

A Game Semantics for System P

Johannes Marti and Riccardo Pinasio

February 15, 2016

Abstract

In this paper we introduce a game semantics for System P, one of the most studied axiomatic systems for non-monotonic reasoning, conditional logic and belief revision. We prove soundness and completeness of the game semantics with respect to the rules of System P, and show that an inference is valid with respect to the game semantics if and only if it is valid with respect to the standard order semantics of System P. Combining these two results leads to a new completeness proof for System P with respect to its order semantics. Our approach allows us to construct for every inference either a concrete proof of the inference from the rules in System P or a countermodel in the order semantics.

Our results rely on the notion of a witnessing set for an inference, whose existence is a concise, necessary and sufficient condition for validity of an inferences in System P.

We also introduce an infinitary variant of System P and use the game semantics to show its completeness for the restricted class of well-founded orders.

1 Introduction

System P is an inference system which formalizes core principles of non-monotonic consequence relations as studied in artificial intelligence [8]. It is also the non-nested fragment of a conditional logic developed in philosophy and linguistics [10, 3, 15].

The standard semantics for System P is based on orders and evaluates a non-monotonic inference or conditional by minimization in the order. A similar order semantics is also used in the theory of AGM belief revision [5], which can be recast in the setting of conditional logic [1].

In this paper we introduce a game semantics for the validity of inferences in System P. The study of logical systems with game-theoretic methods was initiated independently by Lorenzen and Lorenz [11] and Hintikka [6]. Hintikka's approach, known as game theoretic semantics, uses a game to establish the truth of a formula in a given model. Lorenzen and Lorenz developed what is known as dialogical logic. A dialogical game is a game in which two players debate the validity of an inference in a logical system. The main difference between dialogical games and game theoretic semantics is that Lorenzen and Lorenz adopt a proof-theoretical perspective whereas Hintikka presupposes a model theory for the logic. A comparison of the two approaches can be found in [13].

The game semantics for System P developed in this paper is close to Lorenzen and Lorenz' dialogical games. For every inference in System P we define a game

in which the first player, we call her Héloïse, attempts to argue for the validity of the inference against attacks by the second player, whom we call Abélard. The game differs from dialogical games in that Abélard successively chooses from a given domain of objects or possible worlds, whose properties are already fixed. This setup allows us to focus on the semantics of the conditional, since it is already determined in advance whether an object or world instantiates a given Boolean combination of properties.

As an illustration of the game semantics consider the following example:

Example 1. Assume that Héloïse and Abélard have agreed that birds normally fly. Héloïse is now claiming that penguins normally fly. Abélard disputes this. They are having the following dialogue:

ABÉLARD: Look at Pingu! He is a penguin but he doesn't fly.
 HÉLOÏSE: But Pingu is not a good example. He is a bird that doesn't fly and we have agreed that birds normally fly.
 ABÉLARD: Yes, but Pingu is a totally fine penguin. He is just a strange bird. When we talk about birds we would rather think of a sparrow, like Tweety. Tweety can fly.
 HÉLOÏSE: [*has nothing to say*]

In this dialogue Abélard tries to disprove the inference to the conclusion that penguins usually fly by providing an object refuting this conclusion. Héloïse tries to circumvent Abélard's alleged counterexample by pointing out that it does not conform to an agreed upon premise. Abélard defends the counterexample by showing that it is an exception to the premise. He does this by presenting a different object which conforms to the premise and which he claims to be more normal than the object provided before. Héloïse loses since she does not have a premise to dispute the normality of the second object given by Abélard.

The main result of this paper is a proof that for every inference the following statements are equivalent:

1. The inference has a formal proof in System P.
2. The inference is valid in the order semantics.
3. The inference is valid in the game semantics.
4. There exists a witnessing set for the inference.

The notion of a witnessing set mentioned in statement 4 is introduced in this paper. Checking for the existence of a witnessing set is a straightforward method to determine the validity of an inference in System P.

We prove two different instances of the equivalence above, where the details in the four statements vary. One instance is Theorem 19, which concerns the standard System P and validity on the class of all orders and yields a new proof of the completeness result for System P. This result has already been obtained in [8] and for nested conditional logic in [3, 15]. Our approach is however more constructive in that we transform a winning strategy for Héloïse into a concrete proof in System P and a winning strategy for Abélard into a counterexample in the order semantics. The other instance is Theorem 17, which on the syntactic side concerns an extension of System P with an infinitary proof rule introduced

in this paper, and on the semantic side concerns validity on the class of well-founded orders.

The paper is organized as follows. In Section 2 we review the basics of System P and its order semantics. In Section 3 we introduce our game semantics for System P and its infinitary variant. In Section 4 we show that winning strategies for Abélard correspond to counterexamples in the order semantics. In Section 5 we show that winning strategies for Héloïse correspond to proofs in System P. To obtain this result we introduce the notion of a witnessing set. In section 6 we prove compactness for the version of the game semantics that corresponds to provability in standard System P. In Section 7 we put the results from the previous sections together to prove the main theorems of this paper.

2 System P

In this section we introduce the version of System P used in this paper and its standard order semantics.

Let W be any set, whose elements we call *possible worlds*. One might also think of the elements in W as objects, as for instance in Example 1. A *conditional* over W is a pair $(A, C) \in \mathcal{P}W \times \mathcal{P}W$ of subsets of W , written as $A \vdash C$. We call A the *antecedent* and C the *consequent* of the conditional $A \vdash C$. A world w *verifies* a conditional $A \vdash C$ if $w \in A \cap C$. A world w is a *counterexample* to or *falsifies* $A \vdash C$ if $w \in A - C = A \cap C^c$. We write $A - C$ for the set difference of A and C and $C^c = W - C$ for the complement of C relative to W .

Two sets Σ and Γ of conditionals over W are called an *inference*, written as Σ/Γ . We call the elements of Σ the *premises* of the inference and the elements of Γ the *conclusions* of the inference. The conclusions in Γ are understood disjunctively. If $\Gamma = \{A \vdash C\}$ is a singleton set we write $\Sigma/A \vdash C$ instead of Σ/Γ . We also call an inference Σ/Γ a *single-conclusion inference* if Γ is a singleton set, and a *multi-conclusion inference* if we want to stress that Γ need not be a singleton. In this paper we focus on single-conclusion inferences because, as shown in Corollary 20, the completeness result for multi-conclusion inferences follows from that for single-conclusion inferences.

Note the distinction between an inference $\Sigma/A \vdash C$ between conditionals and the non-monotonic inference $A \vdash C$ captured by one conditional. Inferences between conditionals have a classical monotonic semantics, whereas conditionals have a non-monotonic semantics on orders.

System P consists of the rules given in Figure 1. A *proof* of an inference Σ/Γ in system P is a tree which is built by recursive application of the rules of system P , such that the root of the tree is a conclusion $A \vdash C \in \Gamma$ and all leaves which are not instances of (Id) are premises in Σ .

The presentation of system P in [8] includes an additional cut rule, which is shown in Lemma 5.3 to be derivable from the other rules in the system. In the setting of conditional logic an axiomatization analogous to that in Figure 1 was given by [15].

The system P_∞ is obtained from system P by adding the following infinitary rule:

$$\frac{A \vdash C_i \quad \text{for all } i \in I}{A \vdash \bigcap \{C_i \mid i \in I\}} \text{ (And}^\infty\text{)}$$

In our presentation of System P conditionals are over sets of worlds rather

$$\begin{array}{c}
\overline{A \vdash A} \text{ (Id)} \\
\\
\frac{A \vdash C \quad C \subseteq D}{A \vdash D} \text{ (RW)} \qquad \frac{A \vdash C \quad A \vdash D}{A \vdash C \cap D} \text{ (And)} \\
\\
\frac{A \vdash B \quad A \vdash C}{A \cap B \vdash C} \text{ (CM)} \qquad \frac{A \vdash C \quad B \vdash C}{A \cup B \vdash C} \text{ (Or)}
\end{array}$$

Figure 1: System P

than formulas. This simplifies the development of the game semantics in Section 3 and dispenses us from including a rule for replacing logically equivalent antecedents in conditionals, like the rule (Left Logical Equivalence) in [8]. One can recast conditionals over formulas as conditionals over sets of worlds by taking the worlds to be all the maximal consistent sets of formulas, and identifying a formula with the set of maximal consistent sets containing it.

The standard semantics for System P uses posets over sets of worlds, which are reflexive, transitive and antisymmetric relations. All the results of this paper are stated in terms of posets. We however never use antisymmetry, thus our results could be generalized to preorders, which are transitive and reflexive relations.

Let $A \vdash C$ be a conditional over a set of worlds W and $P = (W, \leq)$ a poset with carrier W . The conditional $A \vdash C$ holds in P , written as $P \models A \vdash C$, if it satisfies the following semantic clause:

$$A \vdash C \quad \text{iff} \quad \text{for all } w \in A \text{ there is a } v \in A \text{ with } v \leq w \\
\text{such that } u \in C \text{ for all } u \leq v \text{ with } u \in A.$$

In the context of conditional logic this semantic clause has been given in [3] and is a generalization of the clause in [10, p. 48]. It was later introduced to the setting of non-monotonic consequence relations by [2].

The semantic clause for conditionals is simpler on well-founded orders. A poset $P = (W, \leq)$ is *well-founded* if there is no infinite chain $w_1 \geq w_2 \geq \dots$ where the inequalities are strict, meaning that $w_1 \not\leq w_2 \not\leq \dots$. One can show that for a well-founded poset $P = (W, \leq)$ the above semantic clause is equivalent to

$$A \vdash C \quad \text{iff} \quad w \in C \text{ for all } w \in \text{Min}_{\leq}(A),$$

where $\text{Min}_{\leq}(A) \subseteq W$ is the set of minimal elements of A in \leq , that is

$$\text{Min}_{\leq}(A) = \{m \in A \mid \text{if } w \leq m \text{ then } m \leq w \text{ for all } w \in A\}.$$

This semantic clause on well-founded orders is intuitive if one thinks of the order such that $w \leq v$ if w is more normal than v . The conditional $A \vdash C$ holds in such an order if the most normal instances of A are instances of C .

An inference Σ/Γ , where premises and conclusions are over the set W , is *valid on posets* if for every poset $P = (U, \leq)$ and function $f : U \rightarrow W$ there is

an $A \vdash C \in \Gamma$ such that $P \models f^{-1}[A] \vdash f^{-1}[C]$ whenever $P \models f^{-1}[B] \vdash f^{-1}[D]$ for all $B \vdash D \in \Sigma$.

An inference Σ/Γ , where premises and conclusions are over the set W , is *valid on well-founded posets* if for every well-founded poset $P = (U, \leq)$ and function $f : U \rightarrow W$ there is an $A \vdash C \in \Gamma$ such that $P \models f^{-1}[A] \vdash f^{-1}[C]$ whenever $P \models f^{-1}[B] \vdash f^{-1}[D]$ for all $B \vdash D \in \Sigma$.

In the above definitions of validity the carrier U of the poset P is possibly distinct from the set of worlds W over which the relevant conditionals were defined. The two sets U and W are related by the function f , which we call the *labeling function*. We discuss why labeling functions are needed in Remark 5 at the end of the section.

The following proposition states the soundness of System P for its order semantics.

Proposition 2. *If an inference is provable in system P then it is valid on posets. If an inference is provable in system P_∞ then it is valid on well-founded posets.*

Proof. A routine induction on the complexity of the proof trees. □

We continue with a model-theoretic construction used in Section 7. This construction is used in the proof of Lemma 3 in [9] and is also essential to the complexity results of [4].

Proposition 3. *Let I be an index set and take a poset P_i for every $i \in I$. Consider the disjoint sum $S = \coprod_{i \in I} P_i$ which results from placing all the different P_i next to each other without adding any order that is not already present in some P_i . We write $\iota_i : W_i \rightarrow U$ for the inclusion map of P_i into S . Then it holds for all $A, C \subseteq U$ that*

$$S \models A \vdash C \quad \text{iff} \quad P_i \models \iota_i^{-1}[A] \vdash \iota_i^{-1}[C] \text{ for all } i \in I.$$

Proof. The proof is a routine argument unfolding the semantic clause of \vdash and using the definition of the disjoint sum. □

As observed in [4], one obtains the following disjunction property:

Corollary 4. *A multi-conclusion inference Σ/Γ is valid on posets iff there exists an $A \vdash C \in \Gamma$ such that $\Sigma/A \vdash C$ is valid on posets. The same property holds for validity on well-founded posets.*

Proof. First note that the direction from right-to-left follows immediately from the definition of validity.

We prove the contrapositive of the left-to-right direction. Assume that $\Sigma/A \vdash C$ is not valid on posets for every $A \vdash C \in \Gamma$. Then there is for every $A \vdash C \in \Gamma$ a poset $P_{A,C} = (U_{A,C}, \leq_{A,C})$ and a function $f_{A,C} : U_{A,C} \rightarrow W$ such that $P_{A,C} \models f_{A,C}^{-1}[B] \vdash f_{A,C}^{-1}[D]$ for all $B \vdash D \in \Sigma$ but $P_{A,C} \not\models f_{A,C}^{-1}[A] \vdash f_{A,C}^{-1}[C]$.

We define $P = (U, \leq)$ to be the disjoint sum of all the $P_{A,C}$ for $A \vdash C \in \Gamma$ with inclusions $\iota_{A,C} : U_{A,C} \rightarrow U$. We let $f : U \rightarrow W$ be the unique function such that $f \circ \iota_{A,C} = f_{A,C}$ for all $A \vdash C \in \Gamma$. By Proposition 3 we know that for all $E, F \subseteq W$

$$P \models f^{-1}[E] \vdash f^{-1}[F] \quad \text{iff} \quad P_{A,C} \models f_{A,C}^{-1}[E] \vdash f_{A,C}^{-1}[F] \text{ for all } A \vdash C \in \Gamma. \quad (1)$$

We show that P together with the labeling function f is a counterexample to the validity of Σ/Γ on posets.

That $P \models f^{-1}[B] \vdash f^{-1}[D]$ for all $B \vdash D \in \Sigma$ follows from (1) because $P_{A,C} \models f_{A,C}^{-1}[B] \vdash f_{A,C}^{-1}[D]$ for all $B \vdash D \in \Sigma$ and $A \vdash C \in \Gamma$.

That $P \not\models f^{-1}[A] \vdash f^{-1}[C]$ for all $A \vdash C \in \Gamma$ follows from (1) because $P_{A,C} \not\models f_{A,C}^{-1}[A] \vdash f_{A,C}^{-1}[C]$ for all $A \vdash C \in \Gamma$.

The same construction works in the case of well-founded posets because the disjoint sum of well-founded posets is well-founded. \square

In the following remark we explain why labeling functions are needed in the definition of validity:

Remark 5. The intuitive notion of validity for an inference with premises and conclusions over a set W quantifies only over posets with carrier W . This notion of validity is equivalent to the definition of validity given above when the labeling function is required to be bijective. We sketch an argument showing that this intuitive notion of validity cannot be complete with respect to provability in system P .

We first show that we need to allow for labeling functions that are not surjective. It is easy to see that if $A \subseteq W$ is not empty, then for every surjective function $f : U \rightarrow W$ and every poset (U, \leq) the conditional $f^{-1}[A] \vdash f^{-1}[\emptyset]$ does not hold on (U, \leq) . Hence, the inference

$$\frac{A \vdash \emptyset \quad A \neq \emptyset}{B \vdash D}$$

is valid for the notion of validity with surjective labeling function. However, this inference is not provable in system P ; one can show this using the completeness result of Theorem 19 for the notion of validity with arbitrary labeling functions.

We now argue that we also need to allow for labeling functions that are not injective. An argument for this claim can be found at the end of Section 5.2 on page 193 of [8]. We show here how their reasoning applies to our context. In particular, we provide a multi-conclusion inference that is valid according to the notion of validity with injective labeling functions, but is not provable in system P . Let W be the set $\{x, y, z\}$. Consider the following multi-conclusion inference Σ/Γ with conditionals over W :

$$\frac{\{x, y, z\} \vdash \{y, z\}}{\{x\} \vdash \emptyset \quad \{x, y\} \vdash \{y\} \quad \{x, z\} \vdash \{z\}}$$

One can show that this multi-conclusion inference is valid for the notion of validity where the labeling function is required to be injective. It is however not valid for labeling functions that are not injective. This is witnessed by the poset $P = (U, \leq)$ where U is the set $\{x_0, x_1, y, z\}$ and \leq is as follows:

$$\begin{array}{cc} x_0 & x_1 \\ | & | \\ y & z \end{array}$$

The labeling function $f : U \rightarrow W$ is defined as $x_0, x_1 \mapsto x, y \mapsto y, z \mapsto z$. Hence it follows by Corollary 20 that the inference Σ/Γ is not provable in system P .

One might wonder what rules have to be added to system P to recover completeness with respect to validity with injective labeling functions. The discussion of the coherence condition in Section II.4.1 of [15] suggests that this requires expressive power that goes beyond the language of conditionals.

3 The game

In this section we introduce the semantic games for System P, along with the needed game theoretic notions. We first give the definitions and then provide intuitive explanations and examples.

For every inference $\Sigma/A \vdash C$ we define two games, the *non-well-founded game* $\mathcal{G}_{A,C}^\Sigma$ and the *well-founded game* $\mathcal{F}_{A,C}^\Sigma$. The difference between $\mathcal{G}_{A,C}^\Sigma$ and $\mathcal{F}_{A,C}^\Sigma$ is that the former is about the validity of the inference $\Sigma/A \vdash C$ on posets and its provability in system P , whereas the latter is about the validity of the inference on well-founded posets and its provability in the infinitary system P_∞ . The definitions of $\mathcal{G}_{A,C}^\Sigma$ and $\mathcal{F}_{A,C}^\Sigma$ are mostly the same. Thus we simultaneously define both games for a fixed inference $\Sigma/A \vdash C$ and mention the differences between the two games explicitly.

One can consider the game for an inference $\Sigma/A \vdash C$ to be a graph whose nodes are the *positions* in the game and whose edges are the possible *moves*. Additionally every position is labeled with the player, Héloïse or Abélard, who has to move at that position. A concise specification of this graph is given by the following table:

Pos.	Player	Moves
(R, F)	Abélard	$\{(w, F) \mid w \in R - F\}$
(w, F)	Héloïse	$\{(B \cap D, F \cup (B - D)) \mid B \vdash D \in \Sigma, w \in B - D\} \cup \{(A - C, F) \mid w \in A \cap C\}$

The game contains two types of positions. The first type consists of all pairs of the form $(R, F) \in \mathcal{P}W \times \mathcal{P}W$, where W is the set of worlds that contains the antecedents and consequents of all conditionals in the inference $\Sigma/A \vdash C$. For such a position (R, F) we call R the *required area* and F the *forbidden area* of the position. The second type of positions consists of all pairs of the form $(w, F) \in W \times \mathcal{P}W$. For such a position (w, F) we call w the world of the position and F its forbidden area.

The positions of the form (R, F) belong to Abélard, meaning that he moves next if the game is at such a position. He can choose any world $w \in R - F$ that is in the required area and not in the forbidden area. In this case we say that Abélard plays the move w . When Abélard chooses the world w the game moves to the position (w, F) , where the forbidden area F remains the same.

The positions of the form (w, F) belong to Héloïse. She can play any premise $B \vdash D \in \Sigma$ which is such that $w \in B - D$. This moves the game to the position $(B \cap D, F \cup (B - D))$. In the special case when $w \in A \cap C$, where $A \vdash C$ is the conclusion of the inference $\Sigma/A \vdash C$ that we are playing for, Héloïse has one special move at her disposal which moves the game to the position $(A - C, F)$. We call this move \star and we say that Héloïse plays \star if she chooses this move.

The position $(A - C, \emptyset)$ is defined to be the *starting position* of the game.

A *play* s starting from a position p is a finite sequence $s = p_0, p_1, \dots, p_n$ of positions of the game such that $p_0 = p$ and there is a move from p_i to p_{i+1} for

every $i \in \{0, \dots, n-1\}$. If we call s a *play*, without explicitly mentioning the starting position, then s is assumed to start from the starting position of the game. For two plays s and t starting from the same position we write $t \leq s$ if s is an initial segment of t . It is convenient in the setting of this paper to let \leq be the converse of the initial segment relation, and not the initial segment relation itself. The set of all plays starting from some position is a possibly infinite tree ordered by the converse initial segment relation.

A *match starting from* a position p is a maximal, possibly infinite, branch in the tree of all plays starting from p . If the starting position p is omitted we again assume it to be the starting position of the game.

Note that a finite match is just a play p_0, \dots, p_n such that there is no move in the game leading from p_n to any other position. A player *gets stuck* in the finite match p_0, \dots, p_n if she or he has to move in the position p_n . A finite match is *won* by the player that does not get stuck.

There is no natural way to define the winner of infinite matches. In the game $\mathcal{G}_{A,C}^\Sigma$ the winner of all infinite matches is stipulated to be Abélard. In the game $\mathcal{F}_{A,C}^\Sigma$ the winner of all infinite matches is stipulated to be Héloïse. The winning condition for infinite matches is the only difference between the definitions of $\mathcal{G}_{A,C}^\Sigma$ and $\mathcal{F}_{A,C}^\Sigma$.

A *strategy* S for some player starting from position p is a set of plays starting from p such that:

1. S is closed under initial segments of plays.
2. Whenever a play sp ending in a position p belonging to the player is in S then there is a unique position q reachable from p by a move of the player such that spq is also in S .
3. Whenever a play sp ending in a position p of the opponent is in S then all plays spq such that there is some move of the opponent from p to q are also in S .

If we do not mention the position from which a strategy starts then it should be understood as starting from the starting position of the game.

A strategy for some player starting from some position p is a recipe which tells the player how to continue playing once the game is in position p . The second condition above guarantees that the strategy determines one unique move for the player in every match starting from p in which she or he plays according to the strategy. The third condition guarantees that the strategy covers all possible moves of the opponent.

A strategy S for some player starting from p is a subtree of the tree of all plays starting from p . This subtree is somewhat peculiar since all the nodes ending with a position belonging to the player to whom the strategy belongs have only one child. For this reason we define the *tree* $T = (L, \leq)$ *determined by some strategy* S for some player starting from some position to consist of all plays in S that end in a position that belongs to the opponent, ordered by the converse initial segment relation.

One can check that given a strategy S for Héloïse starting from a position p and a strategy S' for Abélard starting from p there is a unique match starting from p in which Héloïse plays according to S and Abélard plays according to

S' . This match is the unique branch in the subtree of the tree of all plays that corresponds to the intersection of S and S' .

A strategy S for a player starting from p is a *winning strategy* if, for every strategy S' for the opponent starting from p , the unique match in which the player plays according to S and the opponent plays according to S' is winning for the player. Less formally, a strategy for some player is winning if the player wins every match in which he or she plays according to the strategy.

An inference $\Sigma/A \vdash C$ is defined to be *valid in the non-well-founded game semantics* if Héloïse has a winning strategy in the game $\mathcal{G}_{A,C}^\Sigma$. The inferences is defined to be *valid in the well-founded game semantics* if Héloïse has a winning strategy in the game $\mathcal{F}_{A,C}^\Sigma$.

We make use of the fact that winning strategies for some player starting at a given position can be glued together from winning strategies for that player starting from later positions. Assume the position p belongs to the player and there is a move from p to a position q such that the player has a winning strategy S starting from q . Then there is a new winning strategy S' for the player starting from p where the player first moves from p to the position q and then continues playing according to the winning strategy S . We can do a similar thing if the position p belongs to the opponent. Assume that for every move of the opponent leading from p to a position q there is a winning strategy S_q for the player starting from q . Then there is a new winning strategy S' for the player starting from p where she or he waits for the opponent to move to a position q and from then on uses the winning strategy S_q .

The games $\mathcal{G}_{A,C}^\Sigma$ and $\mathcal{F}_{A,C}^\Sigma$ enjoy the property of *determinacy*. Determinacy means that for any such game exactly one of the players has a winning strategy. The determinacy of $\mathcal{G}_{A,C}^\Sigma$ and $\mathcal{F}_{A,C}^\Sigma$ follows from the Gale-Stewart Theorem, a proof of which can be found for instance in [14, sec. 3.5].

We now provide an intuitive explanation of the rules of the game. Consider the game $\mathcal{G}_{A,C}^\Sigma$ or $\mathcal{F}_{A,C}^\Sigma$ for an inference $\Sigma/A \vdash C$. The conditionals in Σ are accepted by both players in advance. Abélard aims to show that conclusion $A \vdash C$ does not follow from Σ by giving a counterexample. Héloïse disputes the relevance of the counterexample by forcing Abélard to admit that there is a more relevant world that verifies the conclusion of the inference. To do so she chooses premises in Σ forcing Abélard to come up with further worlds that verify the chosen premises.

The game starts at position $(A - C, \emptyset)$ which means that in his first move Abélard has to provide a counterexample to the conclusion $A \vdash C$.

After Abélard has chosen some world w the match is at position (w, F) . Héloïse can now show that this world is not relevant by finding a premise $B \vdash D \in \Sigma$ that is falsified by w . This moves the match to the position $(B \cap D, F \cup (B - D))$. Abélard is now required to come up with a more relevant world that verifies the premise $B \vdash D$. Moreover the forbidden area grows such that Abélard is not allowed to choose counterexamples to $B \vdash D$ anymore.

If Abélard moves to a position (w, F) such that w verifies the conclusion $A \vdash C$ then Héloïse can play \star . This moves the game to $(A - C, F)$ forcing Abélard to restart with a new counterexample to the conclusion. It now becomes more difficult for Abélard to find the required worlds because the forbidden area F has grown to include all counterexamples to premises that Héloïse played before.

Let us consider again the dialogue from Example 1 in the introduction.

Formally this dialogue corresponds to the game for the inference $B \vdash F / P \vdash F$ such that there are at least two worlds $p, t \in W$ with $p, t \in B$, $p \in P$, $t \notin P$, $p \notin F$ and $t \in F$. The game starts at the position $(P - F, \emptyset)$. This means that Abélard has to give a counterexample to the conditional $P \vdash F$. He chooses p moving the game to the position (p, \emptyset) . Héloïse tries to neutralize this alleged counterexample by pointing out that it does not conform to the premise $B \vdash F$. So the match moves to the position $(B \cap F, B - F)$. Now Abélard has to come up with an object that is more normal than p in that it verifies $B \vdash F$. He chooses t moving the game to $(t, B - F)$. His counterexample p is thus left intact because t is not in P and hence does not verify $P \vdash F$. Moreover there is no premise that is falsified by t , which means that Héloïse cannot argue that t is not normal in some respect. Hence Abélard wins the match.

We now consider two examples of games that belong to more abstract inferences. It is quite helpful to illustrate the dialectics of simple games by means of Venn diagrams.

Example 6. Consider the following cautious cut rule which is part of the axiomatization of System P in [8]:

$$\frac{A \vdash B \quad A \cap B \vdash C}{A \vdash C} \text{ (CCut)}$$

Take $A, B, C \subseteq W$ to be any subsets of some set of worlds W . Héloïse has a winning strategy in the non-well-founded game for this inference. The strategy can be described as follows. Héloïse's move only depends on the last world w that Abélard has played. If $w \in A \cap B^c$ then Héloïse plays the first premise $A \vdash B$. If $w \in A \cap B \cap C^c$ then she replies with the second premise $A \cap B \vdash C$. If $w \in A \cap B \cap C$ then Héloïse replies with \star . We do not need to specify a reply for the case where $w \in A^c$ because it can be seen to never arise if Héloïse plays the above strategy. One can check that by playing this strategy Héloïse does not get stuck and no infinite match can arise. Hence it is a winning strategy for Héloïse.

The following classical cut rule is however not valid in general:

$$\frac{A \vdash B \quad B \vdash C}{A \vdash C} \text{ (Cut)}$$

Assuming that there are worlds $w_0 \in A \cap B \cap C^c$ and $w_1 \in A^c \cap B \cap C$ we can show that Abélard has a winning strategy in the non-well-founded game for (Cut). In the starting position $(A - C, \emptyset)$ Abélard plays w_0 . Héloïse can then only reply with the second premise $B \vdash C$ moving the match to the position $(B \cap C, B - C)$. Then Abélard picks w_1 and Héloïse gets stuck. Note that the absence of A in the antecedent of the second premise is crucial, since it allows Abélard to escape to a world not in A .

Example 7. This last example demonstrates the difference between the non-well-founded and the well-founded game. Consider the following inference where A_i for all $i \in I$ and C are subsets of any set of worlds W :

$$\frac{A_i \vdash C \quad \text{for all } i \in I}{\bigcup \{A_i \mid i \in I\} \vdash C} \text{ (Or}^\infty\text{)}$$

Héloïse has a winning strategy in the well-founded game for this inference. Whenever Abélard has chosen a world in A_i for some $i \in I$ then Héloïse plays

the premise $A_i \vdash C$. It is guaranteed that all worlds picked by Abélard are in A_i for some $i \in I$. This holds at the beginning of the match because the required area of the starting position is $\bigcup\{A_i \mid i \in I\}$. It remains true later because after Héloïse has played $A_i \vdash C$ Abélard needs to pick a world in $A_i \cap C$ and after she has played \star Abélard needs to pick a world in $\bigcup\{A_i \mid i \in I\} - C$. Playing this strategy Héloïse never gets stuck. So this is a winning strategy for her in the well-founded game for the inference because by definition she wins all infinite matches.

Héloïse does not necessarily have a winning strategy in the non-well-founded game associated to (Or^∞) . To see this consider the following instance of the rule over the set $W = \omega \cup \{\infty\}$ where ω is the set of all natural numbers and ∞ is distinct from all elements in ω :

$$\frac{\{n, \infty\} \vdash \{\infty\} \quad \text{for all } n \in \omega}{\omega \cup \{\infty\} \vdash \{\infty\}}$$

Abélard has a winning strategy in this game because he can enforce an infinite match. If the required area equals ω then Abélard picks the smallest $n \in \omega$ that is not in the forbidden area. Such a world exists because the forbidden area is the finite set containing all k such that Héloïse has played $\{k, \infty\} \vdash \{\infty\}$ before. If the required area is $\{\infty\}$ then Abélard plays ∞ which is never in the forbidden area. In this way Abélard never gets stuck because he has a move for all of his positions. In the starting position and after Héloïse has played \star the required area is ω . After Héloïse has played one of the premises the required area is $\{\infty\}$. A match played according to this strategy is either infinite or a match where Héloïse gets stuck. So it is a winning strategy for Abélard in the non-well-founded game.

4 Abélard orders

In this section we show that a winning strategy for Abélard in the game for some inference yields a countermodel to the inference in the order semantics.

Proposition 8. *If Abélard has a winning strategy in $\mathcal{G}_{A,C}^\Sigma$ then the inference $\Sigma/A \vdash C$ is not valid on posets. If Abélard has a winning strategy in $\mathcal{F}_{A,C}^\Sigma$ then the inference $\Sigma/A \vdash C$ is not valid on well-founded posets.*

Proof. Assume that Abélard has a winning strategy in $\mathcal{G}_{A,C}^\Sigma$. We show that the tree $T = (L, \leq)$ determined by Abélard's winning strategy is a countermodel to the validity of the inference $\Sigma/A \vdash C$. The labeling function $f : L \rightarrow W$ is defined to map a play $s, (w, F)$ ending with a position (w, F) for Héloïse to the world $w \in W$ previously picked by Abélard.

We have to show that $T \models f^{-1}[B] \vdash f^{-1}[D]$ for all $B \vdash D \in \Sigma$ and that $T \not\models f^{-1}[A] \vdash f^{-1}[C]$. To show the former take any $B \vdash D \in \Sigma$ and consider $s \in L$ such that $s \in f^{-1}[B]$. We must find a $t \leq s$ such $t \in f^{-1}[B]$ and for all $u \leq t$ with $u \in f^{-1}[B]$ it holds that $u \in f^{-1}[D]$. We distinguish cases on whether $s \in f^{-1}[D]$.

First assume that $s \notin f^{-1}[D]$. Let $w = f(s)$ be the world from the last position (w, F) in the play s . Because $s \in f^{-1}[B]$ and by assumption $s \notin f^{-1}[D]$ it follows that $w \in B - D$. Then Héloïse can play the premise $B \vdash D$ moving the game to the position $(B \cap D, F \cup (B - D))$. Abélard replies with some

world $v \in B \cap D$ moving to the position $(v, F \cup (B - D))$. This gives a play $t = s, (B \cap D, F \cup (B - D)), (v, F \cup (B - D))$ which is in L because the last position is a position for Héloïse. It is the case that $t \in f^{-1}[B]$ because $f(t) = v$. It remains to be shown that for any $u \leq t$, if $u \in f^{-1}[B]$ then $u \in f^{-1}[D]$. Take any such u whose last position is $(f(u), F')$. Since $u \leq t$, the position $(f(u), F')$ either occurs later in the match than $(v, F \cup (B - D))$ or it is equal to it. Since in a match the forbidden area never decreases, we have that $F \cup (B - D) \subseteq F'$, and therefore $B - D \subseteq F'$. Because Abélard picked $f(u)$ when the forbidden area was F' it follows that $f(u) \notin F'$ and thus $f(u) \notin B - D$. Hence if $u \in f^{-1}[B]$ then also $u \in f^{-1}[D]$.

Consider now the case when $s \in f^{-1}[D]$. We distinguish two further cases. If $r \in f^{-1}[D]$ for all $r \leq s$ with $r \in f^{-1}[B]$ then we can take s to be the witnessing t to satisfy the semantic clause. If on the other hand there is some $r \leq s$ with $r \in f^{-1}[B]$ but $r \notin f^{-1}[D]$ then we can run the argument from the previous paragraph with r for s to find the witnessing $t \leq r \leq s$ such that $t \in f^{-1}[B]$ and for all $u \leq t$ if $u \in f^{-1}[B]$ then $u \in f^{-1}[D]$.

It remains to be shown that $T \not\models f^{-1}[A] \vdash f^{-1}[C]$. We must find an $s \in f^{-1}[A]$ such that for every $t \leq s$ with $t \in f^{-1}[A]$ there exists a $u \leq t$ with $u \in f^{-1}[A]$ but $u \notin f^{-1}[C]$. First note that $f^{-1}[A] \neq \emptyset$, otherwise Abélard would get stuck in the initial position $(A - C, \emptyset)$. Consider then any $s \in f^{-1}[A]$. Take an arbitrary $t \leq s$ such that $t \in f^{-1}[A]$. If $t \notin f^{-1}[C]$ then we can take $u = t$. If $t \in f^{-1}[C]$ then t ends with the position (v, F) where $v \in A \cap C$. Héloïse can reply to this position with \star , to which Abélard's winning strategy must supply a world $z \in A - C$. This moves the game to a position (z, F) , and hence we have a play $u \leq t$ ending with position (z, F) . Thus $f(u) \in A - C$, and so $u \in f^{-1}[A]$ but $u \notin f^{-1}[C]$. \square

Example 9. Consider the following instance of (Or^∞) from Example 7:

$$\frac{\{n, \infty\} \vdash \{\infty\} \quad \text{for all } n \in \omega}{\omega \cup \{\infty\} \vdash \{\infty\}}$$

In Example 7 we show that Abélard has a winning strategy in the non-well-founded game for this inference. The construction from the proof of Proposition 8 transforms this winning strategy into a non-well-founded poset that is a counterexample to the inference above. This poset is displayed in Figure 2.

The converse of Proposition 8 is a consequence of Theorems 17 and 19. One could also show it by a direct proof. The idea is that Abélard obtains a winning strategy in the game for some inference by moving downwards in a poset that is a counterexample to the inference.

5 Héloïse proves

In this section we prove that if Héloïse has a winning strategy in the game for some inference then the inference is provable in System P. An intermediate step of the proof is to show the existence of a witnessing set for the inference, which is defined as follows:

Definition 10. Let Γ be a set of conditionals. The *abnormality area* $U(\Gamma) \subseteq W$ of Γ is defined as:

$$U(\Gamma) = \bigcup \{B - D \mid B \vdash D \in \Gamma\}.$$




Figure 2: Counterexample to (Or^∞)

The set Γ is a *witnessing set* for an inference $\Sigma/A \vdash C$ if $\Gamma \subseteq \Sigma$ and the following conditions are satisfied:

1. $A \subseteq C \cup U(\Gamma)$.
2. $B \cap D \subseteq (A \cap C) \cup U(\Gamma)$ for all $B \vdash D \in \Gamma$.

Proposition 11. *If Héloïse has a winning strategy in the game $\mathcal{F}_{A,C}^\Sigma$ then the inference $\Sigma/A \vdash C$ has a witnessing set.*

Proof. Assume that Héloïse has a winning strategy in the game $\mathcal{F}_{A,C}^\Sigma$. We define the witnessing set $\Gamma \subseteq \Sigma$ as the set of all premises that Héloïse uses in some match played according to her winning strategy. We check that the two conditions from Definition 10 are satisfied.

For the first condition consider any $w \in A$. If $w \in C \subseteq C \cup U(\Gamma)$ we are done. So we can assume that $w \notin C$. Since the game starts in position $(A - C, \emptyset)$ Abélard can then move to (w, \emptyset) in his first move. Héloïse's winning strategy provides her with a reply to this move. The reply cannot be \star because $w \notin A \cap C$. Hence she replies with a premise $B \vdash D \in \Gamma$ from which it follows that $w \in B - D \subseteq U(\Gamma) \subseteq C \cup U(\Gamma)$.

For the second condition consider any $B \vdash D \in \Gamma$. By definition of Γ there is a match played according to Héloïse's winning strategy in which she plays the premise $B \vdash D$. This moves the match to a position of the form $(B \cap D, F)$, where $F \subseteq U(\Gamma)$, because the forbidden area F at a position is the union of the $B' - D'$ of all the premises $B' \vdash D'$ that Héloïse has played so far in the match. Now consider a $w \in B \cap D$ and distinguish cases on whether $w \in F$. If $w \in F$ then also $w \in U(\Gamma)$. Otherwise Abélard can reply with w at the position $(B \cap D, F)$. Héloïse's winning strategy provides her with a reply to this move. If she plays \star then $w \in A \cap C$. If she plays a premise $B \vdash D \in \Gamma$ then $w \in B - D \subseteq U(\Gamma)$. \square

There is also a short direct proof of the converse of Proposition 11. We do not give it here because the converse also follows from Theorem 17.

We now show how one can obtain a formal proof of an inference in System P from a witnessing set for this inference.

Proposition 12. *If the inference $\Sigma/A \vdash C$ has a witnessing set then it is provable in system P_∞ . If the inference $\Sigma/A \vdash C$ has a finite witnessing set then it is provable in system P .*

Proof. We first show that a witnessing set yields a proof in system P_∞ . Thus assume that there is a witnessing set $\Gamma \subseteq \Sigma$ for the inference $\Sigma/A \vdash C$. We index the elements in Γ such that $\Gamma = \{B_i \vdash D_i \mid i \in I\} \subseteq \Sigma$ for some set I . That Γ is a witnessing set means that:

1. $A \subseteq C \cup U(\Gamma)$.
2. $B_i \cap D_i \subseteq (A \cap C) \cup U(\Gamma)$ for all $i \in I$.

We need two consequences of these inclusions.

The first consequence is

$$\bigcup_{i \in I} B_i \cap \bigcap_{i \in I} (B_i^c \cup D_i) \subseteq A \cap C. \quad (2)$$

To see this pick any $w \in \bigcup_{i \in I} B_i \cap \bigcap_{i \in I} (B_i^c \cup D_i)$. Since $w \in \bigcup_{i \in I} B_i$ there exists an $j \in I$ such that $w \in B_j$. Then also $w \in D_j$ because $w \in \bigcap_{i \in I} (B_i^c \cup D_i)$ and hence $w \in B_j^c \cup D_j$. Thus $w \in B_j \cap D_j \subseteq (A \cap C) \cup U(\Gamma)$ by condition 2. It follows that $w \in A \cap C$ because $w \in \bigcap_{i \in I} (B_i^c \cup D_i) = U(\Gamma)^c$.

The second consequence is

$$\left(C \cup \bigcup_{i \in I} B_i \right) \cap A = A. \quad (3)$$

The \subseteq -inclusion is obvious. For the other inclusion we need that $A \subseteq C \cup \bigcup_{i \in I} B_i$. This holds because condition 1 is that $A \subseteq C \cup U(\Gamma)$ and one can verify that $U(\Gamma) \subseteq \bigcup_{i \in I} B_i$.

We now construct the proof of $\Sigma/A \vdash C$. For every $j \in I$ we have the following proof:

$$\frac{\frac{B_j \vdash D_j}{B_j \vdash B_j^c \cup D_j} \text{ (RW)} \quad \frac{\frac{\frac{\bigcup_{i \in I} B_i \cap B_j^c \vdash \bigcup_{i \in I} B_i \cap B_j^c} \text{ (Id)}}{\bigcup_{i \in I} B_i \cap B_j^c \vdash B_j^c \cup D_j} \text{ (RW)}}{\bigcup_{i \in I} B_i \cap B_j^c \vdash B_j^c \cup D_j} \text{ (Or)}}{B_j \cup (\bigcup_{i \in I} B_i \cap B_j^c) \vdash B_j^c \cup D_j}$$

One can check that $B_j \cup (\bigcup_{i \in I} B_i \cap B_j^c) = \bigcup_{i \in I} B_i$. So we have for every $j \in I$ a proof of $\bigcup_{i \in I} B_i \vdash B_j^c \cup D_j$ from premises in Σ . We continue as follows:

$$\frac{\frac{\frac{\bigcup_{i \in I} B_i \vdash \bigcup_{i \in I} B_i} \text{ (Id)}}{\bigcup_{i \in I} B_i \vdash \bigcup_{i \in I} B_i} \text{ (Id)} \quad \frac{\{\bigcup_{i \in I} B_i \vdash B_j^c \cup D_j \mid j \in I\}}{\bigcup_{i \in I} B_i \vdash \bigcap_{j \in I} (B_j^c \cup D_j)} \text{ (And}^\infty\text{)}}{\frac{\bigcup_{i \in I} B_i \vdash \bigcup_{i \in I} B_i \cap \bigcap_{j \in I} (B_j^c \cup D_j)}{\bigcup_{i \in I} B_i \vdash A \cap C} \text{ (RW)}} \text{ (And)}$$

The last application of (RW) is possible because of (2). We use this proof twice to continue as follows:

$$\frac{\frac{\frac{C \cap A \vdash C \cap A}{C \cap A \vdash C} \text{ (RW)}}{\frac{C \cap A \vdash C \cap A}{C \cap A \vdash C} \text{ (Id)}} \quad \frac{\frac{\frac{\bigcup_{i \in I} B_i \vdash A \cap C}{\bigcup_{i \in I} B_i \vdash A} \text{ (RW)}}{\bigcup_{i \in I} B_i \vdash A \cap C} \text{ (RW)}}{\frac{\bigcup_{i \in I} B_i \vdash A \cap C}{\bigcup_{i \in I} B_i \vdash C} \text{ (RW)}} \text{ (CM)}}{\frac{\bigcup_{i \in I} B_i \cap A \vdash C}{(C \cup \bigcup_{i \in I} B_i) \cap A \vdash C} \text{ (Or)}}$$

In the last application of (Or) we use the distributivity of \cap over \cup . By (3) we now have a proof of $A \vdash C$ from premises in Σ .

The same construction can be used to show the second claim. If Σ is finite then so is the index set I , which allows us to replace the application of (And $^\infty$) above by a finite chain of applications of (And). \square

Example 13. Consider again the following cautious cut rule:

$$\frac{A \vdash B \quad A \cap B \vdash C}{A \vdash C} \text{ (CCut)}$$

Using Propositions 11 and 12 we can construct a proof of (CCut) in system P from the winning strategy of Example 6. Both premises $A \vdash C$ and $A \cap B \vdash C$ are played by Héloïse in some match according to this winning strategy. From the proof of Proposition 11 it follows that together they are a witnessing set for the inference (CCut). We can apply the proof of Proposition 12 to this witnessing set to obtain a proof of the inference in system P . After deleting some obvious redundancies we obtain a proof of (CCut) starting as follows:

$$\frac{\frac{\frac{A \cap B \vdash C}{A \cap B \vdash (A \cap B)^c \cup C} \text{ (RW)}}{\frac{A \cap B \vdash C}{A \cap B \vdash (A \cap B)^c \cup C} \text{ (Id)}} \quad \frac{\frac{\frac{A \cap (A \cap B)^c \vdash A \cap (A \cap B)^c}{A \cap (A \cap B)^c \vdash (A \cap B)^c \cup C} \text{ (RW)}}{\frac{A \cap (A \cap B)^c \vdash (A \cap B)^c \cup C} \text{ (Or)}}}{A \vdash (A \cap B)^c \cup C}$$

The above is the rightmost leaf of the following continuation:

$$\frac{\frac{\frac{A \vdash A}{A \vdash A} \text{ (Id)}}{\frac{A \vdash A}{A \vdash A} \text{ (Id)}} \quad \frac{\frac{\frac{\frac{A \vdash B}{A \vdash A^c \cup B} \text{ (RW)}}{A \vdash (A^c \cup B) \cap ((A \cap B)^c \cup C)} \text{ (And)}}{\frac{A \vdash (A \cap B)^c \cup C}{A \vdash (A \cap B)^c \cup C} \text{ (And)}}}{\frac{A \vdash A \cap (A^c \cup B) \cap ((A \cap B)^c \cup C)}{A \vdash C} \text{ (RW)}}$$

The converse of Proposition 12 can be shown by an induction on the complexity of the proof tree of an inference. It also follows from Theorems 17 and 19.

6 Compactness

In this section we show that the semantics of the non-well-founded games is compact. For this we need the following notion:

Definition 14. A subalgebra \mathfrak{A} of \mathcal{PW} is compact if for all elements A of \mathfrak{A} and sets of elements \mathcal{B} of \mathfrak{A} such that $A \subseteq \bigcup \mathcal{B}$ there is a finite $\mathcal{B}' \subseteq \mathcal{B}$ such that $A \subseteq \bigcup \mathcal{B}'$.

Compact subalgebras arise naturally when working with conditionals over formulas in propositional logic. The subalgebra of all sets of maximal consistent sets of formulas containing a given formula is a compact subalgebra of the powerset algebra over the set of all maximal consistent sets of formulas.

Theorem 15. *Let \mathfrak{A} be a compact subalgebra of \mathcal{PW} . If $\Sigma \cup \{A \vdash C\} \subseteq \{B \vdash D \mid B, D \in \mathfrak{A}\}$ and Héloïse has a winning strategy in $\mathcal{G}_{A,C}^\Sigma$ then there is a finite $\Sigma' \subseteq \Sigma$ such that Héloïse has a winning strategy in $\mathcal{G}_{A,C}^{\Sigma'}$.*

Proof. We consider the tree $T = (L, \leq)$ determined by the winning strategy of Héloïse in $\mathcal{G}_{A,C}^\Sigma$ starting from the position $(A - C, \emptyset)$. Let $<$ on L be the strict version of the converse initial segment relation \leq meaning that $t < s$ if $t \leq s$ and not $s \leq t$. The relation $<$ on L is well-founded. If this was not the case then there would be an infinite match in T contradicting the claim that T is the tree of a winning strategy for Héloïse in the game $\mathcal{G}_{A,C}^\Sigma$.

The proof is an induction on the well-founded relation $<$. This means that we show a claim about all elements of L by showing that the claim holds for some $s \in L$ whenever the claim holds for all $t \in L$ with $t < s$. The claim which we show by induction is that Héloïse has a winning strategy in the game $\mathcal{G}_{A,C}^\Gamma$ for a finite $\Gamma \subseteq \Sigma$ starting from the position (R, F) , where (R, F) is the last position in the play s and $R, F \in \mathfrak{A}$. The statement of the theorem then follows from taking s to be the play consisting just of the starting position $(A - C, \emptyset)$ of the game $\mathcal{G}_{A,C}^\Sigma$.

So suppose we have a play $s \in L$ with last position (R, F) where $R, F \in \mathfrak{A}$ such that the claim holds for all $t < s$. From the position (R, F) Abélard can move to any position (w, F) where $w \in R - F$. Héloïse's winning strategy provides a reply $r_w \in \Sigma \cup \{\star\}$ for any such $w \in R - F$. After the reply r_w the game moves to the position p_w such that

$$p_w = \begin{cases} (B \cap D, F \cup (B - D)) & \text{if } r_w = B \vdash D \in \Sigma, \\ (A - C, F) & \text{if } r_w = \star. \end{cases}$$

For every $w \in R - F$ we thus obtain a new play $t_w = s, (w, F), p_w \in L$ to which the induction hypothesis applies because $t_w < s$. This means that for every $w \in R - F$ there is a finite $\Gamma_w \subseteq \Sigma$ such that Héloïse has a winning strategy for the game $\mathcal{G}_{A,C}^{\Gamma_w}$ starting from the position p_w .

We define the area K_w for any $w \in R - F$ such that $K_w = B - D$ if $r_w = B \vdash D$ and $K_w = A \cap C$ if $r_w = \star$. By the rules of game we have that $w \in K_w$ for every $w \in R - F$. Thus we obtain the covering

$$R - F \subseteq \bigcup \{K_w \mid w \in R - F\}.$$

Since all the involved propositions are in \mathfrak{A} it follows by compactness of \mathfrak{A} that there is a finite subcover

$$R - F \subseteq \bigcup \{K_v \mid v \in V\}, \quad (4)$$

where $V \subseteq R - F$ is a finite set.

Define the set

$$\Gamma = \bigcup_{v \in V} \Gamma_v \cup \{B \vdash D \mid B \vdash D = r_v \text{ for some } v \in V\}.$$

The set $\Gamma \subseteq \Sigma$ is finite because Γ_v is finite for every $v \in V$ and V is finite. Note that Héloïse's winning strategy starting from position p_v in $\mathcal{G}_{A,C}^{\Gamma_v}$ is also a winning strategy starting from p_v in the game $\mathcal{G}_{A,C}^{\Gamma}$, since $\Gamma_v \subseteq \Gamma$ and so any move available to her in $\mathcal{G}_{A,C}^{\Gamma_v}$ is also available to her in $\mathcal{G}_{A,C}^{\Gamma}$.

We show that Héloïse has a winning strategy in the game $\mathcal{G}_{A,C}^{\Gamma}$ starting from the position (R, F) . We need to specify a move for Héloïse in any position (w, F) such that $w \in R - F$ is a possible move of Abélard in (R, F) . Consider any such $w \in R - F$. By (4) there is a $v \in V$ such that $w \in K_v$. Hence we can let Héloïse reply r_v in the position (w, F) . This moves the game into position p_v . From there on Héloïse plays according to her winning strategy starting from p_v in the game $\mathcal{G}_{A,C}^{\Gamma}$. \square

Example 16. The compactness result for the non-well-founded game for an inference fails if the antecedents and consequents of the inference are not from a compact subalgebra. To see this consider the following infinitary inference over the set of worlds $W = \{a, b\} \cup \omega$ where we assume a and b to be distinct from each other and from any element of ω :

$$\frac{\{\{a\} \cup \omega \vdash \omega\} \cup \{\{n, b\} \vdash \{b\} \mid n \in \omega\}}{\{a, b\} \vdash \{b\}}$$

Héloïse's winning strategy in the non-well-founded game for this inference is as follows. The starting position is $(\{a, b\} - \{b\}, \emptyset)$. Abélard's first move must thus be a , to which Héloïse replies with the premise $\{a\} \cup \omega \vdash \omega$. This moves the game to the position $(\omega, \{a\})$. Abélard now needs to choose some number $n \in \omega$. Héloïse can then play the corresponding premise $\{n, b\} \vdash \{b\}$ forcing him to play b . Héloïse answers this move with \star and Abélard gets stuck because a is in the forbidden area. In all matches played according to this strategy Abélard gets stuck after a finite number of moves. Hence it is a winning strategy for Héloïse in the non-well-founded game.

There is however no finite subset of the premises for which Héloïse still has a winning strategy. To see this assume that we are playing the game for an inference with the same conclusion as the inference above but only a finite subset of its premises. Then for some $n \in \omega$ the premise $\{n, b\} \vdash \{b\}$ is missing. In this case Abélard has a winning strategy. He starts the match with a . If the premise $\{a\} \cup \omega \vdash \omega$ is not in the finite set of premises available to Héloïse then she loses immediately. Otherwise she plays this premise and forces Abélard to pick a number in ω . He then chooses the number $n \in \omega$ such that $\{n, b\} \vdash \{b\}$ is missing. Now Héloïse is stuck.

7 Completeness of the order semantics

In this section we prove the main theorems of this paper.

For the system P_∞ we obtain the following result:

Theorem 17. *The following are equivalent:*

1. There is a proof of the inference $\Sigma/A \vdash C$ in system P_∞ .
2. The inference $\Sigma/A \vdash C$ is valid on well-founded posets.
3. Héloïse has a winning strategy in the game $\mathcal{F}_{A,C}^\Sigma$.
4. There is a witnessing set for the inference $\Sigma/A \vdash C$.

Proof. The implication from 1 to 2 is given by Proposition 2.

The implication from 2 to 3 is by contraposition. Assume that Héloïse does not have a winning strategy in the game $\mathcal{F}_{A,C}^\Sigma$. By determinacy then Abélard has a winning strategy. From Proposition 8 we obtain a countermodel to the validity of $\Sigma/A \vdash C$ on well-founded posets.

The implication from 3 to 4 is given by Proposition 11.

The implication from 4 to 1 is given by Proposition 12. \square

Using compactness we obtain the version of the main theorem for the system P . We need the following lemma:

Lemma 18. *If Σ is finite then Héloïse has a winning strategy in $\mathcal{G}_{A,C}^\Sigma$ iff she has a winning strategy in $\mathcal{F}_{A,C}^\Sigma$.*

Proof. First recall that the only difference between $\mathcal{F}_{A,C}^\Sigma$ and $\mathcal{G}_{A,C}^\Sigma$ are the winning conditions for infinite matches. In the former infinite matches are won by Héloïse, in the latter by Abélard. So whenever Héloïse has a winning strategy in $\mathcal{G}_{A,C}^\Sigma$ the same strategy is also winning in $\mathcal{F}_{A,C}^\Sigma$. This proves the direction from right to left.

For the direction from left to right we show that if Σ is finite then any match played according to a winning strategy for Héloïse in $\mathcal{F}_{A,C}^\Sigma$ is finite. Hence the strategy is also winning in the game $\mathcal{G}_{A,C}^\Sigma$.

Assume for a contradiction that an infinite match is played. In this match Héloïse has to play infinitely often, and because Σ is finite either some premise $B \vdash D \in \Sigma$ is played infinitely often or otherwise only \star is played infinitely often.

First consider the case that some $B \vdash D \in \Sigma$ is played infinitely often. After Héloïse plays $B \vdash D$ for the first time $B - D$ is contained in the forbidden area. Later in the match the forbidden area only grows so Abélard is never allowed to pick a world in $B - D$. This means that Héloïse can never play $B \vdash D$ again because this move is only available if the match is at a world in $B - D$. Hence $B \vdash D$ cannot be played infinitely often.

If only \star is played infinitely often there is a point in the match from which on Héloïse always plays \star . Hence there is a sequence of plays where Héloïse starts by playing \star , Abélard responds with $w \in W$ and Héloïse replies with \star again. But for w to be a reply to the first \star it needs to hold that $w \in A - C$, which contradicts that Héloïse answers with \star because this move presupposes that $w \in A \cap C$. \square

Theorem 19. *Let \mathfrak{A} be a compact subalgebra of \mathcal{PW} . If $\Sigma \cup \{A \vdash C\} \subseteq \{B \vdash D \mid B, D \in \mathfrak{A}\}$ then the following are equivalent:*

1. There is a proof of the inference $\Sigma/A \vdash C$ in system P .
2. The inference $\Sigma/A \vdash C$ is valid on posets.

3. Héloïse has a winning strategy in the game $\mathcal{G}_{A,C}^\Sigma$.

4. There is a finite witnessing set for the inference $\Sigma/A \vdash C$.

Proof. The implication from 1 to 2 is given by Proposition 2.

The implication from 2 to 3 is by contraposition. Assume that Héloïse does not have a winning strategy in the game $\mathcal{G}_{A,C}^\Sigma$. By determinacy then Abélard has a winning strategy. From Proposition 8 we obtain a countermodel to the validity of $\Sigma/A \vdash C$ on well-founded posets.

For the implication from 3 to 4 assume that Héloïse has a winning strategy in the game $\mathcal{G}_{A,C}^\Sigma$. Then by Theorem 15 there is a finite $\Sigma' \subseteq \Sigma$ such that Héloïse has a winning strategy in the game $\mathcal{G}_{A,C}^{\Sigma'}$. By Lemma 18 it follows that she also has a winning strategy in $\mathcal{F}_{A,C}^{\Sigma'}$. Proposition 11 yields a witnessing set Γ for the inference $\Sigma'/A \vdash C$. By the definition of a witnessing set we have that $\Gamma \subseteq \Sigma'$ and hence Γ is finite because Σ' is finite. Because Γ is a finite witnessing set for $\Sigma'/A \vdash C$ and $\Sigma' \subseteq \Sigma$ one can see from the definition of a witnessing set that Γ is also a finite witnessing set for the inference $\Sigma/A \vdash C$.

The implication from 4 to 1 is given by Proposition 12. \square

Note that if Σ is finite the assumption of compactness in the previous theorem is satisfied because then the subalgebra generated by A , C and all antecedents and consequents of conditionals in Σ is finite and hence compact.

Completeness extends to multi-conclusion inferences as follows:

Corollary 20. *A multi-conclusion inference Σ/Γ is provable in system P_∞ iff it is valid on well-founded posets.*

If $\Sigma \cup \Gamma \subseteq \{A \vdash C \mid A, C \in \mathfrak{A}\}$ for a compact subalgebra \mathfrak{A} of \mathcal{PW} then the multi-conclusion inference Σ/Γ is provable in system P iff it is valid on posets.

Proof. It follows immediately from the definition of provability in system P_∞ that a multi-conclusion inference Σ/Γ is provable in system P_∞ iff there exists a conclusion $A \vdash C \in \Gamma$ such that $\Sigma/A \vdash C$ is provable in system P_∞ . Analogously we have by Corollary 4 that the inference Σ/Γ is valid on well-founded posets iff there exists a conclusion $A \vdash C \in \Gamma$ such that $\Sigma/A \vdash C$ is valid on well-founded posets. Hence the claim follows by Theorem 17.

The second claim follows similarly using Theorem 19. \square

The corollary above yields the following strong completeness result:

Corollary 21. *Let Σ be a set of conditionals over a set of worlds W . Then there is a set of worlds U , a function $f : U \rightarrow W$ and a well-founded poset $P = (U, \leq)$ such that for all $A, C \subseteq W$:*

$$P \models f^{-1}[A] \vdash f^{-1}[C] \quad \text{iff} \quad \Sigma/A \vdash C \text{ is provable in system } P_\infty. \quad (5)$$

Proof. Let Σ be a set of conditionals over W and define Γ to be the following set of conditionals over W :

$$\Gamma = \{A \vdash C \mid \Sigma/A \vdash C \text{ is not provable in } P_\infty\}.$$

Clearly, Σ/Γ is not provable in system P_∞ . It follows by Corollary 20 that it is not valid on well-founded posets. Hence there is a set U , a function $f : U \rightarrow W$ and a well-founded poset $P = (U, \leq)$ such that $P \models f^{-1}[B] \vdash f^{-1}[D]$ for all

$B \vdash D \in \Sigma$ and $P \not\models f^{-1}[A] \vdash f^{-1}[C]$ for all $A \vdash C \in \Gamma$. By the definition of Γ the latter entails the left-to-right direction of (5).

For the right-to-left direction of (5) assume that $\Sigma/A \vdash C$ is provable in system P_∞ . It follows from Proposition 2 that the inference is valid on well-founded posets. From the definition of validity we obtain that $P \models f^{-1}[A] \vdash f^{-1}[C]$ because $P \models f^{-1}[B] \vdash f^{-1}[D]$ for all $B \vdash D \in \Sigma$. \square

Lastly, we obtain a similar result for system P , which is essentially Theorem 5.18 from [8].

Corollary 22. *Let \mathfrak{A} be a compact subalgebra of \mathcal{PW} in the sense of Definition 14. Take some $\Sigma \subseteq \{B \vdash D \mid B, D \in \mathfrak{A}\}$. Then there is a set of worlds U , a function $f : U \rightarrow W$ and a poset $P = (U, \leq)$ such that for all $A, C \in \mathfrak{A}$:*

$$P \models f^{-1}[A] \vdash f^{-1}[C] \quad \text{iff} \quad \Sigma/A \vdash C \text{ is provable in system } P.$$

Proof. This is analogous to the proof of Corollary 21. \square

8 Conclusions and further work

In this paper we introduce a game semantics for System P and use it to provide a new completeness proof with respect to the order semantics.

The game semantics is useful to determine whether a given inference is valid because for many inferences it is hard to find a formal proof in System P. In such cases it is often easier to determine whether Héloïse has a winning strategy in the game for the inference. If one finds such a winning strategy for Héloïse then Propositions 11 and 12 yield a formal proof in System P. If on the other hand one finds a winning strategy for Abélard then one immediately obtains a countermodel using Proposition 8.

The notion of a witnessing set introduced in this paper allows for a concise characterization of validity in System P. One can check whether an inference is valid by searching for a subset of the set of premises that satisfies the condition on a witnessing set from Definition 10. The inference is valid if and only if such a set is found. If the antecedents and consequents in the conditionals are propositional formulas then verifying the two conditions in Definition 10 amounts to checking the validity of two formulas in propositional logic. This problem is in coNP. Thus assuming a coNP oracle our algorithm checks for the validity of an inference by non-deterministically guessing a witnessing set. Hence the algorithm for checking validity is in $\Sigma_2^P = \text{NP}^{\text{coNP}}$. This is theoretically worse than the results from [9, 4, 12] which provide procedures that find countermodels in NP thus demonstrating that validity is in coNP. It would, however, be interesting to compare the performance of the different algorithms in actual applications, since the theoretical differences in complexity rely on non-deterministic Turing machines, which only exist in theory.

In Remark 5 we show that for completeness it is necessary to use labeling functions in the order semantics. This suggests looking for a semantics that does not require labeling functions. We are investigating an approach based on antimatroids [7, ch. 2] which are a generalization of partial orders.

It might also be interesting to adapt the game semantics of this paper to systems of conditional logic that are weaker or stronger than System P.

References

- [1] Alexandru Baltag and Sonja Smets. “Conditional Doxastic Models: A Qualitative Approach to Dynamic Belief Revision”. In: *Electronic Notes in Theoretical Computer Science* 165 (2006), pp. 5–21.
- [2] Craig Boutilier. “Conditional Logics of Normality: A Modal Approach”. In: *Artificial Intelligence* 68.1 (1994), pp. 87–154.
- [3] John Burgess. “Quick Completeness Proofs for Some Logics of Conditionals”. In: *Notre Dame Journal of Formal Logic* 22.1 (1981), pp. 76–84.
- [4] Nir Friedman and Joseph Y. Halpern. “On the Complexity of Conditional Logics”. In: *Principles of Knowledge Representation and Reasoning: Proceedings of the Fourth International Conference (KR’94)*. Morgan Kaufmann, 1994, pp. 202–213.
- [5] Adam Grove. “Two Modellings for Theory Change”. In: *Journal of Philosophical Logic* 17.2 (1988), pp. 157–170.
- [6] Jaakko Hintikka and Gabriel Sandu. “Game-Theoretical Semantics”. In: *Handbook of Logic and Language*. Ed. by Johan van Benthem and Alice ter Meulen. Second Edition. Elsevier, 2010, pp. 415–465.
- [7] Bernhard Korte, László Lovász, and Rainer Schrader. *Greedoids*. Springer, 1991.
- [8] Sarit Kraus, Daniel Lehmann, and Menachem Magidor. “Nonmonotonic Reasoning, Preferential Models and Cumulative Logics”. In: *Artificial Intelligence* 44.1-2 (1990), pp. 167–207.
- [9] Daniel Lehmann and Menachem Magidor. “What does a Conditional Knowledge Base Entail?” In: *Artificial Intelligence* 55.1 (1992), pp. 1–60.
- [10] David Lewis. *Counterfactuals*. Blackwell Publishers, 1973.
- [11] Paul Lorenzen and Kuno Lorenz. *Dialogische Logik*. Wissenschaftliche Buchgesellschaft, 1978.
- [12] Gian Luca Pozzato. *Conditional and Preferential Logics: Proof Methods and Theorem Proving*. Vol. 208. Frontiers in Artificial Intelligence and Applications. IOS Press, 2010.
- [13] Shahid Rahman and Tero Tulenheimo. “From Games to Dialogues and Back”. In: *Games: Unifying Logic, Language, and Philosophy*. Ed. by Ondrej Majer, Ahti-Veikko Pietarinen, and Tero Tulenheimo. Springer, 2009, pp. 153–208.
- [14] Jouko Väänänen. *Models and Games*. Cambridge University Press, 2011.
- [15] Frank Veltman. “Logics for Conditionals”. PhD thesis. University of Amsterdam, 1985.