15
Model Comparison

Logical formulas express properties of semantic structures. Different languages have different expressive strength over models, showing in powers of distinction. A poor language with only "Yes" and "No" distinguishes few situations, while a rich language can distinguish a whole spectrum. Model comparison can be cast as a game between a "duplicator" $\boldsymbol{D}$ who claims that two given models $\boldsymbol{M}$ and $\boldsymbol{N}$ are similar, and a "spoiler" $\boldsymbol{S}$ who claims they are different. In this chapter, we define comparison games for first-order logic, prove their adequacy for model equivalence, give correspondences between players' winning strategies and logical difference formulas or potential isomorphisms, discuss general game-theoretic aspects of the games, and show how to create variations and extensions. These games go back to Fraïssé (1954) and Ehrenfeucht (1961). Thomas (1997) uses them in computer science, and for a mathematics-oriented slant, see Doets (1996) or Väänänen (2011).

15.1 Isomorphism and first-order equivalence

Expressive power and invariances The expressive power of a language shows in its power of distinguishing situations, as we saw in Chapter 1. The notions of transformations and invariants from 19th century geometry make precise sense of this. In logic, this requires two things: a relation of structural *invariance* between models, and a *language* expressing the properties of those models. The analysis aims to show that the invariance matches those differences between models that the language cannot detect. About the most important structural invariance relation is the following widespread notion in mathematics.

DEFINITION 15.1 Isomorphism
Two models $\boldsymbol{M}$ and $\boldsymbol{N}$ are *isomorphic* (written $\boldsymbol{M} \cong \boldsymbol{N}$) if there exists a bijection f between the objects in their domains that preserves all the relevant structure: atomic properties, relations, distinguished objects, and operations. Thus, we have

$R^{\boldsymbol{M}}de$ iff $R^{\boldsymbol{N}}f(d)f(e)$	for all binary predicates R, and objects d, e in $\boldsymbol{M}$
$f(G^{\boldsymbol{M}}(d)) = G^{\boldsymbol{N}}(f(d))$	for all unary functions G, and objects d in $\boldsymbol{M}$

These two clauses show the general pattern of structure preservation. ■

Coarser invariants may be just the right comparison level for some other purpose: witness the notions of bisimulation between process models in Chapters 1 and 11.[179]

First-order expressiveness For convenience, in this chapter, we only use a first-order logic whose vocabulary has finitely many predicate letters and individual constants. The linguistic notion of model comparison is elementary equivalence $\boldsymbol{M} \equiv \boldsymbol{N}$: that is, $\boldsymbol{M}$ and $\boldsymbol{N}$ satisfy the same sentences. How close is this to a structural similarity? Let us look at the basic Isomorphism Lemma to find out.

FACT 15.1 For all models $\boldsymbol{M}$ and $\boldsymbol{N}$, if $\boldsymbol{M} \cong \boldsymbol{N}$, then $\boldsymbol{M} \equiv \boldsymbol{N}$.

Proof An easy induction on first-order formulas φ shows that, for all tuples of objects $\boldsymbol{a}$ in $\boldsymbol{M}$, and any isomorphism f sending the latter to the model $\boldsymbol{N}$, we have that $\boldsymbol{M} \models \varphi[\boldsymbol{a}]$ iff $\boldsymbol{N} \models \varphi[f(\boldsymbol{a})]$. ■

This implication holds for any well-behaved logical language. The converse is by no means true for first-order logic. What does hold is full harmony in the special case of *finite* models.

FACT 15.2 For all finite models, the following two assertions are equivalent:

(a) $\boldsymbol{M}$ is isomorphic with $\boldsymbol{N}$.

(b) $\boldsymbol{M}$ and $\boldsymbol{N}$ satisfy the same first-order sentences.

Proof From (b) to (a). Write a first-order sentence $\delta^{\boldsymbol{M}}$ describing $\boldsymbol{M}$. Let there be k objects. Then quantify existentially over $x_1, \ldots, x_k$, enumerate all true atomic statements about these in $\boldsymbol{M}$, plus the true negations of atoms, and state that no other objects exist. Since $\boldsymbol{N}$ satisfies $\delta^{\boldsymbol{M}}$, it can be enumerated just like $\boldsymbol{M}$. The isomorphism is immediate. ■

179 For much more on invariance and logical definability, see van Benthem (1996, 2002b).

This proof does not extend to infinite models, as first-order logic cannot define finiteness. Nor, for instance, can it tell the rationals $\mathbb{Q}$ apart from the reals $\mathbb{R}$ in their order $<$.

Example 15.1 Natural versus supernatural numbers
Elementary equivalence cannot even distinguish the natural numbers $\mathbb{N}$ from the model $\mathbb{N}+\mathbb{Z}$ that continues with the integers as supernatural numbers:

$\mathbb{N}$	versus	$\mathbb{N}+\mathbb{Z}$
$0, 1, 2, \ldots$		$0, 1, 2, \ \ldots \ \ \ldots \infty+1, \infty, \infty-1, \ldots$

We will see the reason for this indistinguishability in Section 15.5. ■

With a richer vocabulary, however, a language may see differences that used to be invisible. But there is no need to always extend our systems. In fact, weak expressive power can also be a good thing, as it yields transfer of properties across different situations, say, between standard models and nonstandard models.

15.2 Ehrenfeucht-Fraïssé games

The fine structure of the above invariance analysis is brought out by playing a certain type of logic games. These will work for any models, finite or not.

Playing the game Consider two models $\boldsymbol{M}$ and $\boldsymbol{N}$. A player called "duplicator" claims that $\boldsymbol{M}$ and $\boldsymbol{N}$ are similar, while a player called "spoiler" maintains that they are different. Players agree on some finite number k of rounds for the game, the severity of the probe.

Definition 15.2 Comparison games
A *model comparison* game works as follows, packing two moves into one round. *Spoiler* (also written $\boldsymbol{S}$ for brevity) chooses one of the models, and picks an object d in its domain. *Duplicator* (also written $\boldsymbol{D}$ for brevity) then chooses an object e in the other model, and the pair (d, e) is added to the current list of matched objects. After k rounds, the object matching is inspected. If it is a partial isomorphism, duplicator wins; otherwise, spoiler does. Here, a "partial isomorphism" is an injective partial map f between models $\boldsymbol{M}$ and $\boldsymbol{N}$ that is an isomorphism between its own domain and range. ■

This alternating schedule $(DS)^*$ occurs in many games. We now present some sample plays of our comparison games. As in Chapter 14, players may lose, even when they have a winning strategy. We use a language with a binary relation symbol R only, mostly disregarding identity atoms $=$.

EXAMPLE 15.2 Comparing integers and rationals
The linear orders of integers $\mathbb{Z}$ and rationals $\mathbb{Q}$ have different first-order properties: the latter is dense, the former discrete. Here is how this will surface in the game:

$$\begin{array}{ccccccc} & -1 & 0 & +1 & & & \\ \cdots & \circ & \circ & \circ & \cdots & & \mathbb{Z} \\ & & & & & & \\ \cdots\cdots\cdots & \circ\cdots\cdots & \circ\cdots\cdots & \circ\cdots\cdots & \cdots & & \mathbb{Q} \\ & -1 & 0 & +1 & & & \end{array}$$

By choosing objects well, duplicator has a winning strategy for the game over two rounds. But spoiler can always win the game in three rounds. Here is a typical play:

Round 1	***S*** chooses 0 in $\mathbb{Z}$	***D*** chooses 0 in $\mathbb{Q}$
Round 2	***S*** chooses 1 in $\mathbb{Z}$	***D*** chooses 1/3 in $\mathbb{Q}$
Round 3	***S*** chooses 1/5 in $\mathbb{Q}$	any response for ***D*** is losing

These moves will convey the typical strategic flavor of the game. ■

Difference formulas and spoiler's strategies In playing the games, winning strategies for spoiler are correlated with first-order formulas φ that bring out a difference between the models. The correlation is tight. The quantifier syntax of φ triggers the moves for spoiler.

EXAMPLE 15.2, CONTINUED Exploiting definable differences
Spoiler can use the first-order definition of density for a binary order, written as $\forall x \forall y(x < y \rightarrow \exists z(x < z \wedge z < y))$, to distinguish $\mathbb{Q}$ from $\mathbb{Z}$. We spell this out, to show how there is an almost algorithmic derivation of a strategy from a first-order difference formula. For convenience, we use existential quantifiers only. The idea is for spoiler to maintain a difference between the two models, of stepwise decreasing syntactic depth. Spoiler starts by observing that

$$\exists x \exists y(x < y \wedge \neg \exists z(x < z \wedge z < y)) \text{ is true in } \mathbb{Z}\text{, but false in } \mathbb{Q} \qquad \#$$

Spoiler then chooses an integer d for $\exists x$, making $\exists y(d < y \wedge \neg \exists z(d < z \wedge z < y))$ true in $\mathbb{Z}$. Now duplicator can take any rational number d' in $\mathbb{Q}$: the first-order formula $\exists y(d' < y \wedge \neg\exists z(d' < z \wedge z < y))$ will be false for it, by #:

$$\mathbb{Z} \models \exists y(d < y \wedge \neg\exists z(d < z \wedge z < y)), \text{ not } \mathbb{Q} \models \exists y(d' < y \wedge \neg\exists z(d' < z \wedge z < y))$$

In the second round, spoiler continues with a witness e for the new outermost quantifier $\forall y$ in the true existential formula in $\mathbb{Z}$: making $d < e \wedge \neg\exists z(d < z \wedge z < e)$ true there. Again, whatever object e' duplicator now picks as a response in $\mathbb{Q}$, the formula $d' < e' \wedge \neg\exists z(d' < z \wedge z < e')$ will be false there. In the third round, spoiler analyzes the mismatch in truth value. If duplicator kept $d' <' e$ true in $\mathbb{Q}$, then, since $\neg\exists z(d < z \wedge z < e)$ holds in $\mathbb{Z}$, $\exists z(d' < z \wedge z < e')$ holds in $\mathbb{Q}$. Spoiler then switches to $\mathbb{Q}$, chooses a witness for the existential formula, and wins. ■

Thus, even the right model switches for the strategy of spoiler are encoded in the difference formulas. Such switches are mandatory whenever there is a syntactic change from one type of outermost quantifier (existential, universal) to another.[180]

15.3 Adequacy and strategies

As with evaluation games, the interesting information is in players' strategies. In the results to follow, we think of winning strategies for duplicator, although spoiler's strategic point of view will return later. For the sake of brevity, let us write $WIN(\boldsymbol{D}, \boldsymbol{M}, \boldsymbol{N}, k)$ for: "duplicator has a winning strategy against spoiler in the k-round comparison game between the models $\boldsymbol{M}$ and $\boldsymbol{N}$."

Comparison games can start from any initial match of objects in $\boldsymbol{M}$ and $\boldsymbol{N}$. In particular, if models have distinguished objects named by individual constants, these are matched automatically. In the proofs to come, we think of all initial matches in the latter way. We now look at an analogue of the success lemma from Chapter 14.

THEOREM 15.1 For all models $\boldsymbol{M}$ and $\boldsymbol{N}$, and $k \in \mathbb{N}$, the following two assertions are equivalent:

(a) $WIN(\boldsymbol{D}, \boldsymbol{M}, \boldsymbol{N}, k)$: duplicator has a winning strategy in the k-round game.

(b) $\boldsymbol{M}$ and $\boldsymbol{N}$ agree on all first-order sentences up to quantifier depth k.

180 Our examples may also suggest a correlation: "winning strategy for spoiler over n rounds $\sim$ difference formula with n quantifiers altogether." But as we shall soon see, the right measure is different, being the maximum length of a quantifier nesting in a formula.

This improves on the Isomorphism Lemma in two ways. Our adequacy result matches up a language-dependent and a language-independent comparison relation. And it provides fine structure not available before, which helps in applications.

Proof From (a) to (b) is an induction on k. We start with the base step. With 0 rounds, the initial match of objects must have been a partial isomorphism for $\boldsymbol{D}$ to win. So $\boldsymbol{M}$ and $\boldsymbol{N}$ agree on all atomic sentences, and hence on their Boolean combinations, the formulas of quantifier depth 0. We proceed with the inductive step. The inductive hypothesis says that, for any two models, if $\boldsymbol{D}$ can win their comparison game over k rounds, the models agree on all first-order sentences up to quantifier depth k. Now let $\boldsymbol{D}$ have a winning strategy for the $k+1$ round game on $\boldsymbol{M}$ and $\boldsymbol{N}$. Consider any first-order sentence φ of quantifier depth $k+1$. Such a φ is equivalent to a Boolean combination of (i) atoms, (ii) sentences of the form $\exists x\psi$, with ψ of quantifier depth at most k. Thus, it suffices to show that $\boldsymbol{M}$ and $\boldsymbol{N}$ agree on the latter forms.

The essential case is this. Let $\boldsymbol{M} \models \exists x\psi$. Then for some object d, we get $\boldsymbol{M}, d \models \psi$. Think of $(\boldsymbol{M}, d)$ as an expanded model with a distinguished object d to which we assign a new name $\underline{d}$. In this way $(\boldsymbol{M}, d)$ verifies the sentence $\varphi(\underline{d})$. Now, $\boldsymbol{D}$'s winning strategy has a response for whatever $\boldsymbol{S}$ can do in the $k+1$-round game. For instance, let $\boldsymbol{S}$ start with $\boldsymbol{M}$ and object d. Then $\boldsymbol{D}$ has a response e in $\boldsymbol{N}$ to this move such that $\boldsymbol{D}$'s remaining strategy still gives a win in the k-round game played from the given link $d - e$. This yields an expanded model $(\boldsymbol{N}, e)$, with e as its interpretation of the name $\underline{d}$. The remainder is an ordinary k-round game starting from the models $(\boldsymbol{M}, d)$ and $(\boldsymbol{N}, e)$. By the inductive hypothesis, these models agree on all sentences up to quantifier depth k: and hence also on $\varphi(\underline{d})$. Therefore, $\boldsymbol{N}, e \models \varphi(\underline{d})$, and so $\boldsymbol{N} \models \exists x\psi$.

The converse direction from (b) to (a) requires another induction on k. This time we need a small auxiliary result about first-order logic in a finite relational vocabulary, the so-called Finiteness Lemma.

LEMMA Fix variables $x_1, \ldots, x_m$. Up to logical equivalence, there are only finitely many first-order formulas $\varphi(x_1, \ldots, x_m)$ of quantifier depth $\leq k$.[181]

181 The proof is by induction on k, analyzing formulas of quantifier depth $k+1$ in the same way as above, and then using the fact that Boolean combinations of any finite set of formulas are finite modulo logical equivalence.

Now we do the inductive proof from (b) to (a). The base step is trivial: doing nothing is a winning strategy for $\boldsymbol{D}$. As for the inductive step, we give the first move in $\boldsymbol{D}$'s strategy. Let $\boldsymbol{S}$ choose one of the models, say $\boldsymbol{M}$, plus some object d in it. Now, $\boldsymbol{D}$ looks at the set of first-order formulas true of d in $\boldsymbol{M}$, which may refer to distinguished objects available through their names in the language. By the Finiteness Lemma, this set is finite modulo logical equivalence, and so one existential formula $\exists x\psi^d$ true in $\boldsymbol{M}$ summarizes all this information. Now, because the models $\boldsymbol{M}$ and $\boldsymbol{N}$ agree on all first-order sentences of depth $k+1$, and $\exists x\psi^d$ is such a sentence, it also holds in $\boldsymbol{N}$. Therefore, $\boldsymbol{D}$ can choose a witness e for it in $\boldsymbol{N}$. Then the expanded models $(\boldsymbol{M}, d)$, $(\boldsymbol{N}, e)$ agree on all sentences up to quantifier depth k, and so, by the inductive hypothesis, $\boldsymbol{D}$ has a winning strategy in the remaining k-round game between them. $\boldsymbol{D}$'s initial response plus the latter further strategy gives $\boldsymbol{D}$ an overall winning strategy over $k+1$ rounds. ■

15.4 An explicit version: The logic content of strategies

Theorem 15.1 still leaves out the precise match we found earlier between spoiler's winning strategies and first-order formulas. Thus, it displays a phenomenon that we discussed in Chapter 4, under the heading of ∃-*sickness*. In logic (but also elsewhere), one often rushes to formulating notions and results with an existential quantifier when more constructive information would be available if we made the witnesses for that existential quantifier explicit. The symptoms of this disease are overuse of indefinite articles such as "a" or modal affixes such as "-ility." Why have a theory of lovability when we can have one of love?

Fortunately, ∃-sickness can often be cured with a little further effort:[182] The following result is the earlier adequacy theorem made explicit.

Theorem 15.2 There exists an explicit correspondence between

(a) Winning strategies for $\boldsymbol{S}$ in the k-round comparison game for $\boldsymbol{M}$ and $\boldsymbol{N}$.

(b) First-order sentences φ of quantifier depth k with $\boldsymbol{M} \models \varphi$, not $\boldsymbol{N} \models \varphi$.

---

182 Another strain of the disease occurs in standard completeness theorems, that link provability to validity, instead of seeking a more direct match between proofs and semantic verifications. Remedies include the full completeness theorems discussed in Chapter 20.

Proof We first look at the direction from (b) to (a). Every φ of quantifier depth k induces a winning strategy for $\boldsymbol{S}$ in a k-round game between any two models. Each round $k-m$ starts with a match between objects linked so far that differ on some subformula ψ of φ with quantifier depth $k-m$. By Boolean analysis, $\boldsymbol{S}$ then finds some existential subformula $\exists x \bullet \alpha$ of ψ with a matrix formula α of quantifier depth $k-m-1$ on which the models disagree. $\boldsymbol{S}$'s next choice is a witness in that model of the two where $\exists x \bullet \alpha$ holds.

Our next direction is from (a) to (b). Any winning strategy σ for $\boldsymbol{S}$ induces a distinguishing formula of proper quantifier depth. To obtain this, let $\boldsymbol{S}$ make the first choice d in model $\boldsymbol{M}$ according to σ, and now write down an existential quantifier for that object. Our formula will be true in $\boldsymbol{M}$ and false in $\boldsymbol{N}$. We know that each choice of $\boldsymbol{D}$ for an object e in $\boldsymbol{N}$ gives a winning position for $\boldsymbol{S}$ in all remaining $(k-1)$-round games starting from an initial match $d-e$. By the inductive hypothesis, these induce distinguishing formulas of depth $k-1$. By the Finiteness Lemma, only finitely many such formulas exist. Some of these will be true in $\boldsymbol{M}$ (say $A_1, \ldots, A_r$), and others in $\boldsymbol{N}$ (say $B_1, \ldots, B_s$). The total difference formula for strategy σ is then the $\boldsymbol{M}$-true assertion

$$\exists x \bullet (A_1 \wedge \cdots \wedge A_r \wedge \neg B_1 \wedge \cdots \wedge \neg B_s)$$

whose appropriateness is easy to check. ■

Thus, spoiler's winning strategies in a comparison game correspond to formulas, that is, logical objects of independent interest.[183] A similar match exists for the other player. One might call duplicator's strategies analogies of some finite quality measured by the number k. Technically they are cut-off versions of the "potential isomorphisms" that will be defined in Section 15.6.

Remark Explicitness versus computability
We gave an explicit definable match of strategies with other objects, but it need not be computable. Also, strategies in evaluation or comparison games need not be effective. They range from history-free (with all next moves read off from the current state) to dependent on a complete record of the game so far. The strategic invariants in the next section illustrate the kind of memory to be maintained.[184]

183 We have a caveat. The formulation of Theorem 15.2 is still $\exists$-sick. Can you cure it?

184 On the more computational side, Chapter 18 will present an important theorem on the adequacy of history-free strategies in parity games.

15.5 The games in practice: Invariants and special model classes

In practice, comparison games involve not just logic, but also combinatorial analysis of the models involved. Facts 15.1 through 15.3 provide some examples.

FACT 15.3 The rationals $(\mathbb{Q}, <)$ are elementarily equivalent to the reals $(\mathbb{R}, <)$.

It suffices to show that duplicator can win the comparison game for every k. A good method is to identify an *invariant* for duplicator to maintain at intermediate game states. In this particular case, the invariant is simply that all matches so far form a finite partial isomorphism. All further choices of spoiler can then be countered using the unboundedness and density of the orders. More complicated invariants may depend on the number of rounds still to go.

FACT 15.4 $(\mathbb{N}, <)$ is elementarily equivalent with $(\mathbb{N} + \mathbb{Z}, <)$.

This time, if the length of the game is known in advance, duplicator can counter choices of spoiler from the supernatural numbers in $\mathbb{Z}$ by matching them with large natural numbers in $\mathbb{N}$.[185]

Invariants are concrete descriptions of positions where players have a winning strategy. Some descriptions of solutions in the game logics of Part I had this character, witness our brief discussion in Chapter 4.

Finally, comparison games also work on model classes where standard methods of first-order logic fail. Fact 15.3 gives an example of such a negative use of games.

FACT 15.5 Even or odd are not first-order definable on the finite models.

The usual proof for this nondefinability on all models is a compactness argument that fails on finite models. But now, using games, suppose that even size had a first-order definition on finite models, of quantifier depth k. Then any two finite models for which duplicator can win the k-round comparison game are both of even size, or both of odd size. But this is refuted by any two finite models with k versus $k + 1$ objects in their domains.

---

185 The invariant maintains suitable distances between objects. Duplicator makes sure that with k rounds to go, the two sequences $d_1, \ldots, d_m$ in $\mathbb{N}$ and $e_1, \ldots, e_m$ in $\mathbb{N} + \mathbb{Z}$ chosen so far have the following properties: (a) $d_i < d_j$ iff $e_i < e_j$, (b) if d_i, d_j have distance $< 2^k - 1$, then $\mathit{distance}(e_i, e_j) = \mathit{distance}(d_i, d_j)$; else, d_i, d_j and e_i, e_j both have distance $\geq 2^k - 1$ (finite or infinite).

15.6 Game theory: Determinacy, finite and infinite games

Comparison games are two-player zero-sum games of some finite depth k. Therefore, Zermelo's Theorem applies, and either duplicator or spoiler has a winning strategy.

FACT 15.6 Model comparison games are determined.

But these games also have a natural version that goes on forever, say over ω rounds. A natural winning convention in that case is the safety property that duplicator wins the infinite game by not losing at any finite stage, maintaining a partial isomorphism all the time. This is stronger than being able to win all finite-round games. With this understanding, $\mathbb{N}$ and $\mathbb{N} + \mathbb{Z}$ can be distinguished by spoiler in an infinite game: it suffices to start with a supernatural number and keep descending. But when comparing $\mathbb{Q}$ with $\mathbb{R}$, duplicator can hold out indefinitely.

These games still fall under earlier results from Chapter 5. In particular, the winning set for spoiler is open (failure of partial isomorphism always occurs by some finite stage), and hence the Gale-Stewart Theorem applies.

FACT 15.7 The infinite comparison game is determined.

For infinite games, duplicator's winning strategies do correspond to a notion of independent interest.

DEFINITION 15.3 Potential isomorphism
A *potential isomorphism* between two models $\boldsymbol{M}$ and $\boldsymbol{N}$ is a non-empty family I of finite partial isomorphisms between $\boldsymbol{M}$ and $\boldsymbol{N}$ satisfying the following back-and-forth property:

(a) If $f \in I$ and $a \in \boldsymbol{M}$, then there exists a $b \in \boldsymbol{N}$ with $f \cup \{(a, b)\} \in I$.

(b) If $f \in I$ and $b \in \boldsymbol{N}$, then there exists an $a \in \boldsymbol{M}$ with $f \cup \{(a, b)\} \in I$.

This is like a bisimulation, but now for much richer non-modal languages. ■

FACT 15.8 The potential isomorphisms between two models correspond to duplicator's winning strategies in the infinite comparison game.

By contrast, in infinite games, spoiler's winning strategies are methods blocking each attempt at establishing potential isomorphism by some finite stage, guided by a difference formula.

Potential isomorphism implies elementary equivalence. If duplicator can win the infinite game, then duplicator can win every finite cut-off, and the success lemma applies. But the models $\mathbb{N}$ and $\mathbb{N}+\mathbb{Z}$ refuted the converse. It is easy to see directly that the partial isomorphisms in a potential isomorphism I satisfy the same first-order formulas, even with infinite conjunctions and disjunctions added. In fact, Karp's Theorem says that two models are potentially isomorphic iff they satisfy the same sentences in infinitary first-order logic.[186]

15.7 Modifications and extensions

Model comparison games capture a wide variety of logics. In this section we explore some illustrations.

Modal games Restricting players' choices of objects to local successors of currently matched objects leads to basic modal languages, and a link with the notion of bisimulation in Chapter 1.[187] The back-and-forth clauses in a bisimulation between two models strongly suggest a game where one player mentions a challenge, letting one process make an available move, while the other player must then respond with an appropriate simulating move. This might go on forever, but there is also a natural finite variant restricting the number of rounds. More precisely, the fine structure of bisimulation suggests the following games between duplicator and spoiler, comparing successive pairs (m, n) in two models $\boldsymbol{M}$ and $\boldsymbol{N}$.

DEFINITION 15.4 Bisimulation game
Fix a finite number of rounds. In each round of the *bisimulation game*, spoiler chooses a state x in one model that is a successor of the current m or n, and duplicator responds with a matching successor y in the other model. Spoiler wins if, at any stage, x and y differ in atomic properties, or if duplicator cannot choose a matching successor. Infinite bisimulation games have no finite bound, while all other conventions remain the same. ■

186 More on infinitary first-order logic games is found in Barwise & van Benthem (1999), where they are used to prove new kinds of interpolation theorems, as well as a Lindström-type characterization for infinitary modal logic in terms of bisimulation.

187 A more general case of the approach is the Guarded Fragment (Andréka et al. 1998).

Example 15.3 Modal comparison games

Spoiler can win the game between the models depicted below (cf. Example 1.3) starting from their roots:

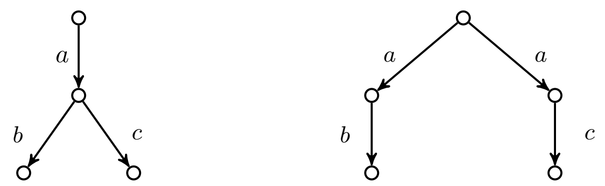

Spoiler needs two rounds, and different strategies do the job. One stays on the left, exploiting the modal difference of depth 2, with three existential modalities:

$$\langle a\rangle(\langle b\rangle\top \wedge \langle c\rangle\top)$$

Another strategy switches models, using the smaller formula

$$[a]\langle b\rangle\top$$

where the type of modality switches between universal and existential. ■

A success lemma can be proved for the finite bisimulation game like for first-order logic (cf. van Benthem 2010b). Analyzing the games further, the following two relevant observations emerge.

Fact 15.9

(a) Spoiler's winning strategies in a k-round game between $(\boldsymbol{M}, s)$ and $(\boldsymbol{N}, t)$ match the modal formulas of operator depth k on which s and t disagree.

(b) Duplicator's winning strategies in an infinite game between $(\boldsymbol{M}, s)$ and $(\boldsymbol{N}, t)$ match the bisimulations between $\boldsymbol{M}$ and $\boldsymbol{N}$ that link s to t.

Clause (b) reveals the close connection between our games and bisimulations.

Pebble games One can also add structure to the games, for instance, in the way that players operate. For instance, to make memory a concern, one can let objects be chosen only by using a finite resource that has been supplied to the players, marking them with one of a finite set of pebbles (Immerman & Kozen 1989). In this setting, duplicator has a winning strategy for the n-round k-pebble game between two models $\boldsymbol{M}$ and $\boldsymbol{N}$ iff $\boldsymbol{M}$ and $\boldsymbol{N}$ agree on all first-order sentences of quantifier depth $\leq k$ that use at most the variables $x_1, \ldots, x_m$ (free or bound), a so-called "finite-variable fragment."[188]

Other languages Other comparison games capture first-order logic with generalized quantifiers (Keenan & Westerståhl 1997), or the first-order fixed point logic LFP(FO) of Chapter 14. This may raise new perspectives. For LFP(FO), for instance, it is not known whether there exists a model comparison game that would be more analogous to the elegant fixed point evaluation game of Chapter 14.

15.8 Connections between logic games

Now that we have seen two major logic games, one for evaluation and one for comparison, questions arise of general architecture, and connections between games. We end with three suggestive observations.

Parallel game operations Comparison games suggest new operations on games in addition to the earlier choices and switch: most obviously, parallel composition. As we will see in Chapter 18, model comparison games are interleaved evaluation games. A general study of parallel game operations will be found in Chapter 20.

Model comparison as evaluation Comparison games sometimes reduce to evaluation games. Through their definition in Chapter 1, bisimulations E may be viewed as non-empty greatest fixed points for a first-order operator between models $\boldsymbol{M}$ and $\boldsymbol{N}$ defined by the equation:

$$\mathsf{E}xy \leftrightarrow \bigwedge_P \big((Px \leftrightarrow Py) \wedge \forall z(R_a xz \to \exists u(R_a yu \wedge \mathsf{E}zu)) \wedge \forall u(R_a yu \to \exists z(R_a xz \wedge \mathsf{E}zu))\big)$$

with the right-hand side taken over all atomic predicates P and actions R_a.

Thus, existence of a bisimulation between states s and t amounts to the truth of some LFP(FO) formula in a suitable disjoint sum model $\boldsymbol{M} + \boldsymbol{N}$. Such a formula can be checked by the fixed point evaluation game of Chapter 14. The latter can be infinite; but so can model comparison games. We will highlight the broader significance of facts such as this in Chapter 18.

---

188 This result is proved in Immerman & Kozen (1989). As an important genre of further results, these authors also show that three pebbles suffice as a working memory for winning all comparison games over linear orders. Finite-variable fragments in general play an important role in finite model theory (cf. Ebbinghaus & Flum 1999, Libkin 2004).

Game equivalence The literature often switches between supposedly equivalent formulations without explanation. For instance, Barwise & van Benthem (1999) define the following comparison game starting with a finite partial isomorphism between two models:

> In each round, duplicator selects a set F^+ of partial isomorphisms satisfying the back-and-forth property that, for every object a in one model, there is an object b in the other with $f \cup \{(a, b)\} \in F^+$; and vice versa. In the same round, spoiler then chooses a match in F^+ again, and so on.

In each round, duplicator offers spoiler a complete panorama of all choices that spoiler could make in the former game, plus duplicator's own responses to them. Spoiler then makes a choice of both spoiler's own move as well as duplicator's prepackaged response, setting the new stage. The two games are obviously power-equivalent in the sense of Chapters 11 and 19, but their internal structure also matches closely. Indeed, the above transformation exemplifies a general turn switch such as we saw in the Thompson transformations that we discussed in Chapter 11.

Chapter 18 has further discussion of equivalence levels for logic games.

15.9 Conclusion

Comparison games are a concrete and powerful way of thinking about the interplay of logic and structure. By now they are widely used for many purposes, and the reader will have understood their appeal, while we have also proved their basic properties. Besides being successful special logical activities, comparison games also demonstrate interesting general features, as we have seen. Just to mention one of these, they perform a striking new sort of parallel combination of evaluation games in different structures.

Further, the games of this chapter have a direct impact on the game logics in Parts I and II of this book. As we already noted in Chapter 1 when discussing invariance, playing comparison games offers a systematic way of adding fine structure to existing notions of simulation between processes or games, revealing further information about invariants. Thus, comparison games can be a useful tool in the general study of the right levels at which to analyze general games, and the design of their languages. In line with the integrative spirit of this book, they can be used as games about games.

15.10 Literature

Two highly accessible textbooks are Doets (1996) and Väänänen (2011). Kolaitis (2001) is an excellent presentation geared more toward computer science.