Voting games

We are now going to introduce an extended example from (Hedges et al. 2015b), which illustrates the use of selection functions that are different from arg max, and argues that selection equilibrium (as supposed to quantifier equilibrium) is the *correct* solution concept for higher order games in normal form.

We will consider an election with three voters and two candidates. The set of candidates is $X = \{a, b\}$. We will consider the election as a game in which the three voters simultaneously²⁴ choose a candidate. The outcome of the game is precisely the candidate who received the most votes. Thus, we take the set of outcomes to also be X , and the outcome function to be the majority²⁵ function $q : X^3 \rightarrow X$.

We will denote the three voters by $i = 1, 2, 3$. The preferences of the voters are defined by multi-valued selection functions $E_i : (X \rightarrow X) \rightarrow \mathcal{P}X$, or by the multi-valued quantifiers \overline{E}_i . We will investigate how the selection and quantifier equilibria vary as we choose different combinations of selection functions.

²³The reason we consider only unilateral deviations to moves, rather than strategies, is that the unilateral continuations are multilinear, and hence are maximised at a vertex of the simplex of strategies, which correspond to moves (or degenerate strategies). It is unclear whether this generalises to more complicated situations than mixed strategies, such as when strategies can mutate a state variable, or to behaviours other than maximisation.

²⁴The term *simultaneous* does not refer to physical time, but to the knowledge of the players at the time of choosing. Elections in the real world are simultaneous in this sense.

²⁵The choice of three voters and two candidates was made precisely so that the majority function is well-defined without needing a tie-breaking rule.

If player i is *rational*, then by definition she has a rational preference relation²⁶ \succeq_i on X . Since X contains two elements, there are precisely three rational preference relations, where respectively a is preferred to b , b is preferred to a , and both are equally preferred. The last case, when player i is indifferent between a and b , is described²⁷ by the multi-valued selection function

$$E_{a \approx b} k = X$$

that always returns the set of all moves. For the remaining two cases, we can fix an ordering $a \succeq b$ and write the corresponding selection functions as $\arg \max$ and $\arg \min$ with respect to this ordering. We can get all possible rational behaviours by choosing E_1 , E_2 and E_3 from $\arg \max$, $\arg \min$ and $E_{a \approx b}$.

As a first example, take the three selection functions to be $E_1 = \arg \max$, $E_2 = \arg \min$ and $E_3 = E_{a \approx b}$, so the first judge prefers a , the second prefers b and the third is indifferent. Consider a strategy profile $\sigma = (\sigma_1, \sigma_2, \sigma_3) : X \times X \times X$. For each of these, we can calculate for $i = 1, 2, 3$ the unilateral contexts $\mathcal{U}_i^q \sigma : X \rightarrow X$, namely

$$\mathcal{U}_1^q \sigma x_1 = q(x_1, \sigma_2, \sigma_3)$$

$$\mathcal{U}_2^q \sigma x_2 = q(\sigma_1, x_2, \sigma_3)$$

$$\mathcal{U}_3^q \sigma x_3 = q(\sigma_1, \sigma_2, x_3)$$

Using the definition in §1.2.3, we see that σ is a selection equilibria iff

$$\sigma_1 \in \arg \max(\mathcal{U}_1^q \sigma) = \arg \max_{x_1: X} q(x_1, \sigma_2, \sigma_3)$$

$$\sigma_2 \in \arg \min(\mathcal{U}_2^q \sigma) = \arg \min_{x_2: X} q(\sigma_1, x_2, \sigma_3)$$

$$\sigma_3 \in E_{a \approx b}(\mathcal{U}_3^q \sigma) = X$$

where the third condition is trivial.

Since there are only 8 strategy profiles, we can simply enumerate them and calculate the equilibria by brute force. We will begin by checking that the obvious strategies (a, b, a) and (a, b, b) , where each player votes for her preferred candidate, are selection equilibria. In this case the unilateral contexts are

$$\mathcal{U}_1^q(a, b, a)x_1 = q(x_1, b, a) = x_1$$

$$\mathcal{U}_2^q(a, b, a)x_2 = q(a, x_2, a) = a$$

Applying the selection functions gives

$$\arg \max(\mathcal{U}_1^q(a, b, a)) = \arg \max_{x_1: X} x_1 = \{a\}$$

$$\arg \min(\mathcal{U}_2^q(a, b, a)) = \arg \min_{x_2: X} a = \{a, b\}$$

Since $a \in \arg \max(\mathcal{U}_1^q(a, b, a))$ and $b \in \arg \min(\mathcal{U}_2^q(a, b, a))$, we have a selection equilibrium. We can similarly verify that (a, b, b) is a selection equilibrium.

²⁶ A total preorder.

²⁷ $E_{a \approx b}$ is the $\arg \max$ operator for the rational preference relation with $a \approx b$, because every element of the order is maximal in every context.

1.2.8 Modelling with selection functions

However, this game has additional equilibria that are less plausible. Consider for example the strategy (a, a, a) . Classically, this is an equilibrium because the second player, who is the only player who might have incentive to deviate, is in fact not pivotal and so cannot increase her utility with any deviation. This is also captured by the selection equilibrium, because the second player's unilateral context is the constant function

$$\mathcal{U}_2^q(a, a, a)x_2 = q(a, x_2, a) = a$$

The minimum value that $q(a, x_2, a)$ can take is a , and this minimum is attained at $x_2 = a, b$. Thus we have $\arg \min(\mathcal{U}_2^q(a, a, a)) = \{a, b\}$, and since $a \in \arg \min(\mathcal{U}_2^q(a, a, a))$ we see that player 2 has no incentive to unilaterally deviate.

If we as a modeller decide that this is implausible, one classical method²⁸ is to modify the preferences of the players. In the initial formulation, the preferences of player 2 are given by a rational preference relation $a \preceq_2 b$ on the type of outcomes $R = X$. However, we can pass from this to a rational preference relation \preceq'_2 on the type of plays $X \times X \times X$, given by

$$(x_1, x_2, x_3) \preceq'_2 (x'_1, x'_2, x'_3) \iff q(x_1, x_2, x_3) \preceq_2 q(x'_1, x'_2, x'_3)$$

Thus, a play is preferred iff it leads to a preferred outcome. In this preference order, we have implausible equivalences such as $(a, a, a) \approx'_2 (a, b, a)$. The solution is to manually modify the preference relation. For example, we could replace \preceq'_2 with a lexicographic relation \preceq''_2 which first takes preferred outcomes, and then if the outcomes are equivalent, uses the preference on moves $b \succeq_2 a$. This models a player who will always vote b , except in the strange context²⁹ in which voting a leads b to win the contest, and voting b leads a to win.

When using selection functions, the idiomatic way to approach this problem is to use selection refinements (§1.1.7). Specifically, we refine player 2's multi-valued selection function $\arg \min : (X \rightarrow X) \rightarrow \mathcal{P}X$ to the single-valued selection function

$$\varepsilon_2 k = \begin{cases} b & \text{if } b \in \arg \min k \\ a & \text{otherwise} \end{cases}$$

Much more detail about this is given in (Hedges et al. 2015b). A particular advantage of higher order game theory that can be exploited is that we never need to assume that the notion of rationality implemented by a selection function is equivalent to maximising over a rational preference relation, so rational and non-rational behaviour can be treated on an equal footing.³⁰

²⁸Another method that is used in practice is simply to ignore equilibria that are considered implausible, making the 'model' a mixture of formal game theory and informal economic intuition.

²⁹This is *tactical voting*, and can never arise in the toy model that we are considering.

³⁰It is often straightforward to 'rationalise' a non-rational selection function by enlarging the set of outcomes and modifying the outcome function, with the consequence that the advantages of selection functions are often only methodological. It is interesting to compare this to the discussion of similar 'rationalisations' of bounded rationality models in (Spiegler 2011).

1.2.9 Coordination and differentiation

The main contribution of (Hedges et al. 2015b) is the demonstration that coordinating and differentiating behaviour of players in the voting game can be modelled by selection functions that do not arise from rational preference relations, specifically fixpoint and non-fixpoint operators. Specifically, coordinating behaviour is modelled by the selection function $E_{\text{fix}} : (X \rightarrow X) \rightarrow \mathcal{P}X$ given by

$$E_{\text{fix}}k = \begin{cases} \{x : X \mid x = kx\} & \text{if nonempty} \\ X & \text{otherwise} \end{cases}$$

and differentiating behaviour by

$$E_{\text{nonfix}}k = \begin{cases} \{x : X \mid x \neq kx\} & \text{if nonempty} \\ X & \text{otherwise} \end{cases}$$

Notice that $\{x : X \mid x \neq kx\}$ is empty iff k is the identity function $X \rightarrow X$.

In the scenario of the voting game, a player modelled by E_{fix} aims to *vote for the winner*, and as such they are a **Keynesian agent**, after (Keynes 1936), chapter 12. Keynes' example of coordinating preferences in economics is as follows: suppose for simplicity that the financial success of a company is dependent only on the number of investors it attracts. An investor, therefore, has no preferences over different companies, but rather must aim to go with the majority of other investors. Similarly, a player modelled by E_{nonfix} is a *punk*, whose only aim is to vote against the majority.

These selection functions are total variants of their refinements

$$E'_{\text{fix}}k = \{x : X \mid x = kx\}$$

$$E'_{\text{nonfix}}k = \{x : X \mid x \neq kx\}$$

which are not total (see §1.1.6). The intuition behind returning X in the exceptional case, rather than \emptyset or some arbitrary default value such as $\{a\}$, is that in this case the player is indifferent and hence should be satisfied with any choice. In practice this has been checked by brute force, by showing that these exact definitions of E_{fix} and E_{nonfix} agree with the informal specification of a player who aims to vote with or against the majority.

Unlike $\arg \max$ and $\arg \min$, the selection functions E_{fix} and E_{nonfix} are not closed (see §1.1.7). To give a specific counterexample, consider the context $k : X \rightarrow X$ given by $kx = a$. The fixpoints of k are $E_{\text{fix}}k = \{a\}$. Then

$$\overline{E_{\text{fix}}k} = \{kx \mid x \in \{a\}\} = \{a\}$$

and

$$\overline{\overline{E_{\text{fix}}k}} = \{x : X \mid kx \in \{a\}\} = \{a, b\}$$

Thus $E_{\text{fix}}k \neq \overline{\overline{E_{\text{fix}}k}}$. Similarly,

$$E_{\text{nonfix}}k = \{b\} \neq \{a, b\} = \overline{\overline{E_{\text{nonfix}}k}}$$

This relates to the choices not being pivotal in the constant context (see §1.2.3).

1.2.10 Illustrating the solution concepts

As a second example, we will take $E_1 = E_{\text{fix}}$, $E_2 = E_{\text{nonfix}}$ and $E_3 = \arg \max$. Thus we have a Keynesian agent, a punk, and a rational player who prefers a . This game has 4 selection equilibria, namely

$$(a, a, b), (a, b, a), (b, b, a), (b, b, b)$$

Of these, (a, a, b) and (b, b, b) are implausible and can be ruled out by refining $\arg \max$ as described in §1.2.8. With (a, b, a) all three voters have achieved their aim: the first is in the majority, the second is in the minority and the third has her preferred outcome. In (b, b, a) player 2 is unable to be in the minority because she is pivotal, and so she is satisfied with the choice of b . Player 2's unilateral context is

$$\mathcal{U}_2^q(b, b, a)x_2 = q(b, x_2, a) = x_2$$

and $\mathcal{U}_2^q(b, b, a)$ has no non-fixpoints, and so $E_{\text{nonfix}}(\mathcal{U}_2^q(b, b, a)) = \{a, b\}$ by definition. This selection equilibrium could be ruled out if desired out by refining E_{nonfix} to return the actual set of non-fixpoints, even if empty.

Now, suppose we define the preferences instead by the multi-valued quantifiers $\overline{E}_1, \overline{E}_2, \overline{E}_3$. The four selection equilibria remain as quantifier equilibria (see §1.2.3), but there are two additional quantifier equilibria which are not selection equilibria, namely (a, a, a) and (b, a, a) . We will focus on the latter as an example. The unilateral context for the first player with this strategy profile is

$$\mathcal{U}_1^q(b, a, a)x_1 = q(x_1, a, a) = a$$

The strategy profile is not a selection equilibrium because player 1's choice is not a fixpoint, in other words,

$$b \notin \{a\} = E_{\text{fix}}(\mathcal{U}_1^q(b, a, a))$$

However, the quantifier equilibrium uses outcomes rather than choices, and the outcome is a fixpoint:

$$q(b, a, a) = a \in \{a\} = \overline{E}_{\text{fix}}(\mathcal{U}_1^q(b, a, a))$$

According to the definition in §1.2.3, taking $\Phi_i = \overline{E}_i$, this makes (b, a, a) a quantifier equilibrium. Alternatively, this can be seen as player 1 not being pivotal (§1.2.3) and deviating from the more plausible equilibrium (a, a, a) .

For a third example, we will take $E_1 = E_2 = E_3 = E_{\text{fix}}$, giving a game in which no players have preferences over the candidates but all aim to vote with the majority.³¹ Intuitively the equilibria of this game should be (a, a, a) and (b, b, b) , and indeed these are exactly the selection equilibria. However, every strategy profile of this game is a quantifier equilibrium. This is essentially a tautology: no matter how the players vote, the majority choice is in the majority. This is an example where the selection equilibrium makes a useful prediction that agrees with intuition, but the quantifier equilibrium makes no prediction.

³¹This scenario is the closest to the original Keynes beauty contest.

For a final example, take $E_1 = E_2 = E_3 = E_{\text{nonfix}}$. This models a population consisting entirely of punks³². Of course, in any particular play a majority of the punks will always fail in their aim to be in the minority. The selection equilibria of this game are precisely the ‘maximally differentiated’ strategy profiles, namely all strategy profiles except for (a, a, a) and (b, b, b) . Once again, every strategy profile is a quantifier equilibrium, because

$$\overline{E_{\text{nonfix}}}k = \begin{cases} \{kx \mid x \neq kx\} & \text{if nonempty} \\ \text{Im } k & \text{otherwise} \end{cases}$$

If we take for example player 1 with the strategy profile (a, a, a) , the unilateral context is $\mathcal{U}_1^q(a, a, a)x_1 = a$, and so

$$\overline{E_{\text{nonfix}}}(\mathcal{U}_1^q(a, a, a)) = \{\mathcal{U}_1^q(a, a, a)b\} = \{q(b, a, a)\} = \{a\}$$

This gives us a second example in which the selection equilibrium makes a useful and intuitive prediction, whereas the quantifier equilibrium makes no prediction.

On the basis of these examples, we choose to use selection equilibria rather than quantifier equilibria as the default solution concept for higher order games. This is essentially the argument of (Hedges et al. 2015b). This approach is more general than (Escardó and Oliva 2011), which (naively) takes the quantifier equilibrium as its solution concept.

1.3 Sequential games

1.3.1 Discussion

The theory of sequential games defined by selection functions was the starting point of higher order game theory. It was first developed in (Escardó and Oliva 2010a), with the presentation in this section based on (Escardó and Oliva 2011). In particular it is not the work of the author, but is included for completeness.

The type of single-valued selection functions, introduced in §1.1.4, carries the structure of a strong monad. Moreover the type of single-valued quantifiers carries a different strong monad structure, namely the continuation monad (see §1.1.3). The key result about higher order games is that the monad operations on these types are compatible with the game-theoretic interpretation, and in particular there is a monoidal product operator that implements the backward induction algorithm (and which moreover extends to unbounded games), which is the standard method in game theory to calculate subgame perfect equilibria of games of perfect information.

We will use the term **sequential game** to refer to games of perfect information with the added restriction that the player choosing at each stage, and the set of moves available to her, is independent of the previously chosen moves. This rules out standard examples such as the market entry game.³³

In a sequential game the players are ordered, with each player being able to observe the choices made by previous players before making her own choice. We

³²Neil Ghani has pointed out that, in fact, *punks don't vote*.

³³Generally speaking, type theoretic approaches to game theory have difficulty with this sort of example. Multi-agent influence diagrams (Koller and Milch 2003), for example, cannot represent the market entry game. This exact issue is the focus of (Botta et al. 2013).

can easily adapt the unilateral continuations from §1.2.2 in order to define Nash equilibria for sequential games. However, sequential games can have implausible Nash equilibria in which players make so-called incredible threats, in which a player can rule out branches of the game tree by contingently making a mutually destructive move. Subgame perfect equilibria rule this out, by forcing play to be rational even in branches that are not reached by playing the strategy. Every sequential game with finitely many players has at least one subgame perfect equilibrium, a result that is proved constructively using backward induction.

The method in this chapter of computing subgame perfect equilibria using the product of selection functions generalises to games with infinitely many players, so long as the outcome function is topologically continuous. These ‘unbounded sequential games’ are crucial to the application of selection functions in proof theory (Powell 2013, Oliva and Powell 2014; 2015), and are discussed from a game-theoretical point of view in (Escardó and Oliva 2012). However, backward induction for unbounded games relies crucially on subtle aspects of higher type computability, namely bar recursion, and for simplicity in this thesis we will focus only on games with finitely many players. See also the discussion in §2.2.1 of infinitely repeated games as coalgebras.³⁴

1.3.2 The category of selection functions

Just as quantifiers form the continuation monad \mathcal{K}_R , so selection functions form the **selection monad** \mathcal{J}_R where

$$\mathcal{J}_R X = (X \rightarrow R) \rightarrow X$$

Explicitly, the unit maps $\eta : X \rightarrow \mathcal{J}_R X$ are given by

$$\eta x = \lambda(k : X \rightarrow R).x$$

and the Kleisli extension of $f : X \rightarrow \mathcal{J}_R Y$ is

$$f^* \varepsilon = \lambda(k : Y \rightarrow R).f(\varepsilon \lambda(x : X).k(f x k))k$$

It is difficult to give a direct proof that \mathcal{J}_R is a monad, and the proof in (Escardó and Oliva 2010a), which also applies to the selection monad transformer in (Hedges 2014), uses the fact that the monad laws are equivalent to the category axioms for the Kleisli category. This proof is simple enough to reproduce here.

Given a locally small category \mathcal{C} and a fixed object R , we will define a new category $\mathcal{J}_R \mathcal{C}$ called a **selection category**. The objects of $\mathcal{J}_R \mathcal{C}$ are exactly the objects of \mathcal{C} , and the hom-sets are given by

$$\text{hom}_{\mathcal{J}_R \mathcal{C}}(X, Y) = \text{hom}_{\mathcal{C}}(Y, R) \rightarrow \text{hom}_{\mathcal{C}}(X, Y)$$

In particular, using the fact that **Set** is enriched in itself gives

$$\text{hom}_{\mathcal{J}_R \mathbf{Set}}(X, Y) = (Y \rightarrow R) \rightarrow (X \rightarrow Y) = X \rightarrow \mathcal{J}_R Y$$

³⁴The relationship between corecursion and bar recursion seems to be an open question. One possibility is that corecursion provides the appropriate general method to verify that a strategy is an equilibrium, whereas bar recursion or products of selection functions provides the appropriate general method to compute equilibria from scratch. Progress towards answering this question was recently made in (Capretta and Uustalu 2016) with the notion of a ‘barred coalgebra’.

The identity on X is given by

$$\text{id}_X k = \text{id}_X$$

where the left hand id_X is the identity in $\text{hom}_{\mathcal{J}_R\mathcal{C}}(X, X)$, and the right hand is in $\text{hom}_{\mathcal{C}}(X, X)$. The composition of $\varepsilon : \text{hom}_{\mathcal{C}}(Y, R) \rightarrow \text{hom}_{\mathcal{C}}(X, Y)$ and $\delta : \text{hom}_{\mathcal{C}}(Z, R) \rightarrow \text{hom}_{\mathcal{C}}(Y, Z)$ is given by

$$(\delta \circ \varepsilon)k = \delta k \circ \varepsilon(k \circ \delta k)$$

The category axioms can be proved very compactly. The left unit law is

$$(\text{id}_Y \circ \varepsilon)k = \text{id}_Y k \circ \varepsilon(k \circ \text{id}_Y k) = \text{id}_Y \circ \varepsilon(k \circ \text{id}_Y) = \varepsilon k$$

The right unit law is

$$(\varepsilon \circ \text{id}_X)k = \varepsilon k \circ \text{id}_X(k \circ \varepsilon k) = \varepsilon k \circ \text{id}_X = \varepsilon k$$

and the associativity law is

$$\begin{aligned} (\gamma \circ (\delta \circ \varepsilon))k &= \gamma k \circ (\delta \circ \varepsilon)(k \circ \gamma k) \\ &= \gamma k \circ \delta(k \circ \gamma k) \circ \varepsilon(k \circ \gamma k \circ \delta(k \circ \gamma k)) \\ &= (\gamma \circ \delta)k \circ \varepsilon(k \circ (\gamma \circ \delta)k) \\ &= ((\gamma \circ \delta) \circ \varepsilon)k \end{aligned}$$

The selection category $\mathcal{J}_R\mathcal{C}$ will appear again in §3.3.5. $\mathcal{J}_R\mathcal{C}$ can also be given a premonoidal product, which we will now introduce.

1.3.3 The product of selection functions

In §1.1.4 we introduced single-valued selection functions as elements of the type

$$\varepsilon : \mathcal{J}_R X = (X \rightarrow R) \rightarrow X$$

In §1.3.2 we proved that \mathcal{J}_R is a monad. In (Escardó and Oliva 2010a) it is moreover proved that \mathcal{J}_R is a strong monad. In particular this means that it is a monoidal monad, in the sense that it has a premonoidal³⁵ product operator

$$\times : \mathcal{J}_R X \times \mathcal{J}_R Y \rightarrow \mathcal{J}_R(X \times Y)$$

This operator is called the binary **product of selection functions**, and is explicitly defined by

$$(\varepsilon \times \delta)k = (a, bx)$$

where

$$\begin{aligned} a &= \varepsilon \lambda(x : X).k(x, bx) \\ bx &= \delta \lambda(y : Y).k(x, y) \end{aligned}$$

This can be folded to any finite number of selection functions $\varepsilon_j : \mathcal{J}_R X_j$, by

$$\bigotimes_{j=i}^N \varepsilon_j = \varepsilon_i \times \bigotimes_{j=i+1}^N \varepsilon_j$$

with the base case $\bigotimes_{j=N}^N \varepsilon_j = \varepsilon_N$.

³⁵The product of selection functions is written \otimes elsewhere in the literature, but we will write it as \times to emphasise that it is noncommutative, and avoid confusion with the tensor product of open games introduced in §2.2.6, which is very different. This notation for a premonoidal product is from (Power and Robinson 1993).

1.3.4 Sequential games

Since our players now choose sequentially, we will number them $1, \dots, N$. Just as for normal form games (§1.2.2), a **sequential game** consists of choice types X_i , and an outcome function

$$q : \prod_{i=1}^N X_i \rightarrow R$$

where R is a type of outcomes.

We will, similarly, model the players by multi-valued selection functions

$$E_i : (X_i \rightarrow R) \rightarrow \mathcal{P}X_i$$

This is the role taken by single-valued quantifiers in (Escardó and Oliva 2010a) and by multi-valued quantifiers in (Escardó and Oliva 2011). Since the same reasoning about coordination and differentiation games applies as in §1.2.10 and (Hedges et al. 2015b), we will use multi-valued selection functions rather than quantifiers.

The first difference between sequential and normal form games is that strategies become nontrivial. When player i makes her choice, she can directly observe the choices made by all players $j < i$. Thus, player i 's strategy is a function

$$\sigma_i : \prod_{j < i} X_j \rightarrow X_i$$

and the type of strategy profiles is

$$\Sigma = \prod_{i=1}^N \left(\prod_{j < i} X_j \rightarrow X_i \right)$$

A play, however, is still a tuple in $\prod_{i=1}^N X_i$.

Given a strategy, we can play it to obtain its **strategic play**. Notice that if $\sigma : \Sigma$ then σ_1 is simply a choice $\sigma_1 : X_1$, because the tuple $\prod_{j < 1} X_j$ is the unit type. Similarly $\sigma_2 : X_1 \rightarrow X_2$, so we can apply σ_2 to σ_1 to obtain the second player's choice, $\sigma_2 \sigma_1 : X_2$. This can be extended by course-of-values recursion:

$$\begin{aligned} \mathbf{P} : \Sigma &\rightarrow \prod_{i=1}^N X_i \\ (\mathbf{P}\sigma)_i &= \sigma_i((\mathbf{P}\sigma)_1, \dots, (\mathbf{P}\sigma)_{i-1}) \end{aligned}$$

This is our second example of a play function (the first being the identity function in §1.2.2), and is again a special case of the play functions introduced in §2.1.4.

1.3.5 Subgame perfection

A **partial play** is defined to be a tuple

$$x_1, \dots, x_{i-1} : \prod_{j < i} X_j$$

for some $1 \leq i \leq N$. Notice that the empty sequence is considered a partial play, but a play is not considered a partial play. Partial plays are in bijection with **subgames**, which are games that can be obtained by replacing some initial segment of players by fixed choices.

Given a strategy profile σ , we can extend a partial play to a play called its **strategic extension** by σ . This is similar to a unilateral deviation, except that instead of a single player i deviating, all players $j \leq i$ deviate. In the strategic extension ν_x^σ of $x = x_1, \dots, x_{i-1}$ by σ , the first $i-1$ players are forced to use the partial play, but subsequent players use σ . The definition is by the course-of-values recursion

$$(\nu_x^\sigma)_j = \begin{cases} x_j & \text{if } j < i \\ \sigma_j((\nu_x^\sigma)_1, \dots, (\nu_x^\sigma)_{j-1}) & \text{if } j \geq i \end{cases}$$

If $i = 1$ then x is the (unique) partial play of length 0, in which case $\mathbf{P}\sigma = \nu^\sigma$.

We will now modify the unilateral continuations introduced in §1.2.2 for use with sequential games. Instead of simply having a single function

$$\mathcal{U}_i^q : \Sigma \rightarrow X_i \rightarrow R$$

for the i th player's unilateral continuation, we must instead have one for each partial play of length $i-1$. This will be defined by

$$\mathcal{U}_x^q \sigma x_i = q(\nu_{x, x_i}^\sigma)$$

for $x = x_1, \dots, x_{i-1}$.

A **subgame perfect equilibrium** is a strategy profile σ that is an equilibrium in every subgame. That is, for every partial play $x = x_1, \dots, x_{i-1}$, the next move played by σ , namely $\sigma_i x$, is rational according to the selection function E_i , in the context in which player i unilaterally deviates in the subgame induced by x . Explicitly, σ must satisfy the conditions

$$\sigma_i x \in E_i(\mathcal{U}_x^q \sigma)$$

for all partial plays $x = x_1, \dots, x_{i-1}$.

Equivalently, a subgame perfect equilibrium is a fixpoint of the best response function $\mathbf{B} : \Sigma \rightarrow \mathcal{P}\Sigma$, where

$$\mathbf{B}\sigma = \{\sigma' : \Sigma \mid \sigma'_i x \in E_i(\mathcal{U}_x^q \sigma) \text{ for all } x = x_1, \dots, x_{i-1}\}$$

1.3.6 Backward induction

We will now prove the theorem that deserves to be called the *fundamental theorem of higher order game theory*, that the product of selection functions computes a play that is rational according to subgame perfect equilibrium.

Suppose each E_i is refined by a single-valued selection function $\varepsilon_i : \mathcal{I}_R X_i$, in the sense of §1.1.7, so $\varepsilon_i k \in E_i k$ for all $k : X_i \rightarrow R$. Then

$$\left(\bigotimes_{i=1}^N \varepsilon_i \right) q$$

is the strategic play of a subgame perfect equilibrium.

1.3. SEQUENTIAL GAMES

This has been proven in (Escardó and Oliva 2010a; 2011), but we will give the proof again here, because the technique is important and will be used again in §3.2.4 and §3.3.4. Specifically, we will prove that the strategy profile σ defined by the course-of-values recursion

$$\sigma_i(x_1, \dots, x_{i-1}) = \varepsilon_i(\mathcal{U}_{x_1, \dots, x_{i-1}}^q \sigma)$$

is subgame perfect, and that its strategic play is $(\bigotimes_{i=1}^N \varepsilon_i) q$. Subgame perfection is immediate, because for every partial play x_1, \dots, x_{i-1} we have

$$\sigma_i(x_1, \dots, x_{i-1}) = \varepsilon_i(\mathcal{U}_{x_1, \dots, x_{i-1}}^q \sigma) \in E_i(\mathcal{U}_{x_1, \dots, x_{i-1}}^q \sigma)$$

For the second part, we will prove the stronger fact that for any partial play x_1, \dots, x_{i-1} ,

$$\left(\bigotimes_{j=i}^N \varepsilon_j \right) q_{x_1, \dots, x_{i-1}} = (\nu_{x_1, \dots, x_{i-1}}^\sigma)_{j=i}^N$$

by strong induction on $N - i$, where

$$q_{x_1, \dots, x_{i-1}} : \prod_{j=i}^N X_j \rightarrow R$$

is defined by

$$q_{x_1, \dots, x_{i-1}}(x_i, \dots, x_N) = q(x_1, \dots, x_{i-1}, x_i, \dots, x_N)$$

The original claim follows, because the empty partial play has the property

$$(\nu^\sigma)_{j=1}^N = \nu^\sigma = \mathbf{P}\sigma$$

This is the characteristic of the proof technique we are using: to prove a property of a sequential game, work by bar induction on the tree of subgames. For games with finitely many players, this is equivalent to fixing a partial play of length $i - 1$ and working by strong induction on $N - i$, ending eventually with the empty partial play, whose induced subgame is the game itself.

In the base case, we have a partial play x_1, \dots, x_{N-1} , and

$$\left(\bigotimes_{j=N}^N \varepsilon_j \right) q_{x_1, \dots, x_{N-1}} = \varepsilon_N q_{x_1, \dots, x_{N-1}} = (\nu_{x_1, \dots, x_{N-1}}^\sigma)_{j=N}^N$$

because by the definition of $(\nu_x^\sigma)_j$ in §1.3.5,

$$\begin{aligned} (\nu_{x_1, \dots, x_{N-1}}^\sigma)_N &= \sigma_N((\nu_{x_1, \dots, x_{N-1}}^\sigma)_1, \dots, (\nu_{x_1, \dots, x_{N-1}}^\sigma)_{N-1}) \\ &= \sigma_N(x_1, \dots, x_{N-1}) \end{aligned}$$

By the definition of σ this is $\varepsilon_N(\mathcal{U}_{x_1, \dots, x_{N-1}}^q \sigma)$. Finally, we have

$$\mathcal{U}_{x_1, \dots, x_{N-1}}^q \sigma = q_{x_1, \dots, x_{N-1}}$$

because for all $x_N : X_N$,

$$\begin{aligned} \mathcal{U}_{x_1, \dots, x_{N-1}}^q \sigma x_N &= q(\nu_{x_1, \dots, x_N}^\sigma) \\ &= q((\nu_{x_1, \dots, x_N}^\sigma)_1, \dots, (\nu_{x_1, \dots, x_N}^\sigma)_N) \\ &= q(x_1, \dots, x_N) \\ &= q_{x_1, \dots, x_{N-1}} x_N \end{aligned}$$

This completes the base case of the proof.

1.3.7 The inductive step

For the inductive step, we take a partial play x_1, \dots, x_{i-1} . Unfolding the product of selection functions once gives

$$\left(\bigotimes_{j=i}^N \varepsilon_j \right) q_{x_1, \dots, x_{i-1}} = \left(\varepsilon_i \times \bigotimes_{j=i+1}^N \varepsilon_j \right) q_{x_1, \dots, x_{i-1}} = (a, ba)$$

where

$$\begin{aligned} a &= \varepsilon_i \lambda(x_i : X_i) \cdot q_{x_1, \dots, x_{i-1}, x_i}(bx_i) \\ bx_i &= \left(\bigotimes_{j=i+1}^N \varepsilon_j \right) q_{x_1, \dots, x_{i-1}, x_i} \end{aligned}$$

The inductive hypothesis gives us

$$bx_i = (\nu_{x_1, \dots, x_{i-1}, x_i}^\sigma)_{j=i+1}^N$$

for all $x_i : X_i$. We must prove that

$$(a, ba) = (\nu_{x_1, \dots, x_{i-1}}^\sigma)_{j=i}^N$$

which is to say that

$$\begin{aligned} a &= (\nu_{x_1, \dots, x_{i-1}}^\sigma)_i \\ ba &= (\nu_{x_1, \dots, x_{i-1}}^\sigma)_{j=i+1}^N \end{aligned}$$

We will first prove that for all $x_i : X_i$,

$$\mathcal{U}_{x_1, \dots, x_{i-1}}^q \sigma x_i = q_{x_1, \dots, x_{i-1}, x_i}(bx_i)$$

The left hand side by definition is

$$\mathcal{U}_{x_1, \dots, x_{i-1}}^q \sigma x_i = q(\nu_{x_1, \dots, x_{i-1}, x_i}^\sigma)$$

and the right hand side, by the inductive hypothesis, is

$$q_{x_1, \dots, x_{i-1}, x_i}(bx_i) = q_{x_1, \dots, x_{i-1}, x_i}((\nu_{x_1, \dots, x_{i-1}, x_i}^\sigma)_{j=i+1}^N)$$

and these are equal because

$$\begin{aligned} q_{x_1, \dots, x_{i-1}, x_i}((\nu_{x_1, \dots, x_{i-1}, x_i}^\sigma)_{j=i+1}^N) &= q(x_1, \dots, x_{i-1}, x_i, (\nu_{x_1, \dots, x_{i-1}, x_i}^\sigma)_{j=i+1}^N) \\ &= q(\nu_{x_1, \dots, x_{i-1}, x_i}^\sigma) \end{aligned}$$

1.3. SEQUENTIAL GAMES

From this we immediately get

$$\begin{aligned}
 (\nu_{x_1, \dots, x_{i-1}}^\sigma)_i &= \sigma_i((\nu_{x_1, \dots, x_{i-1}}^\sigma)_1, \dots, (\nu_{x_1, \dots, x_{i-1}}^\sigma)_{i-1}) \\
 &= \sigma_i(x_1, \dots, x_{i-1}) \\
 &= \varepsilon_i(\mathcal{Q}_{x_1, \dots, x_{i-1}}^q \sigma) \\
 &= \varepsilon_i \lambda(x_i : X_i) \cdot q_{x_1, \dots, x_{i-1}, x_i}(bx_i) \\
 &= a
 \end{aligned}$$

which is the first of the two conditions to be proved.

In order to prove the second condition

$$ba = (\nu_{x_1, \dots, x_{i-1}}^\sigma)_{j=i+1}^N$$

we use the inductive hypothesis to get

$$ba = (\nu_{x_1, \dots, x_{i-1}, a}^\sigma)_{j=i+1}^N$$

and so we will prove by strong induction on $j \geq i$ that

$$(\nu_{x_1, \dots, x_{i-1}}^\sigma)_j = (\nu_{x_1, \dots, x_{i-1}, a}^\sigma)_j$$

(Notice that this also trivially holds for $j < i$.) In the base case we have

$$(\nu_{x_1, \dots, x_{i-1}, a}^\sigma)_i = a = (\nu_{x_1, \dots, x_{i-1}}^\sigma)_i$$

by the previous result. For the inductive step,

$$\begin{aligned}
 (\nu_{x_1, \dots, x_{i-1}, a}^\sigma)_j &= \sigma_j((\nu_{x_1, \dots, x_{i-1}, a}^\sigma)_1, \dots, (\nu_{x_1, \dots, x_{i-1}, a}^\sigma)_{j-1}) \\
 &= \sigma_j((\nu_{x_1, \dots, x_{i-1}}^\sigma)_1, \dots, (\nu_{x_1, \dots, x_{i-1}}^\sigma)_{j-1}) \\
 &= (\nu_{x_1, \dots, x_{i-1}}^\sigma)_j
 \end{aligned}$$

This completes the entire proof.

Chapter 2

The algebra and geometry of games

2.1 Open games

2.1.1 Discussion

In this section we define open games, our objects of study. The barrier to entry is high: this section is quite abstract, and although we focus on examples and intuitions inline with the theory, we must wait until §2.3.9 before giving even the most trivial examples from game theory textbooks. Moreover the definitions we develop here bear no apparent relation whatsoever to the usual, familiar definitions of game due to von Neumann (Leyton-Brown and Shoham 2008), or the higher order games introduced in chapter 1, and a large part of this thesis will be spent developing the reader’s intuition for how these objects behave and how they should be used in mathematical modelling. To some extent the investment in abstraction will not pay off in the scope of this thesis, which is setting the groundwork for serious applications.

There are two key pieces of intuition that should be understood before beginning this section. The first is that a game should be seen as a *process*.¹ In particular a game should be a process that maps observations to choices: a game should input whatever information is observable, compute a decision and then output that decision. We remain agnostic about what is meant by ‘computation’ using the familiar technique of modelling a computation as a morphism of a suitable axiomatically-defined category. In this informal games-as-processes description we can already see that a game may consist of an *aggregate* of players: the game-process may distribute its input to a collection of sub-processes, each of which has only partial access to the true input, and which run in parallel to compute choices.

The idea of viewing games as processes, and hence as morphisms of a suitable category, first appears (to the author’s knowledge) in (Pavlovic 2009). Our

¹This view may already be surprising to readers familiar with game semantics, in which games are objects and (relative) winning strategies are morphisms. See for example (Hyland 1997).

implementation, however, is quite different.² With the theory here we are able to take the logical next step and view sequential and parallel play of games as sequential and parallel composition of processes.

The second key piece of intuition, which is shared with higher order game theory and which we reiterate from §1.1.1, is that all of our computation is done relative to a delimited continuation. A game is a process that runs in a calling environment, and after a process has terminated with a decision the calling environment will use that decision to compute an outcome for the process. This computation taking decisions to outcomes is the process' continuation, and it is one of the most fundamental principles of game theory that this computation is known to the players³, hence the processes run in continuation passing style. We will uniformly use the term 'continuation' to refer to what is variously called a 'utility function', 'payoff function', 'payoff matrix'⁴ or 'outcome function'. However continuations are more general, and include for example contexts in which other players are making rational choices. This ability to reduce multi-player situations to single-player situations, by abstracting away the other players into something akin to an outcome function, is a crucial ingredient of compositional game theory. The outcomes-as-continuations view strongly informed the definitions in this section, and other parts of the theory.

The entirety of chapter 2 is essentially based on (Ghani and Hedges 2016).

2.1.2 The underlying model of computation

We will now introduce our model of computation: **symmetric monoidal categories**. The definitions in this section will be made with respect to an arbitrary symmetric monoidal category, which will be instantiated with particular examples in later sections. This should be seen in the context of premonoidal categories and the computational λ -calculus (Moggi 1989), restricted to commutative side effects.

We will begin with a category with finite products, which for simplicity we will take to be the category **Set** of sets and functions. Given a monad $T : \mathbf{Set} \rightarrow \mathbf{Set}$, we can form the Kleisli category $\mathbf{Kl}(T)$, whose objects are sets and whose morphisms are Kleisli arrows,

$$\mathrm{hom}_{\mathbf{Kl}(T)}(X, Y) = X \rightarrow TY$$

The unit morphisms of $\mathbf{Kl}(T)$ are the units of T , and the composition of morphisms in $\mathbf{Kl}(T)$ is given by Kleisli extension.

A **monoidal monad** is a monad T equipped with a product operator

$$\otimes : TX \times TY \rightarrow T(X \times Y)$$

natural in X and Y . If the diagram

²In that paper the 'processes' are closer to our continuations, while the act of computing a choice is delegated to a fixpoint operator.

³The failure of this is called 'incomplete information', and is crucial to many applications of game theory because it is not a reasonable assumption about the real world. This is not in the scope of this thesis, however.

⁴A payoff matrix represents a linear payoff function, via the linearity of expectation. An unrelated usage of the same term appears in dialogical semantics, in which the 'matrix', or quantifier-free part of a prenexed formula, is another example of a continuation of a game.

$$\begin{array}{ccc}
 TX \times TY & \xrightarrow{\otimes} & T(X \times Y) \\
 \downarrow \sigma_{TX, TY} & & \downarrow T\sigma_{X, Y} \\
 TY \times TX & \xrightarrow{\otimes} & T(Y \times X)
 \end{array}$$

commutes then T is called a **commutative monad**, and this can be used to make $\text{Kl}(T)$ into a symmetric monoidal category (Power and Robinson 1993). The unit object of $\text{Kl}(T)$ is $I = 1$, the set with one element, and the object $X \otimes Y$ is the cartesian product $X \times Y$.

Categories of the form $\text{Kl}(T)$ have some additional structure we will need, namely that the monoidal product is a **weak product**: we have canonical deleting morphisms

$$!_X : \text{hom}_{\text{Kl}(T)}(X, I)$$

and copying morphisms

$$\Delta_X : \text{hom}_{\text{Kl}(T)}(X, X \otimes X)$$

given by composing the unique maps in **Set** with the units of T . These are not true products, but uniqueness is the only condition that fails.⁵ As a consequence, we also get canonical projections

$$\pi_1 : X \otimes Y \xrightarrow{X \otimes !_Y} X \otimes I \xrightarrow{\rho_X} X$$

and

$$\pi_2 : X \otimes Y \xrightarrow{!_X \otimes Y} I \otimes Y \xrightarrow{\lambda_Y} Y$$

The justification and intuition for considering such a category as a ‘model of computation’ is described in §0.7. For the remainder of this chapter, \mathcal{C} is going to refer to an arbitrary category with weak products. One example is **Set**, which translates into game-theoretic terms as pure strategies. Another example is **SRel**, which gives mixed strategies and which we will not describe.

2.1.3 The category of stochastic relations

We begin with the **finitary probability distribution monad** $\mathcal{D} : \text{Set} \rightarrow \text{Set}$. The underlying functor acts on sets by

$$\mathcal{D}X = \left\{ \alpha : X \rightarrow [0, 1] \mid \text{supp}(\alpha) \text{ is finite, } \sum_{x \in \text{supp}(\alpha)} \alpha x = 1 \right\}$$

where $\text{supp}(\alpha) = \{x : X \mid \alpha x \neq 0\}$. We consider values⁶ $\alpha : \mathcal{D}X$ as random variables of type X , and use the notation $\mathbb{P}[\alpha = x]$ for $\alpha x : [0, 1]$. The action

⁵In the language of functional programming, we can delete or copy data in many ways, but there is only one canonical way, namely the way that doesn’t use side effects.

⁶See the statement of Moggi’s thesis in §0.7.

on morphisms $f : X \rightarrow Y$ is given by

$$\mathbb{P}[\mathcal{D}f\alpha = y] = \sum_{\substack{x \in \text{supp}(\alpha) \\ fx=y}} \mathbb{P}[\alpha = x]$$

This makes \mathcal{D} into a functor.

The monad unit $\delta : X \rightarrow \mathcal{D}X$ creates unit mass distributions,

$$\mathbb{P}[\delta x = x'] = \begin{cases} 1 & \text{if } x = x' \\ 0 & \text{otherwise} \end{cases}$$

The Kleisli extension of a function $f : X \rightarrow \mathcal{D}Y$ is given by

$$\mathbb{P}[f^*\alpha = y] = \sum_{x \in \text{supp}(\alpha)} \mathbb{P}[fx = y]$$

This makes \mathcal{D} into a monad on **Set**. Finally, we have a monoidal product $\otimes : \mathcal{D}X \times \mathcal{D}Y \rightarrow \mathcal{D}(X \times Y)$ given by

$$\mathbb{P}[\alpha \otimes \beta = (x, y)] = \mathbb{P}[\alpha = x] \cdot \mathbb{P}[\beta = y]$$

With these operations, \mathcal{D} is a commutative monad on **Set**, and so its Kleisli category is symmetric monoidal. We will refer to the Kleisli category as **SRel**, the category of sets and **stochastic relations**. This is a variant of the usual notion of stochastic relations, which is defined in (Panangaden 1999) using subprobability rather than probability distributions. The distribution monad used here is the same one used in the study of convex sets (Fritz 2009, Jacobs et al. 2015), and is a finitary⁷ version of the Giry monad on the category of measure spaces (Giry 1982). Although subprobability distributions might be interesting in this setting, allowing choices to fail with some probability, we will use true probability distributions to remain close to Nash's original assumptions.

2.1.4 Open games

We can now give the definition of an open game. A **type** of an open game is of the form

$$\mathcal{G} : (X, S) \rightarrow (Y, R)$$

where X, Y, R, S are objects of \mathcal{C} . For now this is purely formal; later we will make open games into the morphisms of a category whose objects are pairs. In §2.3.5 we will introduce an alternative notation for this, namely

$$\mathcal{G} : X \otimes S^* \rightarrow Y \otimes R^*$$

Each of the four objects of \mathcal{C} should be read in a different way. We view X as the type of **observations** that can be made by \mathcal{G} , and Y as the type of **moves** or **choices**⁸, hence a game maps observations to choices. The types R

⁷Recall that Nash's existence theorem uses only finite sets of strategies, so this is not a restriction. We would need to use the Giry monad if we were following the generalisation of Nash's theorem to compact sets of strategies in (Glicksberg 1952).

⁸We continue to use these terms interchangeably.

and S , on the other hand, are ‘dual’ or ‘contravariant’ types. A useful intuition is that a type appearing in the right component of a pair represents a type whose elements are elements of R in the future, and which a rational player is reasoning about. We view R as the type of **outcomes** of \mathcal{G} , that is, the type of values about which the players in \mathcal{G} have preferences. The type S , dually, represents outcomes that are ‘generated’ by \mathcal{G} and returned to the calling environment.

We now finally arrive at the central definition of this thesis, that of an **open game**, which we will often simply call a **game**. An open game of type

$$\mathcal{G} : (X, S) \rightarrow (Y, R)$$

is, by definition, a 4-tuple

$$\mathcal{G} = (\Sigma_{\mathcal{G}}, \mathbf{P}_{\mathcal{G}}, \mathbf{C}_{\mathcal{G}}, \mathbf{B}_{\mathcal{G}})$$

where

- $\Sigma_{\mathcal{G}}$ is a set, called the set of **strategy profiles** of \mathcal{G}
- $\mathbf{P}_{\mathcal{G}} : \Sigma_{\mathcal{G}} \rightarrow \text{hom}_{\mathcal{C}}(X, Y)$ is called the **play function** of \mathcal{G}
- $\mathbf{C}_{\mathcal{G}} : \Sigma_{\mathcal{G}} \rightarrow \text{hom}_{\mathcal{C}}(X \otimes R, S)$ is called the **coplay function** of \mathcal{G}
- $\mathbf{B}_{\mathcal{G}} : \text{hom}_{\mathcal{C}}(I, X) \times \text{hom}_{\mathcal{C}}(Y, R) \rightarrow \Sigma_{\mathcal{G}} \rightarrow \mathcal{P}\Sigma_{\mathcal{G}}$ is the **best response function** of \mathcal{G}

In general we impose no conditions whatsoever on these components. In practice, however, we will restrict to games which are freely generated by the constructions considered in this section and §2.2, which gives some implicit restrictions.⁹

The most straightforward parts of this definition are the first two components. It is intuitive that a game has a set of strategy profiles and that, given a strategy profile and an observation, we can run the strategy profile on the observation to obtain a choice. To give a simple concrete example with $\mathcal{C} = \mathbf{Set}$, suppose $Y = A \times B$, and define

$$\Sigma = (X \rightarrow A) \times (X \times A \rightarrow B)$$

and

$$\mathbf{P}(\sigma_1, \sigma_2)x = (\sigma_1 x, \sigma_2(x, \sigma_1 x))$$

This represents a two-player game of perfect information¹⁰: first the value x is input, then the first player observes this and chooses a , and then the second player observes both x and a and chooses b , and finally the pair (a, b) is output.

Probably the most mysterious part of the definition is the coplay function. The basic idea is that the coplay function takes a utility in the future and transforms it to a utility less far in the future. It is completely unclear, however, why this should depend on an observation and a strategy. The only explanation that will be given is that the dependence on the observation is used only to

⁹Readers who care about foundations can be reassured that there will only be set-many freely generated games of a fixed type, out of class-many arbitrary games.

¹⁰However, the input x that is observed by both players is not expressible in classical game theory.

define the counit game in §2.1.9, which transforms an observation into a utility. The dependence on a strategy profile is a consequence of this, since to define coplay for an aggregate it is necessary to have access to the play functions of the components, which requires a strategy profile.¹¹

2.1.5 The best response function

The final part of the definition of an open game, and the part which needs the most explanation, is the best response function. In classical game theory the best response function of a game is a function $\mathbf{B} : \Sigma \rightarrow \mathcal{P}\Sigma$, where $\sigma' \in \mathbf{B}\sigma$ means that each player i would be satisfied to play her component σ'_i in response to the situation in which every other player j plays the strategy σ_j . When this happens we say that σ' is a ‘best response’ to σ . An equilibrium can be defined as a strategy which is a best response to itself, that is to say, a fixpoint of the multi-valued function \mathbf{B} . The best response functions in §1.2.4 and §1.3.5 behave in this way.

We can quite easily replace \mathbf{B}_G for open games with an ‘equilibrium set function’

$$\mathbf{E}_G : \text{hom}_{\mathcal{C}}(I, X) \times \text{hom}_{\mathcal{C}}(Y, R) \rightarrow \mathcal{P}\Sigma_G$$

by defining $\mathbf{E}_G(h, k)$ to be the set of fixpoints of $\mathbf{B}_G(h, k)$. This is the approach taken in (Ghani and Hedges 2016). If we do this, the remaining definitions in this section can be made correctly, without reference to best responses. However we choose to always carry around the best response function, which after all contains more information, for the reasons given in §1.2.4: it provides a method to prove existence theorems¹² (although see §3.3.7), and it gives a finer notion of equivalence between games (although see §2.2.1). If we were studying ‘abstract game theory’¹³ more seriously we could investigate the idea that the string diagrams in §2.3.5 can be interpreted in two¹⁴ different categories, and the best response function is forgotten by an identity-on-objects functor.

One of the fundamental ideas of this thesis (see §1.1.1 and §2.1.1) is that games should only be defined relative to a continuation. This leads to the idea of allowing a continuation as an additional parameter to the best response function, namely

$$\mathbf{B}_G : \text{hom}_{\mathcal{C}}(Y, R) \rightarrow \Sigma_G \rightarrow \mathcal{P}\Sigma_G$$

Here we see in full generality what is meant by the term ‘continuation’ in this thesis: it is nothing but a function from choices to outcomes.¹⁵ The importance of working relative to a truly *arbitrary* continuation, rather than an arbitrary-but-fixed continuation such as an outcome function, is that when we define $\mathbf{B}_G k \sigma$ in terms of the best response functions of its components, we will use

¹¹More succinctly, the definition is correct because it works (see the remainder of this thesis for details). The exact combination of definitions that works is very subtle and fragile, and is by far the most important contribution of this thesis.

¹²For example, one of the hypotheses of the Kakutani fixpoint theorem, used to prove the Nash existence theorem, is that our multivalued function should have convex values. Best response functions are indeed convex in the appropriate setting, but mere sets of equilibria are badly behaved, consisting only of isolated sets of points in typical examples.

¹³This phrase is from (Pavlovic 2009).

¹⁴The same idea could also work for the Nash and subgame-perfect variants of categorical composition, and possibly many other variant definitions.

¹⁵The simplicity should not diminish the importance of the outcomes-as-continuations idea.

continuations k' that actually depend on the strategy profile σ as well as on k . Thus, allowing continuations to vary is at least as important as allowing strategies to vary.

If we define the best response function in this way we again obtain a logically sound theory. However this definition contains a serious error, because nontrivial examples can only be written if the best response function is also allowed to depend on the observation made. A pair

$$(h, k) : \text{hom}_{\mathcal{C}}(I, X) \times \text{hom}_{\mathcal{C}}(Y, R)$$

will be called a **context**¹⁶ for the game \mathcal{G} , and we can roughly think of h and k as the past¹⁷ and future behaviour of the calling environment. There is another major subtlety here, namely that the best response function does not depend on a pure observation but on the *computation of an observation*, which may have side effects. If T is a commutative monad on the category of sets, and \mathcal{C} is its (symmetric monoidal) Kleisli category, then the observation is a value of type TX rather than X .¹⁸ For example if we have probabilistic choice as a side effect by taking $\mathcal{C} = \mathbf{SRel}$ (see §2.1.3) then we need to consider probability distributions over possible observations, because elements of $\text{hom}_{\mathbf{SRel}}(I, X)$ are random variables of type X . If we try to use pure observations instead, we find that we cannot define categorical composition in §2.2.2.

Finally, we come to a family of variants that are much more important. The definition of best response as a multivalued function and an equilibrium as a fixpoint is classical, and is sufficient for working with pure, mixed and nondeterministic strategies.¹⁹ This idea, however, does not stand up to more complex side effects such as learning.²⁰ If we write the multivalued function instead as

$$\mathbf{B}_{\mathcal{G}} : \text{hom}_{\mathcal{C}}(I, X) \times \text{hom}_{\mathcal{C}}(Y, R) \rightarrow \Sigma_{\mathcal{G}} \rightarrow \Sigma_{\mathcal{G}} \rightarrow \mathbb{B}$$

then the better definition is to replace the booleans \mathbb{B} with some other (in general noncommutative) algebraic structure B . If we only care about Nash equilibria then we can take B to be a monoid, however for reasoning about subgame perfection it seems that B should be some kind of ‘noncommutative complete semilattice’. A typical example would be of the form $B = T'\mathbb{B}$, where T' is another (strong, not necessarily commutative) monad on the category of sets, which can be different to T , the monad with $\mathcal{C} = \text{Kl}(T)$.²¹ Since it is still unclear what the appropriate definition should be we will use the simpler one by default.

¹⁶This generalises contexts in higher order game theory, which consist only of a continuation.

¹⁷The symbol h is a mnemonic for *history*.

¹⁸This is unsurprising when working with monoidal categories but striking when working in Haskell, where it is relatively unusual to see a monad in a contravariant position.

¹⁹The crucial property that these effects share is that the real numbers can be made into an algebra of the corresponding monad: identity for pure strategies, expectation for mixed strategies and supremum/infimum for nondeterministic strategies.

²⁰Since any nontrivial economic model will likely involve at least learning, the simpler definition is suitable for toy examples only.

²¹We can think of T as the effects allowed by the players, such as learning, and T' as the effects allowed by the economist, such as choosing a prior.

2.1.6 Closed games

The reason for the terminology *open game* introduced in §2.1.4 is that an open game is *open* to its environment, in the sense that information can pass back and forth. However, if one of X, Y, R, S in $\mathcal{G} : (X, S) \rightarrow (Y, R)$ is equal to the tensor unit I in \mathcal{C} , then the flow of information is restricted, and if all are equal to I then no information can pass between the game and its environment. A **closed game** is therefore defined²² to be an open game of type $\mathcal{G} : (I, I) \rightarrow (I, I)$.

Looking ahead to the result in §2.2.7, that open games are morphisms of a monoidal category whose objects are pairs of sets, and whose tensor unit is (I, I) , a closed game can equivalently be called an **abstract scalar**. This terminology is from (Abramsky and Coecke 2004, Abramsky 2005), and we will use it interchangeably with ‘closed game’. See also (Kelly and Laplaza 1980).

Closed games are particularly interesting when I is terminal in \mathcal{C} . The categories **Set** and **SRel** both have this property, although **Rel** (which models nondeterminism) does not. In this case, closed games have a very simple formulation: the play and coplay functions become trivial, and the game is described entirely by its set $\Sigma_{\mathcal{G}}$ of strategy profiles, and its best response function, whose type reduces to

$$\mathbf{B}_{\mathcal{G}} : \Sigma_{\mathcal{G}} \rightarrow \mathcal{P}\Sigma_{\mathcal{G}}$$

A game in the intuitive, informal sense can be described by this data (see §1.2.4), and correspondingly there is no notion of information being shared with the environment in classical game theory, which suggests that games in the usual sense should be thought of as closed games.

A closed game \mathcal{G} over any category has a canonical history and a canonical continuation, namely both being given by id_I . Thus we can give a solution concept for closed games: an **equilibrium** of \mathcal{G} is a strategy profile $\sigma : \Sigma_{\mathcal{G}}$ with $\sigma \in \mathbf{B}_{\mathcal{G}}(\text{id}_I, \text{id}_I)\sigma$. This solution concept will be used in §3.1 and §3.2, where we see that it includes pure and mixed Nash and subgame perfect equilibria as special cases.

2.1.7 Decisions

In this section we see our first examples of open games, the **decisions**²³, which are one-player games. Every nontrivial game contains decisions as components, representing the players of the game.

Decisions, by definition, are games of the form

$$\mathcal{D} : (X, I) \rightarrow (Y, R)$$

where the player makes an observation from X and makes a choice from Y , with preferences over outcomes in R . Formally, the requirements for a game with this type to be a decision are

- $\Sigma_{\mathcal{D}} = \text{hom}_{\mathcal{C}}(X, Y)$

²²Games are not doors, or topologies: every closed game is open, although it is reasonable to think of closed games as a degenerate case of open games.

²³This term comes from decision theory, which mathematically is the study of one-player games, but in practice is the applied study of modelling issues that are logically prior to game theory.

- $\mathbf{P}_{\mathcal{D}}\sigma = \sigma : \text{hom}_{\mathcal{C}}(X, Y)$
- $\mathbf{C}_{\mathcal{D}}\sigma = !_X \otimes R : \text{hom}_{\mathcal{C}}(X \otimes R, I)$
- $\mathbf{B}_{\mathcal{D}}(h, k) : \text{hom}_{\mathcal{C}}(X, Y) \rightarrow \mathcal{P} \text{hom}_{\mathcal{C}}(X, Y)$ is constant

Consequently, to define a decision is to define a function

$$\mathbf{R} : \text{hom}_{\mathcal{C}}(I, X) \times \text{hom}_{\mathcal{C}}(Y, R) \rightarrow \mathcal{P} \text{hom}_{\mathcal{C}}(X, Y)$$

Such a function will be called a **rationality function**, and defines the best response function of a decision by

$$\mathbf{B}_{\mathcal{D}}(h, k)\sigma = \mathbf{R}(h, k)$$

Since a decision always consists of a single player and never an aggregate, a strategy is simply a mapping from observations to outcomes. The fact that the strategy is a morphism of \mathcal{C} means that the strategy can have side effects, such as probabilistic choice (see §1.2.6). The condition on $\mathbf{P}_{\mathcal{D}}$ is simply that playing the strategy involves applying it as a function to the observation.

The only possibly unexpected condition is that on $\mathbf{B}_{\mathcal{D}}$. The idea of best responses (see §1.2.4) is that $\mathbf{B}(h, k)\sigma$ should be the set of strategy profiles σ' such that, for each player, playing using σ' is a rational response to the situation in which every other player is playing using σ . Since a decision involves only one player any property of ‘every other player’ is vacuous, and so σ is not used in the definition.²⁴

The definition of a decision is already at the right level of generality²⁵ that an arbitrary multi-valued selection function (§1.1.6) can be considered as a decision in **Game(Set)**. First, notice that a decision $\mathcal{D} : I \rightarrow Y \otimes R^*$ over $\mathcal{C} = \mathbf{Set}$ is defined by a rationality function

$$\mathbf{R} : (Y \rightarrow R) \rightarrow \mathcal{P}Y$$

which is precisely the type of a multi-valued selection function (see §1.1.6). More generally, a multi-valued selection function $E : (Y \rightarrow R) \rightarrow \mathcal{P}Y$ can be converted into a decision $\mathcal{D}_E : X \rightarrow Y \otimes R^*$ for an arbitrary set X , using the response function

$$\mathbf{R}(x, k) = \{\sigma : X \rightarrow Y \mid \sigma x \in Ek\}$$

The resulting decision models a player whose rationality is defined by E .

2.1.8 Preliminary examples of decisions

We will give two example families of decisions representing classically rational players, respectively with pure and mixed strategies. In the first case, we take $\mathcal{C} = \mathbf{Set}$. We will fix a rational preference relation²⁶ \preceq on a set R , representing

²⁴Put another way, in a one-player game there is nothing for a strategy to be a best response to.

²⁵One point of view is that the generalisation from utility maximisation to selection functions was a necessary step en-route to compositional game theory.

²⁶Recall that a rational preference relation is a total preorder, that is, a total order in which we may have $r \neq r'$ with $r \preceq r'$ and $r' \preceq r$. Thus \preceq can be viewed as a total order on equivalence classes.

a player's preferences for different outcomes in R , where $r \succeq r'$ means that the outcome r is considered at least as good as the outcome r' . We will also fix an arbitrary set X and a finite set Y . To define a decision

$$\mathcal{D} : (X, 1) \rightarrow (Y, R)$$

is to define a rationality function

$$\mathbf{R} : X \times (Y \rightarrow R) \rightarrow \mathcal{P}(X \rightarrow Y)$$

We define a particular family of decisions by the response functions

$$\mathbf{R}(x, k) = \{\sigma : X \rightarrow Y \mid k(\sigma x) \succeq ky \text{ for all } y : Y\}$$

We will now explicitly give the data defining \mathcal{D} . Its strategies are $\Sigma_{\mathcal{D}} = X \rightarrow Y$, and its play function $\mathbf{P}_{\mathcal{D}} : (X \rightarrow Y) \rightarrow (X \rightarrow Y)$ is $\mathbf{P}_{\mathcal{D}}\sigma x = \sigma x$. Its coplay function $\mathbf{C}_{\mathcal{D}} : (X \rightarrow Y) \rightarrow (X \times R \rightarrow 1)$ is $\mathbf{C}_{\mathcal{D}}(\sigma, (x, r)) = *$. Its best response function

$$\mathbf{B}_{\mathcal{D}} : X \times (Y \rightarrow R) \rightarrow ((X \rightarrow Y) \rightarrow \mathcal{P}(X \rightarrow Y))$$

is given by

$$\mathbf{B}_{\mathcal{D}}(x, k)\sigma = \{\sigma' : X \rightarrow Y \mid k(\sigma' x) \succeq ky \text{ for all } y : Y\}$$

For our second family of examples we will use $\mathcal{C} = \mathbf{SRel}$ (defined in §2.1.3). For an arbitrary set X and finite set Y we define a decision

$$\mathcal{D} : (X, 1) \rightarrow (Y, \mathbb{R})$$

modelling a player who maximises expected utility (see §1.2.6). To specify such a decision is to specify a response function

$$\mathbf{R} : \mathcal{D}X \times (Y \rightarrow \mathcal{D}\mathbb{R}) \rightarrow \mathcal{P}(X \rightarrow \mathcal{D}Y)$$

where \mathcal{D} is the distribution monad (see §2.1.3). We use the particular function

$$\mathbf{R}(h, k) = \{\sigma : X \rightarrow \mathcal{D}Y \mid \mathbb{E}[k^*(\sigma^*h)] \geq \mathbb{E}[ky] \text{ for all } y : Y\}$$

where $\mathbb{E} : \mathcal{D}\mathbb{R} \rightarrow \mathbb{R}$ is the expectation operator, and $-^*$ is the Kleisli extension of \mathcal{D} . Explicitly, we have the strategy set $\Sigma_{\mathcal{D}} = X \rightarrow \mathcal{D}Y$. The play function

$$\mathbf{P}_{\mathcal{D}} : (X \rightarrow \mathcal{D}Y) \rightarrow (X \rightarrow \mathcal{D}Y)$$

is given by $\mathbb{P}[\mathbf{P}_{\mathcal{D}}\sigma x = y] = \mathbb{P}[\sigma x = y]$. The coplay function

$$\mathbf{C}_{\mathcal{D}} : (X \rightarrow \mathcal{D}Y) \rightarrow (X \times \mathbb{R} \rightarrow \mathcal{D}1)$$

is given by $\mathbb{P}[\mathbf{C}_{\mathcal{D}}\sigma(x, r) = *] = 1$. The best response function

$$\mathbf{B}_{\mathcal{D}} : \mathcal{D}X \times (Y \rightarrow \mathcal{D}\mathbb{R}) \rightarrow (X \rightarrow \mathcal{D}Y) \rightarrow \mathcal{P}(X \rightarrow \mathcal{D}Y)$$

is given by

$$\mathbf{B}_{\mathcal{D}}(h, k)\sigma = \{\sigma' : X \rightarrow \mathcal{D}Y \mid \mathbb{E}[k^*(\sigma'^*h)] \geq \mathbb{E}[ky] \text{ for all } y : Y\}$$

2.1.9 Computations and counit

Besides decisions, the other atomic games we will consider are computations and counits. These components are unable to make strategic decisions, and in particular only have one strategy, which behaves trivially. Their purpose, however, is to control the information flow in a game. Examples of computations include outcome functions, identities, copying and deleting. The counits serve to connect forward-flowing to backward-flowing data, that is, identifying a particular forward-flowing value as the value that some player is reasoning about.

Formally, a **computation** is a game

$$(f, g) : (X, S) \rightarrow (Y, R)$$

defined by morphisms $f : \text{hom}_{\mathcal{C}}(X, Y)$ and $g : \text{hom}_{\mathcal{C}}(R, S)$. After we have defined a symmetric monoidal category of games in §2.2.7, this will give us a symmetric monoidal embedding of $\mathcal{C} \times \mathcal{C}^{\text{op}}$ (see §2.2.9). The definition is given by

- $\Sigma_{(f,g)} = 1 = \{*\}$
- $\mathbf{P}_{(f,g)}^* = f : \text{hom}_{\mathcal{C}}(X, Y)$
- $\mathbf{C}_{(f,g)}^* : \text{hom}_{\mathcal{C}}(X \otimes R, S)$ is given by the composition

$$X \otimes R \xrightarrow{\pi_2} R \xrightarrow{g} S$$

- $\mathbf{B}_{(f,g)}(h, k)^* = \{*\}$ for all h and k

The conditions $\Sigma_{(f,g)} = 1$ and $\mathbf{B}_{(f,g)}(h, k)^* = \{*\}$ implement the idea that (f, g) is ‘strategically trivial’. The fact that $* \in \mathbf{B}_{(f,g)}(h, k)^*$ means that $*$ is always an equilibrium of a computation. The idea behind this is that a strategy for an aggregate game should never *fail* to be an equilibrium because of a computation, because a computation has no preferences. Rather, if a strategy fails to be an equilibrium, it should always be because some *player* has incentive to deviate. We expand on this idea in §2.2.8.

The conditions on $\mathbf{P}_{(f,g)}$ and $\mathbf{C}_{(f,g)}$ determine the information flow, which is explained in more detail in §2.2.11. When $\mathcal{C} = \mathbf{Set}$, the coplay function of a computation is simply

$$\mathbf{C}_{(f,g)}^*(x, r) = gr$$

Next, for each object X of \mathcal{C} we define a **counit** $\varepsilon_X : (X, X) \rightarrow (I, I)$ by

- $\Sigma_{\varepsilon_X} = 1 = \{*\}$
- $\mathbf{P}_{\varepsilon_X}^* = !_X : \text{hom}_{\mathcal{C}}(X, I)$
- $\mathbf{C}_{\varepsilon_X}^* = \rho_X : \text{hom}_{\mathcal{C}}(X \otimes I, X)$
- $\mathbf{B}_{\varepsilon_X}(h, k)^* = \{*\}$ for all h and k

The intuition behind this definition is similar to that for computations. If $\mathcal{C} = \mathbf{Set}$, the coplay function is $\mathbf{C}_{\varepsilon_X}^*(x, *) = x$. Counits will not play a major role in this thesis, although see §2.2.10 and §2.3.6.

2.2 The category of games

2.2.1 Discussion

In §2.1 we introduced many definitions, especially of open games (§2.1.4), decisions (§2.1.7), computations and counits (§2.1.9). In this section we will add algebraic structure, giving two ways to compose games: categorical composition, which is a primitive form of sequential play of games, and tensor product, which is a primitive form of simultaneous play. We will prove that these obey the axioms of a symmetric monoidal category. In §2.3 this algebraic structure will be investigated further.

Readers who are not category theorists can treat these two operations as rather like sequential and parallel composition operators in process algebra²⁷, similar to those in (Blumensath and Winschel 2013), with the axioms we check simply being technical conditions for the string diagram language introduced in §2.3.5 to be sound.

On the other hand, category theorists may wonder about the canonicity of the operations, that is to say, why we use this particular symmetric monoidal category and not some other, and whether a mathematically cleaner, more axiomatic treatment is possible. The answer to this is simply no, our operations are not canonical. In §2.2.3, for example, we discuss alternative definitions of categorical composition, and in §3.3.7 it is suggested that the most basic definitions in §2.1.4 will need to be changed to deal with dynamic games with mixed strategies.

Nevertheless, one important direction that we will not consider in this thesis is to additionally consider morphisms between games, so that we have a monoidal bicategory (Schommer-Pries 2009). This is formally needed in several places where we unreasonably assume that the underlying category \mathcal{C} is strict monoidal. More importantly, when considering infinite games such as repeated games, a natural approach is to define the game as a terminal coalgebra of a functor which precomposes a finite approximation to the repeated game with one additional stage. Because a terminal coalgebra is defined by a universal property of objects, this requires 2-cells between games. There are several possible definitions that could be used, however, and at the present time it is not clear which are useful.²⁸ One possibility that would be elegant, but is only speculation, is that solutions of games (see §3.3.2) are global points, and that backward induction of solutions (§3.3.4) is an instance of horizontal composition of 2-cells. A bicategory of games would also provide a formal theory of equivalences and refinements between games, as discussed in §1.2.4.

2.2.2 Categorical composition of games

Now we come to our first aggregation operation, categorical composition, which models sequential play of games. Given a pair of games $\mathcal{G} : (X, T) \rightarrow (Y, S)$ and $\mathcal{H} : (Y, S) \rightarrow (Z, R)$, we need to define a composition $\mathcal{H} \circ \mathcal{G} : (X, T) \rightarrow (Z, R)$. This is a primitive form of composition in which the intermediate choice at Y is

²⁷An earlier, failed attempt at a compositional game theory by the author took this as a starting point and used series-parallel partial orders (Bechet et al. 2005) as part of the data specifying a game.

²⁸Neil Ghani is currently working on this problem.

2.2. THE CATEGORY OF GAMES

hidden; the more intuitive sequential composition that produces plays of type $X \otimes Y$ will be recovered using tensor product and identities²⁹ in §3.2.2. The game $\mathcal{H} \circ \mathcal{G}$ is an aggregate whose players consist of the players of \mathcal{G} together with the players of \mathcal{H} . This means that a strategy profile for $\mathcal{H} \circ \mathcal{G}$ should consist of a strategy profile for \mathcal{G} together with a strategy profile for \mathcal{H} , that is, $\Sigma_{\mathcal{H} \circ \mathcal{G}} = \Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}}$.

Since the play is sequential, the observation that is made by \mathcal{H} should be the choice that is made by \mathcal{G} . This motivates the definition of the play function of $\mathcal{H} \circ \mathcal{G}$, namely that $\mathbf{P}_{\mathcal{H} \circ \mathcal{G}}(\sigma, \tau)$ should be the composition

$$X \xrightarrow{\mathbf{P}_{\mathcal{G}}\sigma} Y \xrightarrow{\mathbf{P}_{\mathcal{H}}\tau} Z$$

Using the categorical composition in \mathcal{C} appropriately sequences the effects that can be used by the two components. For example if \mathcal{G} and \mathcal{H} both contain players who can make probabilistic choices, the play function of $\mathcal{H} \circ \mathcal{G}$ gives the appropriate probability distribution on Z taking into account the probability distribution on Y .

The coplay function of $\mathcal{H} \circ \mathcal{G}$, on the other hand, is hard or impossible to justify on intuitive grounds. We define $\mathbf{C}_{\mathcal{H} \circ \mathcal{G}}(\sigma, \tau)$ to be the composition

$$X \otimes R \xrightarrow{\Delta_{X \otimes R}} X \otimes X \otimes R \xrightarrow{X \otimes \mathbf{P}_{\mathcal{G}}\sigma \otimes R} X \otimes Y \otimes R \xrightarrow{X \otimes \mathbf{C}_{\mathcal{H}}\tau} X \otimes S \xrightarrow{\mathbf{C}_{\mathcal{G}}\sigma} T$$

(where we are assuming that \mathcal{C} is strict monoidal). When $\mathcal{C} = \mathbf{Set}$, this is given explicitly by

$$\mathbf{C}_{\mathcal{H} \circ \mathcal{G}}(\sigma, \tau)(x, r) = \mathbf{C}_{\mathcal{G}}\sigma(x, \mathbf{C}_{\mathcal{H}}\tau(\mathbf{P}_{\mathcal{G}}\sigma x, r))$$

This will be discussed in §2.2.11.

For now we will also give the best response function without motivation, but it will be discussed in detail in the next section. The definition is

$$\mathbf{B}_{\mathcal{H} \circ \mathcal{G}}(h, k)(\sigma, \tau) = \mathbf{B}_{\mathcal{G}}(h, k_{\tau \circ})\sigma \times \bigcap_{\sigma' : \Sigma_{\mathcal{G}}} \mathbf{B}_{\mathcal{H}}(\mathbf{P}_{\mathcal{G}}\sigma' \circ h, k)\tau$$

where $k_{\tau \circ}$ is the composition

$$Y \xrightarrow{\Delta_Y} Y \otimes Y \xrightarrow{Y \otimes \mathbf{P}_{\mathcal{H}}\tau} Y \otimes Z \xrightarrow{Y \otimes k} Y \otimes R \xrightarrow{\mathbf{C}_{\mathcal{H}}\tau} S$$

With this definition of composition, we obtain a category³⁰ $\mathbf{Game}(\mathcal{C})$ whose objects are pairs of objects of \mathcal{C} and whose morphisms are open games, where the identity on (X, R) is the computation $(\text{id}_X, \text{id}_R)$ and the composition is that given above. This will be proved in §2.2.4 and §2.2.5.

²⁹This is *splitting the atom of game theory*.

³⁰If we are being pedantic then composition is not associative because

$$\Sigma_{\mathcal{I} \circ (\mathcal{H} \circ \mathcal{G})} = \Sigma_{\mathcal{I}} \times (\Sigma_{\mathcal{H}} \times \Sigma_{\mathcal{G}}) \neq (\Sigma_{\mathcal{I}} \times \Sigma_{\mathcal{H}}) \times \Sigma_{\mathcal{G}} = \Sigma_{(\mathcal{I} \circ \mathcal{H}) \circ \mathcal{G}}$$

Rather, $\mathbf{Game}(\mathcal{C})$ should be a bicategory whose 2-cells are an appropriate sort of morphisms between games (see §2.2.1). However this is not a work on category theory, and it would be overcomplicated and unenlightening to use bicategories here, so we will simply *pretend* that the cartesian monoidal category of sets is strict monoidal.

2.2.3 Best response for sequential compositions

In order to give a definition of $\mathbf{B}_{\mathcal{H} \circ \mathcal{G}}(h, k)(\sigma, \tau)$, we need to define what it means for a strategy profile (σ', τ') to be a best response to (σ, τ) , in the context consisting of a history $h : \text{hom}_{\mathcal{C}}(I, X)$ and a continuation $k : \text{hom}_{\mathcal{C}}(Z, R)$.

The idea is that we have a fixed situation in which the strategy profile played is (σ, τ) , and then one player unilaterally deviates to a component of σ' or τ' . This means that if we have a player in \mathcal{G} deviating to a component of σ' we know that every player in \mathcal{H} is playing τ , and vice versa.

The first condition is that σ' should be a best response to σ , in a modified context in which we extend the continuation k backwards in time using the fact that we know that \mathcal{H} will be played using τ . Therefore we define³¹ the extended continuation $k_{\tau \circ}$ to be the composition

$$Y \xrightarrow{\Delta_Y} Y \otimes Y \xrightarrow{Y \otimes \mathbf{P}_{\mathcal{H}} \tau} Y \otimes Z \xrightarrow{Y \otimes k} Y \otimes R \xrightarrow{\mathbf{C}_{\mathcal{H}} \tau} S$$

For the second condition we need that τ' is a best response to τ in an appropriate context. The same continuation k can be used without modification, and the history can be extended forwards using the fact that \mathcal{G} will be played with σ , namely as the composition

$$I \xrightarrow{h} X \xrightarrow{\mathbf{P}_{\mathcal{G}} \sigma} Y$$

This leads to the possible definition

$$\mathbf{B}_{\mathcal{H} \circ \mathcal{G}}(h, k)(\sigma, \tau) = \left\{ (\sigma', \tau') : \Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}} \mid \begin{array}{l} \sigma' \in \mathbf{B}_{\mathcal{G}}(h, k_{\tau \circ})\sigma \text{ and} \\ \tau' \in \mathbf{B}_{\mathcal{H}}(\mathbf{P}_{\mathcal{G}} \sigma \circ h, k)\tau \end{array} \right\}$$

which can be written equivalently as

$$\mathbf{B}_{\mathcal{H} \circ \mathcal{G}}(h, k)(\sigma, \tau) = \mathbf{B}_{\mathcal{G}}(h, k_{\tau \circ})\sigma \times \mathbf{B}_{\mathcal{H}}(\mathbf{P}_{\mathcal{G}} \sigma \circ h, k)\tau$$

Here we face a significant design choice. If we made this definition then very few changes would need to be made to this thesis. Assuming that when a player in \mathcal{H} is deviating to τ' , the players in \mathcal{G} use σ , leads to a theory of games supporting a solution concept based on Nash equilibrium. However our games are dynamic, and it is widely recognised in game theory³² that Nash equilibrium is not a suitable solution concept for dynamic games. In particular, in a Nash equilibrium earlier players can make so-called ‘non-credible threats’ to later players.

The usual solution to this problem is to use subgame-perfect equilibria (see §1.3.5). This is an equilibrium refinement of Nash: every subgame-perfect equilibrium is a Nash equilibrium, but not vice versa.³³ In classical game theory, the usual method is to define ‘subgames’ as subtrees of the game tree, and define a subgame-perfect equilibrium to be a strategy which induces a Nash equilibrium when restricted to any subgame.

³¹This is the only place in which the coplay function is actually used. All of the machinery and the difficult interpretation of coplay begins with the fact that a function of the appropriate type is needed here.

³²This hardly needs a reference, but see for example (Leyton-Brown and Shoham 2008).

³³In particular, the Nash equilibria which are not subgame-perfect are precisely the ones containing non-credible threats.

One possibility, therefore, would be to extend the data specifying a game with a recursively-built collection of subgames, and quantify over this collection when defining the best-response function. However, a simpler solution is possible. An apparent alternative would be to quantify over the type of histories, and say that a best response should be valid for every possible history, not just the ones that arise from the strategy profile. For classical game theory this works³⁴, but for us it is too strong. Consider, for example, a game defined (over the category of sets, for simplicity) as a composition of the form

$$(X, I) \xrightarrow{(\Delta_X, \text{id}_I)} (X \otimes X, I) \xrightarrow{\mathcal{D}} (Y, R)$$

where \mathcal{D} is a decision. Now the earlier component does not contain a player and cannot make strategic decisions, so our intuition about threats is no longer valid. No matter what the starting history is, and no matter what this game is precomposed with, the player making the decisions can only ever observe histories of the form (x, x) . It is therefore too strong to require that the player's strategy should be rational for histories not of this form.

This example, however, suggests the solution: the image of the play function, as the strategy profile of the first component varies, gives precisely those histories for the second component that can possibly arise. Now we can return to the original problem: defining what it means for (σ', τ') to be a best response to (σ, τ) in the context (h, k) . The second condition becomes that τ' should be a best response to τ now in a variety of contexts: all those of the form (h', k) , where h' is of the form $\mathbf{P}_{\mathcal{G}}\sigma'' \circ h$, where σ'' is an arbitrary strategy profile for \mathcal{G} . Putting this together and rewriting it, we get the stated definition

$$\mathbf{B}_{\mathcal{H} \circ \mathcal{G}}(h, k)(\sigma, \tau) = \mathbf{B}_{\mathcal{G}}(h, k_{\tau \circ})\sigma \times \bigcap_{\sigma' : \Sigma_{\mathcal{G}}} \mathbf{B}_{\mathcal{H}}(\mathbf{P}_{\mathcal{G}}\sigma' \circ h, k)\tau$$

2.2.4 The identity laws

We will now prove that $\mathbf{Game}(\mathcal{C})$, as defined in §2.2.2, is a category, beginning with the identity laws. Let $\mathcal{G} : (X, S) \rightarrow (Y, R)$ be a game. We first prove that $(\text{id}_Y, \text{id}_R) \circ \mathcal{G} = \mathcal{G}$.

For the strategy sets,

$$\Sigma_{(\text{id}_Y, \text{id}_R) \circ \mathcal{G}} = \Sigma_{\mathcal{G}} \times \Sigma_{(\text{id}_Y, \text{id}_R)} = \Sigma_{\mathcal{G}} \times 1 = \Sigma_{\mathcal{G}}$$

For the play function,

$$\mathbf{P}_{(\text{id}_Y, \text{id}_R) \circ \mathcal{G}}\sigma = \mathbf{P}_{(\text{id}_Y, \text{id}_R)}^* \circ \mathbf{P}_{\mathcal{G}}\sigma = \text{id}_Y \circ \mathbf{P}_{\mathcal{G}}\sigma = \mathbf{P}_{\mathcal{G}}\sigma$$

For coplay, by definition $\mathbf{C}_{(\text{id}_Y, \text{id}_R) \circ \mathcal{G}}\sigma$ is the composition

$$\begin{aligned} X \otimes R &\xrightarrow{\Delta_X \otimes R} X \otimes X \otimes R \xrightarrow{X \otimes \mathbf{P}_{\mathcal{G}}\sigma \otimes R} X \otimes Y \otimes R \\ &\xrightarrow{X \otimes \mathbf{C}_{(\text{id}_Y, \text{id}_R)}^*} X \otimes R \xrightarrow{\mathbf{C}_{\mathcal{G}}\sigma} S \end{aligned}$$

³⁴In classical game theory a 'history' can be defined as a path starting at the root of the game tree, and such paths are in bijection with subtrees.

2.2. THE CATEGORY OF GAMES

which, expanding the definitions further, is

$$X \otimes R \xrightarrow{\Delta_X \otimes R} X \otimes X \otimes R \xrightarrow{X \otimes \mathbf{P}_G \sigma \otimes R} X \otimes Y \otimes R \xrightarrow{X \otimes \pi_2} X \otimes R \xrightarrow{\mathbf{C}_G \sigma} S$$

The first part of this is the identity on $X \otimes R$, so it is equal to $\mathbf{C}_G \sigma$.

For best response we have by definition

$$\mathbf{B}_{(\text{id}_Y, \text{id}_R) \circ \mathcal{G}}(h, k)\sigma = \mathbf{B}_G(h, k_{*o})\sigma \times \bigcap_{\sigma' : \Sigma_G} \mathbf{B}_{(\text{id}_Y, \text{id}_R)}(\mathbf{P}_G \sigma' \circ h, k)*$$

For each $\sigma' : \Sigma_G$,

$$\mathbf{B}_{(\text{id}_Y, \text{id}_R)}(\mathbf{P}_G \sigma' \circ h, k)* = \{*\}$$

therefore

$$\mathbf{B}_{(\text{id}_Y, \text{id}_R) \circ \mathcal{G}}(h, k)\sigma = \mathbf{B}_G(h, k_{*o})\sigma \times \{*\} = \mathbf{B}_G(h, k_{*o})\sigma$$

under the natural isomorphism $\Sigma_G \times 1 \cong \Sigma_G$. The continuation k_{*o} is given by

$$Y \xrightarrow{\Delta_Y} Y \otimes Y \xrightarrow{Y \otimes \mathbf{P}_{(\text{id}_Y, \text{id}_R)}^*} Y \otimes Y \xrightarrow{Y \otimes k} Y \otimes R \xrightarrow{\mathbf{C}_{(\text{id}_Y, \text{id}_R)}^*} R$$

Expanding the definitions and simplifying, this is

$$Y \xrightarrow{\Delta_Y} Y \otimes Y \xrightarrow{Y \otimes k} Y \otimes R \xrightarrow{\pi_2} R$$

which is equal to k .

For the other identity law, we will prove that $\mathcal{G} \circ (\text{id}_X, \text{id}_S) = \mathcal{G}$. For the strategy sets,

$$\Sigma_{\mathcal{G} \circ (\text{id}_X, \text{id}_S)} = \Sigma_{(\text{id}_X, \text{id}_S)} \times \Sigma_G = 1 \times \Sigma_G = \Sigma_G$$

For the play function,

$$\mathbf{P}_{\mathcal{G} \circ (\text{id}_X, \text{id}_S)} \sigma = \mathbf{P}_G \sigma \circ \mathbf{P}_{(\text{id}_X, \text{id}_S)} = \mathbf{P}_G \sigma \circ \text{id}_X = \mathbf{P}_G \sigma$$

For coplay we have that $\mathbf{C}_{\mathcal{G} \circ (\text{id}_X, \text{id}_S)} \sigma$ is the composition

$$\begin{aligned} X \otimes R &\xrightarrow{\Delta_X \otimes R} X \otimes X \otimes R \xrightarrow{X \otimes \mathbf{P}_{(\text{id}_X, \text{id}_S)}^* \otimes R} X \otimes X \otimes R \\ &\xrightarrow{X \otimes \mathbf{C}_G \sigma} X \otimes S \xrightarrow{\mathbf{C}_{(\text{id}_X, \text{id}_S)}^*} S \end{aligned}$$

Expanding and simplifying, this is

$$X \otimes R \xrightarrow{\Delta_X \otimes R} X \otimes X \otimes R \xrightarrow{X \otimes \mathbf{C}_G \sigma} X \otimes S \xrightarrow{\pi_2} S$$

which is equal to $\mathbf{C}_G \sigma$.

For best response we have

$$\begin{aligned} \mathbf{B}_{\mathcal{G} \circ (\text{id}_X, \text{id}_S)}(h, k)\sigma &= \mathbf{B}_{(\text{id}_X, \text{id}_S)}(h, k_{\sigma o}) * \times \bigcap_{\sigma' : \Sigma_{(\text{id}_X, \text{id}_S)}} \mathbf{B}_G(\mathbf{P}_{(\text{id}_X, \text{id}_S)} \sigma' \circ h, k)\sigma \\ &= \{*\} \times \mathbf{B}_G(\mathbf{P}_{(\text{id}_X, \text{id}_S)}^* \circ h, k)\sigma \\ &= \mathbf{B}_G(\mathbf{P}_{(\text{id}_X, \text{id}_S)}^* \circ h, k)\sigma \\ &= \mathbf{B}_G(\text{id}_X \circ h, k)\sigma \\ &= \mathbf{B}_G(h, k)\sigma \end{aligned}$$

In summary, we have proved that the identity for composition on an object (X, R) is $(\text{id}_X, \text{id}_R)$.

2.2.5 Associativity

Consider games

$$(W, U) \xrightarrow{\mathcal{G}} (X, T) \xrightarrow{\mathcal{H}} (Y, S) \xrightarrow{\mathcal{I}} (Z, R)$$

We will prove that

$$\mathcal{I} \circ (\mathcal{H} \circ \mathcal{G}) = (\mathcal{I} \circ \mathcal{H}) \circ \mathcal{G}$$

This associativity law is not obvious, especially because of the best response function of a composition.

The cases of the strategy sets and play functions are immediate from the associativity of cartesian product and composition:

$$\Sigma_{\mathcal{I} \circ (\mathcal{H} \circ \mathcal{G})} = \Sigma_{\mathcal{H} \circ \mathcal{G}} \times \Sigma_{\mathcal{I}} = \Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}} \times \Sigma_{\mathcal{I}} = \Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{I} \circ \mathcal{H}} = \Sigma_{(\mathcal{I} \circ \mathcal{H}) \circ \mathcal{G}}$$

and

$$\begin{aligned} \mathbf{P}_{\mathcal{I} \circ (\mathcal{H} \circ \mathcal{G})}(\sigma, \tau, \upsilon) &= \mathbf{P}_{\mathcal{I}\upsilon} \circ \mathbf{P}_{\mathcal{H} \circ \mathcal{G}}(\sigma, \tau) \\ &= \mathbf{P}_{\mathcal{I}\upsilon} \circ \mathbf{P}_{\mathcal{H}\tau} \circ \mathbf{P}_{\mathcal{G}}\sigma \\ &= \mathbf{P}_{\mathcal{I} \circ \mathcal{H}}(\tau, \upsilon) \circ \mathbf{P}_{\mathcal{G}}\sigma \\ &= \mathbf{P}_{(\mathcal{I} \circ \mathcal{H}) \circ \mathcal{G}}(\sigma, \tau, \upsilon) \end{aligned}$$

For coplay, by definition $\mathbf{C}_{\mathcal{I} \circ (\mathcal{H} \circ \mathcal{G})}(\sigma, \tau, \upsilon)$ is the composition

$$\begin{aligned} W \otimes R &\xrightarrow{\Delta_{W \otimes R}} W \otimes W \otimes R \xrightarrow{W \otimes \mathbf{P}_{\mathcal{H} \circ \mathcal{G}}(\sigma, \tau) \otimes R} W \otimes Y \otimes R \\ &\xrightarrow{W \otimes \mathbf{C}_{\mathcal{I}\upsilon}} W \otimes S \xrightarrow{\mathbf{C}_{\mathcal{H} \circ \mathcal{G}}(\sigma, \tau)} U \end{aligned}$$

which is

$$\begin{aligned} W \otimes R &\xrightarrow{\Delta_{W \otimes R}} W \otimes W \otimes R \xrightarrow{W \otimes \mathbf{P}_{\mathcal{G}}\sigma \otimes R} W \otimes X \otimes R \\ &\xrightarrow{W \otimes \mathbf{P}_{\mathcal{H}\tau} \otimes R} W \otimes Y \otimes R \xrightarrow{W \otimes \mathbf{C}_{\mathcal{I}\upsilon}} W \otimes S \xrightarrow{\Delta_{W \otimes S}} W \otimes W \otimes S \\ &\xrightarrow{W \otimes \mathbf{P}_{\mathcal{G}}\sigma \otimes S} W \otimes X \otimes S \xrightarrow{W \otimes \mathbf{C}_{\mathcal{H}\tau}} W \otimes T \xrightarrow{\mathbf{C}_{\mathcal{G}}\sigma} U \end{aligned}$$

On the other hand $\mathbf{C}_{(\mathcal{I} \circ \mathcal{H}) \circ \mathcal{G}}(\sigma, \tau, \upsilon)$ is

$$\begin{aligned} W \otimes R &\xrightarrow{\Delta_{W \otimes R}} W \otimes W \otimes R \xrightarrow{W \otimes \mathbf{P}_{\mathcal{G}}\sigma \otimes R} W \otimes X \otimes R \\ &\xrightarrow{W \otimes \mathbf{C}_{\mathcal{I} \circ \mathcal{H}}(\tau, \upsilon)} W \otimes T \xrightarrow{\mathbf{C}_{\mathcal{G}}\sigma} U \end{aligned}$$

which is

$$\begin{aligned} W \otimes R &\xrightarrow{\Delta_{W \otimes R}} W \otimes W \otimes R \xrightarrow{W \otimes \mathbf{P}_{\mathcal{G}}\sigma \otimes R} W \otimes X \otimes R \\ &\xrightarrow{W \otimes \Delta_X \otimes R} W \otimes X \otimes X \otimes R \xrightarrow{W \otimes X \otimes \mathbf{P}_{\mathcal{H}\tau} \otimes R} W \otimes X \otimes Y \otimes R \\ &\xrightarrow{W \otimes X \otimes \mathbf{C}_{\mathcal{I}\upsilon}} W \otimes X \otimes S \xrightarrow{W \otimes \mathbf{C}_{\mathcal{H}\tau}} W \otimes T \xrightarrow{\mathbf{C}_{\mathcal{I}\sigma}} U \end{aligned}$$

and these two morphisms are equal by the comonoid laws for Δ .³⁵ For another perspective on this part of the proof, see §2.2.11.

³⁵Since this is not a thesis on category theory, detailed proofs by diagram chasing will not be included for assertions like this that a complicated pair of morphisms are equal.

For best response we have

$$\begin{aligned}
 & \mathbf{B}_{\mathcal{I} \circ (\mathcal{H} \circ \mathcal{G})}(h, k)(\sigma, \tau, v) \\
 &= \mathbf{B}_{\mathcal{H} \circ \mathcal{G}}(h, k_{v \circ})(\sigma, \tau) \times \bigcap_{(\sigma', \tau') : \Sigma_{\mathcal{H} \circ \mathcal{G}}} \mathbf{B}_{\mathcal{I}}(\mathbf{P}_{\mathcal{H} \circ \mathcal{G}}(\sigma', \tau') \circ h, k)v \\
 &= \mathbf{B}_{\mathcal{G}}(h, (k_{v \circ})_{\tau \circ})\sigma \\
 & \quad \times \bigcap_{\sigma' : \Sigma_{\mathcal{G}}} \mathbf{B}_{\mathcal{H}}(\mathbf{P}_{\mathcal{G}}\sigma' \circ h, k_{v \circ})\tau \times \bigcap_{\substack{\sigma' : \Sigma_{\mathcal{G}} \\ \tau' : \Sigma_{\mathcal{H}}}} \mathbf{B}_{\mathcal{I}}(\mathbf{P}_{\mathcal{H}}\tau' \circ \mathbf{P}_{\mathcal{G}}\sigma' \circ h, k)v
 \end{aligned}$$

and

$$\begin{aligned}
 & \mathbf{B}_{(\mathcal{I} \circ \mathcal{H}) \circ \mathcal{G}}(h, k)(\sigma, \tau, v) \\
 &= \mathbf{B}_{\mathcal{G}}(h, k_{(\tau, v) \circ})\sigma \times \bigcap_{\sigma' : \Sigma_{\mathcal{G}}} \mathbf{B}_{\mathcal{I} \circ \mathcal{H}}(\mathbf{P}_{\mathcal{G}}\sigma' \circ h, k)(\tau, v) \\
 &= \mathbf{B}_{\mathcal{G}}(h, k_{(\tau, v) \circ})\sigma \\
 & \quad \times \bigcap_{\sigma' : \Sigma_{\mathcal{G}}} \left(\mathbf{B}_{\mathcal{H}}(\mathbf{P}_{\mathcal{G}}\sigma' \circ h, k_{v \circ})\tau \times \bigcap_{\tau' : \Sigma_{\mathcal{H}}} \mathbf{B}_{\mathcal{I}}(\mathbf{P}_{\mathcal{H}}\tau' \circ \mathbf{P}_{\mathcal{G}}\sigma' \circ h, k)v \right) \\
 &= \mathbf{B}_{\mathcal{G}}(h, k_{(\tau, v) \circ})\sigma \\
 & \quad \times \bigcap_{\sigma' : \Sigma_{\mathcal{G}}} \mathbf{B}_{\mathcal{H}}(\mathbf{P}_{\mathcal{G}}\sigma' \circ h, k_{v \circ})\tau \times \bigcap_{\substack{\sigma' : \Sigma_{\mathcal{G}} \\ \tau' : \Sigma_{\mathcal{H}}}} \mathbf{B}_{\mathcal{I}}(\mathbf{P}_{\mathcal{H}}\tau' \circ \mathbf{P}_{\mathcal{G}}\sigma' \circ h, k)v
 \end{aligned}$$

Here $k_{v \circ}$ is the composition

$$Y \xrightarrow{\Delta_Y} Y \otimes Y \xrightarrow{Y \otimes \mathbf{P}_{\mathcal{I}}v} Y \otimes Z \xrightarrow{Y \otimes k} Y \otimes R \xrightarrow{\mathbf{C}_{\mathcal{I}}v} S$$

and $(k_{v \circ})_{\tau \circ}$ is the composition

$$X \xrightarrow{\Delta_X} X \otimes X \xrightarrow{X \otimes \mathbf{P}_{\mathcal{H}}\tau} X \otimes Y \xrightarrow{X \otimes k_{v \circ}} X \otimes S \xrightarrow{\mathbf{C}_{\mathcal{H}}\tau} T$$

which expands to

$$\begin{aligned}
 & X \xrightarrow{\Delta_X} X \otimes X \xrightarrow{X \otimes \mathbf{P}_{\mathcal{H}}\tau} X \otimes Y \xrightarrow{X \otimes \Delta_Y} X \otimes Y \otimes Y \\
 & \xrightarrow{X \otimes Y \otimes \mathbf{P}_{\mathcal{I}}v} X \otimes Y \otimes Z \xrightarrow{X \otimes Y \otimes k} X \otimes Y \otimes R \xrightarrow{X \otimes \mathbf{C}_{\mathcal{I}}v} X \otimes S \xrightarrow{\mathbf{C}_{\mathcal{H}}\tau} T
 \end{aligned}$$

On the other hand $k_{(\tau, v) \circ}$ is the composition

$$X \xrightarrow{\Delta_X} X \otimes X \xrightarrow{X \otimes \mathbf{P}_{\mathcal{I} \circ \mathcal{H}}(\tau, v)} X \otimes Z \xrightarrow{X \otimes k} X \otimes R \xrightarrow{\mathbf{C}_{\mathcal{I} \circ \mathcal{H}}(\tau, v)} T$$

which expands to

$$\begin{aligned}
 & X \xrightarrow{\Delta_X} X \otimes X \xrightarrow{X \otimes \mathbf{P}_{\mathcal{H}}\tau} X \otimes Y \xrightarrow{X \otimes \mathbf{P}_{\mathcal{I}}v} X \otimes Z \xrightarrow{X \otimes k} X \otimes R \\
 & \xrightarrow{\Delta_X \otimes R} X \otimes X \otimes R \xrightarrow{X \otimes \mathbf{P}_{\mathcal{H}}\tau \otimes R} X \otimes Y \otimes R \xrightarrow{\mathbf{P}_{\mathcal{I}}v} X \otimes S \xrightarrow{\mathbf{P}_{\mathcal{H}}\tau} T
 \end{aligned}$$

Then $(k_{v \circ})_{\tau \circ} = k_{(\tau, v) \circ}$, and we are done.

2.2.6 Tensor product of games

Besides the categorical composition introduced in §2.2.2, the other aggregation operator we will consider is a tensor (or monoidal) product. Given an arbitrary pair of games

$$\mathcal{G} : (X_1, S_1) \rightarrow (Y_1, R_1)$$

and

$$\mathcal{H} : (X_2, S_2) \rightarrow (Y_2, R_2)$$

we can form their tensor product

$$\mathcal{G} \otimes \mathcal{H} : (X_1 \otimes X_2, S_1 \otimes S_2) \rightarrow (Y_1 \otimes Y_2, R_1 \otimes R_2)$$

Since the weak product \otimes of \mathcal{C} can be extended componentwise to pairs of objects, this will give us a monoidal product on $\mathbf{Game}(\mathcal{C})$, the category defined in §2.2.2. Whereas $\mathcal{H} \circ \mathcal{G}$ behaves like the sequential play of \mathcal{G} and \mathcal{H} , the purpose of $\mathcal{G} \otimes \mathcal{H}$ is to behave like the simultaneous play.

Just as for sequential composition, a strategy profile for an aggregate of two games consists of a strategy profile for each game, thus

$$\Sigma_{\mathcal{G} \otimes \mathcal{H}} = \Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}}$$

The play function can be defined using the tensor product of morphisms in \mathcal{C} , by

$$\mathbf{P}_{\mathcal{G} \otimes \mathcal{H}}(\sigma, \tau) = \mathbf{P}_{\mathcal{G}}\sigma \otimes \mathbf{P}_{\mathcal{H}}\tau$$

Similarly, the coplay function is

$$\mathbf{C}_{\mathcal{G} \otimes \mathcal{H}} : \Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}} \rightarrow \text{hom}_{\mathcal{C}}(X_1 \otimes X_2 \otimes R_1 \otimes R_2, S_1 \otimes S_2)$$

where $\mathbf{C}_{\mathcal{G} \otimes \mathcal{H}}(\sigma, \tau)$ is given by

$$X_1 \otimes X_2 \otimes R_1 \otimes R_2 \xrightarrow{\cong} X_1 \otimes R_1 \otimes X_2 \otimes R_2 \xrightarrow{\mathbf{C}_{\mathcal{G}}\sigma \otimes \mathbf{C}_{\mathcal{H}}\tau} S_1 \otimes S_2$$

To define the best response function

$$\begin{aligned} \mathbf{B}_{\mathcal{G} \otimes \mathcal{H}} : \text{hom}_{\mathcal{C}}(I, X_1 \otimes X_2) \times \text{hom}_{\mathcal{C}}(Y_1 \otimes Y_2, R_1 \otimes R_2) &\rightarrow \Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}} \\ &\rightarrow \mathcal{P}(\Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}}) \end{aligned}$$

essentially the same reasoning applies as in §2.2.3, except that we are now working with simultaneous games, and with selection equilibria (in the sense of §1.2.3), rather than subgame perfect equilibria. Given $h : \text{hom}_{\mathcal{C}}(I, X_1 \otimes X_2)$, we can take the canonical projections of it to form the component histories $h_1 : \text{hom}_{\mathcal{C}}(I, X_1)$ and $h_2 : \text{hom}_{\mathcal{C}}(I, X_2)$. More importantly, given a continuation $k : \text{hom}_{\mathcal{C}}(Y_1 \otimes Y_2, R_1 \otimes R_2)$, a history $h : \text{hom}_{\mathcal{C}}(I, X_1 \otimes X_2)$ and a strategy profile $(\sigma, \tau) : \Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}}$, we can form smaller continuations³⁶ in which each component may deviate from the strategy, while the other stays fixed. The continuation for the first player is denoted $k_{\otimes\tau(h_2)} : \text{hom}_{\mathcal{C}}(Y_1, R_1)$, and is defined by the composition

$$Y_1 \xrightarrow{\cong} Y_1 \otimes I \xrightarrow{Y_1 \otimes h_2} Y_1 \otimes X_2 \xrightarrow{Y_1 \otimes \mathbf{P}_{\mathcal{H}}\tau} Y_1 \otimes Y_2 \xrightarrow{k} R_1 \otimes R_2 \xrightarrow{\pi_1} R_1$$

³⁶These are generalised forms of the unilateral continuations introduced in §1.2.2.

Similarly, the continuation $k_{\sigma(h_1)\otimes} : \text{hom}_{\mathcal{C}}(Y_2, R_2)$ for the second player is

$$Y_2 \xrightarrow{\cong} I \otimes Y_2 \xrightarrow{h_1 \otimes Y_2} X_1 \otimes Y_2 \xrightarrow{\mathbf{P}_{\mathcal{G}} \sigma \otimes Y_2} Y_1 \otimes Y_2 \xrightarrow{k} R_1 \otimes R_2 \xrightarrow{\pi_2} R_2$$

Given these definitions, the best response function for the tensor product is given by

$$\mathbf{B}_{\mathcal{G} \otimes \mathcal{H}}(h, k)(\sigma, \tau) = \mathbf{B}_{\mathcal{G}}(h_1, k_{\otimes \tau(h_2)})\sigma \times \mathbf{B}_{\mathcal{H}}(h_2, k_{\sigma(h_1)\otimes})\tau$$

This definition will be refuted in §3.3.7, however.

2.2.7 The monoidal category axioms

We will now prove that $\mathbf{Game}(\mathcal{C})$ is a symmetric monoidal category. The action of the monoidal product on objects is to pairwise apply the monoidal product of \mathcal{C} , so

$$(X_1, R_1) \otimes (X_2, R_2) = (X_1 \otimes X_2, R_1 \otimes R_2)$$

The action on morphisms, which are games, is exactly the construction given in §2.2.6. The monoidal unit is (I, I) , where I is the monoidal unit of \mathcal{C} . The symmetric monoidal category axioms underly the string diagram language for $\mathbf{Game}(\mathcal{C})$ introduced in §2.3.5 and used in the remainder of this thesis.

In general, proving that a category is monoidal takes a significant amount of work. We must define three families of morphisms, the left and right unitors and the associators, prove their naturality, and then prove the commutativity of two diagrams including the Mac Lane pentagon. To prove that a monoidal category is symmetric monoidal we must additionally define the braiding morphisms, prove their naturality, and prove commutativity of an additional three diagrams.

On the assumption that \mathcal{C} is strict monoidal³⁷, most of this work can be avoided. In that case we have equalities of objects

$$(I, I) \otimes (X, R) = (I \otimes X, I \otimes R) = (X, R) = (X \otimes I, R \otimes I) = (X, R) \otimes (I, I)$$

and so we can take all of the unitors to be the identity morphisms (that is, computations formed of pairs of identities, see §2.2.2), which are automatically natural and satisfy the commutative diagrams, simply by the fact that $\mathbf{Game}(\mathcal{C})$ is a category. Similarly we have equalities

$$\begin{aligned} ((X_1, R_1) \otimes (X_2, R_2)) \otimes (X_3, R_3) &= (X_1 \otimes X_2 \otimes X_3, R_1 \otimes R_2 \otimes R_3) \\ &= (X_1, R_1) \otimes ((X_2, R_2) \otimes (X_3, R_3)) \end{aligned}$$

and so we can also take the associators to be identities.

For the braiding morphisms we take the computation

$$\begin{aligned} s_{(X_1, R_1), (X_2, R_2)} &= (s_{X_1, X_2}, s_{R_2, R_1}) \\ &: (X_1 \otimes X_2, R_1 \otimes R_2) \rightarrow (X_2 \otimes X_1, R_2 \otimes R_1) \end{aligned}$$

For a strict monoidal category the the unit law for a symmetric monoidal category becomes trivial, so we must prove only the associativity and inverse laws.

³⁷In practice, \mathcal{C} is never strict monoidal. The *correct* approach is to consider $\mathbf{Game}(\mathcal{C})$ as a monoidal bicategory (see §2.2.1), but for now we can either appeal to the Mac Lane strictification theorem (Mac Lane 1978), or ignore the problem.

2.2. THE CATEGORY OF GAMES

For this, we will use the result from §2.2.9, that computations respect composition and tensor. This is not circular, because we will only use the part of the result that does not already assume that $\mathbf{Game}(\mathcal{C})$ is a monoidal category, and is really shorthand for copying special cases of that proof into this section.

We will begin with the inverse law. For an arbitrary symmetric monoidal category this is

$$s_{B,A} \circ s_{A,B} = \text{id}_{A \otimes B}$$

We take $A = (X_1, R_1)$ and $B = (X_2, R_2)$, and so this is

$$\begin{aligned} & (s_{X_2, X_1}, s_{R_1, R_2}) \circ (s_{X_1, X_2}, s_{R_2, R_1}) \\ & : (X_1 \otimes X_2, R_1 \otimes R_2) \rightarrow (X_1 \otimes X_2, R_1 \otimes R_2) \end{aligned}$$

Since computation is functorial (§2.2.9), this is

$$(s_{X_2, X_1} \circ s_{X_1, X_2}, s_{R_2, R_1} \circ s_{R_1, R_2})$$

and we can apply the inverse law of \mathcal{C} .

The associativity axiom for a strict symmetric monoidal category is

$$\begin{array}{ccc} A \otimes B \otimes C & \xrightarrow{s_{A,B} \otimes C} & B \otimes A \otimes C \\ \downarrow s_{A,B \otimes C} & & \swarrow B \otimes s_{A,C} \\ B \otimes C \otimes A & & \end{array}$$

In $\mathbf{Game}(\mathcal{C})$, we need to take $A = (X_1, R_1)$, $B = (X_2, R_2)$ and $C = (X_3, R_3)$.

As a lemma, we need the equations

$$(X_2, R_2) \otimes s_{(X_1, R_1), (X_3, R_3)} = (X_2 \otimes s_{X_1, X_3}, R_2 \otimes s_{R_3, R_1})$$

and

$$s_{(X_1, R_1), (X_2, R_2)} \otimes (X_3, R_3) = (s_{X_1, X_2} \otimes X_3, s_{R_2, R_1} \otimes R_3)$$

which are special cases of the fact that computation is a monoidal embedding (§2.2.9).

Now, by functoriality, we have that the computations

$$((X_2, R_2) \otimes s_{(X_1, R_1), (X_3, R_3)}) \circ (s_{(X_1, R_1), (X_2, R_2)} \otimes (X_3, R_3))$$

and

$$((X_2 \otimes s_{X_1, X_3}) \circ (s_{X_1, X_2} \otimes X_3), (s_{R_2, R_1} \otimes R_3) \circ (R_2 \otimes s_{R_3, R_1}))$$

are equal. Therefore we must need only check the equations

$$(X_2 \otimes s_{X_1, X_3}) \circ (s_{X_1, X_2} \otimes X_3) = s_{X_1, X_2 \otimes X_3}$$

and

$$(s_{R_2, R_1} \otimes R_3) \circ (R_2 \otimes s_{R_3, R_1}) = s_{R_2 \otimes R_3, R_1}$$

which both hold because \mathcal{C} is symmetric monoidal.

2.2.8 Strategic triviality

Next, we will formalise some informal remarks that were made in §2.1.9. A game $\mathcal{G} : (X, S) \rightarrow (Y, R)$ will be called **strategically trivial** if it satisfies two conditions. Firstly, there must be only one strategy, so $\Sigma_{\mathcal{G}} = 1 = \{*\}$. Secondly, the strategy $*$ must be trivial, in the sense that it can never fail to be a Nash equilibrium. That is, for all contexts $(h, k) : \text{hom}_{\mathcal{C}}(I, X) \times \text{hom}_{\mathcal{C}}(Y, R)$ we must have

$$* \in \mathbf{B}_{\mathcal{G}}(h, k)*$$

We can now give more explanation for why a strategically trivial game should always be in equilibrium, rather than never: the compositional best response functions in §2.2.3 and §2.2.6 both use cartesian products, and if the set of best responses of a computation was empty, everything would cancel and the entire game would have no equilibria. Put another way, an equilibrium of an aggregate consists of an equilibrium of each component, with suitably modified contexts, and so an equilibrium overall must in particular restrict to an equilibrium on those components that are players.

We directly have that computations and counits (§2.1.9) are strategically trivial. We will now prove that strategically trivial games are closed under both composition and tensor. In both cases, the set of strategies is $1 \times 1 \cong 1$. For tensor products, the best response function is

$$\begin{aligned} \mathbf{B}_{\mathcal{G} \otimes \mathcal{H}}(h, k)(*, *) &= \mathbf{B}_{\mathcal{G}}(h_1, k_{\otimes*(h_2)}) * \times \mathbf{B}_{\mathcal{H}}(h_2, k_{*(h_1) \otimes}) * \\ &= \{*\} \times \{*\} \\ &\cong \{*\} \end{aligned}$$

and for categorical compositions, it is

$$\begin{aligned} \mathbf{B}_{\mathcal{H} \circ \mathcal{G}}(h, k)(*, *) &= \mathbf{B}_{\mathcal{G}}(h, k_{* \circ}) * \times \bigcap_{*:1} \mathbf{B}_{\mathcal{H}}(\mathbf{P}_{\mathcal{G}} * \circ h, k) * \\ &= \{*\} \times \{*\} \\ &\cong \{*\} \end{aligned}$$

This is useful because when reasoning about strategically trivial games, we can focus only on the play and coplay functions, since we already know the strategy profiles and best response function. If $\mathcal{G} : (X, S) \rightarrow (Y, R)$ is strategically trivial we will moreover often write the play function of \mathcal{G} as though it had type $\mathbf{P}_{\mathcal{G}} : \text{hom}_{\mathcal{C}}(X, Y)$ and the coplay function as $\mathbf{C}_{\mathcal{G}} : \text{hom}_{\mathcal{C}}(X \otimes R, S)$, leaving the strategy $*$ implicit.

Categorical composition of a game with a strategically trivial game is particularly simple, because the strategically trivial parts act only by transforming the history and continuation. Suppose we have a composition of games

$$(X', S') \xrightarrow{\mathcal{H}_1} (X, S) \xrightarrow{\mathcal{G}} (Y, R) \xrightarrow{\mathcal{H}_2} (Y', R')$$

where \mathcal{H}_1 and \mathcal{H}_2 are strategically trivial. Then

$$\Sigma_{\mathcal{H}_2 \circ \mathcal{G} \circ \mathcal{H}_1} \cong \Sigma_{\mathcal{G}}$$

and

$$\mathbf{B}_{\mathcal{H}_2 \circ \mathcal{G} \circ \mathcal{H}_1}(h, k) = \mathbf{B}_{\mathcal{G}}(\mathbf{P}_{\mathcal{H}_1} \circ h, k')$$

where

$$k' : Y \xrightarrow{\Delta_Y} Y \otimes Y \xrightarrow{Y \otimes \mathbf{P}_{\mathcal{H}_2}} Y \otimes Y' \xrightarrow{Y \otimes k} Y \otimes R' \xrightarrow{\mathbf{C}_{\mathcal{H}_2}} R$$

which in $\mathbf{Game}(\mathbf{Set})$ is

$$k'y = \mathbf{C}_{\mathcal{H}_2}(y, k(\mathbf{P}_{\mathcal{H}_2}y))$$

As can be seen, \mathcal{H}_1 affects only the history, and \mathcal{H}_2 affects only the continuation.

Similarly for tensor products, if we have games $\mathcal{G} : (X_1, S_1) \rightarrow (Y_1, R_1)$ and $\mathcal{H} : (X_2, S_2) \rightarrow (Y_2, R_2)$ where \mathcal{H} is strategically trivial³⁸, then

$$\mathbf{B}_{\mathcal{G} \otimes \mathcal{H}}(h, k)\sigma = \mathbf{B}_{\mathcal{G}}(h_1, k')\sigma$$

where

$$k' : Y_1 \cong Y_1 \otimes I \xrightarrow{Y_1 \otimes h_2} Y_1 \otimes X_2 \xrightarrow{Y_1 \otimes \mathbf{P}_{\mathcal{H}}} Y_1 \otimes Y_2 \xrightarrow{k} R_1 \otimes R_2 \xrightarrow{\pi_1} R_1$$

In $\mathbf{Game}(\mathbf{Set})$, this is

$$\mathbf{B}_{\mathcal{G} \otimes \mathcal{H}}((h_1, h_2), k)\sigma = \mathbf{B}_{\mathcal{G}}(h_1, \lambda(y_1 : Y_1).k_1(y_1, \mathbf{P}_{\mathcal{H}}h_2))\sigma$$

2.2.9 Computations as a monoidal embedding

We will now prove that computation, defined in §2.1.9, gives us a monoidal embedding

$$(-, -) : \mathcal{C} \times \mathcal{C}^{\text{op}} \hookrightarrow \mathbf{Game}(\mathcal{C})$$

This will be used implicitly many times in §3.1 and §3.2 to simplify calculations by working in \mathcal{C} rather than $\mathbf{Game}(\mathcal{C})$, and it is also needed to justify the syntax for objects in $\mathbf{Game}(\mathcal{C})$ introduced in §2.3.5, and the string diagram notation for computations in §2.3.6.

We will first prove that we have a bifunctor $\mathcal{C} \times \mathcal{C}^{\text{op}} \rightarrow \mathbf{Game}(\mathcal{C})$. In the product category $\mathcal{C} \times \mathcal{C}^{\text{op}}$ the objects are pairs of objects of \mathcal{C} , and the morphisms are pairs of morphisms with the second reversed. The identity morphism on the object (X, R) of $\mathcal{C} \times \mathcal{C}^{\text{op}}$ is $(\text{id}_X, \text{id}_R)$ which, lifted to a computation, is also the identity game on (X, R) (see §2.2.2).

The composition

$$(X, T) \xrightarrow{(f_1, f_2)} (Y, S) \xrightarrow{(g_1, g_2)} (Z, R)$$

in $\mathcal{C} \times \mathcal{C}^{\text{op}}$ is, by definition,

$$(g_1, g_2) \circ (f_1, f_2) = (g_1 \circ f_1, f_2 \circ g_2)$$

and we must prove that the games denoted by these two expressions are equal. Since both are strategically trivial (§2.2.8), we need only check the play and coplay functions. The play functions are

$$\mathbf{P}_{(g_1 \circ f_1, f_2 \circ g_2)} : X \xrightarrow{g_1 \circ f_1} Z$$

and

$$\mathbf{P}_{(g_1, g_2) \circ (f_1, f_2)} : X \xrightarrow{\mathbf{P}_{(f_1, f_2)}} Y \xrightarrow{\mathbf{P}_{(g_1, g_2)}} Z$$

³⁸The case where \mathcal{G} is strategically trivial is symmetric.

which are equal. The coplay functions are

$$\mathbf{P}_{(g_1 \circ f_1, f_2 \circ g_2)} : X \otimes R \xrightarrow{\pi_2} R \xrightarrow{f_2 \circ g_2} T$$

and

$$\begin{aligned} \mathbf{C}_{(g_1, g_2) \circ (f_1, f_2)} : X \otimes R &\xrightarrow{\Delta_X \otimes R} X \otimes X \otimes R \xrightarrow{X \otimes \mathbf{P}_{(f_1, f_2)} \otimes R} X \otimes Y \otimes R \\ &\xrightarrow{X \otimes \mathbf{C}_{(g_1, g_2)}} X \otimes S \xrightarrow{\mathbf{C}_{(f_1, f_2)}} T \end{aligned}$$

The latter is

$$\begin{aligned} X \otimes R &\xrightarrow{\Delta_X \otimes R} X \otimes X \otimes R \xrightarrow{X \otimes f_1 \otimes R} X \otimes Y \otimes R \\ &\xrightarrow{X \otimes \pi_2} X \otimes R \xrightarrow{X \otimes g_2} X \otimes S \xrightarrow{\pi_2} S \xrightarrow{f_2} T \end{aligned}$$

and these are equal³⁹

Next, we must prove that the embedding also respects the monoidal structure. The monoidal unit of the product monoidal category $\mathcal{C} \times \mathcal{C}^{\text{op}}$ is (I, I) , which is also the monoidal unit of $\mathbf{Game}(\mathcal{C})$ (see §2.2.7).

Suppose $(f_1, f_2) : (X_1, S_1) \rightarrow (Y_1, R_1)$ and $(g_1, g_2) : (X_2, S_2) \rightarrow (Y_2, R_2)$ are morphisms of $\mathcal{C} \times \mathcal{C}^{\text{op}}$. Their monoidal product is

$$(f_1 \otimes g_1, f_2 \otimes g_2) : (X_1 \otimes X_2, S_1 \otimes S_2) \rightarrow (Y_1 \otimes Y_2, R_1 \otimes R_2)$$

We must therefore prove the equality of games

$$(f_1, f_2) \otimes (g_1, g_2) = (f_1 \otimes g_1, f_2 \otimes g_2)$$

Again both are strategically trivial, so we need only work with the play and coplay functions. For the play function we have

$$\mathbf{P}_{(f_1, f_2) \otimes (g_1, g_2)} = \mathbf{P}_{(f_1, f_2)} \otimes \mathbf{P}_{(g_1, g_2)} = f_1 \otimes g_1 = \mathbf{P}_{(f_1 \otimes g_1, f_2 \otimes g_2)}$$

For coplay, the former is

$$\begin{aligned} \mathbf{C}_{(f_1, f_2) \otimes (g_1, g_2)} : X_1 \otimes X_2 \otimes R_1 \otimes R_2 &\xrightarrow{\cong} X_1 \otimes R_1 \otimes X_2 \otimes R_2 \\ &\xrightarrow{\mathbf{C}_{(f_1, f_2)} \otimes \mathbf{C}_{(g_1, g_2)}} S_1 \otimes S_2 \end{aligned}$$

which is

$$\begin{aligned} X_1 \otimes X_2 \otimes R_1 \otimes R_2 &\xrightarrow{\cong} X_1 \otimes R_1 \otimes X_2 \otimes R_2 \\ &\xrightarrow{\pi_2 \otimes \pi_2} R_1 \otimes R_2 \xrightarrow{f_2 \otimes g_2} S_1 \otimes S_2 \end{aligned}$$

and the latter is

$$\mathbf{C}_{(f_1 \otimes g_1, f_2 \otimes g_2)} : X_1 \otimes X_2 \otimes R_1 \otimes R_2 \xrightarrow{\pi_2} R_1 \otimes R_2 \xrightarrow{f_2 \otimes g_2} S_1 \otimes S_2$$

which is equal.

Pedantically speaking, to obtain a monoidal functor we also need to prove that the unitors in the two categories are respected, but we will not do so here.

³⁹This can fail if \mathcal{C} is only premonoidal, for example because f_1 can have nontrivial side effects. This may be a very serious problem and block the generalisation to noncommutative effects such as learning (see §0.7 and §0.10).

2.2.10 The counit law

We will now prove a result, which we will return to in §2.3.6, that connects computations with the counit game, also introduced in §2.1.9. Let $f : \text{hom}_{\mathcal{C}}(X, Y)$. Then the following diagram in $\mathbf{Game}(\mathcal{C})$ commutes:

$$\begin{array}{ccc}
 (X, Y) & \xrightarrow{(\text{id}_X, f)} & (X, X) \\
 \downarrow (f, \text{id}_Y) & & \downarrow \tau_X \\
 (Y, Y) & \xrightarrow{\tau_Y} & (I, I)
 \end{array}$$

This will be called the **counit law**. Since both games are strategically trivial, we need only check the behaviour of the play and coplay functions.

The definitions in §2.1.9 give us

$$\begin{array}{ll}
 \mathbf{P}_{(\text{id}_X, f)} : X \xrightarrow{\text{id}_X} X & \mathbf{C}_{(\text{id}_X, f)} : X \otimes X \xrightarrow{\pi_2} X \xrightarrow{f} Y \\
 \mathbf{P}_{(f, \text{id}_Y)} : X \xrightarrow{f} Y & \mathbf{C}_{(f, \text{id}_Y)} : X \otimes Y \xrightarrow{\pi_2} Y \xrightarrow{\text{id}_Y} Y \\
 \mathbf{P}_{\varepsilon_X} : X \xrightarrow{!_X} I & \mathbf{C}_{\varepsilon_X} : X \otimes I \xrightarrow{\pi_1} X \\
 \mathbf{P}_{\varepsilon_Y} : Y \xrightarrow{!_Y} I & \mathbf{C}_{\varepsilon_Y} : Y \otimes I \xrightarrow{\pi_1} Y
 \end{array}$$

Composing these sequentially gives us the play functions

$$\begin{array}{l}
 \mathbf{P}_{\tau_X \circ (\text{id}_X, f)} : X \xrightarrow{\text{id}_X} X \xrightarrow{!_X} I \\
 \mathbf{P}_{\tau_Y \circ (f, \text{id}_Y)} : X \xrightarrow{f} Y \xrightarrow{!_Y} I
 \end{array}$$

which are equal⁴⁰, and the coplay functions

$$\mathbf{C}_{\varepsilon_X \circ (\text{id}_X, f)} : X \otimes I \xrightarrow{\Delta_X \otimes I} X \otimes X \otimes I \xrightarrow{X \otimes \pi_1} X \otimes X \xrightarrow{\pi_2} X \xrightarrow{f} Y$$

and

$$\begin{aligned}
 \mathbf{C}_{\varepsilon_Y \circ (f, \text{id}_Y)} : X \otimes I &\xrightarrow{\Delta_X \otimes I} X \otimes X \otimes I \xrightarrow{X \otimes f \otimes I} X \otimes Y \otimes I \\
 &\xrightarrow{X \otimes \pi_1} X \otimes Y \xrightarrow{\pi_2} Y
 \end{aligned}$$

which are both equal to

$$X \otimes I \xrightarrow{\pi_1} X \xrightarrow{f} Y$$

The game denoted by these two equal expressions is important, because when post-composed with another game it will behave like a continuation, and

⁴⁰As in §2.2.9 this can fail if \mathcal{C} is only premonoidal, because f may have side effects. This relates to the interpretation of this equation in §2.3.6 as a topological move on string diagrams, because string diagrams for premonoidal categories have fewer allowed topological moves (Jeffrey 1997). See §0.10.

we will generally use it when f is an outcome function. Let $\mathcal{G} : (X, S) \rightarrow (Y, R)$ be a game, and let $k : \text{hom}_{\mathcal{C}}(Y, R)$ be a continuation for \mathcal{G} . Consider the game

$$(X, S) \xrightarrow{\mathcal{G}} (Y, R) \xrightarrow{(k, \text{id}_R)} (R, R) \xrightarrow{\varepsilon_R} (I, I)$$

which, by the counit law, can be equivalently written

$$(X, S) \xrightarrow{\mathcal{G}} (Y, R) \xrightarrow{(\text{id}_Y, k)} (Y, Y) \xrightarrow{\varepsilon_Y} (I, I)$$

Then

$$\Sigma_{\varepsilon_R \circ (k, \text{id}_R) \circ \mathcal{G}} \cong \Sigma_{\mathcal{G}}$$

and for any $h : \text{hom}_{\mathcal{C}}(I, X)$ and $\sigma : \Sigma_{\mathcal{G}}$ we have

$$\mathbf{B}_{\mathcal{G}}(h, k)\sigma \cong \mathbf{B}_{\varepsilon_R \circ (k, \text{id}_R) \circ \mathcal{G}}(h, \text{id}_I)\sigma$$

under the same isomorphism. To see this, the right hand side by the definition in §2.2.3 is

$$\begin{aligned} & \mathbf{B}_{\varepsilon_R \circ (k, \text{id}_R) \circ \mathcal{G}}(h, \text{id}_I)\sigma \\ &= \mathbf{B}_{\mathcal{G}}(h, (\text{id}_I)_{*o})\sigma \times \bigcap_{\sigma' : \Sigma_{\mathcal{G}}} \mathbf{B}_{\varepsilon_R \circ (k, \text{id}_R)}(\mathbf{P}_{\mathcal{G}}\sigma' \circ h, \text{id}_I)* \end{aligned}$$

and, since $\varepsilon_R \circ (k, \text{id}_R)$ is strategically trivial, this is

$$\mathbf{B}_{\mathcal{G}}(h, (\text{id}_I)_{*o})\sigma \times \{*\}$$

The final step is to see that $(\text{id}_I)_{*o} = k$. By definition it is

$$(\text{id}_I)_{*o} : Y \xrightarrow{\Delta_Y} Y \otimes Y \xrightarrow{Y \otimes \mathbf{P}_{\varepsilon_R \circ (k, \text{id}_R)}} Y \otimes I \xrightarrow{Y \otimes \text{id}_I} Y \otimes I \xrightarrow{C_{\varepsilon_R \circ (k, \text{id}_R)}} R$$

which reduces to

$$Y \xrightarrow{\Delta_Y} Y \otimes Y \xrightarrow{Y \otimes !Y} Y \otimes I \xrightarrow{\pi_1} Y \xrightarrow{k} R$$

and hence to k .

The equations

$$\mathbf{B}_{\varepsilon_R \circ (k, \text{id}_R) \circ \mathcal{G}}(h, \text{id}_I)\sigma \cong \mathbf{B}_{\mathcal{G}}(h, k)\sigma \cong \mathbf{B}_{\varepsilon_Y \circ (\text{id}_Y, k) \circ \mathcal{G}}(h, \text{id}_I)\sigma$$

are very important because they allow us to move between ‘internal’ and ‘external’ views of the continuation k . We will see this used several times in §2.3, §3.1 and §3.2, when k is an outcome or utility function.

2.2.11 Information flow in games

In §2.1 and §2.2 we have given several definitions whose intuitive justifications have been incomplete at best. This will remain the case until [chapter 3](#), in which we will demonstrate that game theory can be done inside the category $\mathbf{Game}(\mathcal{C})$. However we would also like to have a separate justification of the ‘low level’ operations, especially for readers whose motivations come from category theory. We have argued, for example in §2.2.3, that the definitions are *not*

canonical, in the sense that there are useful variations, and so it would be naive to expect an axiomatic definition of $\mathbf{Game}(\mathcal{C})$, for example using universal properties.

There is, however, one known and very surprising connection to existing mathematics. The data $(\Sigma_{\mathcal{G}}, \mathbf{P}_{\mathcal{G}}, \mathbf{C}_{\mathcal{G}}, \mathbf{B}_{\mathcal{G}})$ specifying a game \mathcal{G} could be divided into two kinds of data: the strategy profiles $\Sigma_{\mathcal{G}}$ and best response functions $\mathbf{B}_{\mathcal{G}}$ are motivated in terms of game theory, whereas the play and coplay functions $\mathbf{P}_{\mathcal{G}}\sigma$, $\mathbf{C}_{\mathcal{G}}\sigma$ applied to a fixed strategy σ are used for *information flow*. These latter are strongly reminiscent of *dialectica categories*.

The dialectica categories are categorical models of intuitionistic logic⁴¹ introduced in (de Paiva 1991), which are based on Gödel’s dialectica interpretation (Avigad and Feferman 1998). The dialectica category⁴² $\mathbf{D}(\mathbf{Set})$ has as objects relations $\varphi \subseteq X \times R$, which in particular specifies a pair⁴³ of sets (X, R) . A morphism from $\varphi \subseteq X \times S$ to $\psi \subseteq Y \times R$ is a pair (f, g) where

$$f : X \rightarrow Y$$

and

$$g : X \times R \rightarrow S$$

such that for all $x : X$ and $r : R$ we have

$$\varphi(x, g(x, r)) \implies \psi(fx, r)$$

This has a notable similarity with $f = \mathbf{P}_{\mathcal{G}}\sigma$ and $g = \mathbf{C}_{\mathcal{G}}\sigma$, if we take the relations to be trivial (either both empty or both full). More importantly, the definition of categorical composition is essentially the same in $\mathbf{Game}(\mathcal{C})$ as in $\mathbf{D}(\mathcal{C})$, so the definition that was claimed in §2.2.2 to not be intuitively justifiable is in fact largely an instance of something already known, and the proofs in §2.2.4 and §2.2.5 partly resemble the corresponding proofs for dialectica categories.

It should be mentioned that dialectica categories (and functional interpretations in proof theory) can be seen in game *semantic* terms, see for example (Blass 1997) and (Hedges 2015a). The relations φ and ψ are seen as games, in the game-semantic sense of two-player win/lose games, with (f, g) being a strategy for the first player (proponent) and (x, r) being a strategy for the second player (opponent), in a relative game $\varphi \rightarrow \psi$. Then the existence of a morphism $\varphi \rightarrow \psi$, which logically relates to provability of the implication, amounts to the existence of a winning strategy for the first player for ψ relative to φ . However, this apparent connection between game theory and game semantics seems to be only coincidental (see §0.8).

⁴¹Dialectica models of linear logic are better known.

⁴²This definition generalises to $\mathbf{D}(\mathcal{C})$ for \mathcal{C} an arbitrary monoidal closed category (de Paiva 2007), and to an even more abstract setting in (Hyland 2002). This suggests it may be possible to give a more axiomatic formulation of $\mathbf{Game}(\mathcal{C})$ if desired.

⁴³There are several constructions in category theory besides dialectica categories that take objects to be pairs, generally to obtain some form of duality. Chu spaces are already known to be related to dialectica categories (de Paiva 2007). The Int-construction (Joyal et al. 1996) is better known, but has a quite different structure. That paper explicitly says that the pair (X, R) should be considered as a formalisation of $X \otimes R^*$, as we will begin to do in §2.3.5. For remarks on duality in $\mathbf{Game}(\mathcal{C})$, see §2.3.6 and §2.3.7.

2.3 String diagrams

2.3.1 Discussion

String diagrams are a graphical calculus that can be used to visualise information flow in monoidal categories. Their earliest appearance may be Penrose’s graphical tensor notation in (Penrose 1971), with another precursor being Girard’s proof nets for linear logic (Girard 1987), and the mathematical foundations were formalised in (Joyal and Street 1991). They became well known through the work of Samson Abramsky, Bob Coecke and others on quantum information theory (Coecke 2011), and later through the work of Bob Coecke, Mehrnoosh Sadrzadeh and others on distributional semantics in linguistics (Heunen et al. 2013). String diagrams are also being applied in other areas such as bialgebra (McCurdy 2012) and computability (Pavlovic 2013).

As is the case in quantum physics and linguistics, string diagrams can be used in game theory to *visualise information flow*. This visualisation is a separate issue to compositionality (although being able to compose string diagrams is a crucial requirement), and is a separate contribution of this thesis. The concept of information flow in games will only be introduced informally and by example, starting in §2.3.9. The effect of the counit game on information flow, described in §2.3.6, is particularly interesting, but again will only be discussed informally.

For readers unfamiliar with string diagrams, §2.3.2 and §2.3.3 introduce them in an informal way that will be sufficient for our purposes. Alternatively, (Baez and Stay 2010) is a good introduction. A survey of the many types of string diagrams is given in (Selinger 2011), although the exact variant we introduce in §2.3.5 does not match any of the usual definitions.

The purpose of string diagrams varies by discipline. In quantum information theory, the emphasis is generally on the ability of string diagrams to reduce complex calculations in tensor calculus to trivial topological deformation, a point made forcibly in (Coecke 2005). In linguistics, there is more emphasis on the use of string diagrams as a device for visualising the logical structure of sentences, whereas the underlying categorical structure is used in a more formal way by considering functorial semantics.

In game theory, we will similarly emphasise string diagrams as a visualisation tool. However, more so than in linguistics, the algebraic expressions denoting even simple games can be quite complicated, and we will make use of string diagrams as a *tool for making definitions*. A good example is in §3.2.2, where we will define a particular game by its (simple) string diagram, and the subsequent work to compute the denotation of the string diagram is quite involved. However, this work could quite easily be automated, as described in §0.10.

2.3.2 String diagrams for monoidal categories

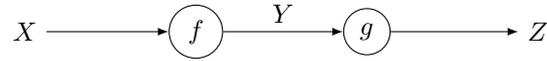
The basic components of string diagrams are **strings** and **beads**. A simple string diagram has the form

$$X \longrightarrow \textcircled{f} \longrightarrow Y$$

This string diagram *denotes* a morphism $f : X \rightarrow Y$ in some monoidal category. The strings in the diagram are labelled by objects of the category, and the beads by morphisms.

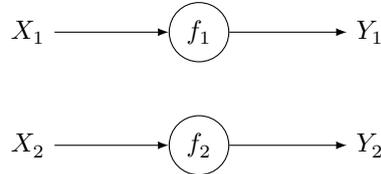
2.3. STRING DIAGRAMS

The composition of the category is denoted by end-to-end composition of string diagrams. If we have morphisms $f : X \rightarrow Y$ and $g : Y \rightarrow Z$, the composition $g \circ f : X \rightarrow Z$ is denoted by

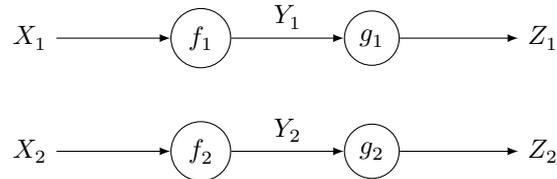


Notice that the associativity of composition is trivialised, because there is only one way to graphically compose three morphisms. This is the simplest example of a powerful fact about string diagrams, that coherence conditions in a category are reduced to graphical identity or topological deformation.

The other way of composing morphisms, namely the tensor product, is denoted by side-by-side composition. If we have morphisms $f_1 : X_1 \rightarrow Y_1$ and $f_2 : X_2 \rightarrow Y_2$, the tensor product $f_1 \otimes f_2 : X_1 \otimes X_2 \rightarrow Y_1 \otimes Y_2$ is denoted by

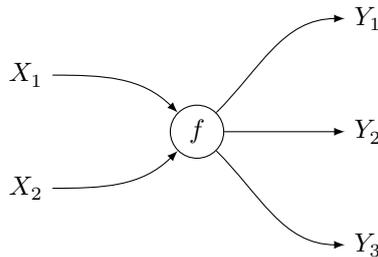


Again, notice that the associator of the tensor product is reduced to graphical identity. According to these rules, the string diagram

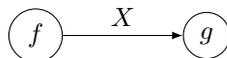


could denote either $(g_1 \circ f_1) \otimes (g_2 \circ f_2)$ or $(g_1 \otimes g_2) \circ (f_1 \otimes f_2) : X_1 \otimes X_2 \rightarrow Z_1 \otimes Z_2$, but these are equal⁴⁴.

A morphism $f : X_1 \otimes X_2 \rightarrow Y_1 \otimes Y_2 \otimes Y_3$ is denoted by a bead with two strings entering on the left and three exiting on the right:



The unit object of the tensor product is denoted by empty space. For example, if we have morphisms $f : I \rightarrow X$ and $g : X \rightarrow I$, the composition $g \circ f : I \rightarrow I$ is denoted by



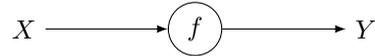
⁴⁴This is the axiom that \otimes is a bifunctor.

The identity morphism on an object X is denoted by simply a string labelled with X . The coherence laws for the identity morphisms and unit object are again reduced to topological deformations.

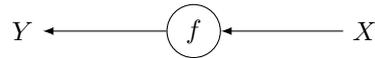
2.3.3 Compact closed categories

An important structure that a monoidal category \mathcal{C} can carry, which has a particularly elegant interpretation in string diagrams, is that of a compact closed category. In a compact closed category we have a **duality**, which is⁴⁵ a monoidal functor $-^* : \mathcal{C} \rightarrow \mathcal{C}^{\text{op}}$. The prototypical example of a compact closed category is the category of finite-dimensional vector spaces and linear maps over a fixed field, with duality given by duality of vector spaces; another well known example is the category of sets and relations, with duality being the identity on sets, and giving inverse relations.

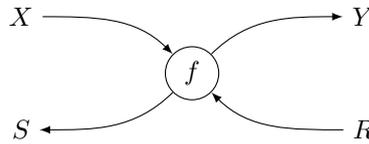
Whereas an object X is denoted by an X -labelled string running from left to right, the dual object X^* is denoted by an X -labelled string running from right to left. Then, given a morphism $f : X \rightarrow Y$ denoted by



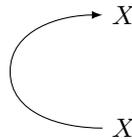
we will denote the dual morphism $f^* : Y^* \rightarrow X^*$ by its rotation:



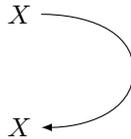
The tensor product is still denoted by side-by-side composition, and so for example a morphism $f : X \otimes S^* \rightarrow Y \otimes R^*$ in a compact closed category would be denoted by



We additionally have unit morphisms $\eta_X : I \rightarrow X \otimes X^*$ denoted by a **cap**⁴⁶



and counits $\varepsilon_X : X \otimes X^* \rightarrow I$ denoted by a **cup**

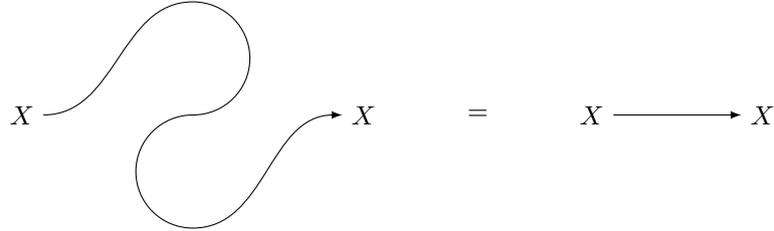


⁴⁵Since this is not a thesis on category theory, this is not intended to be a rigorous introduction to compact closure, but rather to develop intuition that we can use starting in §2.3.5.

⁴⁶The ‘cap’ and ‘cup’ terminology makes more sense if we draw time running from top to bottom rather than right to left.

2.3. STRING DIAGRAMS

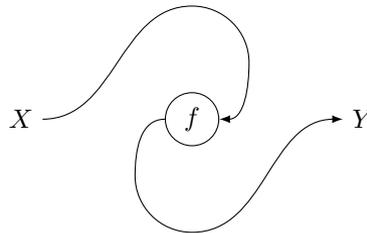
The axioms of a compact closed category specify precisely that the units and counits behave in graphically intuitive ways. For example, a bent string can be straightened (this is called the ‘yanking equation’),



and a bead can be slid around a cap or cup:



Combining these, we can topologically deform a string diagram in arbitrary ways, for example we can equally denote a morphism $f : X \rightarrow Y$ by the string diagram

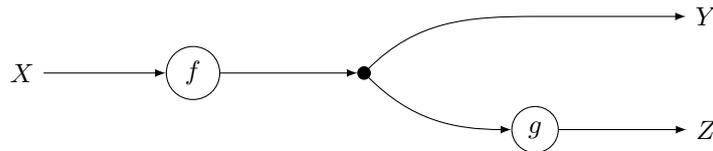


2.3.4 Boxing and compositionality

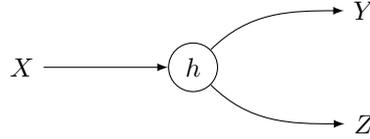
A simple but very important observation is that string diagrams are inherently fully compositional. Given an arbitrary string diagram, we could make it definitionally equal to a single bead, preserving only the strings entering and leaving the diagram. Such strings will be called **open ports**. To give an example, suppose we have morphisms $f : X \rightarrow Y$ and $g : Y \rightarrow Z$, and we define a new morphism

$$(Y \otimes g) \circ \Delta_Y \circ f : X \rightarrow Y \otimes Z$$

which is denoted by the string diagram



This string diagram has an X -labelled open port entering to the left, and Y - and Z -labelled open ports leaving to the right. If we define $h : X \rightarrow Y \otimes Z$ to be this morphism, then we can replace this string diagram with



This amounts to *forgetting* the definition of h and remembering only its type, that is to say, the types, order and variance of its open ports.

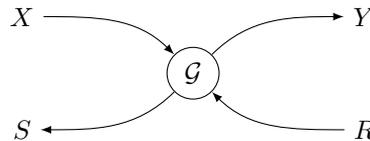
We can imagine that if we *zoom in* to the h -labelled bead in this string diagram, we will come to the previous string diagram, with the open ports being *physically the same as* the strings attached to the h -labelled bead (see §0.10).

Conversely, the act of obtaining the lower string diagram from the upper amounts to drawing a box around the entire diagram, so that the points at which the open strings intersect the box become exactly the points at which the strings enter and leave the h -labelled bead. We will refer to this as **boxing** the string diagram, and it is the graphical analogue of making a definitional equality.

Boxing is the aspect of compositionality that can be used to work in a scalable way. In a more serious example, the morphism h could be a very complicated process, which we can then use and re-use without worrying (or even knowing) how it is defined. In software engineering, the boxed morphism would variously be known as a ‘type’, a ‘signature’ or an ‘interface’.

2.3.5 The geometry of games

The idea behind string diagrams in game theory is that although $\mathbf{Game}(\mathcal{C})$ does not have as much structure as the categories usually considered with string diagrams, we will abuse notation *as though* we had this structure. A game $\mathcal{G} : (X, S) \rightarrow (Y, R)$ will be denoted by the string diagram



That is, we are *pretending* that $\mathbf{Game}(\mathcal{C})$ has a duality $-^*$, and that \mathcal{G} is a morphism $\mathcal{G} : X \otimes S^* \rightarrow Y \otimes R^*$.

This is a reasonable thing to do. If we define an operation on objects by $(X, R)^* = (R, X)$, we find that this operation can be extended in a compatible way to some morphisms, but not all. Thus we can consider $-^*$ to be a partial functor, which is total on objects. Since we have a symmetric monoidal embedding $\mathcal{C} \hookrightarrow \mathbf{Game}(\mathcal{C})$ (§2.2.9), we can write X and R for (X, I) and (R, I) . Thus $R^* = (I, R)$, and so $X \otimes R^* = (X, I) \otimes (I, R) = (X \otimes I, I \otimes R)$ which, treating natural isomorphism as though it is identity, gives $X \otimes R^* = (X, R)$.

2.3. STRING DIAGRAMS

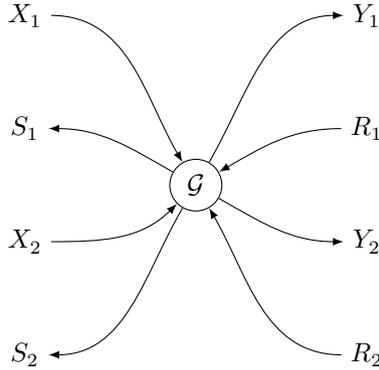
We will extend this notation in several intuitive ways. For example, if we have a game

$$\mathcal{G} : (X_1 \otimes X_2, S_1 \otimes S_2) \rightarrow (Y_1 \otimes Y_2, R_1 \otimes R_2)$$

we may write it as

$$\mathcal{G} : X_1 \otimes S_1^* \otimes X_2 \otimes S_2^* \rightarrow Y_1 \otimes R_1^* \otimes Y_2 \otimes R_2^*$$

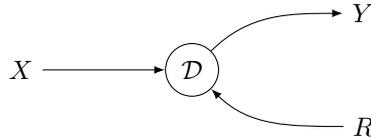
(among many other possibilities, one for each permutation), and the string diagram as



In general, we will be unconcerned but consistent with the order of the types in such situations.

Of course this comes at a cost: not all topological operations on our string diagrams preserve meaning, or even well-formedness. This is something we will need to take extra care about. For example, in a compact closed category there is no real distinction between ‘past’ and ‘future’, but in $\mathbf{Game}(\mathcal{C})$ they are distinct: a game $\mathcal{G} : X \rightarrow Y$ cannot be turned into $\mathcal{G} : I \rightarrow X^* \otimes Y$. An explanation for this is given in §2.3.12.

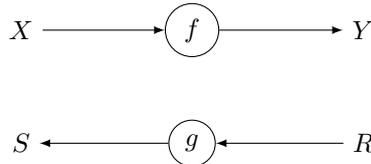
We will omit some strings when some of the types involved are the monoidal unit. For example, if we have a decision $\mathcal{D} : (X, I) \rightarrow (Y, R)$ then we will write $\mathcal{D} : X \rightarrow Y \otimes R^*$, and the string diagram denoting it is



A closed game could be denoted simply by a bead with no strings attached.

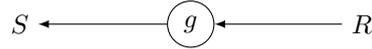
2.3.6 Partial duality

A computation $(f, g) : (X, S) \rightarrow (Y, R)$ will be denoted by the string diagram



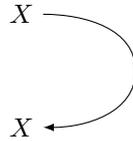
2.3. STRING DIAGRAMS

That is, we are *pretending* that $(f, g) = f \otimes g^*$, where $g^* : S^* \rightarrow R^*$ is the dual of g . This is another case of treating $-^*$ as a partial functor, this time defined on morphisms of \mathcal{C} . From this point onwards we will no longer treat computations formally as pairs (f, g) , but as individual games f and g^* which we tensor together when necessary. This is an extension of the fact that computation is a monoidal embedding (§2.2.9). For example, if $g : R \rightarrow S$ then we can write $g^* = (\text{id}_I, g) : S^* \rightarrow R^*$ and use the string diagram



since id_I is denoted by empty space.

The counit game $\varepsilon_X : (X, X) \rightarrow (I, I)$ can now be written $\varepsilon_X : X \otimes X^* \rightarrow I$, and will be denoted by a cup



The counit law

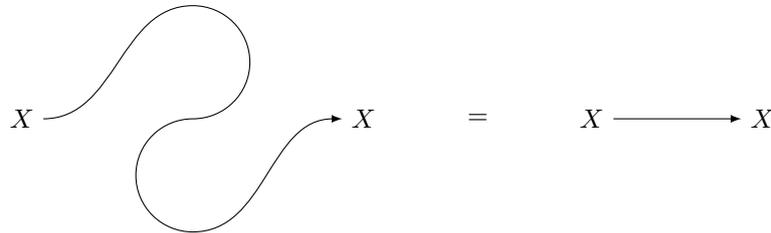
$$\varepsilon_Y \circ (f, \text{id}_Y) = \varepsilon_X \circ (\text{id}_X, f)$$

which we proved in §2.2.10 now appears as a coherence law



telling us that we can slide a computation around the counit like a bead. This takes the same form as the coherence law for the counit of a compact closed category (see §2.3.3), but applies only⁴⁷ to computations.

The partiality of the duality means that, unlike in a compact closed category, we cannot rotate beads. For example, in a compact closed category we have a valid equation

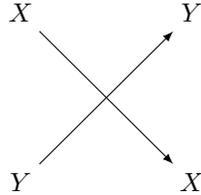


but we cannot do so here because the cap does not denote a morphism of $\mathbf{Game}(\mathcal{C})$, so the left hand side is not a well formed string diagram.

⁴⁷It is unknown whether $\mathbf{Game}(\mathcal{C})$ contains any dualisable morphisms besides the computations.

2.3.7 Covariance, contravariance and symmetries

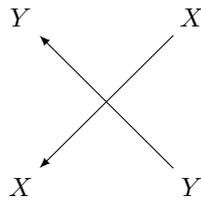
This kind of duality has one strange consequence, concerning how covariant and contravariant (or forward and backward) strings interact. If we have two strings pointing in the same direction which cross, such as



then the denoted morphism is a symmetry $\sigma_{X,Y} : X \otimes Y \rightarrow Y \otimes X$, or more formally, the computation

$$(\sigma_{X,Y}, \text{id}_I) : (X \otimes Y, I) \rightarrow (Y \otimes X, I)$$

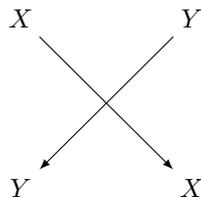
Similarly, in the contravariant direction, the string diagram



denotes $\sigma_{X,Y}^* : Y^* \otimes X^* \rightarrow X^* \otimes Y^*$, or again more formally,

$$(\text{id}_I, \sigma_{X,Y}) : (I, Y \otimes X) \rightarrow (I, X \otimes Y)$$

However, now consider the string diagram



The object on the left is $X \otimes Y^*$, which formally denotes (X, Y) . The object on the right is $Y^* \otimes X$, which again formally denotes (X, Y) . Thus, the denotation σ_{X,Y^*} of this string diagram is an identity, not a symmetry.

More generally, if we have a string diagram which consists only of crossings of covariant and contravariant strings, the denoted game is an identity. We will see this in practice in the coming sections, for example in §2.3.9.

Another example of the same idea is drawing a counit with the opposite orientation than in §2.3.6. This is justified by the topological deformation



where the right hand side is the composition $\varepsilon_X \circ \sigma_{X^*, X}$, because the symmetry $\sigma_{X^*, X}$ is trivial.

2.3.8 Copying and deleting information

The underlying categories \mathcal{C} introduced in §2.1.2 are not compact closed in general, but they do have one piece of additional structure assumed, namely that the tensor product is a weak product. This means that we have canonical morphisms

$$!_X : X \rightarrow I$$

and

$$\Delta_X : X \rightarrow X \otimes X$$

for each object X . Formally, this makes each object X into a cocommutative coalgebra. These morphisms can be lifted into $\mathbf{Game}(\mathcal{C})$ either covariantly, or contravariantly as

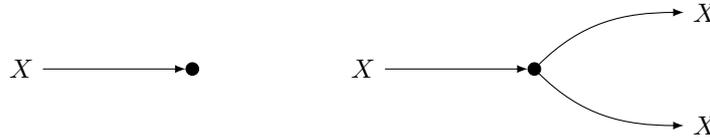
$$!_X^* : I \rightarrow X^*$$

and

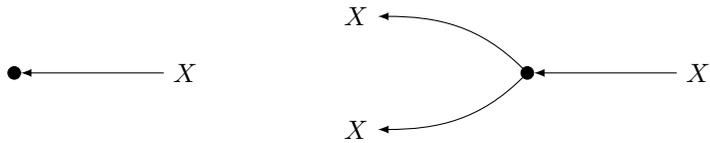
$$\Delta_X^* : X^* \otimes X^* \rightarrow X^*$$

Thus, the covariant objects⁴⁸ of $\mathbf{Game}(\mathcal{C})$ are cocommutative coalgebras, and contravariant objects are commutative algebras.

As is usual with string diagrams, we will denote the algebraic and coalgebraic operators by small filled circles. That is, the coalgebraic operators on a covariant object are denoted by



and the algebraic operators on a contravariant object by their rotations



These operators are not as ubiquitous in game theory as they are, for example, in quantum information theory⁴⁹, but nearly every game will involve copying information at some point in its definition, for example if both a player and a utility function need to be aware of the same value.

⁴⁸A ‘covariant object’ is an object of the form (X, I) , and a ‘contravariant object’ is of the form (I, X) . A general object is the tensor product of a covariant part and a contravariant part, and carries neither an algebra or a coalgebra structure.

⁴⁹See for example the ZX-calculus (Coecke and Duncan 2011).

2.3.9 A bimatrix game

We have now, finally, built up enough theory to give some recognisable, textbook examples of games. We will begin with a bimatrix game, with mixed strategy Nash equilibrium as the solution concept.

Bimatrix games are two player classical games with mixed strategies, introduced in §1.2.6. Thus we have two players, who simultaneously make choices from sets X, Y . For simplicity, we will assume that these are finite. We additionally have utility functions $q_1, q_2 : X \times Y \rightarrow \mathbb{R}$ that give the utility for each player on each play.

The best response function for such a game is

$$\mathbf{B}(\sigma_1, \sigma_2) = \left\{ (\sigma'_1, \sigma'_2) : \mathcal{D}X \times \mathcal{D}Y \left| \begin{array}{l} \mathbb{E}[q_1(\sigma'_1, \sigma_2)] \geq \mathbb{E}[q_1(x', \sigma_2)], \\ \mathbb{E}[q_2(\sigma_1, \sigma'_2)] \geq \mathbb{E}[q_2(\sigma_1, y')], \\ \text{for all } x' : X \text{ and } y' : Y \end{array} \right. \right\}$$

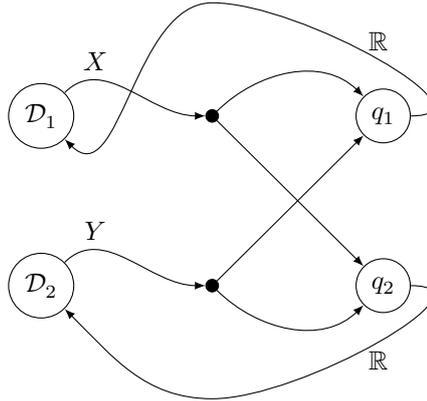
and a mixed strategy Nash equilibrium is a fixpoint of this \mathbf{B} .

A prototypical example of a bimatrix game is matching pennies. In this example, we have $X = Y = \{H, T\}$, with

$$q_1(x, y) = \begin{cases} 1 & \text{if } x = y \\ 0 & \text{if } x \neq y \end{cases}$$

and $q_2(x, y) = 1 - q_1(x, y)$. The unique mixed strategy Nash equilibrium of this game is the strategy profile in which both players choose either H or T with probability $\frac{1}{2}$, with the expected utility for both players being $\frac{1}{2}$. Other well known examples of bimatrix games include the prisoner's dilemma and the chicken game, and can be found in myriad books and lecture notes on introductory game theory.

In order to model a bimatrix game, let \mathcal{G} be the game in $\mathbf{Game}(\mathbf{SRel})$ denoted by the string diagram



This diagram is built from pieces that have been introduced, composed together using the two methods that have been introduced:

- The decisions $\mathcal{D}_1 : I \rightarrow X \otimes \mathbb{R}^*$ and $\mathcal{D}_2 : I \rightarrow Y \otimes \mathbb{R}^*$ are expected utility maximising decisions (§2.1.8), both with histories of type I .
- The two black nodes are comultiplications (§2.3.8).

- The two crossing points are symmetries, one of which is trivial and one of which is nontrivial (§2.3.7).
- The beads labelled q_1 and q_2 are the utility functions considered as co-variant computations (§2.3.6).
- The two counits are drawn with opposite orientations (§2.3.7).

If we calculate the type of strategies of \mathcal{G} , we indeed get $\Sigma_{\mathcal{G}} = \mathcal{D}X \times \mathcal{D}Y$. Because $\mathcal{G} : I \rightarrow I$ is closed game (see §2.1.6), the best response function has type

$$\mathbf{B}_{\mathcal{G}} : \Sigma_{\mathcal{G}} \rightarrow \mathcal{P}\Sigma_{\mathcal{G}}$$

Moreover, $\mathbf{B}_{\mathcal{G}}$ is equal to the best response function \mathbf{B} of the bimatrix game we began with. In particular, the equilibria of \mathcal{G} are exactly the mixed strategy Nash equilibria of the bimatrix game. The proof of this is quite involved, and is mostly given in §3.1.7.

2.3.10 A sequential game

As a second example, we will model a two-player game of perfect information, with players modelled by arbitrary multi-valued selection functions, giving an instance of the sequential games defined in §1.3.4. Let X , Y and R be arbitrary sets, and consider the game defined by the outcome function $q : X \times Y \rightarrow R$ and the multi-valued selection functions

$$E_1 : (X \rightarrow R) \rightarrow \mathcal{P}X$$

$$E_2 : (Y \rightarrow R) \rightarrow \mathcal{P}Y$$

Recall that a subgame perfect equilibrium of this game is a strategy profile

$$(\sigma, \tau) : X \times (X \rightarrow Y)$$

such that

$$\sigma \in E_1(\mathcal{U}_{\langle \rangle}^q(\sigma, \tau))$$

and for all $x : X$,

$$\tau x \in E_2(\mathcal{U}_{\langle x \rangle}^q(\sigma, \tau))$$

where

$$\mathcal{U}_{\langle \rangle}^q(\sigma, \tau)x = q(x, \tau x)$$

and

$$\mathcal{U}_{\langle x \rangle}^q(\sigma, \tau)y = q(x, y)$$

These are instances of the construction in §2.1.7 of a decision from a multi-valued selection function.

Let $\mathcal{D}_1 : I \rightarrow X \otimes R^*$ be the decision in **Game(Set)** defined by the response function

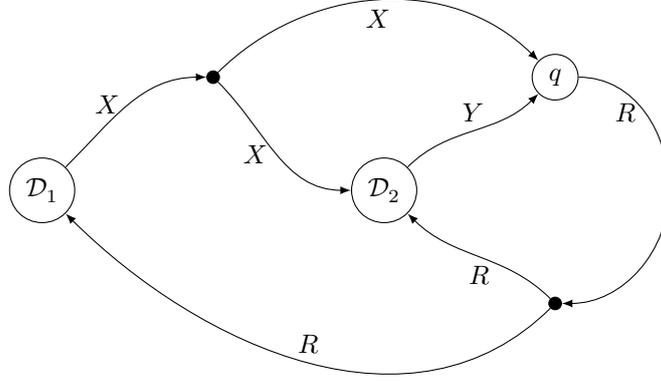
$$\mathbf{R}_{\mathcal{D}_1}(*, k) = E_1 k$$

and let $\mathcal{D}_2 : X \rightarrow Y \otimes R^*$ be the decision defined by

$$\mathbf{R}_{\mathcal{D}_2}(x, k) = \{\tau' : X \rightarrow Y \mid \tau' x \in E_2 k\}$$

2.3. STRING DIAGRAMS

Let $\mathcal{G} : I \rightarrow I$ be the closed game in $\mathbf{Game}(\mathbf{Set})$ denoted by the string diagram



The set of strategy profiles for this game is $\Sigma_{\mathcal{G}} = X \times (X \rightarrow Y)$, and the best response function $\mathbf{B}_{\mathcal{G}} : \Sigma_{\mathcal{G}} \rightarrow \mathcal{P}\Sigma_{\mathcal{G}}$ is precisely the best response function in §1.3.5, whose fixpoints are the subgame perfect equilibria.

This example can also be extended in a graphically intuitive way to games of imperfect information, in which the second player can only observe some function of the first player's choice. The construction will be sketched in §3.2.7.

2.3.11 Coordination and differentiation games

As our final introductory example, we will return to the coordinating and differentiating behaviour of §1.2.9. As in the voting game introduced in §1.2.7, we will let X be finite and consider the outcomes to also be $R = X$. We will consider a pair of decisions

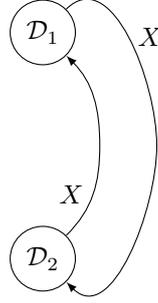
$$\mathcal{D}_{\text{fix}}, \mathcal{D}_{\text{nonfix}} : I \rightarrow X \otimes X^*$$

in $\mathbf{Game}(\mathbf{Set})$, defined by the response functions given exactly by the multi-valued selection functions E_{fix} and E_{nonfix} , so

$$\mathbf{B}_{\mathcal{D}_{\text{fix}}}(*, k)\sigma = E_{\text{fix}}k$$

$$\mathbf{B}_{\mathcal{D}_{\text{nonfix}}}(*, k)\sigma = E_{\text{nonfix}}k$$

Consider the string diagram



where \mathcal{D}_1 and \mathcal{D}_2 are chosen from \mathcal{D}_{fix} and $\mathcal{D}_{\text{nonfix}}$. One way to write this algebraically is

$$\mathcal{G} = (\varepsilon_X \otimes \varepsilon_X) \circ (s_{X,X} \otimes X^* \otimes X^*) \circ (\mathcal{D}_1 \otimes \mathcal{D}_2) : I \rightarrow I$$

The type of strategy profiles is $\Sigma_{\mathcal{G}} = X \times X$. The best response function $\mathbf{B}_{\mathcal{G}} : X \times X \rightarrow \mathcal{P}(X \times X)$ is given by

$$\mathbf{B}_{\mathcal{G}}(\sigma, \tau) = \mathbf{B}_{\mathcal{D}_1}(*, k_1)\sigma \times \mathbf{B}_{\mathcal{D}_2}(*, k_2)\tau$$

where $k_1x = \tau$ and $k_2x = \sigma$ are constant functions. Since $\sigma' \in \mathbf{B}_{\mathcal{D}_{\text{fix}}}(*, k_1)\sigma$ iff $\sigma' = \tau$, and $\tau' \in \mathbf{B}_{\mathcal{D}_{\text{fix}}}(*, k_2)\tau$ iff $\tau' = \sigma$, if we take $\mathcal{D}_1 = \mathcal{D}_2 = \mathcal{D}_{\text{fix}}$ then (σ, τ) is an equilibrium of \mathcal{G} iff $\sigma = \tau$. Similarly, if we take $\mathcal{D}_1 = \mathcal{D}_2 = \mathcal{D}_{\text{nonfix}}$ then (σ, τ) is an equilibrium of \mathcal{G} iff $\sigma \neq \tau$. Thus we again obtain coordinating and differentiating behaviour.

If we take $\mathcal{D}_1 = \mathcal{D}_{\text{fix}}$ and $\mathcal{D}_2 = \mathcal{D}_{\text{nonfix}}$ then \mathcal{G} has no equilibria, because for (σ, τ) to be an equilibrium we must have both $\sigma = \tau$ and $\sigma \neq \tau$. It would be interesting to consider variants of \mathcal{D}_{fix} and $\mathcal{D}_{\text{nonfix}}$ in **Game(SRel)**, because this game⁵⁰ should intuitively have a mixed strategy Nash equilibrium, but this is not free because we cannot take the mixed extension of a selection function (see §1.2.6).

This example demonstrates that open games can model non-classical preferences described by selection functions. It is interesting because the game has no outcome function in a very literal sense, but instead the string diagram visualises the information flow where each player directly reasons about the move of the other. However this issue is orthogonal to compositionality, and we will not discuss it further.

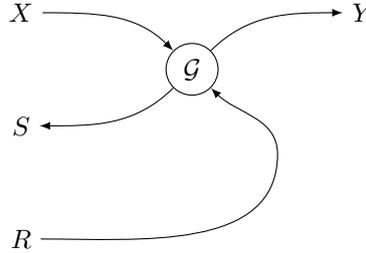
2.3.12 Designing for compositionality

We now return to theoretical considerations, having gained some intuition about the possible examples. The fact that the counit ε_X is not dualisable has some interesting consequences for information flow in **Game(C)**. A game of type $\mathcal{G} : X \otimes S^* \rightarrow Y \otimes R^*$ can be thought of as accepting a value of type X from the past, and a value of type R from the future (see §2.1.4). If we have a value

⁵⁰This game is called the penalty game because it models a penalty shootout. The fixpoint player is the goalkeeper who would like to coordinate, and the non-fixpoint player is the penalty taker who would like to differentiate.

2.3. STRING DIAGRAMS

of type R in the past (for example, because an agent’s utility is determined entirely by events in the past) then we can use the counit to bend the string around, as in



To put it another way: *a value known in the past will still be known in the future.*

On the other hand, if we only have a value of type X in the future, we cannot use it (unlike in a compact closed category) because an agent cannot directly observe the future, but only reason about it (which is the purpose of R). This ‘causality’ point of view serves to explain why $\mathbf{Game}(\mathcal{C})$ has the structure that it does. Indeed, if $\mathbf{Game}(\mathcal{C})$ were compact closed it would be straightforward to construct a paradoxical situation analogous to the grandfather paradox, in which a strategy negates the value chosen by itself. Categories with more structure typically get around this problem using a ‘failure’ state such as the zero vector, the empty relation or the nonterminating computation.

This will have a major impact on *design for compositionality*, which may become an important research topic in applied compositional game theory. In §3.1 and §3.2 we will show that certain known, simple classes of classical games can be faithfully represented as closed games. The purpose of this, as described in §0.4, is to show formally that compositional game theory is indeed game theory. However, a practical use of compositional game theory will likely look quite different.

In §2.3.9 we represented a bimatrix game as a closed game $\mathcal{G} : I \rightarrow I$. Since the game is closed, there is no nontrivial way in which it can interact with its environment, and consequently there is no nontrivial way to compose it with other games. This is unsurprising, because the point was to produce a faithful model of a classical bimatrix game, and classical bimatrix games do not interact with their environment, that is to say, they cannot be composed.

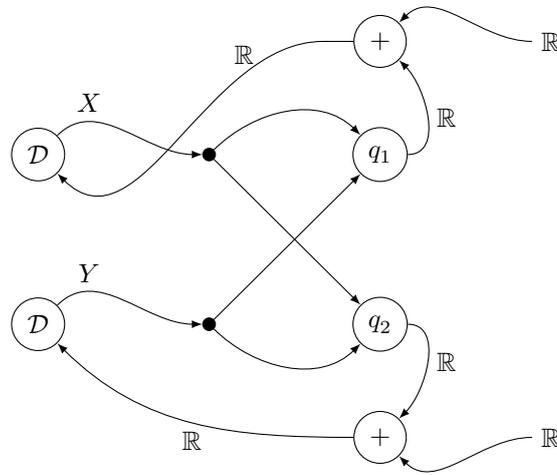
The question of how a bimatrix game should be represented in order to allow compositionality is a nontrivial modelling problem. Motivated by §0.10, we want to model the bimatrix game as an open game so that it could be reused in later, more complicated problems. The simpler half of this problem is what to take for Y and S , the types of values flowing from the game to its environment: a reasonable choice is $Y = X' \otimes Y'$, where X' and Y' are the types of choices made by the two players⁵¹, and $S = I$, because there is no useful sense in which the bimatrix game can generate cointility.

⁵¹In §2.3.9 X' and Y' are called X and Y , and should not be confused with the ‘default’ use of X and Y in §2.1.4 and elsewhere.

2.3. STRING DIAGRAMS

The type X , given suitable features⁵² of the underlying category \mathcal{C} , could vary over all nonempty finite sets, with the game being parametrically polymorphic in X . The decisions, as defined in §2.1.8, do indeed seem to be parametrically polymorphic in X . The non-compositional example in §2.3.9 is recovered by setting $X = I$.

The difficult case is R . The problem is to determine how the utility of two concrete players should be affected by unknown events in the future. The simplest way to deal with this is to take $R = \mathbb{R} \otimes \mathbb{R}$ and use the well-understood ubiquity of real-valued utility as a ‘universal representation’, receiving an additional utility from the future that is simply added to the utility for each player from the game itself. The string diagram corresponding to this description is



One notable drawback of this is that the preferences of the players, which determine the appropriate utilities, become distributed throughout the string diagram rather than being localised. A possible alternative approach would be to have a game that is parametric in both a type R and in a rational preference relation on R . In any particular instantiation, the type and preference relation could be chosen depending on the modelling problem at hand.

⁵²The motivating example is the type system of Haskell (see §0.7). Writing in a parametrically polymorphic way is second nature to a functional programmer.

Chapter 3

Game theory in string diagrams

3.1 Normal form games in string diagrams

3.1.1 Discussion

In [chapter 2](#) of this thesis we have built up an abstract theory, and this will be continued in [§3.3](#) when we study solvability of games. The central argument being laid out (see [§0.10](#)) is that this approach to game theory is *better* than classical approaches, at least insofar as it is more scalable. However a vital piece of the argument is missing: other than the language used in the informal explanations, and the examples without proof at the end of [§2.3](#), we have not demonstrated that open games have any connection *at all* with game theory.

The purpose of this chapter is to show that our theory does agree with classical game theory in certain situations, specifically normal form games with pure and mixed Nash equilibria, and certain extensive form games with subgame perfect equilibria. Given a classical game theoretic situation, we should be able to draw a string diagram that *looks like* the information flow in the target situation, and which moreover denotes *the same* game. In order to discuss sameness of games, without even having a common framework, we will use the approach discussed in [§1.2.4](#): two games will be considered the same if they have the same sets of strategies Σ , and the same best response functions $\mathbf{B} : \Sigma \rightarrow \mathcal{P}\Sigma$. As we saw in [§2.1.6](#), an abstract scalar in $\mathbf{Game}(\mathbf{Set})$ or $\mathbf{Game}(\mathbf{SRel})$ is defined precisely by Σ and \mathbf{B} of this type. Therefore our aim is, given some existing notion of game, to construct an abstract scalar whose strategies and best responses are the same.

Game theory is a large subject, with entire areas devoted to various extensions and special cases of the basic definitions. For this reason there cannot be (and nor should there be) a single theorem that subsumes classical game theory into our new framework. For now, we are going to focus on some simple and common special cases. In this section, we begin with normal form finite games and pure strategies, as introduced in [§1.2](#), and we will prove that every such game can be translated into a string diagram, whose strategies and Nash equilibria are the same. For additional simplicity we will work only with

classical games (see §1.2.5), in order that this section can be understood only in terms of classical game theory, such as is introduced in (Leyton-Brown and Shoham 2008), without requiring knowledge of selection functions. The proofs about **Game(Set)** in this section and §3.2 generalise immediately to arbitrary selection functions, using the translation from multi-valued selection functions to decisions in §2.1.7.

Beginning with §3.1.6, we will also show the same result for normal form games and mixed strategies, which is the setting in which Nash proved his famous existence theorem (Nash 1950).

3.1.2 Tensor products of decisions

Consider a finite normal form classical game, in the sense of §1.2.5. We have $N \geq 1$ players, where the i th player makes a choice from the finite set X_i , and receives a real-valued utility, which she aims to maximise. Since the choices in normal form games take place simultaneously, no player observes any history. Since we are intending to model pure strategies, which involve no side effects, we will work in the category **Game(Set)**. The choice of the i th player will therefore be modelled by the utility-maximising decision

$$\mathcal{D}_i : I \rightarrow X_i \otimes \mathbb{R}^*$$

which is given as the first example in §2.1.8. For completeness, this is defined by the data

- $\Sigma_{\mathcal{D}_i} = X_i$
- $\mathbf{P}_{\mathcal{D}_i} \sigma^* = \sigma$
- $\mathbf{C}_{\mathcal{D}_i} \sigma(x, u) = *$
- $\mathbf{B}_{\mathcal{D}_i} k\sigma = \{\sigma' : X_i \mid k\sigma' \geq kx' \text{ for all } x' : X_i\}$

The purpose of this section is to study the tensor product of finitely many such decisions.

We will work by induction on N , the number of players. Define a game

$$\mathcal{G}_N : I \rightarrow \bigotimes_{i=1}^N X_i \otimes (\mathbb{R}^*)^{\otimes N}$$

where $(\mathbb{R}^*)^{\otimes N}$ denotes the N -fold tensor $\mathbb{R}^* \otimes \cdots \otimes \mathbb{R}^*$, by

$$\mathcal{G}_N = \bigotimes_{i=1}^N \mathcal{D}_i$$

More formally, we recursively define $\mathcal{G}_1 = \mathcal{D}_1$, and $\mathcal{G}_{N+1} = \mathcal{G}_N \otimes \mathcal{D}_{N+1}$.

Our task is to find a closed form description of \mathcal{G}_N . It is simple to see that the set of strategy profiles is a cartesian product

$$\Sigma_{\mathcal{G}_N} = \prod_{i=1}^N X_i$$

The play function, whose type is isomorphic to

$$\mathbf{P}_{\mathcal{G}_N} : \prod_{i=1}^N X_i \rightarrow \prod_{i=1}^N X_i$$

is given by the identity function, because in the inductive case we have

$$\text{id}_{\prod_{i=1}^N X_i} \times \text{id}_{X_{n+1}} = \text{id}_{\prod_{i=1}^{N+1} X_i}$$

Notice that this is equal to the play function \mathbf{P} for normal form games defined in §1.2.2. The coplay function has type isomorphic to

$$\mathbf{C}_{\mathcal{G}_N} : \prod_{i=1}^N X_i \times \mathbb{R}^n \rightarrow 1$$

and so is uniquely defined.

3.1.3 Best response for a tensor of decisions

The best response function has type isomorphic to

$$\mathbf{B}_{\mathcal{G}_N} : \left(\prod_{i=1}^N X_i \rightarrow \mathbb{R}^N \right) \rightarrow \prod_{i=1}^N X_i \rightarrow \mathcal{P} \prod_{i=1}^N X_i$$

and is given by

$$\mathbf{B}_{\mathcal{G}_N} k\sigma = \left\{ \sigma' : \prod_{i=1}^N X_i \mid \begin{array}{l} k_i(\sigma'_i, \sigma_{-i}) \geq k_i(x'_i, \sigma_{-i}) \\ \text{for all } 1 \leq i \leq N, x'_i : X_i \end{array} \right\}$$

where $k_i = \pi_i \circ k$, and the notation (σ_i, σ'_{-i}) was introduced in §1.2.2. We will prove this by induction on N , the number of players.

This is the key step to connect our games with classical normal form games, because it has exactly the same form as the best response function in §1.2.5, and can be stated equivalently as saying that $\mathbf{B}_{\mathcal{G}_N} k\sigma$ is the set of best responses to σ in the N -player normal form game whose outcome function is k .

We begin with the case $N = 1$. Since $\mathcal{G}_1 = \mathcal{D}_1$ is a decision, by definition we directly have

$$\mathbf{B}_{\mathcal{G}_1} k\sigma = \{\sigma' : X_1 \mid k\sigma' \geq kx'_1 \text{ for all } x'_1 : X_1\}$$

To see that this is of the required form, simply note that for $N = 1$ we have $(x_1, \sigma_{-1}) = x_1$, and that the projection $(-)_1$ is the identity.

For the inductive step, by definition of the tensor product,

$$\begin{aligned} \mathbf{B}_{\mathcal{G}_{N+1}} k\sigma &= \mathbf{B}_{\mathcal{G}_N \otimes \mathcal{D}_{N+1}} k(\sigma_{-(N+1)}, \sigma_{N+1}) \\ &= \mathbf{B}_{\mathcal{G}_N} k_{\otimes \sigma_{N+1}} \sigma_{-(N+1)} \times \mathbf{B}_{\mathcal{D}_{N+1}} k_{\sigma_{-(N+1)} \otimes \sigma_{N+1}} \end{aligned}$$

3.1. NORMAL FORM GAMES IN STRING DIAGRAMS

We introduce here the less cluttered notation $k_{\otimes\sigma_{N+1}}$ and $k_{\sigma_{-(N+1)}\otimes}$ as shorthand for $k_{\otimes\sigma_{N+1}(\ast)}$ and $k_{\sigma_{-(N+1)}(\ast)\otimes}$ (see §2.2.6), because all histories are trivial for decisions in normal form games.

Before we proceed further, we must explicitly calculate $k_{\otimes\sigma_N}$ and $k_{\sigma_{-(N+1)}\otimes}$. The original continuation is

$$k : \prod_{i=1}^{N+1} X_i \rightarrow \mathbb{R}^{N+1}$$

which we will write in the isomorphic form

$$k : \prod_{i=1}^N X_i \times X_{N+1} \rightarrow \mathbb{R}^N \times \mathbb{R}$$

The left continuation $k_{\otimes\sigma_{N+1}} : \prod_{i=1}^N X_i \rightarrow \mathbb{R}^n$ is, by the definition in §2.2.6, given by

$$\begin{aligned} \prod_{i=1}^N X_i &\xrightarrow{\cong} \prod_{i=1}^N X_i \times 1 \xrightarrow{\prod_{i=1}^N X_i \times \mathbf{P}_{\mathcal{D}_{N+1}} \sigma_{N+1}} \prod_{i=1}^N X_i \times X_{N+1} \\ &\xrightarrow{k} \mathbb{R}^N \times \mathbb{R} \xrightarrow{\pi_{-(N+1)}} \mathbb{R}^N \end{aligned}$$

Since \mathcal{D}_{N+1} is a decision in **Game(Set)** its play function is

$$\mathbf{P}_{\mathcal{D}_{N+1}} \sigma_{N+1}^\ast = \sigma_{N+1}$$

so this is simply

$$k_{\otimes\sigma_{N+1}} \sigma' = k_{-(N+1)}(\sigma', \sigma_{N+1})$$

Similarly the right continuation $k_{\sigma_{-(N+1)}\otimes} : X_{N+1} \rightarrow \mathbb{R}$ by definition is

$$\begin{aligned} X_{N+1} &\xrightarrow{\cong} 1 \times X_{N+1} \xrightarrow{\mathbf{P}_{\mathcal{G}_N} \sigma_{-(N+1)}^\ast \times X_{N+1}} \prod_{i=1}^N X_i \times X_{N+1} \\ &\xrightarrow{k} \mathbb{R}^N \times \mathbb{R} \xrightarrow{\pi_{N+1}} \mathbb{R} \end{aligned}$$

Since we have already checked in the previous section that

$$\mathbf{P}_{\mathcal{G}_N} \sigma_{-(N+1)}^\ast = \sigma_{-(N+1)}$$

this is explicitly

$$k_{\sigma_{-(N+1)}\otimes} \sigma' = k_{N+1}(\sigma_{-(N+1)}, \sigma')$$

3.1.4 Best response for a tensor of decisions, continued

We can now expand

$$\mathbf{B}_{\mathcal{G}_{N+1}} k\sigma = \mathbf{B}_{\mathcal{G}_N} k_{\otimes\sigma_{N+1}} \sigma_{-(N+1)} \times \mathbf{B}_{\mathcal{D}_{N+1}} k_{\sigma_{-(N+1)}\otimes} \sigma_{N+1}$$

The first term $\mathbf{B}_{\mathcal{G}_N} k_{\otimes\sigma_{N+1}} \sigma_{-(N+1)}$, by the inductive hypothesis, is the set of $\sigma' : \prod_{i=1}^N X_i$ such that for all $1 \leq i \leq N$ and $x'_i : X_i$,

$$(k_{\otimes\sigma_{N+1}})_i(\sigma'_i, (\sigma_{-(N+1)})_{-i}) \geq (k_{\otimes\sigma_{N+1}})_i(x'_i, (\sigma_{-(N+1)})_{-i})$$

Expanding the definition of $k_{\otimes\sigma_{N+1}}$, this is

$$(k_{-(N+1)})_i(\sigma'_i, (\sigma_{-(N+1)})_i, \sigma_{N+1}) \geq (k_{-(N+1)})_i(x'_i, (\sigma_{-(N+1)})_i, \sigma_{N+1})$$

Noticing that the composition of the projections

$$\prod_{i=1}^{N+1} X_i \xrightarrow{\pi_{-(N+1)}} \prod_{i=1}^N X_i \xrightarrow{\pi_i} X_i$$

is itself just a projection, this can be simplified to

$$k_i(\sigma'_i, \sigma_{-i}, \sigma_{N+1}) \geq k_i(x'_i, \sigma_{-i}, \sigma_{N+1})$$

This notation is the obvious extension of that introduced in §1.2.2, namely

$$(x'_i, \sigma_{-i}, \sigma_{N+1})_j = \begin{cases} x'_i & \text{if } j = i \\ \sigma_i & \text{if } 1 \leq j \leq N \text{ and } j \neq i \\ \sigma_{N+1} & \text{if } j = N + 1 \end{cases}$$

Similarly, the second term $\mathbf{B}_{\mathcal{D}_{N+1}} k_{\sigma_{-(N+1)} \otimes \sigma_{N+1}}$ is by definition the set of $\sigma' : X_{N+1}$ such that for all $x'_{N+1} : X_{N+1}$,

$$k_{\sigma_{-(N+1)} \otimes \sigma'} \geq k_{\sigma_{-(N+1)} \otimes x'_{N+1}}$$

and this expands to

$$k_{N+1}(\sigma_{-(N+1)}, \sigma') \geq k_{N+1}(\sigma_{-(N+1)}, x'_{N+1})$$

Putting these together, the cartesian product

$$\mathbf{B}_{\mathcal{G}_{N+1}} k\sigma = \mathbf{B}_{\mathcal{G}_N} k_{\otimes\sigma_{N+1}} \sigma_{-(N+1)} \times \mathbf{B}_{\mathcal{D}_{N+1}} k_{\sigma_{-(N+1)} \otimes \sigma_{N+1}}$$

is equal to the set of $\sigma' : \prod_{i=1}^{N+1} X_i$, such that both of these conditions hold. The first holds for all $i \leq N$, and we notice that the second is equal to the first for $i = N + 1$. This means that the set $\mathbf{B}_{\mathcal{G}_{N+1}} k\sigma$ has exactly the form we required, and we are done.

3.1.5 The payoff functions

It may be that the game \mathcal{G}_N we have constructed is the properly idiomatic representation of a normal form game as an open game, leaving the utility function as a continuation parameter in the best response function. However, if we have in mind some concrete utility function

$$q : \prod_{i=1}^N X_i \rightarrow \mathbb{R}^N$$

then we can use the result proved in §2.2.10, that the closed game

$$\mathcal{H} = \varepsilon_{\mathbb{R}^{\otimes N}} \circ (q \otimes (\mathbb{R}^*)^{\otimes N}) \circ \mathcal{G}_N : I \rightarrow I$$

3.1. NORMAL FORM GAMES IN STRING DIAGRAMS

has the property that

$$\mathbf{B}_{\mathcal{G}_N} q\sigma = \mathbf{B}_{\mathcal{H}} * \sigma$$

Therefore $\Sigma_{\mathcal{H}} = \Sigma_{\mathcal{G}_N}$ and $\mathbf{B}_{\mathcal{H}} : \Sigma_{\mathcal{G}_N} \rightarrow \mathcal{P}\Sigma_{\mathcal{G}_N}$ specify the same strategies and best responses for \mathcal{H} as for the classical normal form game we began with. As discussed in §3.1.1, this is precisely the aim we began with.

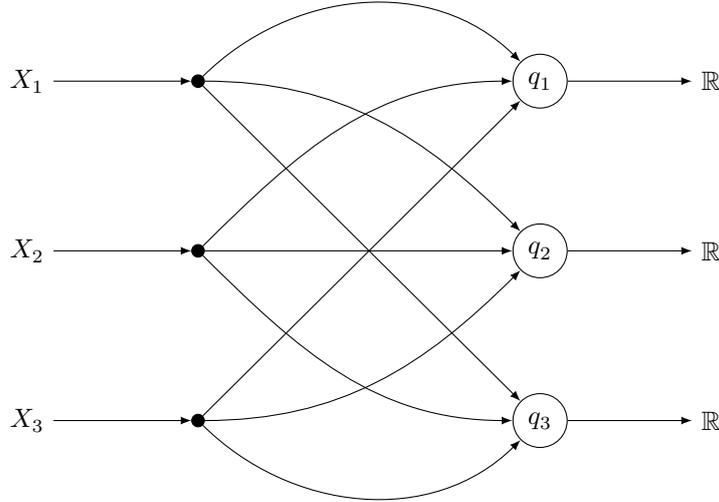
Because computations are a monoidal embedding, we can go further and decompose q , so that the information flow inside q is visible in the string diagram. Typically, q is constructed by beginning with a utility function

$$q_i : \prod_{i=1}^N X_i \rightarrow \mathbb{R}$$

for each player i . The tensor product of these, considered as covariant computations, is

$$\bigotimes_{i=1}^N q_i : \bigotimes_{i=1}^N \bigotimes_{j=1}^N X_j \rightarrow \mathbb{R}^{\otimes N}$$

Since each of the N utility functions requires access to all N choices, we need to copy each choice N times and then braid them into the right order. This is implemented by a copy followed by symmetries. For example, when $N = 3$ we use



In general¹, the copying is implemented by a covariant computation

$$\bigotimes_{i=1}^N X_i \rightarrow \bigotimes_{i=1}^N \bigotimes_{j=1}^N X_j$$

¹Doing this in a somewhat compositional way is possible, but tedious. It is unavoidable to fix the number N , and eventually we will come up against the fact that outcome functions are fundamentally not compositional, because they refer to the totality of all players. Finding ways to avoid this may become a significant research problem in applied compositional game theory.

whose play function is

$$\mathbf{P}_q(x_1, \dots, x_N) = (x_1, \dots, x_N, \dots, x_1, \dots, x_N)$$

Of course, particular features such as common subexpressions can be used to simplify such a string diagram in many examples.

The example game in §2.3.9 has this form, but uses mixed strategies, which we will now develop. In the remainder of this section, and in §3.2, we will work with the open form of a game, leaving the outcome function implicit as a continuation parameter in the best response function.

3.1.6 Stochastic decisions

The result that normal form games can be faithfully represented as open games also holds if we change our solution concept from pure to mixed strategy Nash equilibrium. Mixed strategy Nash equilibria of classical games were introduced in §1.2.6. Following the example in §2.3.9 we now describe the general case, making use of the probability distribution monad \mathcal{D} introduced in §2.1.3.

We will now work in the category $\mathbf{Game}(\mathbf{SRel})$, where \mathbf{SRel} is the Kleisli category of \mathcal{D} . Following §3.1.2, we have a sequence of decisions

$$\mathcal{D}_i : I \rightarrow X_i \otimes \mathbb{R}^*$$

which are defined by the response function

$$\mathbf{R} : (X_i \rightarrow \mathcal{D}\mathbb{R}) \rightarrow \mathcal{P}(\mathcal{D}X_i)$$

given by

$$\mathbf{R}k = \{\sigma : \mathcal{D}X_i \mid \mathbb{E}[k^* \sigma] \geq \mathbb{E}[kx'] \text{ for all } x' : X_i\}$$

These are the expected utility maximising decisions introduced in §2.1.8. To recall, \mathcal{D}_i has the set of strategies

$$\Sigma_{\mathcal{D}_i} = \mathbf{hom}_{\mathbf{SRel}}(I, X_i) = \mathcal{D}X_i$$

Its play function has type

$$\mathbf{P}_{\mathcal{D}_i} : \Sigma_{\mathcal{D}_i} \rightarrow \mathbf{hom}_{\mathbf{SRel}}(I, X_i) = \mathcal{D}X_i \rightarrow \mathcal{D}X_i$$

and is equal to the identity function, and its coplay function

$$\mathbf{C}_{\mathcal{D}_i} : \Sigma_{\mathcal{D}_i} \rightarrow \mathbf{hom}_{\mathbf{SRel}}(I \otimes \mathbb{R}, I) = \mathcal{D}X_i \rightarrow \mathbb{R} \rightarrow \mathcal{D}1$$

and is uniquely defined as the function onto $\mathcal{D}1 = 1$, the only probability distribution on a one-element set. Finally, its best response function is

$$\begin{aligned} \mathbf{B}_{\mathcal{D}_i} &: \mathbf{hom}_{\mathbf{SRel}}(I, I) \times \mathbf{hom}_{\mathbf{SRel}}(X_i, \mathbb{R}) \rightarrow \Sigma_{\mathcal{D}_i} \rightarrow \mathcal{P}\Sigma_{\mathcal{D}_i} \\ &= (X_i \rightarrow \mathcal{D}\mathbb{R}) \rightarrow \mathcal{D}X_i \rightarrow \mathcal{P}(\mathcal{D}X_i) \end{aligned}$$

given by

$$\mathbf{B}_{\mathcal{D}_i}k\sigma = \{\sigma' : \mathcal{D}X_i \mid \mathbb{E}[k^* \sigma] \geq \mathbb{E}[kx'] \text{ for all } x' : X_i\}$$

As in §3.1.2, we will define the game

$$\mathcal{G}_N : I \rightarrow \bigotimes_{i=1}^N X_i \otimes (\mathbb{R}^*)^{\otimes N}$$

by

$$\mathcal{G}_N = \bigotimes_{i=1}^N \mathcal{D}_i$$

By induction we get a closed form for the strategy set

$$\Sigma_{\mathcal{G}_N} = \prod_{i=1}^N \mathcal{D}X_i$$

and also a closed form for the play function

$$\mathbf{P}_{\mathcal{G}_N} : \prod_{i=1}^N \mathcal{D}X_i \rightarrow \mathcal{D} \prod_{i=1}^N X_i$$

given by the monoidal product of the monad \mathcal{D} , in other words,

$$\mathbb{P} [\mathbf{P}_{\mathcal{G}_N} \sigma = x] = \prod_{i=1}^N \mathbb{P} [\sigma_i = x_i]$$

The coplay function is uniquely defined by its type

$$\mathbf{C}_{\mathcal{G}_N} : \prod_{i=1}^N \mathcal{D}X_i \rightarrow \mathbb{R}^N \rightarrow \mathcal{D}1$$

again because $\mathcal{D}1 = 1$.

3.1.7 Best response for a tensor of stochastic decisions

We will prove by induction that the best response function

$$\mathbf{B}_{\mathcal{G}_N} : \left(\prod_{i=1}^N X_i \rightarrow \mathbb{R}^n \right) \rightarrow \prod_{i=1}^N \mathcal{D}X_i \rightarrow \mathcal{D} \prod_{i=1}^N \mathcal{D}X_i$$

has the closed form

$$\mathbf{B}_{\mathcal{G}_N} k\sigma = \left\{ \sigma' : \prod_{i=1}^N \mathcal{D}X_i \mid \mathbb{E}[k_i(\sigma'_i, \sigma_{-i})] \geq \mathbb{E}[k_i(x'_i, \sigma_{-i})] \right. \\ \left. \text{for all } x'_i : X_i, 1 \leq i \leq N \right\}$$

This is exactly the best response function for games with mixed strategies given in §1.2.6. In the case $N = 1$, it is

$$\mathbf{B}_{\mathcal{G}_N} k\sigma = \{ \sigma' : \mathcal{D}X_1 \mid \mathbb{E}[k\sigma'] \geq \mathbb{E}[kx'_1] \text{ for all } x'_1 : X_1 \}$$

which holds by definition.

Note that from this point, we begin to abuse notation by leaving the Kleisli extension of \mathcal{D} implicit. This is standard practice in mathematics, where for example an ordinary non-stochastic function f can be applied to a random variable α , resulting in a random variable $f\alpha$.

By the same reasoning as in §3.1.3, the condition

$$(\sigma'_{-(N+1)}, \sigma'_{N+1}) \in \mathbf{B}_{\mathcal{G}_{N+1}} k(\sigma_{-(N+1)}, \sigma_{N+1})$$

is equivalent to the pair of conditions

$$\sigma'_{-(N+1)} \in \mathbf{B}_{\mathcal{G}_N} k_{\otimes \sigma_{N+1}} \sigma_{-(N+1)}$$

and

$$\sigma'_{N+1} \in \mathbf{B}_{\mathcal{D}_{N+1}} k_{\sigma_{-(N+1)} \otimes \sigma_{N+1}}$$

The former, by the inductive hypothesis, is equivalent to the claim that for all $1 \leq i \leq N$ and $x'_i : X_i$,

$$\mathbb{E}[(k_{\otimes \sigma_{N+1}})_i((\sigma'_{-(N+1)})_i, (\sigma_{-(N+1)})_{-i})] \geq \mathbb{E}[(k_{\otimes \sigma_{N+1}})_i(x'_i, (\sigma_{-(N+1)})_{-i})]$$

The latter, by definition of \mathcal{D}_{N+1} , is equivalent to the claim that for all deviations $x'_{N+1} : X_{N+1}$,

$$\mathbb{E}[k_{\sigma_{-(N+1)} \otimes \sigma'_{N+1}}] \geq \mathbb{E}[k_{\sigma_{-(N+1)} \otimes x'_{N+1}}]$$

We must prove that for all $1 \leq i \leq N+1$ and $x'_i : X_i$,

$$\mathbb{E}[k_i(\sigma'_i, \sigma_{-i})] \geq \mathbb{E}[k_i(x'_i, \sigma_{-i})]$$

In order to obtain this from the inductive hypothesis, it remains to prove that

$$(k_{\otimes \sigma_{N+1}})_i(x'_i, (\sigma_{-(N+1)})_{-i}) = k_i(x'_i, \sigma_{-i})$$

for $1 \leq i \leq N$ and $x'_i : X_i$, and

$$k_{\sigma_{-(N+1)} \otimes x'_{N+1}} = k_{N+1}(\sigma_{-(N+1)}, x'_{N+1})$$

for $x'_{N+1} : X_{N+1}$. Continuing to abuse notation, for the former we can say that

$$\begin{aligned} (k_{\otimes \sigma_{N+1}})_i(x'_i, (\sigma_{-(N+1)})_{-i}) &= (k_{-(N+1)})_i(x'_i, (\sigma_{-(N+1)})_{-i}, \sigma_{N+1}) \\ &= k_i(x'_i, \sigma_{-i}) \end{aligned}$$

and similarly,

$$k_{\sigma_{-(N+1)} \otimes x'_{N+1}} = k_{N+1}(\sigma_{-(N+1)}, x'_{N+1})$$

holds by definition. More carefully, both sides of the former are given by

$$\mathbb{P}[k_i(x'_i, \sigma_{-i}) = u_i] = \sum_{\substack{u_{-i} : \mathbb{R}^N \\ x_{-i} : \prod_{j \neq i} X_j}} \left(\prod_{j \neq i} \mathbb{P}[\sigma_j = x_j] \right) \cdot \mathbb{P}[k(x'_i, x_{-i}) = u]$$

and the latter by

$$\begin{aligned} &\mathbb{P}[k_{N+1}(\sigma_{-(N+1)}, x'_{N+1}) = u_{N+1}] \\ &= \sum_{\substack{u_{-(N+1)} : \mathbb{R}^N \\ x_{-(N+1)} : \prod_{j=1}^N X_j}} \left(\prod_{j=1}^N \mathbb{P}[\sigma_j = x_j] \right) \cdot \mathbb{P}[k(x_{-(N+1)}, x'_{N+1}) = u] \end{aligned}$$

3.2 Extensive form

3.2.1 Discussion

In §3.1 we showed that normal form games can be translated into abstract scalars in **Game(Set)** or **Game(SRel)** in a way that preserves the pure or mixed strategies and Nash equilibria, respectively. In this chapter, we will focus on translating extensive form games into abstract scalars in **Game(Set)** while preserving pure strategies and subgame perfect equilibria.

We give a complete proof for the case of sequential games, that is, perfect information games in which the type of choices of a player may not depend on the values of previous choices (see §1.3.1 and §1.3.4). In §3.2.7 we additionally discuss the representation of imperfect information. Giving a general translation of an arbitrary extensive form game into an abstract scalar of **Game(Set)** would be more work, and is left for the future.

We will see in §3.3.7 that using the same method to translate extensive form games with mixed strategies into **Game(SRel)** *goes wrong*, in the sense that it gives an implausible solution concept, and that the most fundamental definitions in §2.1 and §2.2 will need to be changed as a result. This is why we restrict to pure strategies in this chapter.

This section is, technically speaking, the most difficult of this thesis, with §3.2.4 and §3.2.5 being particularly dense. Verifying these proofs is made more difficult by the fact that it is necessary for readability to leave many isomorphisms of sets implicit, such as when using the notation (x'_i, x_{-i}) for unilateral deviation.² An implementation in a proof assistant such as Coq would be useful, and will be even more important when scaling this style of proof beyond such simple types of games as are considered in this thesis.³

3.2.2 Composition with perfect information

We will begin with a sequence of decisions

$$\mathcal{D}_i : \bigotimes_{j=1}^{i-1} X_j \rightarrow X_i \otimes \mathbb{R}^*$$

in **Game(Set)**. The i th decision represents a player who observes the first $i - 1$ choices. Note that since the empty tensor is the identity, the first of these is $\mathcal{D}_1 : I \rightarrow X_1 \otimes \mathbb{R}^*$. We will define a game

$$\mathcal{G}_N : I \rightarrow \bigotimes_{i=1}^N X_i \otimes (\mathbb{R}^*)^{\otimes N}$$

which behaves like their perfect information composition, by induction on N . Like for simultaneous games, the base case is $\mathcal{G}_1 = \mathcal{D}_1$.

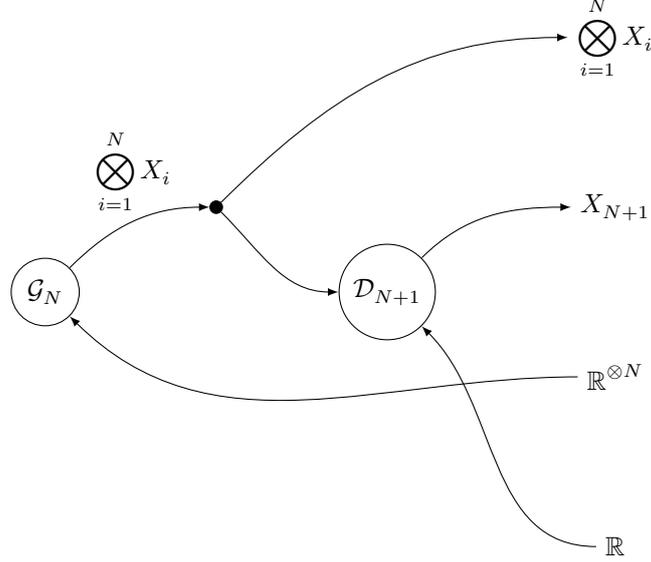
For the inductive step, we are given \mathcal{G}_N and \mathcal{D}_{N+1} , and we must produce \mathcal{G}_{N+1} . In this game, \mathcal{G}_N must play first, and then \mathcal{D}_{N+1} must play observing the choices of \mathcal{G}_N . However, the choices of \mathcal{G}_N must also be preserved in the output

²These are precisely the isomorphisms that need to be explicit in Haskell code, as discussed in §0.7.

³See also comments in §0.2 on certified mathematics in economics.

3.2. EXTENSIVE FORM

of \mathcal{G}_{N+1} , and hence must be copied in parallel with \mathcal{D}_{N+1} . We will define this composition by the string diagram



We will translate this diagram into monoidal category notation in stages. The decision is

$$\mathcal{D}_{N+1} : \bigotimes_{i=1}^N X_i \rightarrow X_{N+1} \otimes \mathbb{R}^*$$

To add the string above it we tensor with the identity

$$\bigotimes_{i=1}^N X_i \otimes \mathcal{D}_{N+1} : \bigotimes_{i=1}^N X_i \otimes \bigotimes_{i=1}^N X_i \rightarrow \bigotimes_{i=1}^{N+1} X_i \otimes \mathbb{R}^*$$

The copying node is the computation

$$\Delta_{\prod_{i=1}^N X_i} : \bigotimes_{i=1}^N X_i \rightarrow \bigotimes_{i=1}^N X_i \otimes \bigotimes_{i=1}^N X_i$$

The top-right part of the diagram is therefore

$$\left(\bigotimes_{i=1}^N X_i \otimes \mathcal{D}_{N+1} \right) \circ \Delta_{\prod_{i=1}^N X_i} : \bigotimes_{i=1}^N X_i \rightarrow \bigotimes_{i=1}^{N+1} X_i \otimes \mathbb{R}^*$$

Before the lower-right crossing, we can tensor with the lower $\mathbb{R}^{\otimes N}$ -labelled contravariant string to get

$$\begin{aligned} & \left(\left(\bigotimes_{i=1}^N X_i \otimes \mathcal{D}_{N+1} \right) \circ \Delta_{\prod_{i=1}^N X_i} \right) \otimes (\mathbb{R}^*)^{\otimes N} \\ & : \bigotimes_{i=1}^N X_i \otimes (\mathbb{R}^*)^{\otimes N} \rightarrow \bigotimes_{i=1}^{N+1} X_i \otimes \mathbb{R}^* \otimes (\mathbb{R}^*)^{\otimes N} \end{aligned}$$

The crossing is the dual of the computation

$$\sigma_{\mathbb{R}^N, \mathbb{R}} : \mathbb{R}^{\otimes(N+1)} \rightarrow \mathbb{R} \otimes \mathbb{R}^{\otimes N}$$

and we must also tensor with the identity above it, to give

$$\bigotimes_{i=1}^{N+1} X_i \otimes \sigma_{\mathbb{R}^N, \mathbb{R}}^* : \bigotimes_{i=1}^{N+1} X_i \otimes \mathbb{R}^* \otimes (\mathbb{R}^*)^{\otimes N} \rightarrow \bigotimes_{i=1}^{N+1} X_i \otimes (\mathbb{R}^*)^{\otimes(N+1)}$$

Finally, we can give the denotation of the entire string diagram as

$$\mathcal{G}_{N+1} = \mathcal{H}_{N+1} \circ \mathcal{G}_N : I \rightarrow \bigotimes_{i=1}^{N+1} X_i \otimes (\mathbb{R}^*)^{\otimes(N+1)}$$

where

$$\mathcal{H} : \bigotimes_{i=1}^N X_i \otimes (\mathbb{R}^*)^{\otimes N} \rightarrow \bigotimes_{i=1}^{N+1} X_i \otimes (\mathbb{R}^*)^{\otimes(N+1)}$$

is given by

$$\mathcal{H} = \left(\bigotimes_{i=1}^{N+1} X_i \otimes \sigma_{\mathbb{R}^N, \mathbb{R}}^* \right) \circ \left(\left(\left(\bigotimes_{i=1}^N X_i \otimes \mathcal{D}_{N+1} \right) \circ \Delta_{\prod_{i=1}^N X_i} \right) \otimes (\mathbb{R}^*)^{\otimes N} \right)$$

3.2.3 Building the composition

Since the strategy sets of the decisions are

$$\Sigma_{\mathcal{D}_i} = \prod_{j=1}^{i-1} X_j \rightarrow X_i$$

it can easily be seen that the strategy sets of the \mathcal{G}_i are isomorphic to the cartesian product of these,

$$\Sigma_{\mathcal{G}_N} = \prod_{i=1}^N \left(\prod_{j=1}^{i-1} X_j \rightarrow X_i \right)$$

The coplay function of \mathcal{G}_N has type

$$\mathbf{C}_{\mathcal{G}_N} : \prod_{i=1}^N \left(\prod_{j=1}^{i-1} X_j \rightarrow X_i \right) \rightarrow 1 \rightarrow 1$$

and hence is uniquely defined.

The play function has type

$$\mathbf{P}_{\mathcal{G}_N} : \prod_{i=1}^N \left(\prod_{j=1}^{i-1} X_j \rightarrow X_i \right) \rightarrow \prod_{i=1}^N X_i$$

We will prove that $\mathbf{P}_{\mathcal{G}_N} = \mathbf{P}$ is the play function for sequential games introduced in §1.3.4. Explicitly, we will prove that

$$(\mathbf{P}_{\mathcal{G}_N} \sigma)_i = \sigma_i((\mathbf{P}_{\mathcal{G}_N} \sigma)_1, \dots, (\mathbf{P}_{\mathcal{G}_N} \sigma)_{i-1})$$

3.2. EXTENSIVE FORM

for all $1 \leq i \leq N$. This should be compared to the defining equation of the play function in §1.3.4. In the base case $\mathcal{G}_1 = \mathcal{D}_1$, the play function is by definition

$$(\mathbf{P}_{\mathcal{D}_1}\sigma)_1 = \mathbf{P}_{\mathcal{D}_1}\sigma = \sigma = \sigma_1$$

which satisfies the equation.

For the inductive step, we need to show that

$$\mathbf{P}_{\mathcal{G}_{N+1}}\sigma = (x, \sigma_{N+1}x)$$

where $x = \mathbf{P}_{\mathcal{G}_N}\sigma_{-(N+1)}$. Since

$$\mathbf{P}_{\mathcal{G}_{N+1}}\sigma = \mathbf{P}_{\mathcal{H}_{N+1}}\sigma_{N+1} \circ \mathbf{P}_{\mathcal{G}_N}\sigma_{-(N+1)}$$

it suffices to prove that

$$\mathbf{P}_{\mathcal{H}_{N+1}}\sigma_{N+1}x = (x, \sigma_{N+1}x)$$

To begin with, we note that

$$\mathbf{P}_{\otimes_{i=1}^N X_i \otimes \sigma_{\mathbb{R}^N, \mathbb{R}}^*} = \text{id}_{\prod_{i=1}^N X_i}$$

and so

$$\begin{aligned} \mathbf{P}_{\mathcal{H}_{N+1}}\sigma_{N+1} &= \mathbf{P}_{\left(\left(\otimes_{i=1}^N X_i \otimes \mathcal{D}_{N+1}\right) \circ \Delta_{\prod_{i=1}^N X_i}\right) \otimes (\mathbb{R}^*)^{\otimes N} \sigma_{N+1}} \\ &= \mathbf{P}_{\left(\otimes_{i=1}^N X_i \otimes \mathcal{D}_{N+1}\right) \circ \Delta_{\prod_{i=1}^N X_i} \sigma_{N+1}} \end{aligned}$$

Then

$$\begin{aligned} \mathbf{P}_{\left(\otimes_{i=1}^N X_i \otimes \mathcal{D}_{N+1}\right) \circ \Delta_{\prod_{i=1}^N X_i} \sigma_{N+1}x} &= \mathbf{P}_{\otimes_{i=1}^N X_i \otimes \mathcal{D}_{N+1}} \sigma_{N+1} \left(\mathbf{P}_{\Delta_{\prod_{i=1}^N X_i}} x \right) \\ &= \mathbf{P}_{\otimes_{i=1}^N X_i \otimes \mathcal{D}_{N+1}} \sigma_{N+1} \left(\Delta_{\prod_{i=1}^N X_i} x \right) \\ &= \mathbf{P}_{\otimes_{i=1}^N X_i \otimes \mathcal{D}_{N+1}} (x, x) \\ &= \left(\mathbf{P}_{\otimes_{i=1}^N X_i} x, \mathbf{P}_{\mathcal{D}_{N+1}} \sigma_{N+1} x \right) \\ &= (x, \sigma_{N+1}x) \end{aligned}$$

Having proved this, we have

$$(\mathbf{P}_{\mathcal{G}_{N+1}}\sigma)_i = \begin{cases} (\mathbf{P}_{\mathcal{G}_N}\sigma_{-(N+1)})_i & \text{if } i \leq N \\ \sigma_{N+1}(\mathbf{P}_{\mathcal{G}_N}\sigma_{-(N+1)}) & \text{if } i = N+1 \end{cases}$$

For $i \leq N$, the inductive hypothesis gives

$$\begin{aligned} (\mathbf{P}_{\mathcal{G}_{N+1}}\sigma)_i &= (\mathbf{P}_{\mathcal{G}_N}\sigma_{-(N+1)})_i \\ &= (\sigma_{-(N+1)})_i((\mathbf{P}_{\mathcal{G}_N}\sigma_{-(N+1)})_1, \dots, (\mathbf{P}_{\mathcal{G}_N}\sigma_{-(N+1)})_{i-1}) \\ &= \sigma_i((\mathbf{P}_{\mathcal{G}_{N+1}}\sigma)_1, \dots, (\mathbf{P}_{\mathcal{G}_{N+1}}\sigma)_{i-1}) \end{aligned}$$

and for $i = N+1$,

$$\begin{aligned} (\mathbf{P}_{\mathcal{G}_{N+1}}\sigma)_{N+1} &= \sigma_{N+1}(\mathbf{P}_{\mathcal{G}_N}\sigma_{-(N+1)}) \\ &= \sigma_{N+1}((\mathbf{P}_{\mathcal{G}_N}\sigma_{-(N+1)})_1, \dots, (\mathbf{P}_{\mathcal{G}_N}\sigma_{-(N+1)})_N) \\ &= \sigma_{N+1}((\mathbf{P}_{\mathcal{G}_{N+1}}\sigma)_1, \dots, (\mathbf{P}_{\mathcal{G}_{N+1}}\sigma)_N) \end{aligned}$$

3.2.4 Best response for the composition

We will now prove that our best response function

$$\mathbf{B}_{\mathcal{G}_N} : \left(\prod_{i=1}^N X_i \rightarrow \mathbb{R}^N \right) \rightarrow \Sigma_{\mathcal{G}_N} \rightarrow \mathcal{P}\Sigma_{\mathcal{G}_N}$$

where

$$\Sigma_{\mathcal{G}_N} = \prod_{i=1}^N \left(\prod_{j=1}^{i-1} X_j \rightarrow X_i \right)$$

is such that $\mathbf{B}_{\mathcal{G}_N} k\sigma$ is the set of all $\sigma' : \Sigma_{\mathcal{G}_N}$ such that, for all players $1 \leq i \leq N$ and all partial plays $x = x_1, \dots, x_{i-1}$ of length $i-1$, and all possible deviations $x_i : X_i$, we have

$$(\mathcal{U}_x^k \sigma(\sigma'_i x))_i \geq (\mathcal{U}_x^k \sigma x_i)_i$$

where \mathcal{U}_x^k is the unilateral deviation operator for sequential games, defined in §1.3.5.

The proof is again by induction on N . In the base case we have $N = 1$ and we need only check $i = 1$ with the empty partial play and some deviation $x_1 : X_1$. Then

$$(\mathcal{U}_x^k \sigma \sigma')_1 = k\sigma'$$

and

$$(\mathcal{U}_x^k \sigma x_1)_1 = kx_1$$

and we are done, because by definition $\mathbf{B}_{\mathcal{G}_1} = \mathbf{B}_{\mathcal{D}_1}$ takes $k : X_1 \rightarrow \mathbb{R}$ and $\sigma : X_1$ to the set of all $\sigma' : X_1$ such that

$$k\sigma' \geq kx'$$

for all $x' : X_1$

Next, by the definition of composition,

$$\begin{aligned} \mathbf{B}_{\mathcal{G}_{N+1}} k\sigma &= \mathbf{B}_{\mathcal{G}_N} k_{\sigma_{N+1} \circ \sigma_{-(N+1)}} \\ &\times \bigcap_{\sigma' : \Sigma_{\mathcal{G}}} \mathbf{B}_{\mathcal{H}_{N+1}} (\mathbf{P}_{\mathcal{G}_N} \sigma', k) \sigma_{N+1} \end{aligned}$$

For $\mathbf{B}_{\mathcal{G}_{N+1}} k\sigma$ to contain some σ' , it must be that

$$\sigma'_{-(N+1)} \in \mathbf{B}_{\mathcal{G}_N} k_{\sigma_{N+1} \circ \sigma_{-(N+1)}}$$

and for all $\sigma''_{-(N+1)}$,

$$\sigma'_{N+1} \in \mathbf{B}_{\mathcal{H}_{N+1}} (\mathbf{P}_{\mathcal{G}_N} \sigma''_{-(N+1)}, k) \sigma_{N+1}$$

The first condition, by the inductive hypothesis, is equivalent to the claim that for all $1 \leq i \leq N$, partial plays $x = x_1, \dots, x_{i-1} : \prod_{j < i} X_j$ and deviations $x_i : X_i$,

$$(\mathcal{U}_x^{k_{\sigma_{N+1} \circ \sigma_{-(N+1)}}} \sigma_{-(N+1)}((\sigma'_{-(N+1)})_i x))_i \geq (\mathcal{U}_x^{k_{\sigma_{N+1} \circ \sigma_{-(N+1)}}} \sigma_{-(N+1)} x_i)_i$$

3.2. EXTENSIVE FORM

We claim that there is an equality of continuations

$$\mathcal{U}_x^{k_{\sigma_{N+1} \circ}} \sigma_{-(N+1)} = \pi_{-(N+1)} \circ \mathcal{U}_x^k \sigma : X_i \rightarrow \mathbb{R}^N$$

which immediately gives us the inductive hypothesis for $N+1$ when $1 \leq i \leq N$. To show this, by definition

$$k_{\sigma_{N+1} \circ} : \prod_{j=1}^N X_j \rightarrow \mathbb{R}^N$$

is given by

$$k_{\sigma_{N+1} \circ} x = k_{-(N+1)}(x, \sigma_{N+1} x)$$

Now we have

$$\begin{aligned} \mathcal{U}_x^{k_{\sigma_{N+1} \circ}} \sigma_{-(N+1)} x_i &= k_{\sigma_{N+1} \circ}(\nu_{x, x_i}^{\sigma_{-(N+1)}}) \\ &= k_{-(N+1)}(\nu_{x, x_i}^{\sigma_{-(N+1)}}, \sigma_{N+1}(\nu_{x, x_i}^{\sigma_{-(N+1)}})) \end{aligned}$$

where ν is defined in §1.3.5, because by definition

$$\mathcal{U}_x^k \sigma x_i = k(\nu_{x, x_i}^\sigma)$$

Therefore it suffices to prove that

$$\nu_{x, x_i}^\sigma = (\nu_{x, x_i}^{\sigma_{-(N+1)}}, \sigma_{N+1}(\nu_{x, x_i}^{\sigma_{-(N+1)}}))$$

By definition

$$(\nu_{x, x_i}^\sigma)_{N+1} = \sigma_{N+1}((\nu_{x, x_i}^\sigma)_1, \dots, (\nu_{x, x_i}^\sigma)_N)$$

so it suffices to prove that

$$(\nu_{x, x_i}^\sigma)_j = (\nu_{x, x_i}^{\sigma_{-(N+1)}})_j$$

for $1 \leq j \leq N$. This is easily proved by strong induction on j .

3.2.5 Best response for the composition, continued

We now return to the second condition for

$$\sigma' \in \mathbf{B}_{\mathcal{G}_{N+1}} k \sigma$$

namely that for all $\sigma''_{-(N+1)} : \Sigma_{\mathcal{G}_N}$,

$$\sigma'_{N+1} \in \mathbf{B}_{\mathcal{H}_{N+1}}(\mathbf{P}_{\mathcal{G}_N} \sigma''_{-(N+1)}, k) \sigma_{N+1}$$

We claim that this is equivalent to

$$\sigma'_{N+1} \in \mathbf{B}_{\mathcal{D}_{N+1}}(\mathbf{P}_{\mathcal{G}_N} \sigma''_{-(N+1)}, k') \sigma_{N+1}$$

where $k' : X_{N+1} \rightarrow \mathbb{R}$ is given by

$$k' x_{N+1} = k_{N+1}(\mathbf{P}_{\mathcal{G}_N} \sigma''_{-(N+1)}, x_{N+1})$$

3.2. EXTENSIVE FORM

We will prove the slightly more general claim where $\mathbf{P}_{\mathcal{G}_N} \sigma''_{-(N+1)}$ is replaced with a general play $x_{-(N+1)}$. To begin, since $\bigotimes_{i=1}^{N+1} X_i \otimes \sigma_{\mathbb{R}^N, \mathbb{R}}^*$ is strategically trivial,

$$\begin{aligned} & \mathbf{B}_{\mathcal{H}_{N+1}}(x_{-(N+1)}, k) \sigma_{N+1} \\ &= \mathbf{B} \left(\left(\bigotimes_{i=1}^N X_i \otimes \mathcal{D}_{N+1} \right) \circ \Delta_{\prod_{i=1}^N X_i} \right) \otimes (\mathbb{R}^*)^{\otimes N} (x_{-(N+1)}, k_{*o}) \sigma_{N+1} \end{aligned}$$

where

$$k_{*o} : \prod_{i=1}^{N+1} X_i \rightarrow \mathbb{R}^{N+1}$$

is the extended continuation defined in §2.2.2, given explicitly by

$$\begin{aligned} k_{*o} x &= \mathbf{C}_{\bigotimes_{i=1}^{N+1} X_i \otimes \sigma_{\mathbb{R}^N, \mathbb{R}}^*} (x, k(\mathbf{P}_{\bigotimes_{i=1}^{N+1} X_i \otimes \sigma_{\mathbb{R}^N, \mathbb{R}}^*} x)) \\ &= \sigma_{\mathbb{R}^N, \mathbb{R}} (k(\text{id}_{\prod_{i=1}^N X_i} x)) \\ &= (k_{N+1} x, k_{-(N+1)} x) \end{aligned}$$

Next, because the identity on $(\mathbb{R}^*)^{\otimes N}$ is strategically trivial,

$$\begin{aligned} & \mathbf{B} \left(\left(\bigotimes_{i=1}^N X_i \otimes \mathcal{D}_{N+1} \right) \circ \Delta_{\prod_{i=1}^N X_i} \right) \otimes (\mathbb{R}^*)^{\otimes N} (x_{-(N+1)}, k_{*o}) \sigma_{N+1} \\ &= \mathbf{B}_{\left(\bigotimes_{i=1}^N X_i \otimes \mathcal{D}_{N+1} \right) \circ \Delta_{\prod_{i=1}^N X_i}} (x_{-(N+1)}, (k_{*o})_1) \sigma_{N+1} \end{aligned}$$

where

$$(k_{*o})_1 x = k_{N+1} x$$

Therefore

$$\begin{aligned} & \mathbf{B}_{\left(\bigotimes_{i=1}^N X_i \otimes \mathcal{D}_{N+1} \right) \circ \Delta_{\prod_{i=1}^N X_i}} (x_{-(N+1)}, k_{N+1}) \sigma_{N+1} \\ &= \mathbf{B}_{\bigotimes_{i=1}^N X_i \otimes \mathcal{D}_{N+1}} (\mathbf{P}_{\Delta_{\prod_{i=1}^N X_i}} x_{-(N+1)}, k_{N+1}) \sigma_{N+1} \\ &= \mathbf{B}_{\bigotimes_{i=1}^N X_i \otimes \mathcal{D}_{N+1}} (\Delta_{\prod_{i=1}^N X_i} x_{-(N+1)}, k_{N+1}) \sigma_{N+1} \\ &= \mathbf{B}_{\bigotimes_{i=1}^N X_i \otimes \mathcal{D}_{N+1}} ((x_{-(N+1)}, x_{-(N+1)}), k_{N+1}) \sigma_{N+1} \\ &= \mathbf{B}_{\mathcal{D}_{N+1}} (x_{-(N+1)}, (k_{N+1})_{*(x_{-(N+1)}) \otimes}) \sigma_{N+1} \end{aligned}$$

where

$$\begin{aligned} (k_{N+1})_{*(x_{-(N+1)}) \otimes} x_{N+1} &= k_{N+1} (\mathbf{P}_{\bigotimes_{i=1}^N X_i} x_{-(N+1)}, x_{N+1}) \\ &= k_{N+1} (x_{-(N+1)}, x_{N+1}) \\ &= k' x_{N+1} \end{aligned}$$

for the definition of k' given earlier.

Now, by definition of $\mathbf{B}_{\mathcal{D}_{N+1}}$, the condition

$$\sigma'_{N+1} \in \mathbf{B}_{\mathcal{D}_{N+1}} (\mathbf{P}_{\mathcal{G}_N} \sigma''_{-(N+1)}, k') \sigma_{N+1}$$

is equivalent to

$$k' (\sigma'_{N+1} (\mathbf{P}_{\mathcal{G}_N} \sigma''_{-(N+1)})) \geq k' x_{N+1}$$

3.2. EXTENSIVE FORM

for all $x_{N+1} : X_{N+1}$, which is

$$k_{N+1}(\mathbf{P}_{\mathcal{G}_N} \sigma''_{-(N+1)}, \sigma'_{N+1}(\mathbf{P}_{\mathcal{G}_N} \sigma''_{-(N+1)})) \geq k_{N+1}(\mathbf{P}_{\mathcal{G}_N} \sigma''_{-(N+1)}, x_{N+1})$$

It remains to show that this, universally quantified over $\sigma''_{-(N+1)} : \Sigma_{\mathcal{G}_N}$, implies the remaining case of the inductive hypothesis for $N + 1$, namely that for all partial plays $x = x_1, \dots, x_N : \prod_{j=1}^N X_j$ and deviations $x_{N+1} : X_{N+1}$,

$$(\mathcal{U}_x^k \sigma(\sigma'_{N+1} x))_{N+1} \geq (\mathcal{U}_x^k \sigma x_{N+1})_{N+1}$$

Notice that

$$(\mathcal{U}_x^k \sigma x_{N+1})_{N+1} = k_{N+1}(\nu_{x, x_{N+1}}^\sigma) = k_{N+1}(x, x_{N+1}) = k' x_{N+1}$$

so we must show

$$k_{N+1}(x, \sigma'_{N+1} x) \geq k_{N+1}(x, x_{N+1})$$

for all partial plays x .

Given such an x , we take the strategy profile

$$(\sigma''_{-(N+1)})_i x' = x_i$$

which ignores its observation, and simply plays the move from x . We claim that $\mathbf{P}_{\mathcal{G}_N} \sigma'' = x$. Using the characterisation of $\mathbf{P}_{\mathcal{G}_N}$ in §3.2.3, for each $1 \leq i \leq N$ we have

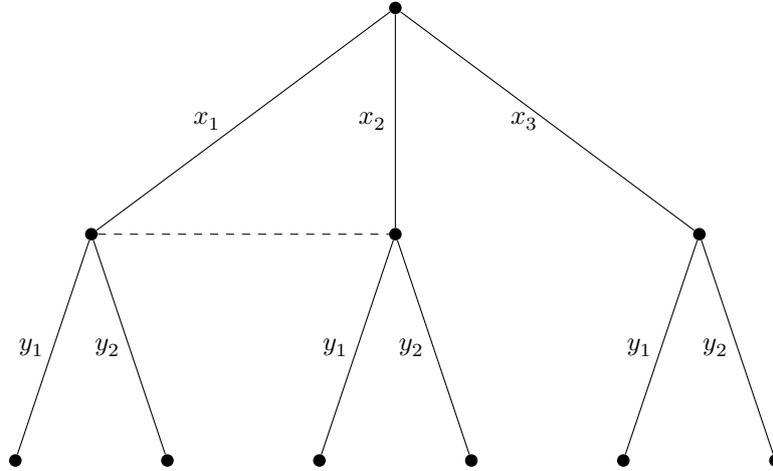
$$\begin{aligned} (\mathbf{P}_{\mathcal{G}_N} \sigma''_{-(N+1)})_i &= (\sigma''_{-(N+1)})_i((\mathbf{P}_{\mathcal{G}_N} \sigma''_{-(N+1)})_1, \dots, (\mathbf{P}_{\mathcal{G}_N} \sigma''_{-(N+1)})_{i-1}) \\ &= x_i \end{aligned}$$

This completes the proof.

3.2.6 Information sets

Information sets are used in (von Neumann and Morgenstern 1944) to give a general theory of imperfect information, including games that are intermediate between sequential and simultaneous. In general, a player may be able to observe partial information about another player's earlier move. That is, they may be able to distinguish some pairs of moves, but not others.

For simplicity, we will begin with a concrete two-player game. Suppose the set of moves for the first player is $X = \{x_1, x_2, x_3\}$ and the set of moves for the second player is $Y = \{y_1, y_2\}$. The game tree for this game is



An **information set** for the second player is simply a partitioning of X , that is to say, an equivalence relation on X . The intention is that two elements of X that are in the same equivalence class cannot be distinguished by the second player, whereas elements in different equivalence classes can be. In this example, the dashed line is the usual notation for specifying that two nodes are in the same equivalence class. Thus, our partitioning of X is $X/\sim = \{\{x_1, x_2\}, \{x_3\}\}$.

For simplicity, we will focus on pure strategies, although sequential equilibrium (Leyton-Brown and Shoham 2008) is a more common solution concept for games with nontrivial information sets (and see §3.3.7). The informal condition that the second player cannot distinguish between x_1 and x_2 can be interpreted formally as a restriction on the allowed strategies, namely that every strategy for player 2 must be compatible with the equivalence relation. Thus, a function $\sigma_2 : X \rightarrow Y$ is an allowed strategy for player 2 iff $\sigma_2 x_1 = \sigma_2 x_2$. This is equivalent to saying that σ_2 is actually a function $\sigma_2 : X/\sim \rightarrow Y$. The set of strategies of this game, therefore, is

$$\Sigma = X \times (X/\sim \rightarrow Y)$$

The general definition of a subgame perfect equilibrium, including with non-trivial information sets, is a strategy that induces a Nash equilibrium on every subgame. For this two-stage game, this is the same as a fixpoint of the best response function $\mathbf{B} : \Sigma \rightarrow \mathcal{P}\Sigma$, where $\mathbf{B}(\sigma_1, \sigma_2)$ is the set of all (σ'_1, σ'_2) satisfying the two conditions⁴

$$q_1(\sigma'_1, \sigma_2[\sigma'_1]) \geq q_1(x', \sigma_2[x'])$$

for all $x' : X$, and

$$q_2(x', \sigma'_2[x']) \geq q_2(x', y')$$

for all $x' : X$ and $y' : Y$. Equivalently, we can give the two conditions as

$$\sigma'_1 = \arg \max_{x':X} q_1(x', \sigma_2[x'])$$

and for all $x' : X$,

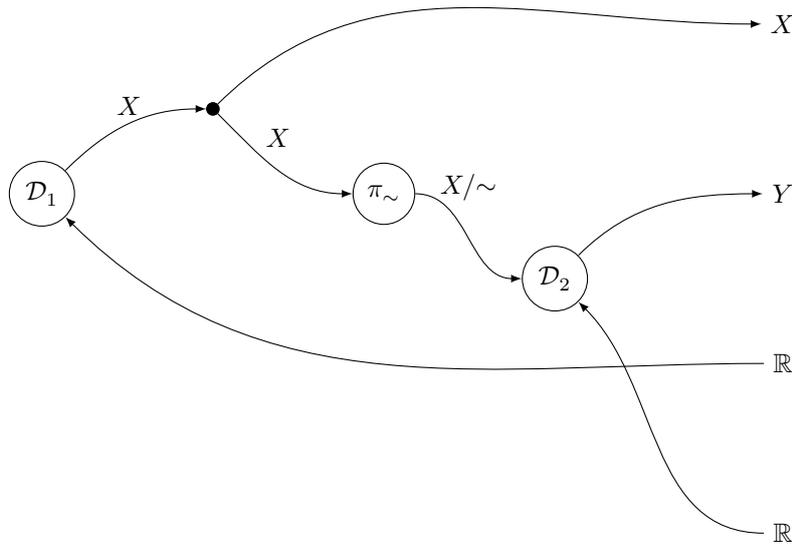
$$\sigma'_2[x'] = \arg \max_{y':Y} q_2(x', y')$$

⁴The notation $[x]$ refers to the information set to which x belongs.

3.2.7 Imperfect information in string diagrams

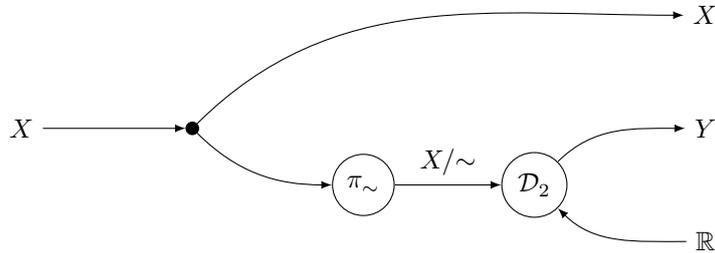
We will end this section by taking a more general two player game involving imperfect information, and showing how it can be visualised as a string diagram. The sets of moves are X and Y , and we have an equivalence relation \sim on X . A strategy profile is a pair $(\sigma_1, \sigma_2) : X \times (X/\sim \rightarrow Y)$. The definition of subgame perfect equilibrium is as before.

We will implement this game using the string diagram



in $\mathbf{Game}(\mathbf{Set})$, where $\pi_{\sim} : X \rightarrow X/\sim$ is the projection onto information sets, $\pi_{\sim}x = [x]$. This should be compared to the string diagrams in §2.3.10 and §3.2.2, which denote perfect information games in which \sim is the identity relation, and so π_{\sim} is the identity on X . Notice how we can explicitly see that the information flowing from \mathcal{D}_1 to \mathcal{D}_2 passes through the projection π_{\sim} , where it is partially hidden.

Since most of the analysis is the same as for the perfect information case, we will focus on the sub-diagram



which is given algebraically by

$$\mathcal{G} = (X \otimes (\mathcal{D}_2 \circ \pi_{\sim})) \circ \Delta_X : X \rightarrow X \otimes Y \otimes \mathbb{R}^*$$

The set of strategy profiles is $\Sigma_{\mathcal{G}} = X/\sim \rightarrow Y$, with play function

$$\mathbf{P}_{\mathcal{G}}\sigma x = (x, \sigma[x])$$

The best response function

$$\mathbf{B}_{\mathcal{G}} : X \times (X \times Y \rightarrow \mathbb{R}) \rightarrow \Sigma_{\mathcal{G}} \rightarrow \mathcal{P}\Sigma_{\mathcal{G}}$$

is calculated as

$$\begin{aligned} \mathbf{B}_{\mathcal{G}}(h, k)\sigma &= \mathbf{B}_{X \otimes (\mathcal{D}_2 \circ \pi_{\sim})}((h, h), k)\sigma \\ &= \mathbf{B}_{\mathcal{D}_2 \circ \pi_{\sim}}(h, k')\sigma \\ &= \mathbf{B}_{\mathcal{D}_2}([h], k')\sigma \end{aligned}$$

where $k'y = k(h, y)$.

Since $\mathbf{B}_{\mathcal{D}_2}$ is a utility-maximising decision, this is

$$\mathbf{B}_{\mathcal{G}}(h, k)\sigma = \{\sigma' : X/\sim \rightarrow Y \mid k(h, \sigma[h]) \geq k(h, y') \text{ for all } y' : Y\}$$

and the condition can be equivalently written

$$\sigma[h] = \arg \max_{y' : Y} k(h, y')$$

This corresponds exactly to the second condition for the best response function in §3.2.6.

3.3 Solvable games

3.3.1 Discussion

In this section we will introduce a suitable solution concept for open games, called solvability. Although the best response function of a game contains all of the information needed to define equilibria, it does not seem to be directly usable as a solution concept, because there is nothing that corresponds directly to an ‘equilibrium of an open game’. A solution of a game, on the other hand, is defined in terms of the best response function and can be directly read as an analogue of Nash or subgame perfect equilibria for open games.

The solvability of a game should be seen as a *witness for an existence theorem*. This is because a solution of an open game implies the existence of equilibria for a class of closed games obtained by varying its continuation, that is, for a class of games whose outcome function varies. For example, as mentioned in §3.3.6 there is a single open game whose solvability is equivalent to the Nash existence theorem for bimatrix games (of fixed size).

The ultimate aim is to show that solvability is respected by the categorical operations used to construct games, which will allow us to prove existence theorems by structural induction, and compute equilibria recursively.⁵ In §3.3.4 we will show that when applied to categorical composition this yields something akin to *backward induction* in classical game theory, and thereby connect solutions with selection functions. In §3.3.6 and §3.3.7, we will discuss how solvability behaves with respect to tensor products, which is more subtle and is ongoing work.

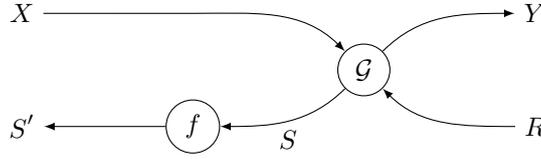
⁵Anything less would shackle the power of compositional game theory.

3.3.2 The definition of solvability

Let $\mathcal{G} : X \otimes S^* \rightarrow Y \otimes R^*$ be a game in $\mathbf{Game}(\mathcal{C})$. We will call \mathcal{G} **solvable** if, for every continuation $k : \text{hom}_{\mathcal{C}}(Y, R)$, there exists a strategy $\sigma : \Sigma_{\mathcal{G}}$, such that for all histories $h : \text{hom}_{\mathcal{C}}(I, X)$, we have $\sigma \in \mathbf{B}_{\mathcal{G}}(h, k)\sigma$.

A **solution** of \mathcal{G} is a function⁶ $s : \text{hom}_{\mathcal{C}}(Y, R) \rightarrow \Sigma_{\mathcal{G}}$ such that sk is a fixpoint of $\mathbf{B}_{\mathcal{G}}(h, k)$ for all contexts $(h, k) : \text{hom}_{\mathcal{C}}(I, X) \times \text{hom}_{\mathcal{C}}(Y, R)$. By the axiom of choice, a game is solvable iff it has a solution.

A first technical issue with solvability is that it is not ‘aware’ of cointility, because the definition of solvability of \mathcal{G} does not refer to S or $\mathbf{C}_{\mathcal{G}}$. Consequently, the solutions of



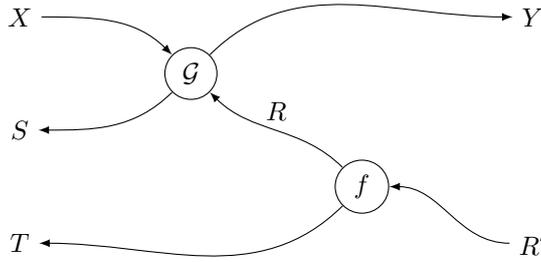
are independent of the computation $f : S \rightarrow S'$.

Notice that all strategically trivial games (§2.2.8) are solvable. Take a strategically trivial game $\mathcal{G} : X \otimes S^* \rightarrow Y \otimes R^*$, so we have $\Sigma_{\mathcal{G}} = 1$, and for every $h : \text{hom}_{\mathcal{C}}(I, X)$ and $k : \text{hom}_{\mathcal{C}}(Y, R)$ we have

$$* \in \mathbf{B}_{\mathcal{G}}(h, k)*$$

Then the unique function $s : \text{hom}_{\mathcal{C}}(Y, R) \rightarrow \Sigma_{\mathcal{G}}$ is trivially a solution.

The following is a simple but useful way to obtain new solvable games from old. Let $\mathcal{G} : X \otimes S^* \rightarrow Y \otimes R^*$ be solvable, and let $f : R' \rightarrow R \otimes T$ be a computation. Then



is also solvable. When solutions are connected with selection functions in §3.3.5, this is analogous to modifying the outcome type of a selection function (§1.1.8). Algebraically, the game denoted by this string diagram is

$$(Y \otimes f^*) \circ (\mathcal{G} \otimes T^*) : X \otimes S^* \otimes T^* \rightarrow Y \otimes R'^*$$

It is obtained from \mathcal{G} by composing and tensoring with strategically trivial games, and so

$$\Sigma_{(Y \otimes f^*) \circ (\mathcal{G} \otimes T^*)} = \Sigma_{\mathcal{G}}$$

⁶A strategy relative to a continuation.

3.3. SOLVABLE GAMES

Its best response function is

$$\mathbf{B}_{(Y \otimes f^*) \circ (\mathcal{G} \otimes T^*)} : \text{hom}_{\mathcal{C}}(I, X) \times \text{hom}_{\mathcal{C}}(Y, R') \rightarrow \Sigma_{\mathcal{G}} \rightarrow \mathcal{P}\Sigma_{\mathcal{G}}$$

given by

$$\mathbf{B}_{(Y \otimes f^*) \circ (\mathcal{G} \otimes T^*)}(h, k)\sigma = \mathbf{B}_{\mathcal{G}}(h, k')\sigma$$

where

$$k' : Y \xrightarrow{k} R' \xrightarrow{f} R \otimes T \xrightarrow{\pi_1} R$$

Let $k : \text{hom}_{\mathcal{C}}(Y, R)$ be a continuation for $(Y \otimes f^*) \circ (\mathcal{G} \otimes T^*)$. By the solvability of \mathcal{G} with the continuation k' , we have $\sigma : \Sigma_{\mathcal{G}}$ such that, for all $h : \text{hom}_{\mathcal{C}}(I, X)$,

$$\sigma \in \mathbf{B}_{\mathcal{G}}(h, k')\sigma$$

Thus $\sigma \in \mathbf{B}_{(Y \otimes f^*) \circ (\mathcal{G} \otimes T^*)}(h, k)\sigma$, as required.

3.3.3 Solvable decisions

We will now provide our first nontrivial examples of a solvable games, namely the decisions in **Game(Set)** defined in §2.1.8 that maximise with respect to a rational preference relation \succeq on a set R . These are

$$\mathcal{D} : X \rightarrow Y \otimes R^*$$

defined by the response function

$$\mathbf{R} : X \times (Y \rightarrow R) \rightarrow \mathcal{P}(X \rightarrow Y)$$

given by

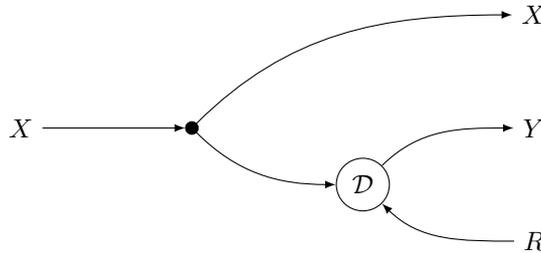
$$\mathbf{R}(x, k) = \{\sigma : X \rightarrow Y \mid k(\sigma x) \succeq ky \text{ for all } y : Y\}$$

We will prove that this decision is solvable, if Y is finite.

Let $k : Y \rightarrow R$ be a continuation. Then since Y is finite and \succeq is a rational preference relation⁷ there is some $y : Y$ for which ky is maximal. Pick the constant⁸ strategy $\sigma : X \rightarrow Y$ given by $\sigma x = y$. Then for all $h : X$ we have $\sigma \in \mathbf{B}_{\mathcal{D}}(h, k)\sigma$ because for all $y' : Y$ we have

$$k(\sigma x) = ky \succeq ky'$$

In **Game(Set)** we can prove a stronger ad-hoc result, namely that



⁷A total preorder.

⁸The fact that we can use a constant strategy here seems to be related to the problem discussed in §3.3.7.

3.3. SOLVABLE GAMES

is also solvable. Algebraically, this game is $(X \otimes \mathcal{D}) \circ \Delta_X$, which is obtained from \mathcal{D} by tensoring and composing with strategically trivial games. Thus, the strategy set is

$$\Sigma_{(X \otimes \mathcal{D}) \circ \Delta_X} = \Sigma_{\mathcal{D}} = X \rightarrow Y$$

Its best response function, by similar reasoning as in §3.2.4, is

$$\mathbf{B}_{(X \otimes \mathcal{D}) \circ \Delta_X}(h, k)\sigma = \{\sigma' : X \rightarrow Y \mid \sigma' \in \mathbf{B}_{\mathcal{D}}(h, k')\sigma\}$$

where $k' : Y \rightarrow R$ is given by $k'y = k(h, y)$. Explicitly, this is

$$\mathbf{B}_{(X \otimes \mathcal{D}) \circ \Delta_X}(h, k)\sigma = \{\sigma' : X \rightarrow Y \mid k(h, \sigma'h) \succeq k(h, y) \text{ for all } y : Y\}$$

Now let $k : X \times Y \rightarrow R$ be a continuation for $(X \otimes \mathcal{D}) \circ \Delta_X$. Define a strategy $\sigma : X \rightarrow Y$ by

$$\sigma x = \arg \max_{y:Y} k(x, y)$$

where $\arg \max$ picks some y that attains the \succeq -maximum. Then for every history $h : X$ we have

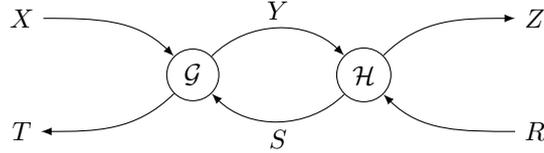
$$k(h, \sigma x) \succeq k(h, y)$$

by construction, and hence $\sigma \in \mathbf{B}_{(X \otimes \mathcal{D}) \circ \Delta_X}(h, k)\sigma$.

3.3.4 Backward induction for open games

We will now prove that solvable games are closed under categorical composition. The proof of this result amounts to *backward induction*, just as does the product of selection functions (§1.3.6). This is the main piece of evidence thus far that solvable games are mathematically natural.

Suppose we have games



with solutions

$$s_G : \text{hom}_{\mathcal{C}}(Y, S) \rightarrow \Sigma_G$$

$$s_H : \text{hom}_{\mathcal{C}}(Z, R) \rightarrow \Sigma_H$$

We define $s_{\mathcal{H} \circ \mathcal{G}} : \text{hom}_{\mathcal{C}}(Z, R) \rightarrow \Sigma_G \times \Sigma_H$ by

$$s_{\mathcal{H} \circ \mathcal{G}}k = (s_G k_{s_H k \circ}, s_H k)$$

where $k_{s_H k \circ} : \text{hom}_{\mathcal{C}}(Y, S)$ is the extended continuation defined in §2.2.2. We must prove that $s_{\mathcal{H} \circ \mathcal{G}}$ is a solution of $\mathcal{H} \circ \mathcal{G}$.

Let $(h, k) : \text{hom}_{\mathcal{C}}(I, X) \times \text{hom}_{\mathcal{C}}(Z, R)$ be a context for $\mathcal{H} \circ \mathcal{G}$. We must prove that

$$s_{\mathcal{H} \circ \mathcal{G}}k \in \mathbf{B}_{\mathcal{H} \circ \mathcal{G}}(h, k)(s_{\mathcal{H} \circ \mathcal{G}}k)$$

This unwinds to the two conditions

$$s_G k_{s_H k \circ} \in \mathbf{B}_G(h, k_{s_H k \circ})(s_G k_{s_H k \circ})$$

and

$$s_{\mathcal{H}}k \in \mathbf{B}_{\mathcal{H}}(\mathbf{P}_{\mathcal{G}}\sigma \circ h, k)(s_{\mathcal{H}}k)$$

for all $\sigma : \Sigma_{\mathcal{G}}$. Both of these follow immediately from the fact that $s_{\mathcal{G}}$ and $s_{\mathcal{H}}$ are solutions of \mathcal{G} and \mathcal{H} respectively.

3.3.5 Selection functions and open games

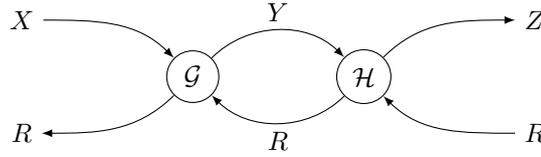
Given a solution $s : \text{hom}_{\mathcal{C}}(Y, R) \rightarrow \Sigma_{\mathcal{G}}$ of a game $\mathcal{G} : X \otimes S^* \rightarrow Y \otimes R^*$, we have an associated selection function

$$\mathbf{P}_{\mathcal{G}} \circ s : \text{hom}_{\mathcal{C}}(Y, R) \rightarrow \text{hom}_{\mathcal{C}}(X, Y)$$

which is a morphism in $\text{hom}_{\mathcal{J}_R\mathcal{C}}(X, Y)$, where $\mathcal{J}_R\mathcal{C}$ is the category of selection functions defined in §1.3.2. We will call ε a selection function for \mathcal{G} if it has the form $\varepsilon = \mathbf{P}_{\mathcal{G}} \circ s$ for a solution s of \mathcal{G} .

We will now for the first time relate the composition in the category of selection functions, or equivalently the Kleisli extension of the selection monad, with game theory.⁹ Unfortunately, this connection does not behave well with respect to coplay; it remains to be seen whether or not it is a useful idea.

Suppose we have games



whose coplay functions are the identity on R . We will prove that if ε is a selection function for \mathcal{G} and δ is a selection function for \mathcal{H} , then $\delta \circ \varepsilon$ is a selection function for $\mathcal{H} \circ \mathcal{G}$, where the composition $\delta \circ \varepsilon$ is taken in the category $\mathcal{J}_R\mathcal{C}$.

By hypothesis, we have a solution $s_{\mathcal{G}} : \text{hom}_{\mathcal{C}}(Y, R) \rightarrow \Sigma_{\mathcal{G}}$ of \mathcal{G} , and a solution $s_{\mathcal{H}} : \text{hom}_{\mathcal{C}}(Z, R) \rightarrow \Sigma_{\mathcal{H}}$ of \mathcal{H} , such that $\varepsilon = \mathbf{P}_{\mathcal{G}} \circ s_{\mathcal{G}}$ and $\delta = \mathbf{P}_{\mathcal{H}} \circ s_{\mathcal{H}}$. The composition of these is

$$(\delta \circ \varepsilon)k = \delta k \circ \varepsilon(k \circ \delta k) = \mathbf{P}_{\mathcal{H}}(s_{\mathcal{H}}k) \circ \mathbf{P}_{\mathcal{G}}(s_{\mathcal{G}}(k \circ \mathbf{P}_{\mathcal{H}}(s_{\mathcal{H}}k)))$$

The solution $s_{\mathcal{H} \circ \mathcal{G}}$ of $\mathcal{H} \circ \mathcal{G}$ is defined by

$$s_{\mathcal{H} \circ \mathcal{G}}k = (s_{\mathcal{G}}k_{s_{\mathcal{H}}k \circ}, s_{\mathcal{H}}k)$$

The selection function associated to this solution is

$$\mathbf{P}_{\mathcal{H} \circ \mathcal{G}}(s_{\mathcal{H} \circ \mathcal{G}}k) = \mathbf{P}_{\mathcal{H}}(s_{\mathcal{H}}k) \circ \mathbf{P}_{\mathcal{G}}(s_{\mathcal{G}}(k_{s_{\mathcal{H}}k \circ}))$$

We get that $\delta \circ \varepsilon$ is precisely this, and so it is a selection function for $\mathcal{H} \circ \mathcal{G}$. Here we need to use the restricted form of the coplay function of \mathcal{H} in order to get that

$$k_{s_{\mathcal{H}}k \circ} = k \circ \mathbf{P}_{\mathcal{H}}(s_{\mathcal{H}}k)$$

⁹The connection between the game theory and the product of selection functions is well understood. However the product of selection functions is a derived operator in terms of the more primitive Kleisli extension and strength operators, neither of which have any previously known meaning in terms of game theory.

rather than (in pure strategies)

$$k_{s_{\mathcal{H}}k}y = \mathbf{C}_{\mathcal{H}}(s_{\mathcal{H}}k)(y, k(\mathbf{P}_{\mathcal{H}}(s_{\mathcal{H}}k)y))$$

This provides us with a high level way to think about solvability. Given an object R of \mathcal{C} , we consider the subcategory $\mathbf{Game}_R(\mathcal{C})$ whose objects are of the form $X \otimes R^*$ for X an object of \mathcal{C} , and whose morphisms are games whose coplay function is constant. Then we can think in terms of functors from $\mathbf{Game}_R(\mathcal{C})$ to the selection category $\mathcal{J}_R(\mathcal{C})$. Although it is possible to do a certain amount of game theory in $\mathbf{Game}_R(\mathcal{C})$ when R is suitably chosen, it is unlikely that this way of thinking has enough benefit to offset the restriction of not having coplay available, which limits compositionality in practice. One way in which this is useful, however, is to understand the relationship between selection functions and open games.

3.3.6 Tensor does not preserve solvability

In general, the class of solvable games is not closed under tensor. This is a high level way of saying that simultaneous games do not have equilibria for every solution concept, and in particular they do not have pure strategy Nash equilibria.

We will give a counterexample using matching pennies, which is the simplest game with no pure strategy equilibrium. In this game two players each simultaneously choose H ('heads') or T ('tails'). If both players make the same choice then player 1 wins, and if the choices are different player 2 wins. This game has no pure strategy Nash equilibrium, because when playing a strategy in which both choices are the same player 2 has incentive to deviate, and when playing a strategy in which the choices are different player 1 has incentive to deviate. (This game has a unique mixed strategy equilibrium, in which both players mix with equal probability.)

We will formalise this game by working over $\mathbf{Game}(\mathbf{Set})$. Let $X = \{H, T\}$, and let R be the poset $\{0 < 1\}$, where 1 indicates winning and 0 indicates losing. The maximising decision

$$\mathcal{D} : I \rightarrow X \otimes R^*$$

was proved to be solvable in §3.3.3.

We will now prove that the tensor product $\mathcal{D} \otimes \mathcal{D}$ is not solvable. To show this, we must exhibit a particular continuation $k : X \times X \rightarrow R \times R$ with the property that

$$\mathbf{B}_{\mathcal{D} \otimes \mathcal{D}}k : X \times X \rightarrow \mathcal{P}(X \times X)$$

has no fixpoints. This is equivalent to saying that the game with outcome function k has no pure strategy Nash equilibrium. Therefore, we take k to be the outcome function of matching pennies, namely

$$k(x, y) = \begin{cases} (1, 0) & \text{if } x = y \\ (0, 1) & \text{if } x \neq y \end{cases}$$

This gives

$$\mathbf{B}_{\mathcal{D} \otimes \mathcal{D}}k(x, y) = \{(y, \bar{x})\}$$

where $\bar{\cdot} : X \rightarrow X$ interchanges H and T . This function indeed has no fixpoint.

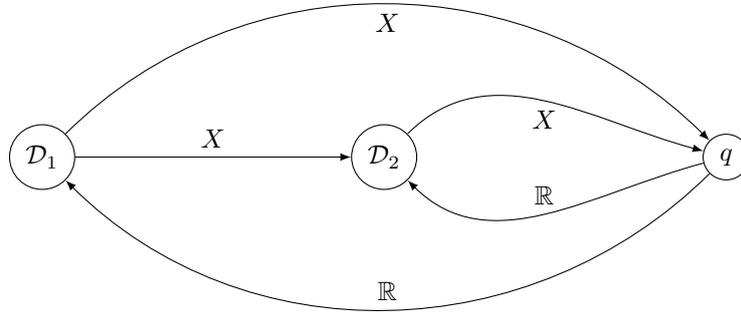
If we reasoned in **Game(SRel)** instead of **Game(Set)** then this does not happen: by Nash’s theorem, the tensor product of finitely many decisions over finite sets of choices in **Game(SRel)** is solvable.

3.3.7 Failure of compositional Nash’s theorem

An early motivation for introducing solvable games, besides connecting open games and selection functions (§3.3.5), was to prove a suitable generalisation of Nash’s theorem for open games. The fact that solvable games are closed under composition (§3.3.4), but fail to be closed under tensor because pure strategy Nash equilibria need not exist (§3.3.6), strongly suggests the conjecture that in **Game(SRel)**, *all freely generated games are solvable*.

This fails, however, because it is not the case that arbitrary solvable games in **Game(SRel)** are closed under tensor product. This appears to result from a much deeper issue with the definitions in §2.1 and §2.2. In §3.1 it was shown that normal form games with mixed strategies Nash equilibria can be faithfully represented by open games, and similarly for extensive form games and pure strategy subgame perfect equilibria in §3.2. However, if the two are combined, we end up with a uselessly implausible solution concept.¹⁰ The previous conjecture, that all freely generated games are solvable, can be taken as a requirement for a repaired theory.

We will end with a specific counterexample to the conjecture. Consider the game in **Game(SRel)** denoted by



or $\mathcal{G} = q \circ (X \otimes \mathcal{D}_2 \otimes \mathbb{R}^*) \circ \mathcal{D}_1$, where $X = \{H, T\}$. \mathcal{D}_2 is an ordinary decision that maximises expected utility, and \mathcal{D}_1 also maximises expected utility but chooses a pair from $X \times X$ with a mixed strategy. The strategically trivial game q (which could be decomposed into computations and counits) implements the outcome function, which is again that of matching pennies. Specifically, the

¹⁰In retrospect this is unsurprising, because the equivalently naive generalisation of mixed strategies to extensive form in classical game theory also fails. The standard solution concept in that situation is sequential equilibrium (Leyton-Brown and Shoham 2008), which involves beliefs (probability distributions on information sets giving finer information about the position in the game tree), that undergo Bayesian updating. Because open games are inherently dynamic games, a repaired theory of open games with mixed strategies must similarly have Bayesian reasoning built into the definitions.

3.3. SOLVABLE GAMES

coplay function $\mathbf{C}_q : X \times X \rightarrow \mathcal{D}(\mathbb{R} \times \mathbb{R})$ is given by

$$\mathbf{C}_q(x, y) = \begin{cases} \delta(0, 1) & \text{if } x = y \\ \delta(1, 0) & \text{if } x \neq y \end{cases}$$

where δ is the unit of the monad \mathcal{D} (see §2.1.3).

Since \mathcal{G} is a scalar, its best response function has the form

$$\mathbf{B}_{\mathcal{G}} : \Sigma_{\mathcal{G}} \rightarrow \mathcal{P}\Sigma_{\mathcal{G}}$$

Despite the fact that we are working over **SRel** and so we have mixed strategies available, and the fact that \mathcal{G} has been built from only expectation-maximising decisions, computations and counits using composition and tensor, we will prove that $\mathbf{B}_{\mathcal{G}}$ has no fixpoints.

The set of strategies is

$$\Sigma_{\mathcal{G}} = \Sigma_{\mathcal{D}_1} \times \Sigma_{\mathcal{D}_2} = \mathcal{D}(X \times X) \times (X \rightarrow \mathcal{D}X)$$

Geometrically, this is the product of a tetrahedron and a square, and hence is closed, convex and finite-dimensional, satisfying the hypotheses of the Kakutani fixpoint theorem.

To begin with, since q is strategically trivial we have

$$\mathbf{B}_{\mathcal{G}}(*, *) (\sigma, \tau) = \mathbf{B}_{(X \otimes \mathcal{D}_2 \otimes \mathbb{R}^*) \circ \mathcal{D}_1}(*, k) (\sigma, \tau)$$

where $k : X \times X \rightarrow \mathcal{D}(\mathbb{R} \times \mathbb{R})$ is the continuation $k = *_{*o} = q$. This is

$$\mathbf{B}_{\mathcal{D}_1}(*, k_{\tau o}) \sigma \times \bigcap_{\sigma' : \mathcal{D}(X \times X)} \mathbf{B}_{X \otimes \mathcal{D}_2 \otimes \mathbb{R}^*}(\mathbf{P}_{\mathcal{D}_1} \sigma' \circ *, k) \tau$$

We will only need to focus on the second part of this. Since \mathcal{D}_1 is a decision we have $\mathbf{P}_{\mathcal{D}_1} \sigma' \circ * = \sigma'$. Then

$$\mathbf{B}_{X \otimes \mathcal{D}_2 \otimes \mathbb{R}^*}(\sigma', k) \tau = \mathbf{B}_{X \otimes \mathcal{D}_2}(\sigma', k_1) \tau$$

where $k_1 : X \times X \rightarrow \mathcal{D}\mathbb{R}$ is given by

$$k(x_1, x_2) = \begin{cases} \delta 0 & \text{if } x_1 = x_2 \\ \delta 1 & \text{if } x_1 \neq x_2 \end{cases}$$

The game $X \otimes \mathcal{D}_2 : X \otimes X \rightarrow X \otimes X \otimes \mathbb{R}^*$ is essentially the component of \mathcal{G} that causes the problems. Its best response function is

$$\mathbf{B}_{X \otimes \mathcal{D}_2}(h, k_1) \tau = \mathbf{B}_{\mathcal{D}_2}(h_1, k') \tau$$

The continuation $k' = (k_1)_{*(h_1) \otimes} : X \rightarrow \mathcal{D}\mathbb{R}$ is given by $k' x_2 = k_1(h_1, x_2)$, where we are abusing notation as in §3.1.7. Explicitly, using the particular k we defined,

$$\mathbb{P}[k' x_2 = 1] = \mathbb{P}[h_1 \neq x_2]$$

Now using the definition of $\mathbf{B}_{\mathcal{D}_2}$,

$$\mathbf{B}_{X \otimes \mathcal{D}_2}(h_1, k_1) \tau = \{\tau' : X \rightarrow \mathcal{D}X \mid \mathbb{E}[k'(\tau' h_1)] \geq \mathbb{E}[k' x_2] \text{ for } x_2 = H, T\}$$

3.3. SOLVABLE GAMES

Putting this together, if we had a fixpoint (σ, τ) of $\mathbf{B}_{\mathcal{G}}$, then for every possible history $h : \mathcal{D}(X \times X)$ we would have

$$\mathbb{E}[k_1(h_1, \tau h_1)] \geq \mathbb{E}[k_1(h_1, x_2)]$$

for $x_2 = H, T$. We will now prove that this is impossible.

Pick some $x_1 \in \text{supp}(\tau H)$, so $\mathbb{P}[\tau H = x_1] > 0$. We will take $h = \delta x_1 \otimes \delta H$, that is, h is the probability distribution with $\mathbb{P}[h = (x_1, H)] = 1$. Then

$$k'(\tau h_1) = k'(\tau x_1) = k_1(x_1, \tau H)$$

and since $\mathbb{P}[\tau H = x_1] > 0$ we have $\mathbb{E}[k_1(x_1, \tau H)] < 1$. This means that player 2 has incentive to unilaterally deviate to $x_2 \neq x_1$, which has

$$\mathbb{E}[k'x_2] = \mathbb{E}[k'(x_1, x_2)] = 1$$

Bibliography

- Samson Abramsky. Abstract scalars, loops, and free traced and strongly compact closed categories. In *Proceedings of CALCO'05*, volume 3629 of *Lectures notes in computer science*, pages 1–29, 2005. doi: 10.1007/11548133_1.
- Samson Abramsky. Arrow’s theorem by arrow theory. In *Logic without borders: essays on set theory, model theory, philosophical logic and philosophy of mathematics*, volume 5 of *Ontos mathematical logic*, pages 15–30. De Gruyter, 2015.
- Samson Abramsky and Bob Coecke. A categorical semantics of quantum protocols. In *Proceedings of the 19th Annual IEEE Symposium on Logic in Computer Science: LICS 2004*, 2004.
- Samson Abramsky and Viktor Winschel. Coalgebraic analysis of subgame-perfect equilibria in infinite games without discounting. To appear in *Mathematical Structures in Computer Science*, 2015.
- Samson Abramsky, Alexander Kurz, Pierre Lescanne, and Viktor Winschel. Coalgebraic semantics of reflexive economics (Dagstuhl seminar 15042). *Dagstuhl reports*, 5(1), 2015. doi: 10.4230/DagRep.5.1.197.
- Michael Atiyah. Advice to a young mathematician. In Timothy Gowers, editor, *The Princeton companion to mathematics*. Princeton university press, 2008.
- Jeremy Avigad and Solomon Feferman. Gödel’s functional (“Dialectica”) interpretation. In S. Buss, editor, *Handbook of proof theory*, volume 137 of *Studies in logic and the foundations of mathematics*, pages 337–405. North Holland, Amsterdam, 1998.
- John Baez and Mike Stay. Physics, topology, logic and computation: a Rosetta stone. In Bob Coecke, editor, *New structures for physics*, pages 95–172. Springer, 2010.
- Denis Bechet, Philippe de Groote, and Christian Retoré. A complete axiomatisation for the inclusion of series-parallel partial orders. In *Rewriting techniques and applications*, Lecture Notes in Computer Science, pages 230–240. Springer, 2005. doi: 10.1007/3-540-62950-5_74.
- Andreas Blass. Some semantical aspects of linear logic. *Logic journal of the interest group in pure and applied logic*, 5(4):487–503, 1997. doi: 10.1093/jigpal/5.4.487.

BIBLIOGRAPHY

- Achim Blumensath and Viktor Winschel. A coalgebraic framework for games in economics. Working paper, 2013.
- Nicola Botta, Antoine Mandel, Cezar Ionescu, Mareen Hofmann, Daniel Lincke, Sibylle Schupp, and Carlo Jaeger. A functional framework for agent-based models of exchange. *Applied mathematics and computation*, 218(8):4025–4040, 2011. doi: 10.1016/j.amc.2011.08.051.
- Nicola Botta, Cezar Ionescu, and Edwin Brady. Sequential decision problems, dependently typed solutions. In *Proceedings of PLMMS13*, 2013.
- Pierre Boudes, Antoine Kaszycz, and Luc Pellisser. Monetary economics simulation: Stock-flow consistent invariance, monadic style. HAL ID: hal-01181278, 2015.
- Venanzio Capretta and Tarmo Uustalu. A coalgebraic view of bar recursion and bar induction. To appear in *Proceedings of FoSSaCS'16*, 2016.
- Bob Coecke. Kindergarten quantum mechanics. ArXiv, 2005.
- Bob Coecke. Introducing categories to the practicing physicist. In *What is category theory? Advanced studies in mathematics and logic*, volume 30, pages 45–74. Polimetrica, 2006.
- Bob Coecke, editor. *New structures for physics*. Springer, 2011.
- Bob Coecke and Ross Duncan. Interacting quantum observables: Categorical algebra and diagrammatics. *New Journal of Physics*, 13(4), 2011. doi: 10.1088/1367-2630/13/4/043016.
- Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. *STOC '06 Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 71–79, 2006.
- Valeria de Paiva. The dialectica categories. Technical report, University of Cambridge, 1991.
- Valeria de Paiva. Dialectica and Chu constructions: Cousins? *Theory and applications of categories*, 17(7):127–152, 2007.
- Martin Erwig and Steve Kollmansberger. Probabilistic functional programming in Haskell. *Journal of functional programming*, 16(1):21–34, 2006. doi: 10.1017/S0956796805005721.
- Martin Escardó. Synthetic topology of data types and classical spaces. *ENTCS*, 87:21–156, 2004.
- Martin Escardó and Paulo Oliva. Selection functions, bar recursion and backward induction. *Mathematical structures in computer science*, 20(2):127–168, 2010a. doi: 10.1017/S0960129509990351.
- Martin Escardó and Paulo Oliva. What sequential games, the Tychonoff theorem and the double-negation shift have in common. In *Proceedings of MSFP'10*, 2010b. doi: 10.1145/1863597.1863605.

BIBLIOGRAPHY

- Martin Escardó and Paulo Oliva. Sequential games and optimal strategies. *Proceedings of the Royal Society A*, 467:1519–1545, 2011. doi: 10.1098/rspa.2010.0471.
- Martin Escardó and Paulo Oliva. Computing Nash equilibria of unbounded games. *Proceedings of the Turing centenary conference*, 2012.
- Martin Escardó and Paulo Oliva. The Herbrand interpretation of the double negation shift. Submitted, 2015.
- Tobias Fritz. Convex spaces I: Definitions and examples. ArXiv, 2009.
- Drew Fudenberg and Jean Tirole. *Game theory*. MIT Press, 1991.
- Neil Ghani and Jules Hedges. A compositional approach to economic game theory. Submitted, 2016.
- Jean-Yves Girard. Linear logic. *Theoretical computer science*, 1987.
- Jean-Yves Girard. Locus solum: from the rules of logic to the logic of rules. *Mathematical structures in computer science*, 11(3):301–506, 2001. doi: 10.1017/S096012950100336X.
- Michèle Giry. A categorical approach to probability theory. *Categorical aspects of topology and analysis*, pages 68–85, 1982. doi: 10.1007/BFb0092872.
- Irving Glicksberg. A further generalisation of the Kakutani fixed point theorem, with application to Nash equilibrium points. *Proc. Amer. Math. Soc.*, 3(1): 170–174, February 1952.
- Jules Hedges. A generalisation of Nash’s theorem with higher-order functionals. *Proceedings of the Royal Society A*, 469, 2013. doi: 10.1098/rspa.2013.0041.
- Jules Hedges. Monad transformers for backtracking search. In *Proceedings of MSFP’14*, EPTCS, pages 31–50, 2014. doi: 10.4204/EPTCS.153.3.
- Jules Hedges. Dialectica categories and games with bidding. In *Post-proceedings of TYPES’14*, volume 39 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 89–110, 2015a. doi: 10.4230/LIPIcs.TYPES.2014.89.
- Jules Hedges. The selection monad as a CPS translation. ArXiv, 2015b.
- Jules Hedges and Mehrnoosh Sadrzadeh. A generalised quantifier theory of natural language in categorical compositional distributional semantics with bialgebras. Submitted, 2016.
- Jules Hedges, Paulo Oliva, Evguenia Sprints, Viktor Winschel, and Philipp Zahn. Higher-order decision theory. ArXiv, 2015a.
- Jules Hedges, Paulo Oliva, Evguenia Sprints, Viktor Winschel, and Philipp Zahn. Higher-order game theory. Submitted, 2015b.
- Chris Heunen, Mehrnoosh Sadrzadeh, and Edward Grefenstette, editors. *Quantum physics and linguistics*. Oxford University Press, 2013.
- Doug Hoyte. *Let over lambda*. lulu.com, 2008.

BIBLIOGRAPHY

- Martin Hyland. Game semantics. In Andrew Pitts and Peter Dybjer, editors, *Semantics and logics of computation*, pages 131–184. Cambridge University Press, 1997.
- Martin Hyland. Proof theory in the abstract. *Annals of pure and applied logic*, 114(1-3):43–78, 2002. doi: 10.1016/S0168-0072(01)00075-6.
- Cezar Ionescu and Patrik Jansson. Dependently-typed programming in scientific computing: examples from economic modelling. In *Implementation and application of functional languages*, Lectures notes in computer science, pages 140–156. Springer, 2013. doi: 10.1007/978-3-642-41582-1_9.
- Bart Jacobs, Bas Westerbaan, and Bram Westerbaan. States of convex sets. In *Proceedings of FoSSaCS’15*, Lectures notes in computer science, pages 87–101, 2015. doi: 10.1007/978-3-662-46678-0_6.
- Alan Jeffrey. Premonoidal categories and a graphical view of programs. Unpublished manuscript, 1997.
- André Joyal and Ross Street. The geometry of tensor calculus I. *Advances in mathematics*, 88(1):55–112, 1991. doi: 10.1016/0001-8708(91)90003-P.
- André Joyal, Ross Street, and Dominic Verity. Traced monoidal categories. *Mathematical proceedings of the Cambridge philosophical society*, 119:447–468, 1996. doi: 10.1017/S0305004100074338.
- Shizuo Kakutani. A generalisation of brouwer’s fixed point theorem. *Duke Math. J.*, 8(3):457–459, 1941. doi: 10.1215/S0012-7094-41-00838-4.
- G. M. Kelly and M. L. Laplaza. Coherence for compact closed categories. *Journal of pure and applied algebra*, 19(193–213), 1980. doi: 10.1016/0022-4049(80)90101-2.
- John Maynard Keynes. *The general theory of employment, interest and money*. Macmillan, 1936.
- Aleks Kissinger, Alex Merry, and Matvey Soloviev. Pattern graph rewrite systems. *Proceedings of DCM’12*, 143:54–66, 2014. doi: 10.4204/EPTCS.143.5.
- Daphne Koller and Brian Milch. Multi-agent influence diagrams for representing and solving games. *Games and economic behaviour*, 45(1):181–221, 2003. doi: 10.1016/S0899-8256(02)00544-4.
- David Kreps. *A course in microeconomic theory*. Princeton university press, 1990.
- David Kreps. *Game theory and economic modelling*. Clarendon lectures in economics. Oxford University Press, 1991.
- Steven M. LaValle. *Planning algorithms*. Cambridge University Press, 2006.
- Stephane le Roux. From winning strategy to Nash equilibrium. *Mathematical logic quarterly*, 60(4–5):354–371, 2014. doi: 10.1002/malq.201300034.

BIBLIOGRAPHY

- Stéphane le Roux and Arno Pauly. Infinite sequential games with real-valued payoffs. In *Proceedings of CSL-LICS'14*, pages 62:1–62:10, 2014. doi: 10.1145/2603088.2603120.
- Pierre Lescanne and Matthieu Perrinel. On the rationality of escalation. ArXiv, 2010.
- Pierre Lescanne and Matthieu Perrinel. “Backward” coinduction, Nash equilibrium and the rationality of escalation. *Acta Informatica*, 49(3):117–137, 2012.
- Kevin Leyton-Brown and Yoav Shoham. *Essentials of game theory: a concise, multidisciplinary introduction*. Morgan and Claypool, 2008.
- Saunders Mac Lane. *Categories for the working mathematician*. Springer, 1978.
- Adrian Mathias. A term of length 4,523,659,424,929. *Synthese*, 133:75–86, 2002.
- Micah McCurdy. Graphical methods for Tannaka duality of weak bialgebras and weak Hopf algebras. *Theory and applications of categories*, 26(9):233–280, 2012.
- Eugenio Moggi. Computational lambda calculus and monads. In *Proceedings of fourth annual symposium on logic in computer science*, 1989.
- Eugenio Moggi. Notions of computation and monads. *Information and Computation*, 93:55–92, 1991.
- Lawrence Moss and Ignacio Iglizzo. Harsanyi type spaces and final coalgebras constructed from satisfied theories. In *Proceedings of CMCS'04*, volume 106 of *ENTCS*, pages 279–295, 2004. doi: 10.1016/j.entcs.2004.02.036.
- A. Mostowski. On a generalisation of quantifiers. *Fundamenta mathematicae*, 44:12–36, 1957.
- John Nash. Equilibrium points in n -person games. *Proceedings of the national academy of sciences*, 36(1):48–49, 1950.
- John Nash. Non-cooperative games. *Annals of mathematics*, 54(2):286–295, 1951.
- Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay Vazirani, editors. *Algorithmic game theory*. Cambridge University Press, 2007.
- Paulo Oliva and Thomas Powell. A constructive interpretation of Ramsey’s theorem via the product of selection functions. *Mathematical structures in computer science*, 2014. doi: 10.1017/S0960129513000340.
- Paulo Oliva and Thomas Powell. A game-theoretic computational interpretation of proofs in classical analysis. In *Gentzen’s centenary: The quest for consistency*, pages 501–531. Springer, 2015. doi: 10.1007/978-3-319-10103-3_18.
- Prakash Panangaden. The category of Markov kernels. *Proceedings of the first international workshop on probabilistic methods in verification*, 22:171–187, 1999. doi: 10.1016/S1571-0661(05)80602-4.

BIBLIOGRAPHY

- Ross Paterson. A new notation for arrows. In *Proceedings of the sixth ACM SIGPLAN international conference on functional programming*, pages 229–240. ACM, 2001. doi: 10.1145/507635.507664.
- Dusko Pavlovic. A semantical approach to equilibria and rationality. In *Algebra and coalgebra in computer science*, volume 5728 of *Lectures notes in computer science*, pages 317–334. Springer, 2009. doi: 10.1007/978-3-642-03741-2_22.
- Dusko Pavlovic. Monoidal computer I: Basic computability by string diagrams. *Information and Computation*, 226:94–116, 2013. doi: 10.1016/j.ic.2013.03.007.
- Roger Penrose. Applications of negative dimensional tensors. In D. J. A. Welsh, editor, *Combinatorial mathematics and its applications*, pages 221–244. Academic Press, 1971.
- Gordon Plotkin and John Power. Notions of computation determine monads. In *Proceedings of FoSSaCS’02*, volume 2303 of *Lectures notes in computer science*, pages 342–356. Springer, 2002.
- Thomas Powell. *On Bar Recursive Interpretations of Analysis*. PhD thesis, Queen Mary University of London, 2013.
- John Power and Edmund Robinson. Premonoidal categories and notions of computation. *Mathematical structures in computer science*, 1993.
- John Power and Hayo Thielecke. Closed Freyd and κ -categories. In *Proceedings of the 26th International Colloquium on Automata, Languages and Programming*, pages 625–634. Springer-Verlag, 1999.
- Chris Schommer-Pries. *The classification of two-dimensional extended topological field theories*. PhD thesis, University of California, Berkeley, 2009.
- Ulrich Schwalbe and Paul Walker. Zermelo and the early history of game theory. *Games and economic behaviour*, 34:123–137, 2001.
- Peter Selinger. A survey of graphical languages for monoidal categories. In Bob Coecke, editor, *New structures for physics*, pages 289–355. Springer, 2011.
- Ron Spiegler. “But can’t we get the same thing with a standard model?” Rationalising bounded-rationality models. *Economics and philosophy*, 27(1):23–43, 2011. doi: 10.1017/S0266267110000453.
- Daniele Varacca and Glynn Winskel. Distributing probability over nondeterminism. *Mathematical structures in computer science*, 16(1):87–113, 2006. doi: 10.1017/S0960129505005074.
- John von Neumann and Oskar Morgenstern. *Theory of games and economic behaviour*. Princeton university press, 1944.