A new game equivalence, its logic and algebra

Johan van Benthem, Nick Bezhanishvili, Sebastian Enqvist

Abstract

We present a new notion of game equivalence that captures basic powers of interacting players. We provide a representation theorem, a complete logic, and a new game algebra for basic powers. In doing so, we establish connections with imperfect information games and epistemic logic. We also identify new open problems concerning logic and games.

1 Introduction

Games are a basic model for interactive agency, but how much structure do we need? Game theory offers strategic form games and extensive games, representing different levels of detail. Logicians have studied further levels, such as powers of players over the outcomes of a game. Each of these levels comes with a notion of invariance between structures that matches a logical language – and as in many fields, the search for new invariances is ongoing. In this paper we offer a new notion bridging between game theory and logic: basic power equivalence, that uses powers encoding a sort of qualitative equilibria. We determine its properties in a new representation theorem, find a complete associated logic, and explore a new game algebra for basic powers that eventually forces us to change from functional to relational strategies. Moreover, we establish interesting connections with imperfect information games and epistemic logic.

This paper fits with a body of earlier work. Our approach is partly inspired by the computational literature on process equivalences, ranging from coarser trace equivalence to more fine-grained notions of bisimulation [**?**]. Even more central to us is the notion of power equivalence, implicit in the game algebra of Parikh [**?**], which also links with the set-theoretic forms for games in [**?**]. A precursor inside game theory is the celebrated transformation analysis of equivalent games with imperfect information by Thompson [**?**], refined in [**?**]. Game theory also has comparative discussions of the information available in extensive forms and in strategic normal forms [**?**], a style of invariance analysis that remains to be connected to our logic-based approach.

Finally, our analysis has clear limitations in what it takes on board. More delicate intuitions of game equivalence emerge once we consider players' preferences, or their types, or when we focus on correlations between available equilibria in the games being compared. These richer settings are beyond the scope of this paper, but they pose a natural challenge to logic-based approaches.

2 Equivalence of games, old and new answers

The question when two games are considered equivalent is fundamental to any game theory. And as in many areas of mathematics, there is no unique answer: different natural candidates exist. Here is an example from van Benthem [**?**].

Example 2.1. Are the following two games the same?

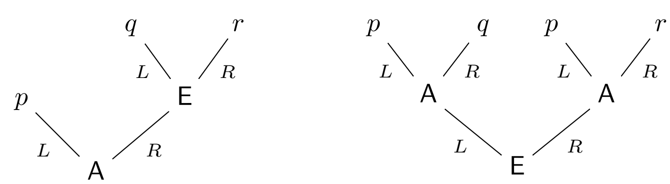

Figure 1: Equivalent games?

There are two players A, E, whose turns are indicated at nodes in the extensive game trees, and the proposition letters p, q, r mark the outcomes at the end of the game. Here are two different natural ways of answering the stated question.

2.1 Bisimulation

The two games are clearly not equivalent when we zoom in on turns and player's available choices. For instance, in $\mathcal{G}_1$ there is a choice point where E can decide whether the game ends in q or in r, but no such choice point occurs in $\mathcal{G}_2$. At this level, a good notion of game equivalence is modal bisimulation. The two games $\mathcal{G}_1$, $\mathcal{G}_2$ are not bisimilar, and this difference shows in the matching propositional modal logic that can define many properties of our games in detail. So, at this level of structure, standard modal logic is a good language for games.

2.2 Power equivalence

Next, we abstract away from specific moves and choices, and focus on players' powers for controlling outcomes. Then the two games may be considered the same. In $\mathcal{G}_1$, A has a strategy "left" which forces the game to end with outcome p, something we can write as the power $\{p\}$. A also has a strategy "right" which forces the set of outcomes $\{q, r\}$. A can do no better than this: the eventual choice between q and r is up to player E. As for the other player, E's strategy "left" forces $\{p, q\}$ (since A can always go left at the start), and E's strategy "right" forces $\{p, r\}$. Now, if we compute powers of the players, we will find the same powers in $\mathcal{G}_2$. This is obvious for player E, but since player A has four strategies now, we need a bit of work. The power $\{p\}$ comes from A's strategy "left, left", the power $\{q, r\}$ from "right, right". The remaining strategies "left right" and "right left" give powers $\{p, r\}$ and $\{q, p\}$, but these can be considered derived, being weakenings of the power $\{p\}$ that A already has. Here we use monotonicity: powers of a player are closed under taking supersets. Formal definitions of games and powers of players can be found in Appendix A.

The power level, too, comes with a perspicuous language for describing game structure, viz. a modal logic over neighborhood models whose relations $N^i sX$ record that player i has a power X in the relevant game. Naturally, this modal language describes less detail of the extensive game tree. But the power perspective also has another logic connection. The above two game trees mirror a syntax for logical formulas where A controls conjunctions and E disjunctions. And then the equivalence of the above games qua powers reflects the distributive law of propositional logic. Indeed, most game semantics for logics operate at the level of powers, often associated with 'winning strategies'.

A further theory for both approaches, including some themes that we will list below, can be found in van Benthem [**?**]. However, we are satisfied with neither answer, and believe that there is yet a third analysis of the question whether the above two games are equivalent, intermediate between the preceding two, and we believe it may be the better answer from a game-theoretic perspective.

2.3 Basic powers

Monotonicity is appealing from the perspective of one individual player, but it ignores what other players can achieve. For instance, in the game $\mathcal{G}_2$, the power $\{p, q\}$ of player A corresponds to something that can be realized in the following sense: A can decide to narrow down the possible outcomes to $\{p, q\}$, but leave the choice of which of these outcomes to player E. But this is not a power A has in game $\mathcal{G}_1$: to allow for the outcome p, he needs to go "left", but this precludes the possibility that E ever gets to choose q. A *basic power* of player P in an extensive game is a set of outcomes X associated with some strategy σ for i as follows: X is the set of all endpoints that can be reached when P plays according to σ, while the choices of the other player $\overline{\mathsf{P}}$ are free. When the game is infinite, the outcomes meant here are identified with complete histories. A more formally precise definition follows. Here and in what follows we denote the set of games for players A with outcomes in O as $\mathbb{G}(A, O)$. For two-player games we call the players by E (Eloise) and A (Abelard). We set $\overline{\mathsf{E}} = \mathsf{A}$ and $\overline{\mathsf{A}} = \mathsf{E}$.

Definition 2.2. Let $\mathcal{G}$ be any game in $\mathbb{G}(A, O)$, let $a \in A$. A power $P \subseteq O$ is said to be a *basic power* for a in $\mathcal{G}$ if there is a strategy σ for a in $\mathcal{G}$ such that $P = \{o(m) \mid m \in \mathsf{Match}(\sigma)\}$. The set of all basic powers of a in $\mathcal{G}$ is denoted by $B_a(\mathcal{G})$.

Example 2.3. Example 2.1 revisited with basic powers.

The basic powers of player A in game $\mathcal{G}_1$ are $\{p\}, \{q, r\}$, while those of E are $\{p, q\}, \{p, r\}$. In game $\mathcal{G}_2$, the basic powers of E are the same as in game $\mathcal{G}_1$, but those of A are $\{p\}, \{p, q\}, \{p, r\}, \{q, r\}$.

In terms of basic powers then, *the two games in Figure 1 are not the same.*

This motivates a new notion of game equivalence:

Definition 2.4. Two games are said to be *basic power equivalent* if each player has the same basic powers in both.

In this paper, we study this stricter game equivalence with logical techniques. We modify the theory of power equivalence to obtain new representation theorems, complete game logics, and algebras of game constructions for basic powers. In the process, we find new themes as well, in particular, a pervasive role for imperfect information and the need for an extended notion of relational strategy.

Basic powers have a game-theoretic slant, moving from the egocentric winning strategies in logical games to 'qualitative equilibria' that depend on what all players do. Even so, few game-theoretic notions are needed for our study. We assume familiarity with the notions of extensive game tree, strategy, outcome, games in strategic form, perfect information, and imperfect information. However, we will point out further connections with game theory as we go.

3 Logical perspectives on game equivalence

Before starting our study of basic powers, we recall some notions and results about standard powers in extensive games, [?]. First, here are three properties of powers of players A, E, viewed as subsets of a total set of outcomes $\mathcal{O}$:

- *Non-emptiness:* Each player has at least one power.
- *Consistency*: If X is a power of player A, and Y a power of player E, then $X \cap Y \neq \emptyset$.
- *Monotonicity*: If X is a power of a player, then so is any superset $Y \supseteq X$.
- *Determinacy*: If X is not a power of a player, then $\mathcal{O} - X$ is a power of the other player.

A simple representation result shows that these conditions are characteristic:

Theorem 3.1. *The following are equivalent for families $\mathcal{X}$, $\mathcal{Y}$ of subsets of $\mathcal{O}$:*

1. *Consistency, Monotonicity, and Determinacy hold for $\mathcal{X}$, $\mathcal{Y}$,*
2. *$\mathcal{X}$, $\mathcal{Y}$ are the powers of two players in some game of perfect information.*

There is also a characterization for the more general case of imperfect information games that carry equivalence relations $\sim$ for both players, giving their uncertainty about where they are in the game tree. Equivalence classes of these relations are often called 'information sets'. It is assumed that $\sim$-connected points have the same available moves, and importantly, the only allowed strategies are those that assign the same move throughout information sets for the player. Apart from this, the definition of powers is the same as before.

Theorem 3.2. *The following are equivalent for families $\mathcal{X}$, $\mathcal{Y}$ of subsets of $\mathcal{O}$:*

1. *Consistency and Monotonicity hold for $\mathcal{X}$, $\mathcal{Y}$,*
2. *$\mathcal{X}$, $\mathcal{Y}$ are the powers of two players in some game of imperfect information.*

Now we generalize sameness of powers to allow for comparison across different games. Indeed, going further, we move from game trees altogether to more abstract state spaces connected with games, with abstract relations on states that indicate powers. These are the proper models for our logic to come.

Definition 3.3. Fix a set Gm of labels for games. A *game state model* $\mathfrak{M} = (S, N, V)$ has a set of states S, a valuation V for proposition letters, and abstract neighborhood relations $N^g_{\mathsf{P}} \subseteq S \times \mathcal{P}S$ for each game $g \in \mathsf{Gm}$ and each player $\mathsf{P} \in \{\mathsf{A}, \mathsf{E}\}$. A game state model is a *power model* if for all games g and $s \in S$, the pair of families $N^g_{\mathsf{A}}[s], N^g_{\mathsf{E}}[s]$ satisfies Consistency and Monotonicity.

The relation $N^g_{\mathsf{P}} sX$ says that, with game g started in state s, player P has a strategy guaranteeing that g will end in some state in X as its outcome. One can view states in S as positions associated with nodes in a game tree, not necessarily one-to-one, as in graph or board games. Much information can be captured at such a level, including invariants that determine players' powers.

Definition 3.4. Let $\mathfrak{M} = (W, N, V)$, $\mathfrak{M}' = (W', N', V')$ be game state models. A binary relation $B \subseteq W \times W'$ is a *power bisimulation* if the following conditions hold for all sBs' and each game g and player P:

Harmony s, s' satisfy the same atomic proposition letters.

Forth For all X with $N^g_{\mathsf{P}} sX$, there exists an X' such that $N'^g_{\mathsf{P}} s'X'$ and also:

Forth-Back For all $v' \in X'$ there is some $v \in X$ such that vBv'.

Back For all X' with $N'^g_{\mathsf{P}} s'X'$, there is some X such that $N^g_{\mathsf{P}} sX$ and also:

Back-Forth For all $v \in X$ there is some $v' \in X'$ such that vBv'.

States in two pointed game models are *power bisimilar*, written $\mathfrak{M}, w \leftrightarrows \mathfrak{N}, v$, if there is a power bisimulation B between $\mathfrak{M}$ and $\mathfrak{N}$ such that wBv.

There is a modal language matching this level of game description. The formulas of the basic modal logic $\mathtt{PL}$ of powers are defined as follows:

$$\varphi := p \in \mathsf{Prop} \mid \varphi \wedge \varphi \mid \neg\varphi \mid \langle g, \mathsf{E}\rangle\varphi \mid \langle g, \mathsf{A}\rangle\varphi$$

Semantics of formulas in game state models are defined by the usual recursion, where the clause for the modal operators reads: $\mathfrak{M}, s \Vdash \langle g, \mathsf{P}\rangle\varphi$ iff there is some $X \subseteq S$ with $N^g_{\mathsf{P}} sX$ and $\mathfrak{M}, v \Vdash \varphi$ for all $v \in X$. Informally, this says that "player P can force the condition φ to hold for the outcomes of game g".

Fact 3.5. *Formulas of the language of* $\mathtt{PL}$ *are invariant for power bisimulation.*

Con: $\langle g, \mathsf{P}\rangle p \to \neg\langle g, \overline{\mathsf{P}}\rangle \neg p$	NonEm: $\langle g, \mathsf{P}\rangle \top$
Mon: $\dfrac{\varphi \to \psi}{\langle g, \mathsf{P}\rangle\varphi \to \langle g, \mathsf{P}\rangle\psi}$	

Figure 2: Axioms and rules for the logic PL

Fact 3.6. *Together with a complete system of axioms for classical propositional logic and the usual rules of modus ponens and uniform substitution, the axioms and rules in Figure 2 are sound and complete for validity over game state models. Furthermore,* PL *is decidable and has the finite model property.*

Game state models for perfect information games also satisfy Determinacy, and they are complete by strengthening the axiom Con to an equivalence:

$$\langle g, \mathsf{P}\rangle p \leftrightarrow \neg\langle g, \overline{\mathsf{P}}\rangle \neg p.$$

One can extend PL with natural game constructions of choices for the two players: $\mathcal{G}_1 \cup \mathcal{G}_2$, $\mathcal{G}_1 \cap \mathcal{G}_2$, game dual $\mathcal{G}^d$ switching roles, and sequential composition $\mathcal{G}_1 \circ \mathcal{G}_2$. These validate a game algebra whose basic part for choice and dual is a 'De Morgan algebra', while $\circ$ is an associative binary operation that is monotonic in both arguments, and distributive over its left, but not its right argument. This algebra can be added in the format of a dynamic logic [**?**], and the semantic interpretation of this logic on game state models relies on the fact that each of the operations is safe for standard neighborhood bisimulations.

Fact 3.7. *The dynamic logic of the operations $\cup, \cap, \circ$ over game state models is decidable and completely axiomatizable.*

In what follows, we want to preserve the methodology behind these results, and it turns out that we can. But there will be surprises in making things work.

4 Representation theorem for basic powers

4.1 Imperfect information games

We start with basic powers in extensive games of imperfect information. These are defined just as before, but keeping in mind that we only consider uniform strategies for each player. Consider the following three conditions:

- *Non-Emptiness*: Each player has at least one basic power.
- *Consistency*: If X is a basic power for player A, and Y is a basic power for player E, then $X \cap Y \neq \emptyset$.

– *Exhaustiveness*: For each basic power X of player P and each $x \in X$, there is a basic power Y for player $\overline{\mathsf{P}}$ such that $x \in Y$.

The proof of the following fact can be found in Appendix B:

Fact 4.1. *Basic powers satisfy Non-Emptiness, Consistency, Exhaustiveness.*

Our first main result is a representation theorem for basic powers.

Theorem 4.2. *The following are equivalent for families $\mathcal{X}$, $\mathcal{Y}$ of subsets of $\mathcal{O}$:*

1. *Non-Emptiness, Consistency and Exhaustiveness hold for $\mathcal{X}$, $\mathcal{Y}$,*
2. *$\mathcal{X}$, $\mathcal{Y}$ are the sets of basic powers of the two players in some game of imperfect information whose possible outcomes are the set $\mathcal{O}$.*

Proof. One direction is the content of Fact 4.1. For the converse direction from (1) to (2), we define a game $\mathcal{G}$ as follows. $\mathcal{G}$ has two steps: first A makes a move, then E makes a move. Moreover, A does not know which move player E has played: the game $\mathcal{G}$ has imperfect information.

Formally, we define $\mathcal{G} = (\mathcal{T}, t, o, \Pi)$ as follows: we fix an enumeration $t_0, ..., t_{n-1}$ of all triples (X, i, x), where $X \in \mathcal{X}$, $x \in X$ and $i \in \{0, 1\}$. Next we fix an enumeration $s_0, ..., s_{m-1}$ of all triples (Y, j, f) where $Y \in \mathcal{Y}$, $j \in \{0, 1\}$ and f is a function from $n = \{0, ..., n-1\}$ to $\mathcal{O}$ subject to the following constraints for each $k < n$ where $t_k = (X, i, x)$:

1. $f(t_k) \in X \cap Y$.
2. In addition, if $i = j$ and $x \in X \cap Y$, then $f(t_k) = x$.

Now define $\mathcal{T}$ by closing the set of words $\{kl \mid k < n \;\&\; l < m\}$ of length 2 under prefixes. For Π we pick some partition that contains the set of positions $\{0, ..., n-1\}$ (each of which is a word of length 1) as a partition cell. We set $t(\varepsilon) = \mathsf{A}$ and we set $t(k) = \mathsf{E}$ for $k < n$. For a branch kl with $k < n, l < m$ and $s_l = (Y, j, f)$ we set $o(kl) = f(t_k)$.

Note that in this two-step game, the strategies for the first player A are just the moves for A. Because of the imperfect information for the second player E, strategies for E can also be identified with moves, i.e. indices corresponding to triples (Y, j, f), since the same choice must be made at each turn.

By Consistency, for any $X \in \mathcal{X}$ and $Y \in \mathcal{Y}$, the intersection $X \cap Y$ is non-empty, which means that the set of choice functions f satisfying conditions 1 and 2 above with respect to any given X, Y, i, j, x is non-empty too. Hence the Non-Emptiness constraint guarantees that both players have non-empty sets of moves at their turns, so we have a well-defined game $\mathcal{G}$ of imperfect information.

The theorem is proved if we can establish the following:

Claim 4.3. *The basic powers of the players in $\mathcal{G}$ are $\mathcal{X}$ for A and $\mathcal{Y}$ for E.*

Proof. First we consider the case of player A.

Subclaim 1. The basic power corresponding to the move k with $t_k = (X, i, x)$ is the set X.

Proof. Note that the set of all outcomes of branches beginning with k is equal to the set of all $f(X, i, x) \in \mathcal{O}$ such that for some $Y \in \mathcal{Y}$, some $j \in \{0, 1\}$ and some $l < m$, $s_l = (Y, j, f)$. (a) Inclusion $\supseteq$: By the definition of the set of triples $s_l = (Y, j, f)$ for $l < m$, any outcome of the form $f(X, i, x)$ for some (Y, j, f) among $p_0, ..., p_{m-1}$ belongs to $X \cap Y \subseteq X$. (b) Inclusion $\supseteq$: Consider any $x' \in X$. By Exhaustiveness, there is some $Y \in \mathcal{Y}$ with $x' \in X \cap Y$. Let f be a choice function defined as follows: given $(Z, r, z) \in \{t_1, ..., t_n\}$ set $f(Z, r, z) = x'$ if $r = i$ and $x' \in Y \cap Z$, set $f(Z, r, z) = z$ if $r = |i - 1|$ and $z \in Y \cap Z$, and set $f(Z, r, z)$ to be some arbitrary element of $Y \cap Z$ otherwise (which is guaranteed to exist by Consistency). Then the triple $(Y, |i - 1|, f)$ is a legitimate choice for E, i.e. it is equal to s_l for some $l < m$. Furthermore, we have $o(k, l) = f(X, i, x) = x'$. □

Thus, the basic powers of A in G equal the family $\mathcal{X}$. Next, consider player E.

Subclaim 2. For any move $l < m$ with $s_l = (Y, j, f)$, its corresponding basic power is $Y \in \mathcal{Y}$.

Proof. (a) Inclusion $\subseteq$: By definition, for all $k < n$ with $t_k = (X, i, x)$ we must have $o(kl) = f(X, i, x) \in X \cap Y \subseteq Y$. (b) Inclusion $\supseteq$: We show that all objects $y \in Y$ will in fact be chosen somewhere by f. Consider any $y \in Y$. By Instantiatedness, there is an $X \in \mathcal{X}$ with $y \in X \cap Y$. For that X, (X, j, y) corresponds to a possible move for A i.e. it is equal to t_k for some $k < m$. By the second constraint on moves for E, we must have $o(kl) = f(X, j, y) = y$.

Since for any Y we can pick $i \in \{0, 1\}$, say $i = 0$, and find an appropriate choice function f such that $(Y, 0, f)$ is a legitimate move for E (using Consistency), it follows that the basic powers of E in $\mathcal{G}$ are exactly the family $\mathcal{Y}$. □

Claim 4.3 now follows from the two Subclaims. □

This concludes the proof of the representation theorem for basic powers. □

4.2 Two game-theoretic angles

Our approach relates to game theory in various ways. Here are two connections.

Strategic forms Our representation theorem also characterizes games in strategic form. To see the connection, we note an analogy between basic powers in game trees and in the rows and columns of two-player matrix games.

Example 4.4. Strategic form games and basic powers.

Revisiting Example 2.1, the strategic forms of the two games are as follows, with rows encoding strategies for A and columns strategies of E:

p	p
q	r

p	p
q	p
p	r
q	r

Now one can read off the basic powers of the two players directly from the rows and columns, and it is clear that the two games are not basic power equivalent.

Given a two-player game in strategic form, let the *yield* of a row or column be the unordered set of outcomes appearing in it.

Theorem 4.5. *The conditions Non-Emptiness, Consistency and Exhaustiveness are necessary and sufficient for representation of yields of rows and columns in a strategic form two-player game.*

Proof. This follows directly from the proof of Theorem 4.2. The basic powers of players in our two-step extensive game are identical with the yields of rows and columns in the corresponding strategic normal form game. □

Even so, basic power equivalence for games is not the same as having the same strategic form (up to permutation and duplication of rows and columns).

Example 4.6. Strategic form equivalence versus basic power equivalence.

Both players have the same basic powers in the following games:

0	1	0
1	0	0
0	0	0

1	1	0
0	0	0

However, there is clearly no way to transform one game into the other simply by duplicating and switching the order of rows and columns. Moreover, viewed as models for a cylindric modal logic $\mathbf{S5}^2$ with the usual "horizontal" and "vertical" accessibility relations, [**?**], and $1, 0$ regarded as truth values of some fixed propositional variable, the two games are not related by any modal bisimulation.

Perfect information games An important special case are games with perfect information, where all information sets are singletons. In this case, additional properties hold, such as this Determinacy condition for standard powers:

> *Determinacy′* For each set $X \subseteq \mathcal{O}$, either one player has a basic power contained in X, or the other player has a basic power contained in $\mathcal{O} - X$.

We have no generalization of our result for powers with perfect information.

Conjecture 4.7. *Consistency, Exhaustiveness, and Determinacy′ are necessary and sufficient for representing given powers in a game of perfect information.*

This result may be harder than the earlier one. For families of sets satisfying Consistency, Monotonicity and Determinacy, one can find a two-step game of perfect information inducing exactly these sets as the standard monotonic powers for the players. But with basic powers, more rounds may be essential.

Example 4.8. Basic powers in perfect information games.
Consider the following extensive game:

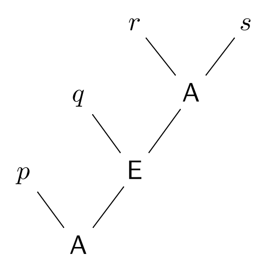

Figure 3: A three-step game

Basic powers for A are $\{p\}, \{q, r\}, \{q, s\}$, and for E: $\{p, q\}$, $\{p, r, s\}$. It is easy to see that no two-step perfect information game produces just these basic powers.

5 Modal logic for basic powers

5.1 Instantial neighborhood logic for basic powers

To find a logical perspective on basic powers, note that basic power equivalence can be generalized to game state models as before. Again, the invariance issue then shifts to: When are two game state models the same at the level of *basic* powers? The answer is in the following definition.

Definition 5.1. Let $\mathfrak{M} = (W, N, V)$, $\mathfrak{M}' = (W', N', V')$ be game state models. The relation $B \subseteq W \times W'$ is a *basic power bisimulation* if the following conditions hold for all sBs' and each game g and player P:

Harmony s, s' satisfy the same atomic proposition letters.

Forth For all X with $N^g_{\mathsf{P}} sX$, there is some X' such that $N'^g_{\mathsf{P}} s'X'$ and also:

Forth-Forth For all $v \in X$, there is a $v' \in X'$ such that vBv'.

Forth-Back For all $v' \in X'$, there is a $v \in X$ such that vBv'.

Back For all X' with $N'^g_{\mathsf{P}} s'X'$, there is some X such that $N^g_{\mathsf{P}} sX$ and also:

Back-Forth For all $v \in X$, there is a $v' \in X'$ such that vBv'.

Back-Back For all $v' \in X'$, there is a $v \in X$ such that vBv'.

We say that two states in two given pointed game state models are *basic power bisimilar*, written $\mathfrak{M}, w \leftrightarrows \mathfrak{N}, v$, if there is a basic power bisimulation B between $\mathfrak{M}$ and $\mathfrak{N}$ such that wBv.

The language that matches this notion of invariance is a richer version of the standard modal language for neighborhood models, which was introduced under the name instantial neighborhood logic in [**?**]. We denote it here by `BPL`, for "basic power logic". Its formulas are defined by the following grammar:

$$\varphi := p \in \mathsf{Prop} \mid \varphi \wedge \varphi \mid \neg\varphi \mid \langle g, \mathsf{P}\rangle(\Psi, \varphi)$$

where Ψ is a finite set of `BPL`-formulas, called the *instantial* formulas. As a notational convention, if $\Psi = \{\psi_1, ..., \psi_n\}$, then we may write $\langle g, \mathsf{P}\rangle(\psi_1, ..., \psi_n, \varphi)$ rather than $\langle g, \mathsf{P}\rangle(\Psi, \varphi)$. However, $\psi_1, ..., \psi_n$ should then still be understood as a set rather than a tuple of formulas.

That this is in fact the right modal language for basic powers was observed, although in different terminology, in [**?**], where the following results were proved:

Fact 5.2. *All formulas of* `BPL` *are invariant for basic power bisimulations.*

Fact 5.3 (Hennessy-Milner property)**.** *Two finite, pointed game state models are basic power bisimilar if, and only if, they satisfy the same formulas in* `BPL`*.*

6 Axiomatization

In this section we axiomatize the valid formulas of `BPL`, thus pinning down the modal logic of basic powers. Our system is a gentle modification of instantial neighborhood logic, and the key step of the completeness proof is to verify that the conditions characterizing basic powers can be captured by suitable axioms.

As axioms and rules for `BPL`, we take all propositional tautologies, modus ponens, uniform substitution plus the following principles:

Theorem 6.1. *The displayed axioms and rules for* `BPL` *are sound and complete.*

Soundness is an easy check. The completeness proof is a normal form argument following [**?**], but the adaptation to basic powers is not trivial, as we must deal with the new frame constraints of Non-emptiness, Consistency and Exhaustiveness. The main contribution here is to prove that the model construction satisfies these constraints. We outline the key parts of the proof below.

Definition 6.2. The *modal depth* of a formula is defined inductively by:

$$d(p) = 0, \quad d(\neg\varphi) = d(\varphi), \quad d(\varphi \wedge \psi) = \mathtt{max}(d(\varphi), d(\psi)),$$

$$d(\langle g, \mathsf{P}\rangle(\Gamma; \varphi)) = \mathtt{max}(d[\Gamma \cup \{\varphi\}]) + 1.$$

$$\text{Weak:}\quad \langle g, \mathsf{P}\rangle(\Phi; p) \to \langle g, \mathsf{P}\rangle(\Phi'; p) \qquad (\Phi' \subseteq \Phi)$$

$$\text{Un:}\quad \langle g, \mathsf{P}\rangle(q_1, ..., q_n; p) \to \langle g, \mathsf{P}\rangle(q_1 \wedge p, ..., q_n \wedge p; p)$$

$$\text{Lem:}\quad \langle g, \mathsf{P}\rangle(\Phi; p) \to \langle g, \mathsf{P}\rangle(\Phi \cup \{q\}; p) \vee \langle g, \mathsf{P}\rangle(\Phi; \neg q \wedge p)$$

$$\text{Bot:}\quad \neg\langle g, \mathsf{P}\rangle(\bot; p)$$

$$\text{Mon:}\ \frac{p \to p' \qquad q_i \to q_i{}' \quad (1 \leq i \leq n)}{\langle g, \mathsf{P}\rangle(q_1, ..., q_n; p) \to \langle g, \mathsf{P}\rangle(q_1', ..., q_n'; p')}$$

Figure 4: Axioms and rules for instantial neighborhood logic

$$\text{Non-Em:}\quad \langle g, \mathsf{P}\rangle\top \qquad\qquad \text{Cons:}\quad \langle g, \mathsf{P}\rangle p \to \neg\langle g, \overline{\mathsf{P}}\rangle\neg p$$

$$\text{Inst:}\quad \langle g, \mathsf{P}\rangle(p; \top) \leftrightarrow \langle g, \overline{\mathsf{P}}\rangle(p; \top)$$

Figure 5: The additional axioms for `BPL`

Definition 6.3. Given a finite set of propositional variables Q, a `BPL` formula φ is a Q-*formula* if all propositional variables appearing in φ belong to Q.

Given $k \in \omega$ and a finite set Q of propositional variables, a (Q, k)-*description* is a consistent Q-formula φ of modal depth $\leq k$, such that for any Q-formula θ of depth $\leq k$, we have $\varphi \vdash \theta$ or $\varphi \vdash \neg\theta$.

By a standard argument, there are at most finitely many Q-formulas of depth $\leq k$ up to logical equivalence, given that Q is finite.

The key lemma for the completeness proof is the following result from [**?**].

Lemma 6.4. *Let* $\langle g, \mathsf{P}\rangle(\Gamma; \varphi)$ *be a formula with* $\max(d[\Gamma \cup \{\varphi\}]) \leq k$ *and* Q *a finite set of propositional variables containing all variables appearing in this formula. Then* $\langle g, \mathsf{P}\rangle(\Gamma; \varphi)$ *is provably equivalent to a disjunction of the form:*

$$\bigvee_{i \in I} \langle g, \mathsf{P}\rangle(\Theta_i; \bigvee \Theta_i)$$

with I *finite, and for each* $i \in I$, Θ_i *a finite set of* (Q, k)*-descriptions such that:*

- *every member of* Θ_i *provably entails* φ*, and*
- *every member of* Γ *is provably entailed by some member of* Θ_i*.*

Fix a finite set of propositional variables Q. Given a Q-formula φ, $\widehat{\varphi}$ is its equivalence class under provable equivalence. For a finite set of formulas Γ, set

$$\widehat{\Gamma} = \{\widehat{\varphi} \mid \varphi \in \Gamma\}$$

We construct a neighborhood model $\mathfrak{M} = (W, N, V)$ as follows:

- $W = \{(\widehat{\varphi}, k) \mid \varphi$ is a (Q, k)-description and $k < \omega\}$
- For any player P, let N_{P} be the union of the sets

$$\{((\widehat{\varphi}, k+1), \widehat{\Gamma} \times \{k\}) \in W \times \mathcal{P}(W) \mid \varphi \vdash \langle g, \mathsf{P}\rangle(\widehat{\Gamma}; \bigvee \widehat{\Gamma})\}$$

  and

$$\{((\widehat{\varphi}, 0), W) \mid (\widehat{\varphi}, 0) \in W\}$$

- Finally, for any propositional variable p, set $V(p) = \{\widehat{\varphi} \mid \varphi \vdash p\}$ if $p \in Q$, $V(p) = \emptyset$ otherwise.

Note that this is well defined, i.e. whether $(\widehat{\varphi}, \widehat{\Gamma}) \in N_{\mathsf{P}}$ is independent of the choice of witnesses φ, Γ of the equivalence classes. The following lemma can be proved exactly as in [**?**], and we refer to that paper for the details:

Lemma 6.5 (Truth lemma)**.** *Let $\mathfrak{M}$ be constructed as above, and let ψ be any basic formula of modal depth $\leq k$ whose propositional variables all belong to Q, and which is such that all game terms appearing in ψ belong to τ. Then for every (Q, τ, k)-description φ, we have:*

$$\mathfrak{M}, (\widehat{\varphi}, k) \Vdash \psi \text{ iff } \varphi \vdash \psi$$

The addition we need for present purposes is the following lemma:

Lemma 6.6. *The structure $\mathfrak{M}$ constructed above is a power model, that is, it satisfies the Non-emptiness, Consistency and Exhaustiveness constraints.*

Proof. First, note that all the conditions hold for the image of each relation on an element of W of the form $(\widehat{\varphi}, 0)$. So we can focus on the images of relations of the form N_{P} on states of the form $(\widehat{\varphi}, k+1)$ for some k.

The Non-emptiness condition is straightforward from the axiom (Non-Em) of BPL. Next, for Instantiatedness, suppose that

$$((\widehat{\varphi}, k+1), \widehat{\Theta} \times \{k\}) \in N_{\mathsf{A}}$$

By definition, we get $\varphi \vdash \langle g, \mathsf{E}\rangle(\Theta; \bigvee \Theta)$. Pick an element $(\widehat{\theta}, k) \in \Theta \times \{k\}$. By (Weak) and (Mon) we get $\langle g, \mathsf{E}\rangle(\Theta; \bigvee \Theta) \vdash \langle g, \mathsf{E}\rangle(\theta; \top)$, so $\varphi \vdash \langle g, \mathsf{E}\rangle(\theta; \top)$. By the axiom (Inst) we get $\varphi \vdash \langle g, \mathsf{A}\rangle(\theta; \top)$ as well. Since φ is a $(Q, k+1)$-description, it follows from Lemma 6.4 that there is a set Ψ of (Q, k)-descriptions with $\varphi \vdash \langle g, \mathsf{A}\rangle(\Psi; \bigvee \Psi)$ and such that there exists some $\psi \in \Psi$ with $\psi \vdash \theta$. But since ψ, θ are both (Q, k)-descriptions, this means that $\widehat{\theta} = \widehat{\psi}$, so $(\widehat{\theta}, k) = (\widehat{\psi}, k)$.

But then we get $((\widehat{\varphi}, k+1), \widehat{\Psi} \times \{k\}) \in N_{\mathsf{E}}$ and $(\widehat{\theta}, k) \in \widehat{\Psi} \times \{k\}$ as required. The converse direction is proved in the same manner.

For Consistency, let $((\widehat{\varphi}, k+1), \widehat{\Theta} \times \{k\}) \in N_{\mathsf{A}}$ and $((\widehat{\varphi}, k+1), \widehat{\Theta'} \times \{k\}) \in N_{\mathsf{E}}$. Clearly, since Θ and Θ' are both sets of (Q,k)-descriptions, that $\widehat{\Theta} \times \{k\}$ does not intersect $\widehat{\Theta'} \times \{k\}$, and so $\bigvee \Theta' \to \neg \bigvee \Theta$. But we have $\varphi \vdash \langle g, \mathsf{E}\rangle(\Theta; \bigvee \Theta)$, hence $\varphi \vdash \langle g, \mathsf{E}\rangle \bigvee \Theta$ by the axiom schema (Weak). Furthermore we have:

$$\varphi \vdash \langle g, \mathsf{A}\rangle(\Theta'; \bigvee \Theta') \vdash \langle g, \mathsf{A}\rangle \bigvee \Theta' \vdash \langle g, \mathsf{A}\rangle \neg \bigvee \Theta$$

But then $\varphi \vdash \langle g, \mathsf{E}\rangle \bigvee \Theta \wedge \langle g, \mathsf{A}\rangle \neg \bigvee \Theta$, and it follows from the axiom schema (Cons) that the formula φ cannot be consistent, which contradicts our assumption that φ was a $(Q, k+1)$-description. □

Combining Lemmas 6.6 and 6.5 with the easy observation that any consistent basic formula of depth $\leq k$, variables in Q and atomic games among τ is provably entailed by some (Q,k)-description[1], we obtain Theorem 6.1.

As a corollary to this proof, we get a further property of our new game logic:

Theorem 6.7. BPL *is decidable and has the effective finite model property.*

6.1 Special cases and variations

Logics of specific games Our move to abstract game state models has one consequence that should be noted. In the original example motivating this paper, powers were sets of endpoints of some concrete game. In that special case, iterating power modalities is almost trivial, as powers of players at endpoints s are just the singleton sets $\{s\}$. But because of this, the logic of game trees will validate further principles that were not in our BPL logic of games.

Fact 6.8. *The formula* $[i]([i]\varphi \leftrightarrow \varphi)$ *is valid on extensive games.*

This trivializes iterated play of the current game, while in a general game logic, we can play a game to its end in some state s in the state space, and then restart the game from there. We leave axiomatizating concrete games as an open problem, but we conjecture that adding the preceding principle suffices.

Intermediate powers Strategies in game trees are not just directed toward endpoints: players can also make sure that play has to pass through relevant intermediate positions. Accordingly, we can define powers that allow both intermediate nodes and endpoints as outcomes. (Cf. the intermediate forcing in Grossi and Turrini [**?**].) The earlier modal neighborhood language can still describe such powers, but now it gets closer to a standard modal action language over game trees, as powers can also be singletons of states reached after one move. We leave the logic of this game equivalence as one more open problem.

[1] This follows in a standard way from Lindenbaum's lemma applied to at most finitely many formulas of depth $\leq k$, variables in Q and game terms among τ up to provable equivalence.

7 Algebra of game constructions

Our modal neighborhood language describes the structure of games from an "internal" perspective, where formulas are evaluated at a given point in the associated state space of a game. But we can also describe game structure "externally" by means of an algebra of game constructions.

This approach, too, has connections with logic. We start by considering some natural operations on games, and then develop this perspective with games of imperfect information. This may be surprising, as these games have no compositional tree-like structure in general (epistemic uncertainty links can cross between subtrees), but we shall see that we can still come a long way.

7.1 Propositional operations

We will use some basic concepts of universal algebra, see for example [**?**]. For simplicity, we restrict attention to finite games, so that $\mathbb{G}(\{\mathsf{A}, \mathsf{E}\}, O)$ is now the set of finite games with outcomes in O. Thus the outcome map of a game $\mathcal{G}$ can be viewed a map o from the leaves in $\mathcal{G}$ into O.

Consider a family of extensive games on a fixed set of outcomes O. We define operations in a standard manner, with binary $\cup, \cap$ corresponding to choice for A, E respectively, and a unary operation $-$ for game dual ('role switch').

Definition 7.1. Let $\mathcal{G}_1 = (\mathcal{T}_1, t_1, o_1, \Pi_1)$ and $\mathcal{G}_2 = (\mathcal{T}_2, t_2, o_2, \Pi_2)$. The *choice games* $\mathcal{G}_1 \cup \mathcal{G}_2$ ($\mathcal{G}_1 \cap \mathcal{G}_2$) are defined as follows. We first construct a tree $\mathcal{T}'$ by adding a new root r with two successors, where the left successor is the root of a subtree isomorphic with $\mathcal{T}_1$ via a fixed isomorphism i_1, and the right successor the root of a subtree isomorphic with $\mathcal{T}_2$ via a fixed i_2. The turn function t' is defined by setting $t'(r) = \mathsf{A}$ ($t'(r) = \mathsf{E}$). For a node u in the subtree for the left successor of the root r we set $t'(u) = t_1(i_1(u))$, and for a node u in the subtree for the right successor of r, we set $t'(u) = t_2(i_2(u))$. The outcome map o' is defined by setting $o'(l) = o_1(i_1(l))$ for a leaf in the subtree corresponding to the left successor of r, and $o'(l) = o_2(i_2(l))$ for a leaf in the subtree corresponding to the right successor of r. We define a partition Π' by setting

$$\Pi' = \{\{r\}\} \cup \{i_1^{-1}[Z] \mid Z \in \Pi_1\} \cup \{i_2^{-1}[Z] \mid Z \in \Pi_2\}.$$

The game $\mathcal{G}_1 \cup \mathcal{G}_2$ ($\mathcal{G}_1 \cap \mathcal{G}_2$) is then defined as $(\mathcal{T}', t', o', \Pi')$.

The definition of $-\mathcal{G}$ is much simpler, it merely changes the turn assignment by switching players at each position, otherwise keeping everything the same.

Definition 7.2. Let $\mathcal{G} = (\mathcal{T}, t, o, \Pi)$. Then $\underline{-\mathcal{G}}$ is defined to be the structure $(\mathcal{T}, t^-, o, \Pi)$, where t^- is defined by $t^-(u) = \overline{t(u)}$ for each $u \in \mathcal{T}$.

Remark 7.3. These operations make strong assumptions. In particular, the imperfect information stays inside the games connected by $\cup$, $\cap$, since the initial choice introduces no uncertainty (in other words, the first round of a choice game has perfect information). We will return to this issue later on.

Basic powers in composite games can be straightforwardly computed from the basic powers of their components, as is shown by the following fact:

Fact 7.4. *Basic powers of players in these games obey the following conditions:*

1. *Basic powers for* A *in* $\mathcal{G}_1 \cup \mathcal{G}_2$ *are the union of* E*'s basic powers in the games* $\mathcal{G}_1$*,* $\mathcal{G}_2$ *separately,*

2. *Basic powers of* E *in* $\mathcal{G}_1 \cup \mathcal{G}_2$ *are unions of non-empty basic powers of* A *in* $\mathcal{G}_1$*,* $\mathcal{G}_2$ *separately.*

3. *Analogously for powers in* $\mathcal{G}_1 \cap \mathcal{G}_2$*.*

4. *Basic powers of* E *in* $-\mathcal{G}$ *are basic powers of* A *in* $\mathcal{G}$*, and likewise for* A*.*

Proof. By a simple verification. □

As a direct corollary we get the following proposition, which opens the door to a game algebra of basic power equivalence:

Proposition 7.5. *Basic power equivalence is a congruence on the algebra:*

$$\langle \mathbb{G}(\{\mathsf{A}, \mathsf{E}\}, O), \cup, \cap, -\rangle.$$

This motivates the following definition:

Definition 7.6. The *game algebra of basic power equivalence* $\mathfrak{G}$ (with outcomes O) is defined to be the quotient:

$$\langle \mathbb{G}(\{\mathsf{A}, \mathsf{E}\}, O), \cup, \cap, -\rangle / \simeq$$

An algebraic equation in the signature $\langle \cup, \cap, -\rangle\rangle$ is said to be *valid* on $\mathfrak{G}$ if it belongs to the equational theory of the algebra $\mathfrak{G}$.

As we did with the logic for basic powers, we now generalize from extensive game trees to arbitrary game state models $\mathfrak{M} = (S, N, V)$, where the neighborhood relations $N^{i,G}sX$ satisfy abstract analogues of the preceding observations, for all game terms G. In doing so, we take a dynamic view of games as state-transforming processes in the style of dynamic game logic (cf. [**?**], [**?**], [**?**]). This leads to an extended algebraic perspective.

Definition 7.7. A *dynamic two-player game* over a set X (of "states") is a map $g : X \to \mathbb{G}(\{\mathsf{A}, \mathsf{E}\}, X)$, assigning a game with outcome set X to each state in X. We denote the set of dynamic two-player games over X by $\mathbb{D}(\{\mathsf{A}, \mathsf{E}\}, X)$.

The operations $\cup, \cap$ and $-$ are naturally lifted to dynamic games in a component-wise manner, and so is the relation $\simeq$ of basic power equivalence. The following observation shows we get an algebra once more:

Proposition 7.8. *Basic power equivalence is a congruence on the algebra:*

$$\langle \mathbb{D}(\{\mathsf{A}, \mathsf{E}\}, O), \cup, \cap, -\rangle.$$

Definition 7.9. The *dynamic game algebra of basic power equivalence* $\mathfrak{D}$ (with outcomes O) is defined to be the quotient:

$$\langle \mathbb{D}(\{\mathsf{A}, \mathsf{E}\}, O), \cup, \cap, - \rangle / \simeq$$

Equational validity for $\mathfrak{D}$ is again defined in the standard manner.

The following observations about valid equations are easy to verify.

Fact 7.10. *The following are valid in the game algebra of basic powers:*

$$\begin{array}{ll} x \cup y = y \cup x & x \cap y = y \cap x \\ x \cup (y \cup z) = (x \cup y) \cup z & x \cap (y \cap z) = (x \cap y) \cap z \\ - - x = x & \\ -(x \cup y) = -x \cap -y & -(x \cap y) = -x \cup -y. \end{array}$$

These are all valid in propositional logic. But there are also striking failures:

Fact 7.11. *The following equations are not valid in the algebra of basic powers:*

$$\begin{array}{ll} x \cup x = x & x \cap x = x \\ x \cup (y \cap z) = (x \cup y) \cap (x \cup z), & x \cap (y \cup z) = (x \cap y) \cup (x \cap z). \end{array}$$

That distribution fails was observed already in our introduction. We only retain a weaker principle which follows from the above valid equations:

$$x \cup (y \cap x) = (x \cup y) \cap x$$

Most striking, perhaps, is the failure of the idempotence laws in our algebra:

Example 7.12. Failure of idempotence.

Consider a game $\mathcal{G}$ with outcomes in $O = \{p, q\}$, where Player A moves first, and gets a choice between outcomes p or q, with no moves at all for Player E. The basic powers of Player A in $\mathcal{G}$ are just $\{p\}$ and $\{q\}$. But in the game $\mathcal{G} \cap \mathcal{G}$, starting with a choice between two copies of $\mathcal{G}$ for Player E, Player A clearly also has the basic power $\{p, q\}$, so it is not the case that $\mathcal{G} \cap \mathcal{G} \simeq \mathcal{G}$. The two games are displayed in Figure 6, with $\mathcal{G}$ to the left, $\mathcal{G} \cap \mathcal{G}$ to the right.

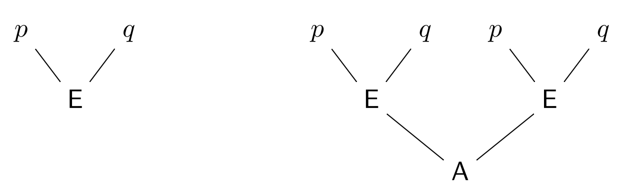

Figure 6: Failure of idempotence.

This makes our game algebra a weak propositional logic sharing a failure of distribution with quantum logic, and a failure of idempotence with linear logic.

Open problem 7.13. *Axiomatize the complete game algebra of basic powers.*

An interesting feature of this game algebra is its use as a tool for studying game equivalence under basic powers. Valid laws of propositional logic that interchange conjunctions and disjunctions correspond to transformations that change order of turns in games. Normal forms yield normal forms for games: say, a distributive normal form shows that we can let one player start (all consecutive moves can then be collected into one), then the other, and let the game end there. However, our weaker game algebra for basic powers does not support standard propositional normal forms, a topic to which we return below.

7.2 Composition and relational strategies

There are further fundamental operations on games. Clearly, games compose to form sequential games, and the corresponding operation is as follows.

Definition 7.14. Given dynamic games $\mathcal{G}_1, \mathcal{G}_2 : X \to \mathbb{G}(\{\mathsf{A}, \mathsf{E}\}, X)$, we can define the *sequential composition* $\mathcal{G}_1 \circ \mathcal{G}_2$ by letting $\mathcal{G}_1 \circ \mathcal{G}_2(u)$ be constructed by replacing each leaf l in $\mathcal{G}_1(u)$ by a copy of the game tree $\mathcal{G}_2(o_1(l))$, where o_1 is the outcome map associated with $\mathcal{G}_1(u)$.

Remark 7.15. This is just one definition of sequential game composition, as it does not allow non-trivial imperfect information links across sequential games. But for instance, if endpoints of G can have epistemic links, we might want to let this continue to the games placed at these. On the other hand game theory often assumes a mysterious act of "revelation" at the end of the game.

However, there is a more immediate problem to deal with, even for this simple notion. It turns out that basic powers are an inherently "global" property of games, and fail to be compositional with respect to sequential composition.

Fact 7.16. *Basic power equivalence is not generally a congruence with respect to the operation of sequential composition on the algebra* $\mathbb{D}(\{\mathsf{A}, \mathsf{E}\}, O)$.

Proof. Let $O = \{p, q\}$. The two perfect information games displayed in Figure 7 can obviously be split up as sequential compositions of pairwise basic power equivalent games. The games are not basic power equivalent themselves, since player E has a basic power $\{p, q\}$ in the game to the right, but not to the left.

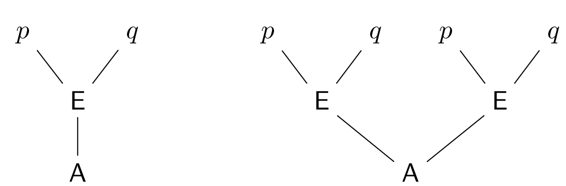

Figure 7: Failure of safety

□

Why does this happen? An obvious culprit is the functional nature of strategies. If a model gets inflated modulo bisimulation, creating copies of earlier moves, then new functions can create new basic powers that were not there before. One solution, then, is to change the functional nature of strategies.

To remedy this situation, we will now widen the notion of a strategy to allow *non-determinism*, so that strategies may constrain the moves of a player, but not determine them uniquely. This move is not altogether new: mixed strategies in game theory can be interpreted in a relational way – and the use of relations rather than functions has also been defended for the broader notion of a 'plan' in [**?**] leading to a better model theory of strategies.

Definition 7.17. A *relational strategy* for player P in a game $\mathcal{G} = (\mathcal{T}, t, o, \Pi)$ is a binary relation σ over $\mathcal{T}$ such that:

- $\sigma[u] \neq \emptyset$ whenever $u \in t^{-1}[\mathsf{P}]$, and
- $\sigma[u] = \sigma[v]$ if u, v are in the same partition cell in Π.

The set $\mathsf{Match}(\sigma)$ of non-deterministic matches guided by a strategy σ is defined in the obvious way. We say that $P \subseteq O$ is a *relational basic power* of P if there is a relational strategy σ for P in $\mathcal{G}$ such that $P = \{o(m) \mid m \in \mathsf{Match}(\sigma)\}$. The set of relational basic powers of P is denoted by $R_{\mathsf{P}}(\mathcal{G})$. We say that $\mathcal{G}_1, \mathcal{G}_2$ are *relational basic power equivalent* if $R_{\mathsf{P}}(\mathcal{G}_1) = R_{\mathsf{P}}(\mathcal{G}_2)$, for each $\mathsf{P} \in \{\mathsf{A}, \mathsf{E}\}$. Finally, we write $\mathcal{G}_1 \equiv \mathcal{G}_2$ when $\mathcal{G}_1, \mathcal{G}_2$ are relational basic power equivalent.

Relational basic power equivalence can be lifted to an equivalence relation between dynamic games in the same component-wise manner as before.

Our first observation is that the new notion solves our problem.

Fact 7.18. *Relational basic power equivalence over* $\mathbb{D}(\{\mathsf{A}, \mathsf{E}\}, O)$ *is a congruence* $\equiv$ *with respect to the operations* $\cup, \cap, -$ *plus sequential game composition* $\circ$.

This motivates the following definition:

Definition 7.19. The *dynamic game algebra* $\mathfrak{R}$ *of relational basic power equivalence* (with outcome set O) is the quotient

$$\langle \mathbb{D}(\{\mathsf{A}, \mathsf{E}\}, O), \cup, \cap, -, \circ \rangle / \equiv .$$

But of course, this change in our set-up for defining basic powers needs to be checked for its compatibility with what went on before. We briefly state a number of results, whose proofs involve straightforward verification:

Theorem 7.20. *The following conditions capture representability of families of sets as relational basic powers in a game of imperfect information: Non-Emptiness, Consistency, Exhaustiveness, and one new property:*

Union Closure*: If all members of a family F are relational basic powers of a player, then so is the union $\bigcup F$.*

Proof. With Union Closure for two families of sets $\mathcal{X}$ and $\mathcal{Y}$, the basic powers of the two-player game constructed in the proof of Theorem 4.2 are closed under unions, which clearly means the basic powers of each player in that game coincide with their *relational* basic powers. Since the basic powers of the two players are the members of $\mathcal{X}$ and $\mathcal{Y}$ respectively, by Theorem 4.2, the result follows. □

We can also easily adapt our complete axiomatization of the modal logic of basic powers, so that it fits relational basic powers:

Theorem 7.21. *Basic* INL *game logic is also complete for relational powers, provided we add the axiom:*

$$\langle g, \mathsf{P}\rangle(\alpha; \varphi_1) \wedge \langle g, \mathsf{P}\rangle(\beta; \varphi_2) \rightarrow \langle g, \mathsf{P}\rangle(\alpha, \beta; \varphi_1 \vee \varphi_2)$$

That is, the logic obtained by adding this axiom is sound and complete for validity over neighborhood models satisfying the extra constraint of Union Closure.

Proof. The construction of a game in the proof of Theorem 6.1 stays exactly as before, but one has to check that the new axiom enforces Union Closure. □

The basic game algebra changes in this new setting. We retain all valid equations from before, since we are dealing with a *looser* game equivalence when working with relational, rather than functional strategies. But the algebra remains weaker than the algebra of standard power equivalence.

Fact 7.22.

1. *All validities mentioned above for functional strategies are also valid with basic powers computed with relational strategies.*
2. *Propositional distribution is still not valid, but idempotence is.*

Next, we turn to valid principles for composition of games, the motivation for making the change to relational strategies in the first place.

Fact 7.23.

1. *Game composition validates associativity, and left-distribution over choice:*

$$(x \cup y) \circ z = (x \circ z) \cup (y \circ z)$$

2. *Dual distributes over composition:*

$$-(x \circ y) = -x \circ -y$$

3. *Game composition does not validate right-distribution over choice.*

Proof. The first two items are immediate. Failure of right-distribution follows by the known fact that right-distribution fails already in the game algebra of standard power equivalence [**?**, **?**]. Since standard power equivalence is a coarser equivalence notion than relational basic power equivalence, any equation that fails in the game algebra of standard power equivalence also fails now. □

Open problem 7.24. *Axiomatize the complete game algebra of the operations* $\{\cup, \cap, -, \circ\}$ *based on relational powers.*

8 Dynamic game logic and a new game algebra

Various extensions are possible for the algebraic setting we have found so far. We mention two directions with further attractive tools for analyzing games.

8.1 Dynamic game logic

Putting instantial game logic together with game algebra yields the following system in the style of dynamic game logic, [**?**], [**?**]. We will then have instantial modalities describing basic powers of player P in the game $\mathcal{G}$:

$$\langle \mathcal{G}, \mathsf{P} \rangle (\Psi; \varphi)$$

This formalism can be interpreted on our neighborhood models, provided these satisfy the constraint of Union Closure as well as the Non-Emptiness, Consistency and Exhaustiveness constraints. The crucial point here is that, with the earlier obstacle to compositionally overcome, we can define the power relation for a composite game of the form $\mathcal{G}_1 \circ \mathcal{G}_2$ in the following inductive manner:

$(u, Z) \in R^{\mathsf{P}}_{\mathcal{G}_1 \circ \mathcal{G}_2}$ iff $Z = \bigcup F$, for some family $F \subseteq \mathcal{P}W$ and some $Y \subseteq W$ with $(u, Y) \in R^{\mathsf{P}}_{\mathcal{G}_1}$ and $(Y, F) \in \widetilde{R}^{\mathsf{P}}_{\mathcal{G}_2}$,

where $(Y, F) \in \widetilde{R}^{\mathsf{P}}_{\mathcal{G}_2}$ means that the following back-and-forth conditions hold:

for all $y \in Y$ there is some $S \in F$ with $(y, S) \in R^{\mathsf{P}}_{\mathcal{G}^2}$, and vice versa,
for all $S \in F$ with $(y, S) \in R^{\mathsf{P}}_{\mathcal{G}^2}$ there is some $y \in Y$ with $(y, S) \in R^{\mathsf{P}}_{\mathcal{G}^2}$.

This semantics validates principles that allow for recursive reasoning about basic game operations in a dynamic logic format. We merely display a few:

Reduction axioms:

- $[\pi_1 \cup \pi_2, E](\Psi; \varphi) \leftrightarrow [\pi_1, E](\Psi; \varphi) \vee [\pi_2, E](\Psi; \varphi)$
- $[\pi_1 \cap \pi_2, E](\Psi; \varphi) \leftrightarrow \bigvee \{[\pi_1, E](\Theta_1; \varphi) \wedge [\pi_2, E](\Theta_2; \varphi) \mid \Psi = \Theta_1 \cup \Theta_2\}$
- $[\pi_1 \circ \pi_2, E](\psi_1, ..., \psi_n; \varphi) \leftrightarrow [\pi_1, E]([\pi_2, E](\psi_1; \varphi), ..., [\pi_2, E](\psi_n; \varphi); [\pi_2, E]\varphi)$

When added to our base logic BPL, these lead to the following result.

Theorem 8.1. *The dynamic game logic of the operations* $\{\cup, \cap, -, \circ\}$ *with relational basic powers is completely axiomatizable.*

In the companion paper [**?**], we explore this perspective in more depth, in the form of a PDL-like dynamic logic over instantial neighborhood semantics with a new computational interpretation. This logic includes the game operations defined here, except for unrestricted game dual, but it also adds a further natural construction of finite iteration of games, which calls for new techniques.

8.2 New game constructions and extended game algebra

The move to relational strategies solved a design problem, and led to an interesting game algebra and game logic. Even so, it is not the only possible approach. In particular, as noticed several times already, we have not deeply analyzed the role of imperfect information, and what repertoire of algebraic operations would best fit with that. Here is an example showing how a greater focus on imperfect information provides an alternative analysis of our earlier problems.

Example 8.2. Restoring distribution through imperfect information.

Consider our initial example of failure of the propositional distribution law with basic powers in Section 2. We insert one uncertainty link in the game to the right, as depicted in Figure 8. Now, keeping in mind that strategies in imperfect information games must choose the same moves at epistemically linked points, it is easy to see that basic powers are the same in the two games depicted:

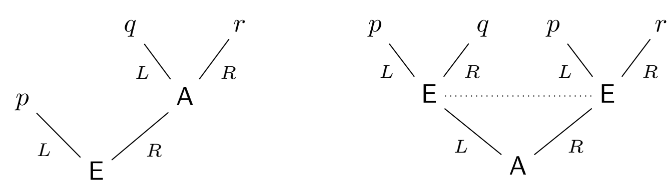

Figure 8: Distribution revisited

This simple observation suggests a natural new choice operation that combines imperfect information games, namely, one that introduces uncertainty:

Let $\mathcal{G}_1 \underline{\cup}\, \mathcal{G}_2$ be the game where A starts by choosing either $\mathcal{G}_1$ or $\mathcal{G}_2$, but this choice is not observed by A – and likewise for $\mathcal{G}_1 \underline{\cap}\, \mathcal{G}_2$.

The game algebra with 'partly invisible choices' that now results has operations

$$\{-, \cup, \underline{\cup}, \cap, \underline{\cap}\}, \circ.$$

Interestingly, it now has valid equations in the extended vocabulary that take the place of the invalid distribution law, such as

$$x \cap (y \cup z) = (x \cap y) \underline{\cup} (x \cap z).$$

This new game algebra provides a much tighter fit for our earlier game transformations. Its extended propositional logic has strong analogies with the 'Thompson transformations' for imperfect information games in game theory, cf. [**?**]. But we do not have a full-fledged theory to offer at this stage.

Open problem 8.3. *Axiomatize the complete algebra of* $\{-, \cup, \underline{\cup}, \cap, \underline{\cap}, \circ\}$.

This game algebra can be viewed as a generalized propositional logic. It has even intriguing stronger principles than the one we noted. Here is one very strong distributivity principle without a classical propositional counterpart.

Fact 8.4. *The following equivalence is valid:* $(x \cap y) \underline{\cup} (z \cap u) = (x \cup z) \underline{\cap} (y \cup u)$.

Elaborating the syntax of this propositional logic takes a good deal of care, since in stating valid principles, we need to 'balance' formulas, and mark subformulas (viewed as possible moves) in a way that respects the uniformity conditions on strategies in imperfect information games (indistinguishable positions should have the same moves available, strategies should make the same choice at indistinguishable positions). We will not pursue this system here, but note that there are analogies with 'IF-logic' of imperfect information, as pointed out in van Benthem [**?**, Chapter 21]. The latter source also suggests adding even one more operation to the algebra, namely parallel composition of games.

As before, we should also add sequential composition to the signature of this algebra. In that case, another suggestive observation can be made. Imperfect information can play a similar role to our introduction of relational strategies.

Example 8.5. Safety revisited.

Consider the example of failure of safety given in the proof of Fact 7.16, which blocked compositionality for basic powers with functional strategies. This time, however, assume that in the duplicated model, there is an uncertainty link between the two copies of the midlevel point. Now the only admissible strategies are those that make the same choice at both midlevel points, and accordingly, the resulting basic powers are the same in both models.

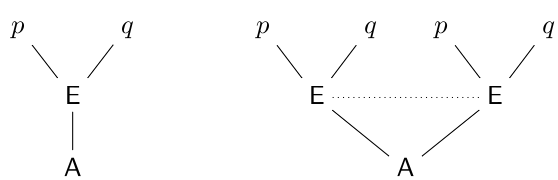

Figure 9: Safety revisited

We have not been able to determine precisely to which extent imperfect information, made explicit in the manner proposed above, can play the same role in the design of game algebra and dynamic game logic as our move from functional to relational strategies. Both approaches have their own motivation, but we leave a deeper comparison of the two as an open problem.

Finally, we started this section by saying that imperfect information games are a typical challenge to compositional algebraic thinking because of the arbitrary crossing of uncertainty lines between subtrees, making the notion of a subgame problematic. The way we got around this was by having algebraic operations that only introduce epistemic uncertainty in very controlled ways. We will discuss a few further aspects of this setting in Section 9.1 below.

9 Further directions

In this final section, we collect a few further perspectives on our motivating issue of natural game equivalences that form a background to what we have done in this paper. Our discussion will be light, and we mainly raise new questions.

9.1 Equivalence levels and translating logics

We have seen various natural levels for studying games, defined by invariance notions, and we gave matching logics. The resulting total picture raises questions. How are the various logics that we considered connected? Can we switch between levels in a systematic manner, perhaps via translations of logics?

We look at these issues for finite game trees with perfect information. Imperfect information will be discussed separately later on. Here is a known result.

Fact 9.1. *The neighborhood modality $\langle g, \mathsf{P}\rangle\varphi$ for standard powers is definable in the modal (fixpoint) language of actions by the following equivalence, where standard modalities $\Diamond$, $\Box$ run over all available moves for the players* A *and* E*:*

$$\langle g, \mathsf{P}\rangle\varphi \leftrightarrow \mu x.\varphi \vee (\mathsf{turn}_\mathsf{P} \wedge \Diamond x) \vee (\mathsf{turn}_{\overline{\mathsf{P}}} \wedge \Box x)$$

This logical pattern of alternating existential and universal modalities over moves is characteristic of many game-theoretic arguments. It follows that standard powers are invariant for standard modal bisimulations of finite game-trees. However, with basic powers, this changes: these are not invariant for bisimulation. In fact, interpreting the example shown in Figure 7 in a different way, we can see it as showing precisely that. The two games are bisimilar as game trees, but do not have the same basic powers. Stated in other terms, the INL-formula

$$\langle g, \mathsf{A}\rangle(p, q; p \vee q)$$

interpreted over these models is not invariant for standard bisimulation. Hence our BPL-language cannot be defined in the modal base language of games. One symptom of this is the impossibility of generalizing the earlier recursion for neighborhood modalities to the INL-modality with instantial information.

However, things change with basic powers based on relational strategies.

Fact 9.2. *Relational basic powers are invariant for standard bisimulation.*

Proof. It is easy to see that any two bisimilar game trees can be turned into *isomorphic* trees by simply creating sufficiently many copies of the moves of each player. This construction obviously does not change the relational basic powers of a game (though it does change the basic powers), and the result follows. □

We can provide an explicit definition of the instantial neighborhood modality for relational basic powers in the modal μ-calculus via the following definition:

$$\langle g, \mathsf{P}\rangle(\psi_1, ..., \psi_n; \varphi) \;\leftrightarrow\; \bigwedge_{1\leq i\leq n} \mu x.(\psi_i \wedge \varphi) \vee (\mathsf{turn}_\mathsf{P} \wedge \Diamond x) \vee (\mathsf{turn}_{\overline{\mathsf{P}}} \wedge \Diamond x \wedge \Box\langle g, \mathsf{P}\rangle\varphi)$$

where we recall that $\langle g, \mathsf{P}\rangle\varphi$ was definable in the μ-calculus as noted before.

Coda: logics for strategic forms

Finally, we mention one more level of looking at games, the strategic forms that were linked to basic power analysis in Section 4.2. These matrix games look like simple games of imperfect information, and this is a valid analogy. Indeed, they have a simple bimodal logic where one modality ranges over choices of a player i (i's 'freedom') and the other over choices of the other player (that is, i's uncertainty), cf. [**?**]. How is this new logic related to our earlier ones? This issue is more delicate. Points in the strategic form are strategy profiles, higher-order objects from the perspective of extensive game trees. Therefore, we do not expect simple translations between the logics of extensive games and strategic form games, and we leave a complete analysis for future study.

9.2 Imperfect information and epistemic logic

From a logical point of view, imperfect information games combine two relations: actions and uncertainty links. This suggests extending our earlier languages to a richer bimodal logic that can talk about action, or at a coarser level: powers, via suitable modalities, say in INL, but that also has an epistemic operator $K\varphi$ for knowledge. This interplay is important because we can then formulate non-trivial statements about the interplay of knowledge and action, such as

- $K_{\mathsf{A}}\langle g, \mathsf{A}\rangle\varphi$: player A knows E has a power for achieving the truth of φ,
- $\langle g, \mathsf{E}\rangle K_{\mathsf{E}}\varphi$: player A has a power for making player E know that φ.

Example 9.3. Natural axioms formulated in such a bimodal language express special properties one can assume about players, such as (in simplified form)

$$K_{\mathsf{P}}\langle g, \mathsf{P}\rangle\varphi \to \langle g, \mathsf{P}\rangle K_{\mathsf{P}}\varphi.$$

This axiom says essentially that actions that can be perfectly observed do not increase uncertainty, a version of Perfect Recall for players, cf. [**?**] for details. The converse implication also has a clear game-theoretic meaning:

$$\langle g, \mathsf{P}\rangle K_{\mathsf{P}}\varphi \to K_{\mathsf{P}}\langle g, \mathsf{P}\rangle\varphi.$$

This expresses essentially that learning cannot take place spontaneously without a difference in action, a property of players sometimes called No Miracles.

Combined logics of action and knowledge are well-understood in terms of expressive power and of computational complexity. In particular, [**?**] ties Perfect Recall and No Miracles to a style of play where uncertainty links in a game tree arise by an update mechanism from dynamic-epistemic logic. Both players can observe each other's moves wholly or partially, which creates horizontal uncertainty links only, resulting in a synchronous view of the game. This is precisely the view of games encoded in our earlier game algebra of $\{-, \cup, \underline{\cup}, \cap, \underline{\cap}\}$.

We will not pursue this epistemic logic, but it relates to many earlier topics. For instance, strategies in games of imperfect information are like 'knowledge programs', where players' choices are guided only by what they know. Also, issues of definability between levels of representing games get more complicated.

But perhaps the major problem left open here concerns our initial topic of game equivalence. Suppose we are presented with two extensive games annotated with uncertainty links indicating the information available to players at various stages of the game. When do we consider these games to be the same? This is not just a matter of comparing actions or powers, it also depends on what we assume about the type of players inhabiting these games. Perhaps we need to refine our earlier intuitions then to equivalence of games "as played".

9.3 Preference and rationality

As a final topic, we mention one more crucial feature of actual games: players' preferences. Behavior in games involves a balance between available actions, information, and players' preferences between outcomes. A standard game-theoretic solution method showing this mixture at work is this procedure:

Example 9.4. Backward Induction.

To show the preferences in the games depicted, outcomes have been marked numerically at the three endpoints y, u, v in the order (value for A, value for E).

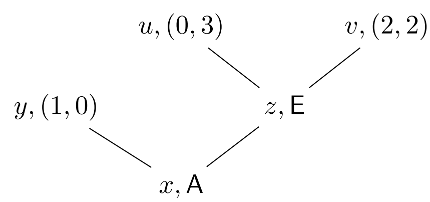

Figure 10: Backward induction

Here, a rational player E will go left at his turn, since he prefers the outcome 3 over the 2 to the right. Then player A, knowing (or believing) this, will go left as well, as she prefers outcome 1 to 0. Thus, the Backward Induction strategy for A is "left", and so is that for E – though the latter strategy never gets played.

Backward Induction has long been a benchmark for logical studies of games, and there is an extensive theory of this and other solution methods. In particular, the reasoning leading to the above prediction about the game has been analyzed in fixed-point logics for relational, rather than functional strategies, a move similar to that in Section 9.2. The Backward Induction reasoning can also be captured in a bimodal logic of actions with an added preference modality.

We conclude this preliminary discussion by connecting with our theme of powers. The above computation can also be seen as a process of "pruning powers", more precisely, of intermediate powers in the sense of Section 4.2.

Example 9.5. Backward Induction and powers.

At point z in the game tree of Figure 10, E has powers $\{u\}$ and $\{v\}$. But these powers stand in a preference relation $\{u\} >_{\mathsf{E}} \{v\}$ in an obvious sense. We can lift players' preference relation to any two sets X, Y as follows:

$$X >_{\mathsf{P}} Y$$

if all points in X are preferred to all points in Y. There are other ways of lifting relations on points to relations on sets, but this one suffices for our purposes.

The rationality embodied in the Backward Induction algorithm assumes essentially that "players never play a power for which they have a better one" in this lifted sense. Continuing in this way along the above reasoning steps, a subtle change takes place. Player A now needs to compare, not $\{y\}$ and $\{u, v\}$, but $\{y\}$ and $\{u\}$, where $\{u\}$ is the prediction or a belief about best play further on. But then, since $\{x\} >_{\mathsf{A}} \{u\}$, her best power is $\{x\}$.

Thus we must extend our earlier analysis of games in terms of players' basic powers to include their best or optimal powers. We do not yet know how to do this in a satisfying manner, so we leave it as an open problem.

We end with an observation connecting to our earlier concerns.

Example 9.6. Failure of distribution with preference.

Applying our running example of the propositional distribution law to the preceding game, yields the new extensive game depicted in Figure 11:

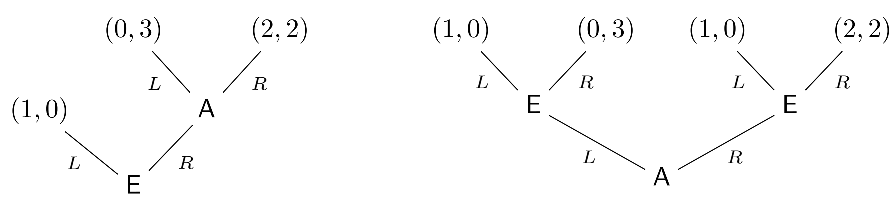

Figure 11: Distribution with preference

Now compare the Backward Induction outcomes for these two games. In the game to the right, it is given by the bold-face links, and clearly, this is not the same prediction as for the game on the right.

So, distribution fails for games with preference, but for other reasons than we had before. The game algebra of games under preference seems a good place for studying the effects of considering players' preferences explicitly.

10 Conclusion

We have defined a new game equivalence in terms of basic powers, and developed its modal logic and game algebra. In the process, we found analogues of results for other game equivalences, but also a number of non-trivial new features that suggest new lines of investigation on logic and games, such as the game algebra of operations that take on board imperfect information in various ways.

Even so, our style of analysis has clear limitations. We considered two-player games only, and we mainly looked at finite instead of infinite games where the histories themselves are the output of the game. And we encountered a serious barrier to our style of analysis when discussing games with preferences. We do not see these limitations as definitive, but clearly, much needs to be done before we can assess the true merits of the logical perspective pursued in this paper.

A Game theory definitions and notation

Definition A.1. A *tree* $\mathcal{T}$ is a prefix closed subset of $\mathbb{N}^*$, subject to the condition that if $w \cdot j \in \mathcal{T}$ and $i < j$ then $w \cdot i \in \mathcal{T}$ as well. The empty word ε is the *root* of the tree.

Standard concepts like branches, leaves and roots of trees are defined as usual.

Definition A.2. An *extensive game* $\mathcal{G}$ for a finite set of players A with outcomes in the set O is a tuple $(\mathcal{T}, t, o, \Pi)$ where $\mathcal{T}$ is a finite tree, t a map from non-leaf nodes of $\mathcal{T}$ to A, o a map from branches of $\mathcal{T}$ to O, and Π a partition of $\mathcal{T}$ subject to the following condition: for any pair w, v within the same partition cell of Π, w and v have the same number of children in $\mathcal{T}$, and furthermore $t(n) = t(n')$. If all partition cells of Π are singletons we call $\mathcal{G}$ a game of perfect information, and we omit Π.

Maximal branches of $\mathcal{T}$ will also be called *full matches*, and prefixes of maximal branches are called *partial matches*.

A *strategy* for player $a \in A$ is a map $\sigma : t^{-1}[a] \to \mathbb{N}$ where $w \cdot \sigma(w)$ is a child of w for each w with $t(w) = a$, and $\sigma(w) = \sigma(w')$ whenever w, w' are in the same partition cell in Π. A *strategy profile* is a tuple $(\sigma_a)_{a \in A}$ of one strategy for each player in A. A strategy profile p determines a unique full match guided by the strategies of each player, so we can speak of the outcome of the profile p and denote it by $o(p)$. Generally, we say that a full match m of $\mathcal{G}$ is *guided* by the strategy σ for a if for every prefix w of m such that $t(w) = a$, $\sigma(w)$ is also a prefix of m. $\mathsf{Match}(\sigma)$ is the set of σ-guided matches.

We denote the set of games for players A with outcomes in O as $\mathbb{G}(A, O)$. For two-player games we call the players by E (Eloise) and A (Abelard). We set $\overline{\mathsf{E}} = \mathsf{A}$ and $\overline{\mathsf{A}} = \mathsf{E}$.

Note that we have not attributed payoffs to matches in a game or preferences over the outcomes, but rather (and more generally) simply outcomes from some fixed chosen set. In this sense we are dealing with *game forms* rather than proper games. We return to the issue of preferences in Section 7.1.

Definition A.3. Let $\mathcal{G} = (\mathcal{T}, t, o, \Pi)$ be a game with outcomes in O. A set $P \subseteq O$ is a *power* of player $a \in A$ in the game $\mathcal{G}$ if there is a strategy σ for a in

G such that $o(m) \in P$ for every σ-guided match m. Given a player $a \in A$ we let $P_a(\mathcal{G})$ denote the set of powers of a in $\mathcal{G}$.

Two games $\mathcal{G}_1, \mathcal{G}_2 \in \mathbb{G}(A, O)$ are *power equivalent* if for all $a \in A$: $P_a(\mathcal{G}_1) = P_a(\mathcal{G}_2)$. We denote this by $\mathcal{G}_1 \sim \mathcal{G}_2$. If $P_a(\mathcal{G}_1) = P_a(\mathcal{G}_2)$ for some specific $a \in A$, we write $\mathcal{G}_1 \sim_a \mathcal{G}_2$.

We use the convention of writing $\overline{\mathsf{E}} = \mathsf{A}$ and $\overline{\mathsf{A}} = \mathsf{E}$.

B Omitted proofs

Proof of Fact 4.1. For Non-Emptiness, pick a player P. By definition of an extensive game, the set of strategies available to each player is non-empty (in particular, if $t^{-1}[\mathsf{P}] = \emptyset$ then the empty function is the unique available strategy for P). So pick an arbitrary strategy σ for P. Then the set $\{o(m) \mid m \in \mathsf{Match}(\sigma)\}$ is a basic power for P.

For Consistency, let P be a basic power for E and let Q be a basic power for A. Then there are strategies σ, τ for E, A respectively such that $P = \{o(m) \mid m \in \mathsf{Match}(\sigma)\}$ and $Q = \{o(m) \mid m \in \mathsf{Match}(\tau)\}$. Hence $o(p) \in P \cap Q$ for the strategy profile $p = (\sigma, \tau)$.

For Exhaustiveness, suppose that P is a basic power of Player P and $x \in P$. Pick a strategy σ for P such that $P = \{o(m) \mid m \in \mathsf{Match}(\sigma)\}$. Since $x \in P$ there is a σ-guided match m such that $o(m) = x$. We define a strategy τ for Player $\overline{\mathsf{P}}$ as follows: if a position $u \in t^{-1}[\overline{\mathsf{P}}]$ belongs to m then u cannot be a leaf since $t(u)$ is defined, so since m is a full match there must be some unique child v of u such that v also belongs to m. Set $\tau(u) = v$. If $u \in t^{-1}[\overline{\mathsf{P}}]$ does not belong to m then define $\tau(u)$ arbitrarily. Then m is a τ-guided match, so $x = o(m) \in Q$ where $Q = \{o(m) \mid m \in \mathsf{Match}(\tau)\}$, which is a basic power of $\overline{\mathsf{P}}$. □