Chapter 8

Infinite Horizon Extensive Form Games, Coalgebraically

Matteo Capucci
Neil Ghani
Clemens Kupke
Jérémy Ledent
Fredrik Nordvall Forsberg
Mathematically Structured Programming group,
Department of Computer and Information Sciences,
University of Strathclyde

Contents

1. Introduction

Game theory is the study of how agents make decisions in order to maximise their outcomes [18, 20]. A strategy profile describes how each agent will play the game, and is said to be a *Nash equilibrium* if no player has any incentive to deviate from their strategy; it is called *subgame perfect* if it is a

Nash equilibrium in every subgame of the game. In a series of papers [5–8], Douglas Bridges investigated constructive aspects of the theory of games where players move simultaneously (so-called normal form games), and their preference relations. This article is concerned with a constructive treatment of games where players move sequentially.

A common way to model sequential games is using their *extensive form*: a game is represented as a tree, whose branching structure reflects the decisions available to the players, and whose leaves (corresponding to a complete 'play' of the game) are decorated by payoffs for each player. When the number of rounds in the game is infinite (e.g. because a finite game is repeated an infinite number of times, or because the game may continue forever), the game tree needs to be infinitely deep. One way to handle such infinite trees is to consider them as the metric completion of finite trees, after equipping them with a suitable metric [19]. However, as a definitional principle, this only gives a method to construct functions into other complete metric spaces, and the explicit construction as a quotient of Cauchy sequences [4, §4.3] can be unwieldy to work with. Instead, we prefer to treat the infinite as the *dual* of the finite, in the spirit of category theory and especially the theory of coalgebras [21].

We are not the first to attack infinite extensive form games using coalgebraic methods. Lescanne [15, 16], Lescanne and Perrinel [17] and Abramsky and Winschel [1] define infinite two-player games coalgebraically, and show that coinductive proof methods can be used to constructively prove properties of games. However, their definition only assigns utility to finite plays. For that reason, they restrict attention to *strongly convergent* strategies, i.e., strategy profiles that always lead to a leaf of the tree in a finite number of steps. This restriction rules out infinitely repeated games, where utility could be assigned using discounted sums or limiting averages — both methods crucially making use of the entire infinite history of the game. Building on our on work on infinitely repeated open games [11], we extend Lescanne's and Abramsky and Winschel's coalgebraic framework to not necessarily convergent strategies.

The *one-shot deviation principle* is a celebrated theorem of classical game theory. It asserts that a strategy is a subgame perfect equilibrium if and only if there is no profitable one-shot deviation in any subgame. While this principle holds for all finite games, in the case of infinite trees, it requires an extra assumption called *continuity at infinity* (see e.g. Fudenberg and Tirole [10, Chapter 4.2]). Essentially, this property says that the actions taken in the distant future have a negligible impact on the current payoff.

In the coalgebraic setting, Abramsky and Winschel [1] claim to prove the one-shot principle without continuity assumptions — we argue that this is not entirely the case. Indeed, they show that the natural coalgebraic equilibrium concept (which they call "SPE") satisfies the one-shot deviation principle. However they do not discuss how this coalgebraic concept relates to the traditional notion of subgame perfect Nash equilibria. As we show in Theorem 4.12, these two notions are indeed equivalent, but only assuming continuity of the utility function. In that regard, the predicate "SPE" of Abramsky and Winschel (called $\Box\mathsf{Unimprov}$ in our work) is in fact closer to a coalgebraic version of the one-shot equilibrium.

Our proof of the one-shot deviation principle extends the previous ones in several ways. Compared to the one of Abramsky and Winschel [1], it applies to games where infinite plays are possible; and it relates to the more standard definition of subgame-perfect Nash equilibrium, $\Box\mathsf{Nash}$. Additionally, our theorem applies to *any* coalgebra of the extensive-form tree functor, whereas Abramsky and Winschel only work with the final coalgebra. Compared to the usual proofs found in the game theory literature, we carefully analyse the constructivity of the proof. The only extra assumption that we require is decidable equality on the set of players (which is typically finite), and decidability of the order relation on the set of payoffs (typically, the set of rational numbers). Moreover, continuity at infinity is usually expressed using uniform continuity; we remark that pointwise continuity suffices.

Structure of the paper We recall the basics of coalgebras for endofunctors in Section 2. In Section 3, we define infinite extensive form games as final coalgebras, and use properties of coalgebras to define notions such as strategies, moves, payoffs and equilibria in a game. We then relate our coalgebraic notions with the existing notions from the literature in Section 4. Throughout the paper, we demonstrate how coinductive proof principles can be used to reason constructively about infinite games.

Notation We use $\mathcal{P} : \mathsf{Set} \to \mathsf{Set}$ for the covariant powerset functor mapping a set to its set of subsets. Given a set-indexed collection of sets $Y : I \to \mathsf{Set}$ and $i \in I$, we write interchangably Y_i or $Y(i)$ for the ith component of the collection. The dependent sum $\sum_{i \in I} Y_i$ is the disjoint union of all of the sets Y_i in the collection, with elements pairs (i, y) where $i \in I$ and $y \in Y_i$, while the dependent function space $\prod_{i \in I} Y_i$ is the set of functions mapping an input $i \in I$ to an element of Y_i. We write $A + B$ for the disjoint union of two sets A and B, with injections $\mathsf{inl} : A \to A + B$

and $\mathsf{inr} : B \to A + B$. We denote the set of natural numbers by $\mathbb{N}$, the set of positive natural numbers by $\mathbb{N}^+$, and write $[n] = \{0, \ldots, n-1\}$ for a canonical n-element set. We also write $\mathbf{1} = [1]$, $\mathbf{2} = [2]$ and so on for fixed small finite sets.

2. Coalgebraic Preliminaries

We assume familiarity with basic category theory.

2.1. *Final coalgebras*

Let $\mathcal{C}$ be a category and $F : \mathcal{C} \to \mathcal{C}$ an endofunctor. An F-coalgebra is a pair (A, α), where A is an object of $\mathcal{C}$, and $\alpha : A \to FA$ is a morphism. An F-coalgebra homomorphism from (A, α) to (B, β) is a morphism $f : A \to B$ preserving the coalgebra structure, i.e. such that the following diagram commutes:

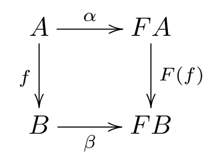

F-coalgebras and F-coalgebra homomorphisms form a category. If F is well behaved (e.g. finitary), this category will have a final object, called the *final F-coalgebra*, and denoted $(\nu F, \mathsf{out})$. Its universal property is a *corecursion* or *coinduction principle*: for every F-coalgebra (A, α), there exists a unique coalgebra homomorphism $\mathsf{unfold} : (A, \alpha) \to (\nu F, \mathsf{out})$. We will make use of Lambek's Lemma, which says that for a final coalgebra $(\nu F, \mathsf{out})$, the map $\mathsf{out} : \nu F \to F(\nu F)$ is an isomorphism.

2.2. *Coinductive families and predicates*

Let I be a set. The category Set^I of *I-indexed families of sets* is the category whose objects are functors from I (viewed as a discrete category) to Set, and whose morphisms are natural transformations. In other words, given two I-indexed families P and Q, a morphism $f : P \to Q$ in Set^I is a family of functions $f_i : P_i \to Q_i$ for each $i \in I$.

A *coinductive family* indexed by I is the final coalgebra νG of an endofunctor G on Set^I. Its corresponding "coinduction principle" says that for every G-coalgebra (P, g), there exists a unique morphism $\mathsf{unfold} : P \to$

νG making the following diagram in Set^I commute:

$$\begin{array}{ccc} P & \xrightarrow{g} & G(P) \\ {\scriptstyle \mathsf{unfold}}\downarrow & & \downarrow{\scriptstyle G(\mathsf{unfold})} \\ \nu G & \xrightarrow[\mathsf{out}]{} & G(\nu G) \end{array}$$

Let us spell out this principle. It says that for every I-indexed family P, if there is a family of functions $g_i : P_i \to G(P)_i$, then there is a unique family of morphisms $\mathsf{unfold}_i : P_i \to \nu G_i$ commuting with the coalgebra maps.

In particular, we will be interested in coinductive families indexed by the carrier X of a coalgebra (X, γ) of a functor $F : \mathsf{Set} \to \mathsf{Set}$, in which case there is a canonical way to obtain coinductive families via predicate liftings of F, as we now explain.

A *(proof-relevant) predicate lifting* of a functor $F : \mathsf{Set} \to \mathsf{Set}$ is a natural transformation $\{\varphi_X : \mathsf{Set}^X \to \mathsf{Set}^{FX}\}_{X \in \mathsf{Set}}$. Given an F-coalgebra $(X, \gamma : X \to FX)$ and a predicate lifting φ of F, we can define an endofunctor on Set^X by

$$\mathsf{Set}^X \xrightarrow{\varphi_X} \mathsf{Set}^{FX} \xrightarrow{-\circ\gamma} \mathsf{Set}^X$$

and consider its final coalgebra. Similarly, in Section 3.3, we also use *proof-irrelevant* predicate liftings $\{\varphi_X : \mathcal{P}(X) \to \mathcal{P}(FX)\}_{X \in \mathsf{Set}}$.

3. Infinite Extensive Form Games

In the game theory literature [13, 18, 22], extensive form games are typically defined using a non-recursive formulation. We take advantage of a more categorical presentation, as it is more compact, supports (co-)recursive function definitions and (co-)inductive reasoning, and smoothly generalises to richer semantic domains, e.g. metric, probabilistic and topological spaces. Throughout this section, let P be a finite set of players.

In this work we separate the description of an extensive form game into its 'tree' part, representing the dynamical structure of the game, and its 'payoff' part, classically given as a decoration on the leaves. The reason we do so stems from our goal of treating infinite games, where plays may be infinite. In that case, the correspondence between leaves and paths in the tree breaks down, making decorations on the leaves an ill-suited device to specify payoff functions.

Below we define first the tree part of a game (Definition 3.1), then the family of payoffs over a given extensive form tree (Definition 3.10), and finally an "extensive form game" as a combination a given tree part with a payoff over it (Definition 3.13). During our journey to the latter definition, we also define the set of *moves* of a game as paths in its tree (Definition 3.7), and the set of *strategy profiles* as choices of a branch for each node of the tree (Definition 3.4). These are fundamental ingredients in the description and analysis of a game.

Definition 3.1. The set ETree^∞ of *infinite extensive form game trees with players* P is the final coalgebra $(\mathsf{ETree}^\infty, \mathsf{out}_{\mathsf{ETree}^\infty})$ of the functor $F_{\mathsf{ETree}} : \mathsf{Set} \to \mathsf{Set}$ defined by

$$F_{\mathsf{ETree}}(X) = \mathbf{1} + P \times \sum_{n \in \mathbb{N}^+} X^n.$$

This supports the Haskell-like data type

$$\text{data } \mathsf{ETree}^\infty = \mathsf{Leaf} \mid \mathsf{Node}\ P\ (n : \mathbb{N}^+)\ ([n] \to \mathsf{ETree}^\infty)$$

Concretely a tree $T \in \mathsf{ETree}^\infty$ is either a leaf, indicating no further plays are possible, or an internal node labeled with a player $p \in P$ who is to play at that point in the game, and an *arity* $n \in \mathbb{N}^+$ representing the number of different moves available, followed by n subtrees. Crucially, being a final coalgebra, ETree^∞ includes paths of infinite depth.

Example 3.2 (Dollar Auction)**.** The Dollar auction is an infinite game introduced by Shubik [24] to exemplify a situation of 'rational escalation'. The game has two players, A and B, bidding over a dollar bill. Player A bids first and then players alternate turns. At each turn, a player chooses between two actions:

- *quit*, in which case the game ends and the other player wins the \$1.
- *bid*, which costs \$0.1, and yields the turn to the other player.

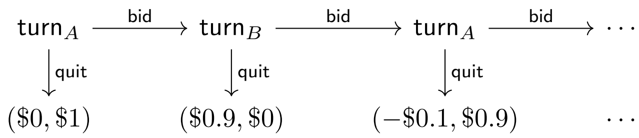

Notice that when players bid, they immediately pay and are not refunded in case they lose the auction. The tree part of this game can be defined by mutual corecursion:

$$\mathsf{Dollar}_A = \mathsf{Node}\ A\ 2\ (\mathsf{Leaf}, \mathsf{Dollar}_B)$$
$$\mathsf{Dollar}_B = \mathsf{Node}\ B\ 2\ (\mathsf{Leaf}, \mathsf{Dollar}_A)$$

Then $\mathsf{Dollar} := \mathsf{Dollar}_A$, as A moves first. As an F_{ETree}-coalgebra, Dollar is defined starting from the coalgebra (D, δ), where $D = \{\mathsf{turn}_A, \mathsf{turn}_B, \mathsf{end}\}$, $P = \{A, B\}$ and δ is defined by:

$$\begin{aligned} &\delta : D \longrightarrow \mathbf{1} + P \times \sum_{n \in \mathbb{N}^+} D^n \\ &\delta\ \mathsf{turn}_A = \mathsf{inr}\ (A, 2, (\mathsf{end}, \mathsf{turn}_B)) \\ &\delta\ \mathsf{turn}_B = \mathsf{inr}\ (B, 2, (\mathsf{end}, \mathsf{turn}_A)) \\ &\delta\ \mathsf{end} = \mathsf{inl}\ * \end{aligned}$$

The coalgebra (D, δ) can be represented by the automaton below, where the two elements of $\mathbf{2}$ are named quit and bid:

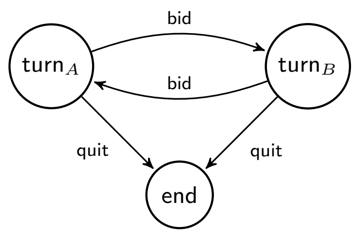

By terminality of $(\mathsf{ETree}^\infty, \mathsf{out}_{\mathsf{ETree}^\infty})$, there is a unique map $\mathsf{unfold}_{(D,\delta)} : (D, \delta) \to (\mathsf{ETree}^\infty, \mathsf{out}_{\mathsf{ETree}^\infty})$, and we define $\mathsf{Dollar} := \mathsf{unfold}_{(D,\delta)}(\mathsf{turn}_A)$. Thus, Dollar is the following infinite tree:

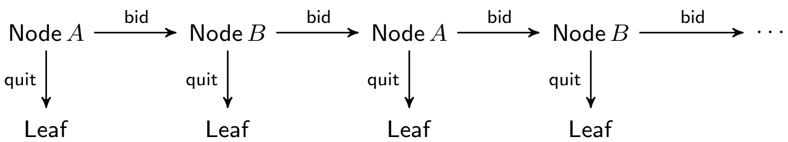

The payoff decoration making this into *the* dollar auction game will be discussed later in Example 3.11.

Example 3.3 (Repeated games)**.** Let T be a finite, perfect-information, extensive-form game, with set of players P. Such games (without utility information) are represented as the elements of the initial algebra of F_{ETree} (see Capucci, Ghani, Ledent, and Nordvall Forsberg [9, Section 2]). Any such tree can be converted to an F_{ETree}-coalgebra given by the automaton whose states and transitions correspond, respectively, to nodes and branches of T.

If we now identify the final states (given by leaves of T) with the initial state (given by the root of T) of the automaton, we get another F_{ETree}-coalgebra (Rep_T, ρ_T): the *repeated game coalgebra.*

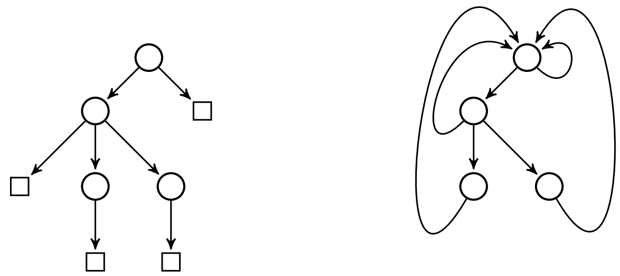

By terminality of $(\mathsf{ETree}^\infty, \mathsf{out}_{\mathsf{ETree}^\infty})$, there is a unique map $\mathsf{unfold}_{(\mathsf{Rep}_T, \rho_T)} : (\mathsf{Rep}_T, \rho_T) \to (\mathsf{ETree}^\infty, \mathsf{out}_{\mathsf{ETree}^\infty})$, and we define

$$T^\infty := \mathsf{unfold}_{(\mathsf{Rep}_T, \rho_T)}(\mathsf{root}).$$

One concrete example is the Market Entry game [23], a game with players $P = \{A, B\}$ described by the extensive-form tree M (here with payoff-labeled leaves):

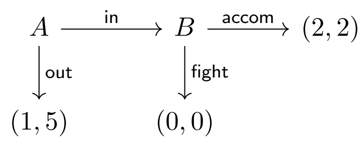

Player A decides whether to enter a new market or not. If staying out, the game ends, but if A enters then player B has to decide whether to accommodate or fight the incumbent. In this case (Rep_T, ρ_T) corresponds to the automaton:

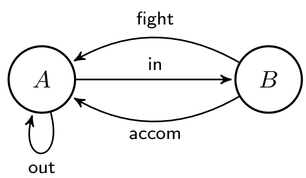

The induced infinite tree M^∞ is then given by

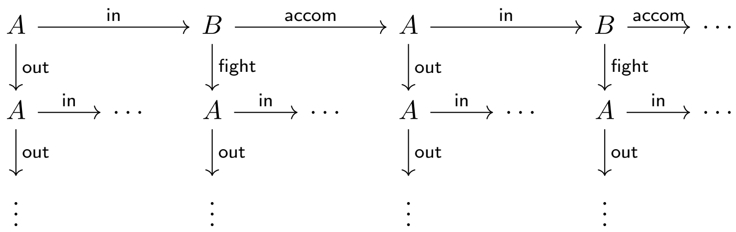

3.1. *Strategies and moves*

Throughout the section, let (X, γ) be an F_{ETree}-coalgebra.

3.1.1. *Strategy profiles*

A strategy profile for the coalgebra (X, γ) at state $x \in X$ consists of a choice of an action at each node in the game tree induced by (X, γ). Recall that an element of $F_{\mathsf{ETree}}(X)$ is either a leaf (of the form $\mathsf{inl}\,*$) or an internal node of the form $\mathsf{inr}\,(q, n, f)$, where q is a player, $n \in \mathbb{N}^+$, and $f : [n] \to X$.

Definition 3.4. We define the family of *strategy profiles* $\mathsf{prof}_{(X,\gamma)} : X \to \mathsf{Set}$ as the final coalgebra associated with the predicate lifting

$$\begin{aligned}
&\varphi_{\mathsf{prof},X} : (X \to \mathsf{Set}) \to F_{\mathsf{ETree}}(X) \to \mathsf{Set}\\
&\varphi_{\mathsf{prof},X}\ P\ (\mathsf{inl}\,*) = \mathbf{1}\\
&\varphi_{\mathsf{prof},X}\ P\ (\mathsf{inr}\,(q, n, f)) = [n] \times \prod_{a \in [n]} P_{f(a)}
\end{aligned}$$

i.e. we define $\mathsf{prof}_{(X,\gamma)}$ as the final coalgebra of the functor $F_{\mathsf{Prof}} : \mathsf{Set}^X \to \mathsf{Set}^X$ defined by $F_{\mathsf{Prof}}(P) = \varphi_{\mathsf{prof},X}(P) \circ \gamma$.

That $\mathsf{prof}_{(X,\gamma)}$ is the final coalgebra implies that for every $x \in X$, there is an isomorphism $s_x : \mathsf{prof}_{(X,\gamma)}(x) \to \varphi_{\mathsf{prof},X}(\mathsf{prof}_{(X,\gamma)})(\gamma(x))$. If $\gamma(x) = \mathsf{inr}\,(q, n, f)$, we thus have

$$s_x(\sigma) \in [n] \times \prod_{a \in [n]} \mathsf{prof}_{(X,\gamma)}(f\,a)$$

and we write $s_x(\sigma) = (\mathsf{now}\,\sigma, \mathsf{next}\,\sigma)$. Intuitively, $\mathsf{now}\,\sigma \in [n]$ is the branch chosen by player q in the current node x of the game tree; and $\mathsf{next}\ \sigma$ is a function assigning, for every branch $a \in [n]$ — not only the one chosen by player q — a strategy profile $\mathsf{next}\,\sigma\,a \in \mathsf{prof}_{(X,\gamma)}(f(a))$ starting at node $f(a)$.

Example 3.5 (Dollar Auction (continues from Example 3.2))**.** For the Dollar game of Example 3.2, we would expect the set of strategy profiles to be isomorphic to $\mathbf{2}^{\mathbb{N}}$, since a strategy profile selects, for every node of the game, an action in $\mathbf{2} \cong \{\mathsf{bid}, \mathsf{quit}\}$.

Formally, we check that $\mathsf{prof}_{(D,\delta)} : D \to \mathsf{Set}$ is the following family of sets, where (D, δ) is the coalgebra that defines the Dollar game in Example 3.2.

$$\mathsf{prof}_{(D,\delta)}(x) \cong \begin{cases} \mathbf{1} & \text{if } x = \mathsf{end} \\ \mathbf{2}^{\mathbb{N}} & \text{if } x = \mathsf{turn}_A, \mathsf{turn}_B \end{cases}$$

Indeed, a function $\sigma \in \mathbf{2}^{\mathbb{N}}$ contains exactly the data of a strategy profile in $\mathsf{prof}_{(D,\delta)}(\mathsf{turn}_A)$, since we can define

$$\begin{aligned} &\mathsf{now}\,\sigma = \sigma(0) \in \mathbf{2} \\ &\mathsf{next}\ \sigma\ \mathsf{quit} = * \in \mathbf{1} \cong \mathsf{prof}_{(D,\delta)}(\mathsf{end}) \\ &\mathsf{next}\ \sigma\ \mathsf{bid} = (n \mapsto \sigma(n+1)) \in \mathbf{2}^{\mathbb{N}} \cong \mathsf{prof}_{(D,\delta)}(\mathsf{turn}_B) \end{aligned}$$

and similarly for a profile $\sigma \in \mathsf{prof}_{(D,\delta)}(\mathsf{turn}_B)$. It is straightforward (but cumbersome) to check that this satisfies the universal property of the final coalgebra of F_{Prof}.

Example 3.6 (Repeated games (continues from Example 3.3))**.** For a finite game T, we have defined in Example 3.3 its repeated game coalgebra (Rep_T, ρ_T), whose unfolding is the infinitely repeated game T^∞. A strategy profile for (Rep_T, ρ_T) with initial state x is the greatest solution to

$$\mathsf{prof}_{(\mathsf{Rep}_T,\rho_T)}(x) \cong \{\text{strategy profiles of } T|_x\} \times \prod_{\ell \in \mathsf{leaves}\ x} \mathsf{prof}_{(\mathsf{Rep}_T,\rho_T)}(\mathsf{root})$$

where $T|_x$ is the subtree of T starting at $x \in \mathsf{Rep}_T$, root is the state corresponding to the root of the tree T, and $\mathsf{leaves}\ x$ denotes the set of leaves in the subtree $T|_x$. In the concrete case of the market entry game, this becomes (where we put $\mathsf{prof}_M = \mathsf{prof}_{(\mathsf{Rep}_M,\rho_M)}$ to ease notation):

$$\mathsf{prof}_M(A) \cong \{\mathsf{in}, \mathsf{out}\} \times \{\mathsf{fight}, \mathsf{accom}\} \times \underbrace{\mathsf{prof}_M(A) \times \mathsf{prof}_M(A) \times \mathsf{prof}_M(A)}_{\text{3 leaves accessible from } A}$$

$$\mathsf{prof}_M(B) \cong \{\mathsf{fight}, \mathsf{accom}\} \times \underbrace{\mathsf{prof}_M(A) \times \mathsf{prof}_M(A)}_{\text{2 leaves accessible from } B}$$

Therefore $\mathsf{prof}_M(A)$ is the final coalgebra of the functor

$$X \mapsto \{\mathsf{in}, \mathsf{out}\} \times \{\mathsf{fight}, \mathsf{accom}\} \times X^3$$

3.1.2. *Moves*

The set of moves in the game is the set of paths in the tree, which is another coinductive family.

Definition 3.7. We define the family of *moves* $\mathsf{moves} : X \to \mathsf{Set}$ as the final coalgebra associated with the predicate lifting

$$\begin{aligned}
&\varphi_{\mathsf{moves},X} : (X \to \mathsf{Set}) \to F_{\mathsf{ETree}}(X) \to \mathsf{Set} \\
&\varphi_{\mathsf{moves},X}\ P\ (\mathsf{inl}\ *) = \mathbf{1} \\
&\varphi_{\mathsf{moves},X}\ P\ (\mathsf{inr}\ (q,n,f)) = \sum_{a\in[n]} P_{f(a)}
\end{aligned}$$

i.e. we define $\mathsf{moves}_{(X,\gamma)}$ as the final coalgebra of the functor $F_{\mathsf{moves}} : \mathsf{Set}^X \to \mathsf{Set}^X$ defined by $F_{\mathsf{moves}}(P) = \varphi_{\mathsf{moves},X}(P) \circ \gamma$.

Again, for every $x \in X$, we have an isomorphism $m_x : \mathsf{moves}_{(X,\gamma)}(x) \to \varphi_{\mathsf{moves},X}(\mathsf{moves}_{(X,\gamma)})(\gamma(x))$. If $\gamma(x) = \mathsf{inr}\ (q, n, f)$, we have

$$m_x(\pi) \in \sum_{a\in[n]} \mathsf{moves}_{(X,\gamma)}(f\ a)$$

Note that, as m_x is iso, if $\gamma(x) = \mathsf{inr}\,(q, n, f)$ then for each $a \in [n]$ and $\pi' \in \mathsf{moves}_{(X,\gamma)}(f\ a)$ there is a unique element $\mathsf{cons}(a, \pi') \in \mathsf{moves}_{(X,\gamma)}(x)$ such that $m_x(\mathsf{cons}(a, \pi')) = (a, \pi')$. Intuitively, $\mathsf{cons}(a, \pi')$ represents a path in the game tree starting at node x, where the first chosen branch is $a \in [n]$, and the following branches are chosen according to the path π'.

Example 3.8 (Dollar Auction (continues from Example 3.5))**.** The moves of Dollar are given by the final coalgebra of $X \mapsto \mathbf{1} + X$, i.e.

$$\mathsf{moves}(\mathsf{Dollar}) \cong \mathbf{1} + \mathsf{moves}(\mathsf{Dollar})$$

The final coalgebra of this functor is known as the *conatural numbers* $(\mathbb{N}^\infty, \mathsf{pred})$, which include all finite natural numbers and an 'infinite' number ω. The map pred sends $0 \in \mathbb{N}^\infty$ to $\mathsf{inl}\ *$ and every other natural number to the right injection of its predecessor. The predecessor of ω is itself, $\mathsf{pred}\ \omega = \mathsf{inr}\,\omega$. Note that it is not decidable if a given conatural number x is finite or infinite; however, by applying pred a finite number of times, we can decide if $x \geq n$ for any finite natural number n.

We interpret $n \in \mathsf{moves}(\mathsf{Dollar})$ as the path starting from the root and ending at the n-th leaf, i.e. the play where players bid n times before one of them[a] decides to quit. The unique infinite play ω corresponds to infinite escalation, with players never quitting.

[a]The player who quits can be determined from the parity of n.

Similarly, $\mathsf{moves}_{(D,\delta)}$ is given by

$$\mathsf{moves}_{(D,\delta)}(x) \cong \begin{cases} \mathbf{1} & \text{if } x = \mathsf{end} \\ \mathbb{N}^\infty & \text{if } x = \mathsf{turn}_A, \mathsf{turn}_B \end{cases}$$

Example 3.9 (Repeated games (continues from Example 3.6))**.** For any finite extensive-form game tree T, one has

$$\mathsf{moves}_{(\mathsf{Rep}_T, \rho_T)}(x) \cong (\mathsf{leaves}\ x) \times (\mathsf{leaves\ root})^{\mathbb{N}}.$$

In the specific instance of the market entry game M, moves are three: $1 : A \overset{\mathsf{out}}{\to} *$, $2 : A \overset{\mathsf{in}}{\to} B \overset{\mathsf{accom}}{\to} *$ and $3 : A \overset{\mathsf{in}}{\to} B \overset{\mathsf{fight}}{\to} *$, forming the set $\mathbf{3}$. These are all accessible from A, therefore M^∞ has set of moves specified by the final coalgebra of

$$X \mapsto \mathbf{3} \times X$$

which is readily seen to be $\mathbf{3}^{\mathbb{N}}$. Indeed, a move of the repeated game is a move for every stage game. On the other hand, only moves 2 and 3 are accessible from B, therefore we get

$$\mathsf{moves}_{(\mathsf{Rep}_M, \rho_M)}(x) \cong \begin{cases} \mathbf{3} \times \mathbf{3}^{\mathbb{N}} & \text{if } x = A \\ \mathbf{2} \times \mathbf{3}^{\mathbb{N}} & \text{if } x = B \end{cases}$$

3.2. *Evaluating strategies*

In order to compare strategies, we need a way to assign a payoff to them. This is done in two steps: the *play function* turns a strategy profile into a sequence of moves; and the *payoff function* explains how outcomes turn into rewards for the players. This will allow us, in Section 3.3, to define several equilibrium concepts, i.e., predicates on strategy profiles that express when all players are happy with their given rewards.

3.2.1. *The Play Function*

We can use the universal property of final coalgebras to define a play function $\mathsf{play}_{(X,\gamma)} : \mathsf{prof}_{(X,\gamma)} \to \mathsf{moves}_{(X,\gamma)}$ which computes the sequence of moves generated by playing according to a strategy profile.

To define $\mathsf{play}_{(X,\gamma)} : \mathsf{prof}_{(X,\gamma)} \to \mathsf{moves}_{(X,\gamma)}$, we use the finality of $\mathsf{moves}_{(X,\gamma)}$, i.e., we want to give $\mathsf{prof}_{(X,\gamma)}$ a F_{moves}-coalgebra structure. It is sufficient to give, for $Q \in \mathsf{Set}^X$, a natural transformation $p_Q : \varphi_{\mathsf{prof},X}(Q) \to \varphi_{\mathsf{moves},X}(Q)$ in $\mathsf{Set}^{F_{\mathsf{ETree}}(X)}$, which we can do as follows.

$$p_Q\ (\mathsf{inl}\ *)\ * = *$$
$$p_Q\ (\mathsf{inr}\ (q, n, f))\ (a, \sigma) = (a, \sigma(a))$$

Instantiating at component $Q = \mathsf{prof}_{(X,\gamma)}$ and composing with the isomorphism s_x gives $\mathsf{prof}_{(X,\gamma)}$ a F_{moves}-coalgebra structure, as required. Hence there is a unique family of functions $\mathsf{play}_{(X,\gamma)}(x) : \mathsf{prof}_{(X,\gamma)}(x) \to \mathsf{moves}_{(X,\gamma)}(x)$, which, up to the isomorphisms s_x and m_x, satisfies the following definition.

$$\begin{aligned}
&\mathsf{play} \in \prod_{x\in X} \left(\mathsf{prof}_{(X,\gamma)}(x) \to \mathsf{moves}_{(X,\gamma)}(x)\right) && \\
&\mathsf{play}(x, *) = * && \text{when } \gamma(x) = (\mathsf{inl}\, *) \\
&\mathsf{play}(x, (a, \sigma)) = (a, \mathsf{play}(f(a), \sigma(a))) && \text{when } \gamma(x) = \mathsf{inr}\,(q, n, f)
\end{aligned}$$

3.2.2. *Payoff functions and the game coalgebra*

From now on, let R be a set of 'rewards' or 'payoffs'. Common choices are $R = \mathbb{Q}$, the rational numbers or $R = \mathbb{R}$, the real numbers, but dealing with infinite plays quickly leads to contemplating infinite rewards (negative and positive) as well. We assume R is ordered, and eventually, we will need to assume that this order is trichotomous. Recall that we denote by P the set of players. We write R^P for the set of functions $R \to P$, representing choices of payoffs for each of the players.

Definition 3.10. A payoff function for an F_{ETree}-coalgebra (X, γ) at state $x \in X$ is a function

$$u : \mathsf{moves}_{(X,\gamma)}(x) \to R^P.$$

The set of payoff functions for each $x \in X$ is denoted by $\mathsf{pay}_{(X,\gamma)}(x)$.

Example 3.11 (Dollar Auction (continues from Example 3.8))**.** Recall that $\mathsf{moves}_{(D,\delta)}(\mathsf{turn}_p)$ is given by the conatural numbers. The payoff function for the Dollar Auction game (where $R = [-\infty, +\infty)$[b]) can be defined in two steps. First, we coinductively define a map into colists[c] of payoffs $\mathsf{List}^\infty\ R^P$

[b]We assume that $-\infty + m = -\infty$ for every $m \in R$.
[c]Colists of A are 'possibly infinite lists', i.e. terms of the final coalgebra of $X \mapsto 1 + A \times X$ for a given A. We denote $\mathsf{inl}\, *$ as Empty and $\mathsf{inr}\,(a, x)$ as $a :: x$.

(which we think of as 'ledgers'):

$$\mathsf{led} \in \prod_{x \in D} \left(\mathsf{moves}_{(D,\delta)}(x) \to \mathsf{List}^{\infty}\ R^P\right)$$

$$\mathsf{led}\ \mathsf{end}\ * = \mathsf{Empty}$$

$$\mathsf{led}\ \mathsf{turn}_A\ m = \begin{cases} (A \mapsto 0, B \mapsto 1) :: \mathsf{Empty} & \mathsf{pred}\ m = \mathsf{inl}\ * \\ (A \mapsto -0.1, B \mapsto 0) :: (\mathsf{led}\ \mathsf{turn}_B\ n) & \mathsf{pred}\ m = \mathsf{inr}\ n \end{cases}$$

$$\mathsf{led}\ \mathsf{turn}_B\ m = \begin{cases} (A \mapsto 1, B \mapsto 0) :: \mathsf{Empty} & \mathsf{pred}\ m = \mathsf{inl}\ * \\ (A \mapsto 0, B \mapsto -0.1) :: (\mathsf{led}\ \mathsf{turn}_A\ n) & \mathsf{pred}\ m = \mathsf{inr}\ n \end{cases}$$

Then the actual utility function is given by summing up componentwise all the payoffs collected by the players during the game:

$$u_{\mathsf{Dollar}}\ m = \sum_{n=0}^{+\infty} p_i, \quad \text{where } p = \mathsf{led}\ \mathsf{turn}_A\ m$$

Here p_i is defined to be zero when i is greater than the length of p.

In the case of $m = \omega$, this will unfold into an infinite sum where the summands alternate between $(A \mapsto -0.1, B \mapsto 0)$ and $(A \mapsto 0, B \mapsto -0.1)$, therefore yielding the payoff vector $(A \mapsto -\infty, B \mapsto -\infty)$.

Example 3.12 (Repeated games (continues from Example 3.9))**.** The payoff function of an infinitely repeated game is obtained similarly to the previous example: 'partial' payoffs are summed at each iteration of the stage game. Unlike the Dollar Auction however, in an infinitely repeated game all plays are infinite, therefore *discounting* has to be adopted. This means that at each stage of the game, payoff vectors are uniformly scaled by a discount factor $0 < \delta < 1$. Discounting reflects the real-world tendency to value future payoffs less than present ones.

Thus let $v_T \in \prod_{x \in \mathsf{Rep}_T} \left(\mathsf{leaves}(x) \to \mathbb{R}^P\right)$ be the utility function of the (finite) stage game T (such as the one represented by the diagram of M in Example 3.3). For a given discount factor δ, we get a payoff function $u^{\delta}_{T^\infty} \in \prod_{x \in \mathsf{Rep}_T} \left(\mathsf{moves}_{(\mathsf{Rep}_T, \rho_T)}(x) \to \mathbb{R}^P\right)$ by setting (recall that $\mathsf{moves}_{(\mathsf{Rep}_T, \rho_T)}(x) = (\mathsf{leaves}\ x) \times (\mathsf{leaves}\ \mathsf{root})^{\mathbb{N}}$):

$$u^{\delta}_{T^\infty}\ x\ (m_0, ms) := (v_T\ x\ m_0) + \sum_{i=0}^{+\infty} \delta^{i+1} \cdot (v_T\ x\ (ms\ i)).$$

The assumption of $|\delta| < 1$ guarantees the convergence of such a sum.

We are now ready to define the "game coalgebra" — an F_{ETree}-coalgebra which will enable us to study equilibria. We do this by collecting all the information needed: the current state of the game, a strategy profile, and a payoff function.

Definition 3.13. Let (X, γ) be an F_{ETree}-coalgebra. The *game coalgebra* $(Z_{(X,\gamma)}, \Gamma)$ is the F_{ETree}-coalgebra with carrier set

$$Z_{(X,\gamma)} = \sum_{x \in X} \mathsf{prof}_{(X,\gamma)}(x) \times \mathsf{pay}_{(X,\gamma)}(x)$$

and dynamics given by the map Γ defined by

$$\Gamma(x, \sigma, u) = \begin{cases} \mathsf{inl}\, * & \text{if } \gamma(x) = \mathsf{inl}\, * \\ \mathsf{inr}\ (q, n, a \mapsto (f\, a, \mathsf{next}\, \sigma\, a, u_a)) & \text{if } \gamma(x) = \mathsf{inr}\ (q, n, f) \end{cases}$$

where $u_a(\pi') = u(\mathsf{cons}(a, \pi'))$, for $a \in [n]$ and $\pi' \in \mathsf{moves}_{(X,\gamma)}(f\, a)$.

3.3. *Equilibrium concepts*

We are now going to introduce the equilibrium concepts for our games: a 'one-shot' equilibrium concept that expresses that no player can improve their payoff by unilaterallty deviating from their strategy in exactly one state, and the usual Nash equlibrium that expresses that none of the players can change their strategy unilaterally (at possibly arbitrarily many places) such that their overall payoff increases. Before we are able to introduce the 'one-shot' equilibrium concept, we will first introduce a modal operator that allows to specify properties that should hold everywhere in the game tree.

3.3.1. *The 'Everywhere' modality*

Notions from game theory such as subgame perfection require a predicate to hold at every node of a tree (i.e., in every subgame). This is achieved by the so-called 'everywhere' or 'henceforth' modality. Using standard techniques from coalgebra, we can construct it using a proof-irrelevant predicate lifting as follows:

Definition 3.14. Let (X, γ) be an F_{ETree}-coalgebra. Consider the predicate lifting $\varphi_\Box : \mathcal{P}(X) \to \mathcal{P}(F_{\mathsf{ETree}} X)$ defined by

$$\varphi_\Box(Q) = \{\mathsf{inl}\, *\} \cup \{\mathsf{inr}(q, n, f) \mid (\forall a \in [n])\, f(a) \in Q\}$$

and for a predicate $P \in \mathcal{P}(X)$, define $\Box P$ to be the greatest fixpoint of

$$F_{\Box,P}(U) = P \cap \gamma^{-1}(\varphi_\Box(U)).$$

A detailed discussion of this operator can be found e.g. in Jacobs [12] where $\Box$ is referred to as the "henceforth" operator. Intuitively, given a predicate P, $\Box P$ holds at a state x of a coalgebra if x itself and all its descendants satisfy P, i.e., P is invariably true in the game starting at x. The $\Box$-operator satisfies the properties one would expect from basic modal logic.

Lemma 3.15. *The modality $\Box$ is monotone, i.e. if P implies Q then $\Box P$ implies $\Box Q$. Furthermore $\Box P \subseteq P$ and $\Box P \subseteq \Box\Box P$.*

Proof. Assume $P \subseteq Q$. To show $\Box P \subseteq \Box Q$, we use the finality of $\Box Q$ to conclude $\Box P \subseteq \Box Q$ by showing that $\Box P \subseteq F_{\Box,Q}(\Box P)$. This follows since $\Box P \subseteq F_{\Box,P}(\Box P)$ and $P \subseteq Q$. In the same way, $\Box P \subseteq P$ and $\Box P \subseteq \Box\Box P$. $\Box$

3.3.2. *Unimprovability*

A very simple equilibrium concept is the following: at each node of the game, the current player cannot improve their payoff by changing their action at the current node. We call this the 'one-shot' equilibrium concept. Formally, we can encode it as follows. First we define a predicate Unimprov which verifies that, at a node in $Z_{(X,\gamma)}$, the current strategy is unimprovable for the current player. We then ask that this predicate holds everywhere in the tree using the 'everywhere' modality.

Definition 3.16. We define the predicate Unimprov on $Z_{(X,\gamma)}$ by:

$$
\begin{aligned}
&(x, \sigma, u) \in \mathsf{Unimprov} \\
&\qquad \text{if} \qquad \gamma(x) = \mathsf{inl}\ * \\
&\qquad \text{or} \qquad \gamma(x) = \mathsf{inr}\ (q, n, f) \\
&\qquad\qquad \text{and } \mathsf{now}\,\sigma \in \mathsf{argmax}(a \mapsto \pi_q(u_a(\mathsf{play}\ (f\,a)\ (\mathsf{next}\,\sigma\,a))))
\end{aligned}
$$

The 'one-shot' equilibrium concept can now be defined as $\Box$Unimprov.

In words, $(x, \sigma, u) \in$ Unimprov holds if x is a leaf, or if x is an internal node and $\mathsf{now}\,\sigma$ gives the best payoff for the current player q at the node according to the utility function u, assuming that all players continue to play according to σ. This equilibrium concept also occurs in Lescanne and Perrinel [17] and Abramsky and Winschel [1], who call it "subgame perfect equilibria". We prefer to reserve that name for the predicate $\Box$Nash that we will define in the next section.

3.3.3. *Nash Equilibria and Subgame Perfection*

The predicate Unimprov from the previous section says that a player cannot improve their payoff by changing their action at the current node only. In contrast, Nash equilibria are concerned with deviations where a player may change their action at several nodes simultaneously. The only restriction is that all such nodes must belong to the same player. So, we first define a predicate $\equiv_p$ which characterises when two strategy profiles are the same, except for deviations by one player p. Since we want to allow an infinite number of deviations by player p, we define this as a coinductive predicate.

Definition 3.17. For each player $p \in P$ and state $x \in X$, we define a family of relations

$$\equiv_p \subseteq \mathsf{prof}_{(X,\gamma)}(x) \times \mathsf{prof}_{(X,\gamma)}(x)$$

as the maximal family such that for all $\sigma, \sigma' \in \mathsf{prof}_{(X,\gamma)}(x)$ and $x \in X$, we have $\sigma \equiv_p \sigma'$ if and only if one of the following is satisfied:

(1) $\gamma(x) = \mathsf{inl}\ *$, or
(2) $\gamma(x) = \mathsf{inr}\ (p, n, f)$ and $\mathsf{next}\,\sigma\, a \equiv_p \mathsf{next}\,\sigma'\, a$ for all $a \in [n]$, or
(3) $\gamma(x) = \mathsf{inr}\ (q, n, f)$ with $q \neq p$, $\mathsf{now}\,\sigma = \mathsf{now}\,\sigma'$ and $\mathsf{next}\,\sigma\, a \equiv_p \mathsf{next}\,\sigma'\, a$ for all $a \in [n]$.

We can use the universal property of $\equiv_p$ to deduce the following:

Lemma 3.18. *Assume the set of players P has decidable equality. For each player $p \in P$, the relation $\equiv_p$ is reflexive.* □

Using $\equiv_p$ to talk about deviations, we can now formulate the Nash equilibrium concept. Unlike the 'one-step' equilibrium, the definition of a Nash equlibrium relies on a global quantification over all possible alternative strategies of a player rather than a local quantification over transitions in the game coalgebra. Consequently, Nash is neither an inductive nor a coinductive predicate.

Definition 3.19. In the game coalgebra $(Z_{(X,\gamma)}, \Gamma)$ we define

$$\begin{aligned} &(x, \sigma, u) \in \mathsf{Nash} \\ &\quad \text{if } \forall p \in P.\ \forall \sigma' \in \mathsf{prof}_{(X,\gamma)}(x). \\ &\qquad (\sigma \equiv_p \sigma') \rightarrow (\pi_p u(\mathsf{play}\ x\ \sigma) \geq \pi_p u(\mathsf{play}\ x\ \sigma')) \end{aligned}$$

In other words, given an initial node x and payoff function u, the strategy profile σ is a Nash equilibrium if every possible deviation σ' by player p yields a payoff for p which is smaller than the one of σ.

We can now succinctly define the solution concept of subgame perfect Nash equilibria simply as $\Box\mathsf{Nash}$ — a strategy profile is subgame perfect if it is a Nash equilibrium in every subgame of the tree.

4. Relating Unimprovability and Subgame Perfect Nash Equilibria

In this section, we relate the coalgebraic subgame perfect Nash equilibria $\Box\mathsf{Nash}$ and the one-deviation equilibrium $\Box\mathsf{Unimprov}$, thus connecting our coalgebraic treatment with the standard notions from game theory. One direction is almost immediate:

Lemma 4.1. *Assume the set of players P has decidable equality. Let (x, σ, u) be a state of the game coalgebra $(Z_{(X,\gamma)}, \Gamma)$. If $(x, \sigma, u) \in \Box\mathsf{Nash}$ then $(x, \sigma, u) \in \Box\mathsf{Unimprov}$.*

Proof. Since $\Box$ is monotone by Lemma 3.15, it is sufficient to show that $(x, \sigma, u) \in \mathsf{Nash}$ implies $(x, \sigma, u) \in \mathsf{Unimprov}$. If $\gamma(x) = \mathsf{inl}\,*$, this is trivial, so we concentrate on the case when $\gamma(x) = \mathsf{inr}\,(q, n, f)$. By definition, we have to show that

$$\pi_q(u_{\mathsf{now}\,\sigma}(\mathsf{play}\ (f(\mathsf{now}\,\sigma))\ (\mathsf{next}\,\sigma\ (\mathsf{now}\,\sigma)))) \geq \pi_q(u_a(\mathsf{play}\ (f\ a)\ (\mathsf{next}\,\sigma\ a)))$$

for every $a \in [n]$. For each a, let σ_a be the strategy profile with $\mathsf{now}\,\sigma_a = a$ and $\mathsf{next}\,\sigma_a = \mathsf{next}\,\sigma$. By Lemma 3.18, $\sigma \equiv_q \sigma_a$, and the conclusion follows from the assumption that $(x, \sigma, u) \in \mathsf{Nash}$. □

For the other direction, we need to assume that the utility function is suitably well behaved; this is known as *continuity at infinity* in the game theory literature [10, Chapter 4.2]. We formulate it more generally for arbitrary F_{moves}-coalgebras.

4.1. *Continuity at infinity*

To formally define continuity at infinity we assume that the set of payoffs R is a metric space and R^P is the P-fold product of this metric space obtained via taking the maximum. To obtain a metric on a F_{moves}-coalgebra we use the projections into the terminal sequence of F_{moves}. This technique can be formulated for arbitrary functors on indexed families of sets.

Definition 4.2. Let $H : \mathsf{Set}^X \to \mathsf{Set}^X$ a functor, and (A, γ) an H-coalgebra. Recall that $\top_X(x) = \mathbf{1}$ is the terminal object in Set^X. We define a family of natural transformations $(\gamma^i : A \to H^i(\top_X))_{i\in\mathbb{N}}$ inductively by:

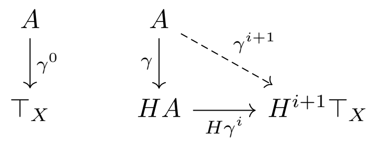

where $\gamma^0 := !_A$ is the unique morphism from A into the terminal object. We call states $a, a' \in A(x)$ *n-step equivalent*, and we write $a \sim_n a'$, when $\gamma^n_x(a) = \gamma^n_x(a')$. This induces a pseudometric on $A(x)$ by putting $d^x_{(A,\gamma)}(a, a') = 2^{-m}$, where $m = \mathsf{sup}\{n \mid a \sim_n a'\}$.

If H is finitary, i.e. determined by its action on finitely presentable objects [2], then if two states in an H-coalgebra (A, γ) agree for all finite observations, they are equal. Hence in this case $d^x_{(A,\gamma)}$ is actually a metric:

Lemma 4.3. *Let $H : \mathsf{Set}^X \to \mathsf{Set}^X$ be a finitary functor. For each H-coalgebra (A, γ) and $x \in X$, $d^x_{(A,\gamma)}$ is a metric on $A(x)$.*

The lemma is a straightforward consequence of a similar result for Set-functors [3, 25]. We apply the above lemma to the functor $F_{\mathsf{moves}} : \mathsf{Set}^X \to \mathsf{Set}^X$ from Definition 3.7, which is finitary. As a result, we are now ready to define continuity at infinity coalgebraically.

Definition 4.4. Let (X, γ) be an F_{ETree}-coalgebra. We call $u \in \mathsf{pay}_{(X,\gamma)}(x)$ continuous at infinity if $u : \mathsf{moves}_{(X,\gamma)}(x) \to R^P$ is *uniformly continuous* as a map between metric spaces, i.e., if

$$\forall\varepsilon > 0.\ \exists\delta > 0.\ \forall m, m'.\ d^x_{\mathsf{moves}_{(X,\gamma)}}(m, m') < \delta \ \to\ d_{R^P}(u(m), (u\ m')) < \varepsilon$$

Remark 4.5. This generalises the usual formulation of continuity at infinity from the game theory literature (see e.g. Fudenberg and Tirole [10, Def. 4.1]) to coalgebras. We observe that the weaker assumption of *pointwise* continuity would be sufficient to prove Theorem 4.12 (or the corresponding traditional statement [10, Theorem 4.2]). Classically, $\mathsf{moves}_{(X,\gamma)}(x)$ is compact [14], and the distinction disappears, but this is of course not constructively valid.

Spelling out the definition of $d^x_{\mathsf{moves}_{(X,\gamma)}}$ and d_{R^P}, we arrive at the following concrete definition of continuity at infinity:

Proposition 4.6. *Let (X, γ) be an F_{ETree}-coalgebra. A payoff function $u \in \mathsf{pay}_{(X,\gamma)}(x)$ is continuous at infinity if and only if*

$$\forall\varepsilon > 0.\ \exists n \in \mathbb{N}.\ \forall m, m'.\ (m \sim_n m') \ \to\ \forall p \in P.\ d_R(\pi_p(u\ m), \pi_p(u\ m')) < \varepsilon$$

□

Example 4.7 (Dollar Auction (continues from Example 3.11))**.** We claim the payoff function for Dollar is continuous at infinity. It will suffice to focus on one component, say $u_A = \pi_A \circ u$, since $\pi_B \circ u$ is the same up to a shift. Let us begin by specifying a metric on $R = [-\infty, +\infty)$:[d]

$$d_R(r, r') = |\arctan\, r - \arctan\, r'|$$

This choice of metric makes R into a bounded space, since evidently $\mathrm{diam}(R) = \pi$. In particular, $d_R((u_A\ m), (u_A\ m'))$ is finite for every $m, m' \in \mathsf{moves}_{(D,\delta)}(x)$.

By applying tan at both sides[e] of $d_R((u_A\ m), (u_A\ m')) < \varepsilon$, our thesis becomes

$$\forall \varepsilon > 0.\ \exists n \in \mathbb{N}.\ \forall m, m'.\ m \sim_n m' \ \to\ \frac{|(u_A\ m) - (u_A\ m')|}{1 + (u_A\ m)(u_A\ m')} < \tan\, \varepsilon. \quad (1)$$

Observe that, when $m, m' \to +\infty$,

$$g((u_A\ m), (u_A\ m')) \longrightarrow 0 \quad \text{where } g(x, y) = \frac{|x - y|}{1 + xy}.$$

since $(u_A\ m), (u_A\ m') \to -\infty$ by the definition of u_A and $g(x, y) \to 0$ as $x, y \to -\infty$. This convergence gives an $n \in \mathbb{N}$ for each chosen $\tan\, \varepsilon > 0$. Suppose now $m \sim_n m'$. In this particular example, this can happen if and only if either $m, m' < n$ and $m = m'$, or $m, m' \geq n$. In the first case, $d_{R^P}((u_A\ m), (u_A\ m')) = 0 < \tan\, \varepsilon$. In the second case, we have chosen n to satisfy $d_{R^P}((u_A\ m), (u_A\ m')) < \tan\, \varepsilon$. Thus we conclude that (1) is satisfied.

Example 4.8 (Repeated games (continues from Example 3.12))**.** The utility function of a repeated game with discounting is almost immediately seen to be continuous. Setting $v = \pi_p \circ (v_T\ A)$ (using notation from Example 3.12), we see that we are tasked to prove:

$$\begin{aligned}
&\forall \varepsilon > 0.\ \exists n \in \mathbb{N}.\ \forall (m_0, ms), (m'_0, ms').\\
&\quad (m_0, ms') \sim_n (m'_0, ms') \ \to\\
&\qquad \left| (v(m_0) - v(m'_0)) + \sum_{i=0}^{+\infty} \delta^{i+1} (v(ms\ i) - v(ms'\ i)) \right| < \varepsilon
\end{aligned}$$

[d]We cannot use the traditional Euclidean metric on $R = \mathbb{R}$ since $|-\infty - r| = +\infty$ for finite $r \in \mathbb{R}$. Any well-defined metric that makes R into a bounded space is equivalent.
[e]The function tan, by virtue of being monotone on the domain $(-\pi/2, +\pi/2)$, preserves inequalities for small enough ε (and, by previous considerations on the diameter of (R, d_R), for every value of $d_R((u_A\ m), (u_A\ m'))$).

In this case, $(m_0, ms') \sim_n (m'_0, ms')$ holds exactly when m_0 and m'_0 agree and (if $n > 0$) if ms and ms' also agree on their first n entries. When this happens, the first $n+1$ terms of the series cancel out. By convergence of said series (easily obtainable by comparison with a geometric series), we can make that quantity as small as we need to by eliding enough leading terms.

4.2. *The one-shot deviation principle*

The one-shot equilibrium concept states that there is no profitable single-node deviation. As an intermediate step towards subgame perfect Nash equilibria, we can also consider profitable deviations in a finite number of nodes. Following Lescanne [15, §5], this concept can be formalised as an *inductive* definition as follows:

Definition 4.9. Let $p \in P$ be a player, $x \in X$ a state, and (X, γ) an F_{ETree}-coalgebra. We define a family of relations $\equiv_p^{\text{fin}} \subseteq \mathsf{prof}_{(X,\gamma)}(x) \times \mathsf{prof}_{(X,\gamma)}(x)$ inductively as the least family such that for all $\sigma, \sigma' \in \mathsf{prof}_{(X,\gamma)}(x)$ and $x \in X$ we have $\sigma \equiv_p^{\text{fin}} \sigma'$ iff one of the following is satisfied

(1) $\sigma = \sigma'$, or
(2) $\gamma(x) = \mathsf{inl} \ast$, or
(3) $\gamma(x) = \mathsf{inr}\; p\; n\; f$ and $\mathsf{next}\,\sigma\, a \equiv_p^{\text{fin}} \mathsf{next}\,\sigma'\, a$ for all $a \in [n]$, or
(4) $\gamma(x) = \mathsf{inr}\; q\; n\; f$ with $q \neq p$, $\mathsf{now}\,\sigma = \mathsf{now}\,\sigma'$ and $\mathsf{next}\,\sigma\, a \equiv_p^{\text{fin}} \mathsf{next}\,\sigma'\, a$ for all $a \in [n]$.

Thus strategies with $\sigma \equiv_p^{\text{fin}} \sigma'$ can differ in their choice of action $\mathsf{now}\,\sigma \neq \mathsf{now}\,\sigma'$ at p-nodes; since the definition is inductive, this can only happen a finite number of times before reaching a base case. Given two strategies σ and σ', we can "truncate" σ after n rounds by replacing it with σ', resulting in a new strategy $\lceil \sigma \rceil_n^{\sigma'}$.

Lemma 4.10. *If $\sigma \equiv_p \sigma'$, then $\sigma \equiv_p^{fin} \lceil \sigma' \rceil_n^{\sigma}$ for any $n \in \mathbb{N}$. Conversely, if the set of players P has decidable equality, then $\sigma \equiv_p^{fin} \sigma'$ implies $\sigma \equiv_p \sigma'$.*

Proof. The first statement follows by induction on n. Since $\lceil \sigma' \rceil_0^{\sigma} = \sigma$, we have by $\sigma \equiv_p^{\text{fin}} \lceil \sigma' \rceil_0^{\sigma}$ by (1). For the step case, consider the form of the proof $\sigma \equiv_p \sigma'$: if $\sigma \equiv_p \sigma'$ because $\gamma(x) = \mathsf{inl} \ast$, we immediately have $\sigma \equiv_p^{\text{fin}} \lceil \sigma' \rceil_{n+1}^{\sigma}$; otherwise $p = q$ or $p \neq q$, and in either case $\mathsf{next}\,\sigma\, a \equiv_p \mathsf{next}\,\sigma'\, a$ for all $a \in [m]$, hence by the induction hypothesis $\mathsf{next}\,\sigma\, a \equiv_p^{\text{fin}} \lceil \mathsf{next}\,\sigma'\, a \rceil_n^{\mathsf{next}\,\sigma\, a}$ for every a, and by either (3) or (4) and how $\lceil \sigma \rceil_n^{\sigma'}$ is defined, we have $\sigma \equiv_p^{\text{fin}} \lceil \sigma' \rceil_{n+1}^{\sigma}$ as required.

The second statement follows easily by induction on the proof that $\sigma \equiv_p^{\text{fin}} \sigma'$, using decidable equality on P to invoke Lemma 4.1 for the base case (1). $\square$

Although allowing a finite number of deviations might seem like a stronger notion of equilibrium than allowing just one, they turn out to be equivalent. This is because the one-shot equilibrium concept is quantified on every subgame: assuming a player can improve their payoff with a finite number of deviations, we can find a single profitable deviation by restricting to the subgame starting at the last deviation. Recall that an order relation $<$ is *trichotomous* if, for every pair of elements x, y, it is decidable whether $x < y$ or or $x = y$ or $x > y$.

Lemma 4.11. *Let (x, σ, u) be a state of the game coalgebra $(Z_{(X,\gamma)}, \Gamma)$. Assume that the order relation $<$ on R is trichotomous. If there is a player p and a strategy σ' such that $\sigma \equiv_p^{fin} \sigma'$ and $\pi_p(u(\mathsf{play}\, x\, \sigma')) > \pi_p(u(\mathsf{play}\, x\, \sigma))$ then $(x, \sigma, u) \in \neg\Box\mathsf{Unimprov}$.*

Proof. By induction on the proof that $\sigma \equiv_p^{\text{fin}} \sigma'$, we can find a strategy σ'' that differs from σ (and agrees with σ') a minimum number of times, whilst still being a profitable deviation. In addition, there is a deepest node where σ'' differs from σ; let σ''' be the strategy that agrees with σ everywhere but at this node, where it agrees with σ'. By trichotomy, we either have $\pi_p(u(\mathsf{play}\, x\, \sigma''')) > \pi_p(u(\mathsf{play}\, x\, \sigma))$ or $\pi_p(u(\mathsf{play}\, x\, \sigma''')) \leq \pi_p(u(\mathsf{play}\, x\, \sigma))$. If the former, then this contradicts $\mathsf{Unimprov}$ at this node as required, so we only need to show that the latter is impossible. This is so, because the latter case violates the assumption that σ'' is minimal. $\square$

Armed with this lemma, we can now tackle the difficult direction of the one-shot deviation principle, assuming the payoff function is continuous. For simplicity, we only state the theorem for $R = \mathbb{Q}$ with the standard metric $d_{\mathbb{Q}}(x, y) = |x - y|$. Note that the order on $\mathbb{Q}$ certainly is trichotomous.

Theorem 4.12. *Assume the set of players P has decidable equality. Let (x, σ, u) be a state of the game coalgebra $(Z_{(X,\gamma)}, \Gamma)$ with rewards $R = \mathbb{Q}$. If u is continuous at infinity and $(x, \sigma, u) \in \Box\mathsf{Unimprov}$, then $(x, \sigma, u) \in \Box\mathsf{Nash}$.*

Proof. Since $\Box P \subseteq \Box\Box P$ and $\Box$ is monotone by Lemma 3.15, it is sufficient to show that $(x, \sigma, u) \in \Box\mathsf{Unimprov}$ implies $(x, \sigma, u) \in \mathsf{Nash}$. Hence, assume a player p and a strategy σ' with $\sigma \equiv_p \sigma'$ are given; by trichotomy of $<$, it is enough the show that

$$u_p(\mathsf{play}\ x\ \sigma) < u_p(\mathsf{play}\ x\ \sigma')$$

where we write $u_p(m) = \pi_p(u(m))$, is impossible. By continuity of u with $\varepsilon = u_p(\mathsf{play}\ x\ \sigma') - u_p(\mathsf{play}\ x\ \sigma)$, we find n such that if $m \sim_n \mathsf{play}\ x\ \sigma'$ then $|u_p(\mathsf{play}\ x\ \sigma') - u_p(m)| < \varepsilon$. Consider the history $m = \mathsf{play}\ x\ \sigma''$, where $\sigma'' = \lceil\sigma'\rceil_n^\sigma$ is σ' up to depth n, and σ thereafter. By construction, $m \sim_n \mathsf{play}\ x\ \sigma'$. We claim that σ'' is still a strictly profitable deviation, i.e. $u_p(\mathsf{play}\ x\ \sigma) < u_p(\mathsf{play}\ x\ \sigma'')$, by trichotomy of $<$ again: if $u_p(\mathsf{play}\ x\ \sigma) \geq u_p(\mathsf{play}\ x\ \sigma'')$ then both $|u_p(\mathsf{play}\ x\ \sigma') - u_p(\mathsf{play}\ x\ \sigma'')| \geq \varepsilon$ and $|u_p(\mathsf{play}\ x\ \sigma') - u_p(\mathsf{play}\ x\ \sigma'')| < \varepsilon$, which is absurd. Hence we have found a profitable finite deviation, which by Lemmas 4.10 and 4.11 contradicts the assumption that $(x, \sigma, u) \in \Box\mathsf{Unimprov}$. Therefore, we must instead have $u_p(\mathsf{play}\ x\ \sigma) \geq u_p(\mathsf{play}\ x\ \sigma')$ as required. □

5. Conclusions and future work

In this paper, we have built on Lescanne's and Abramksy and Winschel's coalgebraic treatment of infinite extensive form games to also consider plays that go on forever, rather than just convergent strategies that eventually lead to a terminal node. We are thus able to treat games such as infinitely repeated games, in addition to the games previously considered. We also connected the coalgebraic and traditional notions of equilibria by proving them equal under the assumption that the payoff function is continuous—a well-known result in the game theory literature, here extended to more general coalgebras. In future work, we plan to exploit techniques from coalgebra and automata theory to use our framework to *solve* various infinite-horizon games, i.e., to compute the set of equilibria for some given notion of equilibrium — Nash, sub game perfect equilibria, etc. We also hope to extend our recent translation of finite extensive form games into the framework of open games [9] to infinite games; since open games are well-suited for software implementation, this might point to another approach for computing the equilibria of infinite extensive form games in practice.

Acknowledgements We would like to thank Pierre Lescanne and Philipp Zahn for discussions about coinduction and the one-step deviation principle, respectively. Thanks also to the anonymous reviewer, whose comments and suggestions helped improve the paper.

References

[1] S. Abramsky and V. Winschel, Coalgebraic analysis of subgame-perfect equilibria in infinite games without discounting, *Math. Struct. Comput. Sci.* **27**(5), 751–761 (2017). doi:10.1017/S0960129515000365.

[2] J. Adamek and J. Rosicky, *Locally Presentable and Accessible Categories*. London Mathematical Society Lecture Note Series, Cambridge University Press (1994). doi:10.1017/CBO9780511600579.

[3] M. Barr, Terminal coalgebras in well-founded set theory, *Theoretical Computer Science*. **114**(2), 299–315 (1993). ISSN 0304-3975. doi:10.1016/0304-3975(93)90076-6.

[4] E. Bishop and D. Bridges, *Constructive Analysis*. Springer (1985).

[5] D. S. Bridges, Preference and utility: A constructive development, *Journal of Mathematical Economics*. **9**(1), 165–185 (1982). doi:10.1016/0304-4068(82)90025-8.

[6] D. S. Bridges, The constructive theory of preference relations on a locally compact space, *Indagationes Mathematicae*. **92**(2), 141–165 (1989). doi:10.1016/S1385-7258(89)80023-X.

[7] D. S. Bridges, The constructive theory of preference relations on a locally compact space II, *Mathematical Social Sciences*. **27**(1), 1–9 (1994). ISSN 0165-4896. doi:10.1016/0165-4896(94)00718-7.

[8] D. S. Bridges, First steps in constructive game theory, *Mathematical Logic Quarterly*. **50**(4-5), 501–506 (2004).

[9] M. Capucci, N. Ghani, J. Ledent, and F. Nordvall Forsberg. Translating extensive form games to open games with agency. In *Applied Category Theory '21* (2021).

[10] D. Fudenberg and J. Tirole, *Game Theory*. MIT Press, Cambridge, MA (1991).

[11] N. Ghani, C. Kupke, A. Lambert, and F. N. Forsberg, A compositional treatment of iterated open games, *Theor. Comput. Sci.* **741**, 48–57 (2018). doi:10.1016/j.tcs.2018.05.026.

[12] B. Jacobs, *Introduction to Coalgebra: Towards Mathematics of States and Observation*. Cambridge Tracts in Theoretical Computer Science, Cambridge University Press (2016). doi:10.1017/CBO9781316823187.

[13] H. W. Kuhn. Extensive games and the problem of information. In eds. H. W. Kuhn and A. W. Tucker, *Contributions to the theory of games, Volume II (Annals of Mathematics Studies, 28)*, pp. 193–216. Princeton University Press (1953).

[14] A. Kurz and D. Pattinson, Definability, canonical models, compactness for finitary coalgebraic modal logic, *Electronic Notes in Theoretical Computer Science*. **65**(1), 135–155 (2002). ISSN 1571-0661. doi:10.1016/S1571-0661(04)80363-3.

[15] P. Lescanne. A simple case of rationality of escalation. In eds. R. Heckel and S. Milius, *CALCO '13*, vol. 8089, *LNCS*, pp. 191–204, Springer (2013). doi:10.1007/978-3-642-40206-7_15.

[16] P. Lescanne, Dependent types for extensive games, *J. Formaliz. Reason.* **11**

(1), 1–17 (2018). doi:10.6092/issn.1972-5787/7517.

[17] P. Lescanne and M. Perrinel, "Backward" coinduction, Nash equilibrium and the rationality of escalation, *Acta Informatica*. **49**(3), 117–137 (2012). doi:10.1007/s00236-012-0153-3.

[18] K. Leyton-Brown and Y. Shoham, *Essentials of Game Theory: A Concise Multidisciplinary Introduction*. Morgan & Claypool (2008).

[19] J. Mycielski and W. Taylor, A compactification of the algebra of terms, *Algebra Universalis*. **6**, 159–163 (1976).

[20] M. Osbourne and A. Rubinstein, *A course in game theory*. MIT Press (1994).

[21] J. Rutten, Universal coalgebra: a theory of systems, *Theoretical Computer Science*. **249**(1), 3–80 (2000). doi:10.1016/S0304-3975(00)00056-6.

[22] R. Selten, Reexamination of the perfectness concept for equilibrium points in extensive games, *International Journal of Game Theory*. **4**, 25–55 (1975). ISSN 0020-7276. doi:10.1007/BF01766400.

[23] R. Selten, The chain store paradox, *Theory and decision*. **9**(2), 127–159 (1978). doi:10.1007/BF00131770.

[24] M. Shubik, The dollar auction game: A paradox in noncooperative behavior and escalation, *Journal of conflict Resolution*. **15**(1), 109–111 (1971).

[25] J. Worrell. *On coalgebras and final semantics*. PhD thesis, University of Oxford (2000).