



HAL
open science

On the Semantics of Disjunctive Logic Programs

Athanasios Tsouanas

► **To cite this version:**

Athanasios Tsouanas. On the Semantics of Disjunctive Logic Programs. Logic in Computer Science [cs.LO]. Ecole Nationale Supérieure de Lyon, 2014. English. NNT: . tel-02139060

HAL Id: tel-02139060

<https://theses.hal.science/tel-02139060>

Submitted on 24 May 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

THÈSE

en vue de l'obtention du grade de

**Docteur de l'Université de Lyon,
délivré par l'École Normale Supérieure de Lyon**

discipline : Informatique

Laboratoire de l'Informatique du Parallélisme (LIP)

École Doctorale Infomaths

présentée et soutenue publiquement le 2 juillet 2014 par

Athanasios TSOUANAS

ON THE SEMANTICS OF DISJUNCTIVE LOGIC PROGRAMS

Directeur de thèse :

Olivier LAURENT

Rapporteurs :

Guy MCCUSKER

Dale MILLER

Devant la commission d'examen formée de :

Christophe FOUQUERÉ (président)

Ekaterina KOMENDANTSKAYA

Olivier LAURENT

Guy MCCUSKER

ON THE SEMANTICS OF
DISJUNCTIVE LOGIC PROGRAMS

by Thanos Tsouanas

a PhD dissertation supervised by Olivier Laurent

*To miláki,
for tudo.*

Acknowledgements

This is the acknowledgements part, where I thank and give credit to those who I feel played a crucial and helpful rôle in all the work and research that lead up to this thesis.¹ Thus it is very easy to choose where to begin: my supervisor, Olivier Laurent, for all his support (scientific and non-scientific alike), his constructive guidance, and his patience, and for putting up with various peculiarities of mine, including my night-owl sleeping patterns. It has been a really enlightening experience working under his supervision and attending his seminars. My work would have been tremendously harder without his input. I had the pleasure to be part of the wonderful Plume team of the LIP lab here in ENS de Lyon, and I would especially like to thank Alexandre Miquel for many annoyingly intriguing discussions, and Colin Riba and Patrick Baillot for all their help. Filippo Bonchi, Damien Pous and Fabio Zanasi have all assisted me in times that I needed tools from their fields of expertise. Thanks also goes to Daniel Hirschhoff for being ἀστῆρος, and for providing a lot of entertainment with his band, among other things. Credit is definitely due to our secretariat staff members for all their hard work and kindness; I would especially like to thank Catherine Desplanches, Damien Séon & Sèverine Morin, for all they have done. Outstanding technical support has been provided by Serge Torres; thanks!

The MALOA project has funded the vast majority of my PhD studies.² I'm grateful to them, and glad that I met many interesting and friendly people in their events, especially Dugald Macpherson and Zoé Chatzidakis. As dictated by my scholarship, I had to spend some time working in different universities: I am thankful to Luke Ong in the University of Oxford, and Dale Miller in École Polytechnique, for giving me the chance to work in their teams, which assisted me a lot with my research. One more reason to thank Dale, is for many enlightening conversations, constructive feedback, and for taking the time to thoroughly read this manuscript. The same goes for Guy McCusker whom I had the pleasure to meet in plenty occasions. Christophe Fouqueré and Katya Komendantskaya were also very kind to accept to be examiners in the jury of my defense, besides my very short notice. John Power's continuous

¹Some similarities with the acknowledgements of my MSc thesis are merely a good indication of the fact that some values withstand time. Of that, and also of my laziness to come up with new text.

²The European Community's Seventh Framework Programme FP7/2007-2013 under grant agreement #238381 (the Marie Curie Initial Training Network in Mathematical Logic).

encouragement and expertise aided me tremendously multiple times that I felt desperately stuck in my work. I owe a lot to him. Also, Paulo and Isabel Oliva offered me priceless support; thank you both!

I'm indebted to the most influential professors I've had in my undergraduate and post-graduate studies, especially to Yiannis Moschovakis whose set theory and recursion theory courses made me realize my passion for mathematical logic. I would also like to thank Aristides Katavolos for his brilliant mathematical analysis courses, as well as Nikolaos Papaspyrou, Costas Koutras, and Dimitrios Thilikos, who played a crucial part in my progress. Martín Escardó's fascinating course introduced me to denotational semantics, for which I am grateful. Panos Rondogiannis, my MSc thesis supervisor, is in fact the one to blame for my interest in the semantics of logic programming and game semantics in particular. He has always been available for discussing any problems or doubts, each of the many times that I needed him during my PhD studies.

Assuming you don't speak Ελληνικά, you should also thank my English teacher, Mike Pozantzis, for being excellent at what he does; I owe him a lot. I feel lucky to have attended the French courses of Yasmine Daâs, even though I did a terrible job as a student—désolé!

Of course I owe a great lot to my family: Maroui, my mother, who is one of the smartest and kindest persons I know, for all her endless love and priceless lessons; my beloved sister, Eva, for being supportive and next to me since a few months before I was even born; and my father, Nikos, an old school hacker to whom I definitely owe—assuming this is a good thing—my geekiness with computers, since he introduced me to them when I could barely speak. I would also like to mention my brother-in-law, Spyridon Samothrakis, for many interesting conversations, scientific or not. My girlfriend, Ramile Carvalho da Silva, has gone through everything that I have throughout all this time that it took to complete my thesis, and even more. It has been a blessing to know that her love and comfort awaits for me after many stressful and disappointing moments, and that I have her support, taking care of me and of all those things that I seem to be incapable of dealing with on my own. Tons of hard parts were only made manageable thanks to her smile or effort. *Eu te amo!* ♡

It is great to have friends one can rely upon, and with whom one can have a great time as well. I want to thank my brother Yiannis Tzi Tselekounis, Harris Chanielidis, Matoula Petrolia, Georgios Zoulou Kapetanakis, Marilou Karanika, Irimi Petroula, Christos Tarantilis, Iraklis Choraitis, and Stathis Darkorang Lytras, for their love and support. Michail Pshlos Dimakos, Panos Tsaknias, and Alexandros Enecon Batsis have in addition provided mathematical insight and help more times than I could possibly remember; and if it hadn't been for Pshlos, I wouldn't even have a BSc degree. Mary Chlouveraki deserves to be singled out, as I feel I owe it to her, my sudden realization of interest in pure mathematics. A brilliant person and dearest friend, probably the only one whose advice I do not automatically ignore—and for good reason. I also thank Barbara Petit and João Aidar for standing by my side in numerous crazy occasions, and together with Rafael Padilha, Fernando Thiers, Nikos Konstantinides, Yannis Arapis, Emília Santos, Rodrigo Pereira, and Joel Creppy, for many relaxing basketball games and fun times in general, especially in periods of great pressure; and café 203 for providing sufficient amounts of much-needed, Bourganel beer. I am grateful to Tainá Lima, for keeping me company while working on the last details of my thesis and defense, and for not

letting me fall too hard when I know I would have. Myla Medina, you talked some sense into me when I was in need of; obrigado! Settling in a bureaucratic place without speaking the language is a challenging and tedious task. Christina Apostolopoulou & Florent Murat, Lionel Rieg, Jean-Marie Madiot, and Barbara gave me invaluable help, time and time again. I'm obliged to Irina Kalliamvakou for being a solid rock by my side for more than half of my life so far, and for so many things that still cannot be expressed in any language I know. Giselle Reis has been extremely kind and always offering help, even at times that I wouldn't even think of asking. I owe you ω !

I thank all the friendly and welcoming people in UFMG, especially Loïc Cerf, Fernando Pereira, and Mário Alvim, and while I'm at it, the ICEx cafeteria for all the pães de queijo and açai. I am also grateful to Ronaldo Silva & Luzia Carvalho for their generous hospitality, and for offering me a welcoming place where I could concentrate and work on my thesis while I was in Espírito Santo, Brazil. The same goes for Leo Valadão and the schools Uptime Linhares and Yázigí São Mateus. The children's computers school on Filellinon Street, Aghios Nikolaos, Crete, let me use their scanners and printers on a very crucial minutes-before-the-deadline moment, getting stressed on my behalf. Thanks!

I am also indebted to Spyros Lyberis who introduced me to the world of Unix back in the summer of '99, instantly changing my relationship with computers. I have become dependant on all those computer tools that make my life so much easier, so I feel I should definitely give credit to the Vim text editor, Don Knuth for T_EX, Theo de Raadt and all the OpenBSD developers, Simon Peyton Jones and the entire Haskell community, xmonad and darcs, and of course, freenode and Wikipedia.

One cannot be productive in research, if one is dead. I'm therefore also thankful to the splendid person who shot me in the back while I was in Brazil, for failing to kill me. I hope they have given up this hobby by now, although I doubt it. Ramile managed to convince three police officers to get me to a hospital—thank you all!—when they were thinking of leaving me behind, with a bullet in my back. I'm grateful to the doctors and nurses of Roberto Silveiras hospital in São Mateus, and to the ones in Hôpital Pierre Wertheimer in Bron, especially Dr. Eurico Freitas Olim and Dr. Cédric Barrey. I owe all these people my (bullet-free) life.

Finally, I hail Crete—I was on decks of Cretan ferries when I began developing my own ideas on semantics for disjunctive logic programs. And if I have forgotten to include your name while I should have, feel free to write it in the margin below. There's plenty of space.³

Thank you!

³Similarly, if you did not like being mentioned, take a pen and cross your name out. Sorry.

Contents

List of Tables	xiii
List of Figures	xiii
I Basic material	1
1 Introduction	3
1.1 What is a logic program?	3
1.2 Approaches to semantics	4
1.3 Outline	5
1.4 Contributions	6
1.5 Related works	6
1.6 Notation	7
2 Logic (and) languages	9
2.1 The language \mathcal{L}_0 of propositional logic	9
2.2 The language \mathcal{L}_1 of first-order, predicate logic	9
2.3 Truth value spaces	10
2.4 The spaces \mathbb{V}_κ	10
3 The syntax of logic programs	13
3.1 An informal look	13
3.2 Terminology	14
3.3 Four logic programming languages	17
3.4 Clean programs	18
4 Model-theoretic declarative semantics	21
4.1 Definitions and notation	21
4.2 LP—the least Herbrand model	22
4.3 DLP—the minimal models	22
4.4 LPN—the well-founded model	24
4.5 DLPN—the infinite-valued minimal models	27
4.6 First-order programs vs. infinite propositional ones	28

II From non-disjunctive to disjunctive semantics	31
5 An abstract framework for semantics	33
5.1 The framework	33
5.2 Semantics of LP	34
5.3 Semantics of DLP	34
5.4 Semantics of LPN	35
5.5 Semantics of DLPN	35
6 Disjunctive operations	37
6.1 Restricting	37
6.2 Splitting	38
6.3 Combinations	39
6.4 Definite instantiations and D-sections	41
7 The abstract transformation $(-)^{\vee}$	43
7.1 Definitions and theory	43
7.2 Applications and examples	46
III Game semantics	49
8 The LPG semantics	51
8.1 Introduction to game semantics	51
8.2 The LPG game	52
8.3 Example plays	53
8.4 Benefit of the doubt	54
8.5 Semantics from games	54
8.6 Soundness and completeness	55
9 The LPNG semantics	57
9.1 The LPNG game	57
9.2 Example plays	58
9.3 Game semantics	58
9.4 Soundness and completeness	60
10 The DLPG semantics	61
10.1 The simplified DLP game	61
10.2 The DLPG game	63
10.3 Plays	68
10.4 Strategies	75
10.5 Game semantics	82
10.6 Soundness and completeness	83
11 Applications of $(-)^{\vee}$ on game semantics	87
11.1 A different game semantics for DLP	87
11.2 An encoding of DLP to LP	88
11.3 A first game semantics for DLPN	89

IV Appendices	91
A Lattices and Heyting algebras	93
References	95
Index of symbols	101
General index	103

List of Tables

3.1	Summary of representative rules for each language	18
8.1	Development of game semantics for logic programming.	51
10.1	How to read believer and doubter statements.	64

List of Figures

1.1	The four logic programming languages studied in this thesis.	4
10.1	Commutative diagram for Property 10.22.	77
10.2a	The penultimate believer move of $\text{proj}_{\sigma_2}(\pi)$ is a stalling.	79
10.2b	The penultimate believer move of $\text{proj}_{\sigma_1}(\pi)$ is a stalling.	79
10.2c	Otherwise.	79

Part I

Basic material

Introduction

1.1 What is a logic program?

No matter where one meets logic programming, or what the specific features of the underlying language under investigation are, a logic program is always some sort of set of rules of the form

$$\mathbf{this} \leftarrow \mathbf{that},$$

read as “*this* holds, if *that* holds”, or “I can solve *this*, if I know how to solve *that*”. Depending on what restrictions we impose on **this** (the *head* of the rule) and **that** (the *body*), we enable or disable features of the resulting programming language.

In its simplest form, a rule looks like this:

$$\mathbf{a} \leftarrow \mathbf{b}_1, \dots, \mathbf{b}_m, \tag{LP}$$

where the commas on the right stand for conjunctions. The standard denotational or declarative semantics for this kind of programs is provided by a specific two-valued model, the so-called *least Herbrand model*. We will briefly review this in Section 4.2, in an attempt to be self-contained; consult [vEK76] or [Llo87] for further information. One extension is to allow *negations* to appear in bodies of rules:

$$\mathbf{a} \leftarrow \mathbf{b}_1, \dots, \mathbf{b}_m, \sim \mathbf{c}_1, \dots, \sim \mathbf{c}_k. \tag{LPN}$$

By negation, we mean *negation-as-failure* ([Cla78]); the semantics we have in mind here is supplied by the many-valued *well-founded model*, defined in [VGRS91]. But the extension in which we are mostly interested in this text is the appearance of *disjunctions* in heads:

$$\mathbf{a}_1 \vee \dots \vee \mathbf{a}_n \leftarrow \mathbf{b}_1, \dots, \mathbf{b}_m. \tag{DLP}$$

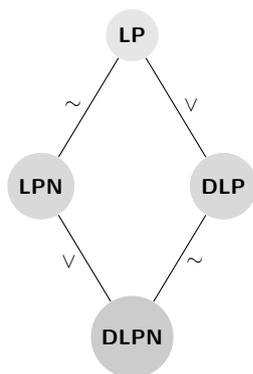
This enables us to express uncertainty and to derive ambiguous information. Instead of a single least model, in this case, we use a *set of minimal models* for the semantics, as defined in [Min82]; we introduce this in Section 4.3. Disjunctive logic programs are extensively studied in [LMR92]. Finally, one can consider both extensions simultaneously, by allowing *both* negations in

bodies *and* disjunctions in heads:¹

$$\mathbf{a}_1 \vee \cdots \vee \mathbf{a}_n \leftarrow \mathbf{b}_1, \dots, \mathbf{b}_m, \sim \mathbf{c}_1, \dots, \sim \mathbf{c}_k. \quad (\text{DLPN})$$

A satisfactory, infinite-valued, model-theoretic semantics for this extension was recently defined in [CPRW07].

Unfortunately, in the logic programming literature, terminology is not as stable as one could hope. To contribute to this dismay, we introduce in the following figure four abbreviations that will hopefully help the reader:



LP: the so-called *definite* logic programs
(with neither negation nor disjunctions);

DLP: disjunctive logic programs;

LPN: logic programs with negation;

DLPN: disjunctive logic programs with negation.

Figure 1.1: The four logic programming languages studied in this thesis.

1.2 Approaches to semantics

The model-theoretic approach to semantics outlined above, is not the only one that has proven itself worthy:

Fixpoint semantics. Frequently, to construct the model-theoretic semantics of logic programming languages, we use an *immediate consequence operator* (traditionally denoted by $T_{\mathcal{P}}$) associated with each program \mathcal{P} , and look at its fixpoints; see [Llo87]. In this thesis we will not concern ourselves much with the $T_{\mathcal{P}}$ operator. An excellent survey of fixpoint semantics for logic programming is [Fit99].

Procedural or operational semantics. The actual implementation of each of the above languages is usually given by *refutation* processes. Given a goal, the system tries to disprove it by constructing a counterexample: a proof that the program, together with the goal is an inconsistent set of rules. Traditionally, such proofs make use of some inference rule based on *resolution*. This might be, for example, SLD resolution in the case of LP, and SLI resolution for DLP. In this work, we do not touch this operational side of semantics either; see [Apt90] for the non-disjunctive and [LMR92] for the disjunctive cases.

Before turning to the next approach, game semantics, one should have a clear understanding of the nature of the aforementioned methods. On one side,

¹What about disjunctions in bodies, or conjunctions in heads? It is an easy exercise to show that this does not affect the programming language in any meaningful way; it only makes the programmer happier. See also Section 3.4.

we have the denotational, model-theoretic semantics and their fixpoint characterizations. These provide us with a notion of *correctness* for every possible answer to a goal that we might give to our program. On the operational side, the procedural semantics provide a construction of an answer to our question (the so-called *computed answer*), and this answer has to be correct. Conversely, such a procedure is expected to be able to derive all of the answers that the denotational semantics considers correct. We then say that the procedural semantics is sound and complete with respect to the denotational one. In Chapter 5 we formally define an abstract framework for semantics of logic programming languages, thus eliminating any kind of ambiguity from the term “semantics”.

Game semantics. Here we adopt a more anthropomorphic point of view, and treat each program as a set of rules for a game, in which two players compete against each other with respect to the truth of a given goal. One player, the so-called “Doubter” doubts the goal’s truthness, while the other player, the “Believer”, believing that the goal is true, tries to defend his stance. To get a meaningful semantics out of such games, we look at the *winning strategies* of the players, and depending on their existence, we assign an actual truth value to the goal. A game semantics may have a denotational or an operational flavor, or lie somewhere in-between the two. In [DCLN98], for example, they stay close to the procedural side of semantics, while the game semantics that we investigate here are more of a denotational nature. Part III is entirely devoted to this kind of semantics. This game-theoretic approach to semantics is influenced by Lorenzen’s dialogue games for logic (see [Lor61]). For an encyclopædic treatment of the use of games in logic, consult [Hod09].

Coalgebraic semantics. This is a very recent approach to semantics for logic programming, which involves using *coalgebraic* methods. Logic programs and their derivation strategies can be modelled as coalgebras on a certain category of presheaves. We will not use coalgebraic semantics in this treatment; for the interested reader, some works on this subject are [BM09], [KMP10], [KP11], [KPS13], and [BZ13]; other examples of approaches that also use categorical methods, include [CM92], [KP96], and [ALM09].

1.3 Outline

Part I introduces all the basic material that we will need for our development. In Chapter 2 we fix a propositional language \mathcal{L}_0 on which we will build our logic programs, we define the notion of a truth value space, and introduce the family \mathbb{V}_κ of such spaces. In the following two chapters we study the syntax and the model-theoretic semantics of the four logic programming languages.

In Part II we build our toolkit to deal with disjunctive programs. First, we introduce an abstract framework for semantics in Chapter 5 and show how the semantics that we have seen so far can be placed under this framework. In the next chapter we present restrictions, splittings, combinations, definite instantiations, and D-sections, which are the main tools we will use when dealing with disjunctive programs. In Chapter 7 we define an operator that acts on any semantics of a non-disjunctive logic programming language, and transforms it

into an “analogous” semantics for the corresponding disjunctive language. We also demonstrate its use by showing some of its applications.

Part **III** is devoted to game semantics. With the exception of DLPN, each language is directly given a game semantics (LP in Chapter **8**, DLP in **10**, and LPN in **9**). Finally, in Chapter **11**, we show some more applications of the semantic transformation defined in Chapter **7**, this time on game semantics.

1.4 Contributions

The contributions of this dissertation can be summarized as follows:

- An *abstract framework for logic programming semantics* is defined and all semantic approaches that we study are placed within this framework. We introduce the general notion of a *truth value space*, on which we evaluate formulæ. As expected, the booleans form the canonical example of a truth value space, but we need to consider much more general ones when dealing with negation-as-failure.
- The first game semantics for LP was given in [vE86] and [DCLN98]; this was later extended in [GRW08] for the case of LPN programs. Here a *game semantics for DLP programs* is developed in full detail; we prove that it is sound and complete with respect to the standard, minimal model semantics of [Min82]. Even though the game itself can be seen as an extension of the LP game, the formalization and notation we have used is influenced by the games used in functional programming instead.²
- We define a *semantic operator* which transforms any given abstract semantics of a non-disjunctive language to a semantics of the “corresponding” disjunctive one. We exhibit the correctness of this transformation by proving that it preserves equivalences of semantics, and we present some applications of it, obtaining some new semantics for DLP and DLPN. In particular, two novel semantics for infinite, propositional DLPN programs are constructed: one model-theoretic and one game-theoretic.

1.5 Related works

Some recent treatments benefit by incorporating tools from areas such as proof theory and linear logic into logic programming:³ the proof search is often presented in terms of *sequent calculi*, formulæ are not necessarily restricted to Horn clauses of first-order logic, linear connectives are taken into account, etc. [MN12], [MNPS91], [AP91], and [And92], are some works along this line of research, just to mention a few of them. What is more, games have been used successfully in such settings as well: in [PR05] and [MS06], for example, two different approaches for game semantics are presented in the context of *computation-as-proof-search*.

A different school of negation in logic programming, namely *stable model semantics* (see [GL88]) gave rise to a relatively new kind of declarative programming: *answer set programming*, or ASP (see [Gel08]). It allows for disjunctions

²This work was published in [Tso13].

³This one, does not.

(besides negations), among other things. Unsurprisingly, a game semantics approach is being investigated as well: in [NV07], a game is briefly outlined for programs with single-headed clauses, i.e., for the non-disjunctive fragment of ASP.

1.6 Notation

The Greek letters ϕ , ψ , and ρ will stand for rules, as well as for the corresponding logic formulæ. Programs will usually be denoted by calligraphic capital letters like \mathcal{P} , \mathcal{Q} , and \mathcal{R} . Lowercase letters such as a , d , f , and p , will always denote atoms, while uppercase such symbols will stand for sets of atoms. For instance, D could be the set $\{d_1, \dots, d_n\}$ which, as you shall shortly see, is identified with the disjunction $d_1 \vee \dots \vee d_n$. We will use a monospaced font when we show their appearances in programs. We use capital “script” letters for families or sequences of sets of atoms: \mathcal{D} could stand for $\langle D_1, \dots, D_n \rangle$, each D_i being a set of atoms $\{d_1^i, \dots, d_{k_i}^i\}$. The *truth values true* and *false* are written as **T** and **F** respectively; logical equivalence as \equiv .

Given a set X , its powerset is $\mathcal{P}(-)$, and we subscript it with n to refer to the set of all subsets of X of cardinality n : e.g., $\mathcal{P}_1(X)$ is the set of all singleton-subsets of X . For the set of all finite subsets of X we use $\mathcal{P}_f(X)$, so that $\mathcal{P}_f(X) = \bigcup_{i \in \omega} \mathcal{P}_i(X)$.

We will work with sequences a lot; we use ϵ or $\langle \rangle$ for the empty sequence, $\#$ for concatenation and $|\cdot|$ for length. We shall write $s \sqsubseteq s'$ to indicate that the sequence s is a prefix of s' , decorating it with an “e” in case s is of even length: $s \sqsubseteq_e s'$ (note that $\sqsubseteq_e \subseteq \sqsubseteq$). Proper (even) prefixes will be shown as \sqsubset (\sqsubset_e). $s \upharpoonright_n$ stands for the sequence of the first n elements of s , and it is equal to the whole sequence if $|s| \leq n$. We will frequently need to extract the longest, even, proper prefix of a sequence s ; we therefore introduce the notation s^- for this, with the convention that it leaves the empty sequence unaltered: $\langle \rangle^- \triangleq \langle \rangle$. Influenced by lists in programming languages, we use $::$ for the *cons* operator:

$$x :: s \triangleq \langle x \rangle \# s.$$

Products of posets are equipped with the product order by default:

$$(s_1, s_2) \sqsubseteq (s'_1, s'_2) \iff s_1 \sqsubseteq s'_1 \text{ and } s_2 \sqsubseteq s'_2.$$

As has been just demonstrated, \triangleq and \iff are used to introduce the definition of a function or symbol, while $:=$ is used to “let-bind” the variables appearing on its left to the corresponding expressions that appear on its right (or the other way around). Two abbreviations extensively found in texts of mathematics and logic that we will use are “iff” for “if and only if”, and “wff” for “well-formed formula”.

Further notational conventions will be introduced as soon as it is sensible to do so. Thus far, we have what we need to begin.

Logic (and) languages

2.1 The language \mathcal{L}_0 of propositional logic

In this section we describe the propositional logic language which we denote by \mathcal{L}_0 . Foremost we assume a countably infinite set \mathcal{A}_{t_0} whose elements we denote by a, b, c, \dots and we call *atoms*. We will use the binary connectives \vee , \wedge , and \rightarrow , and the unary connective \sim , which is meant to stand for *negation-as-failure*.

A (*well-formed*) *formula* (wff) is defined inductively: every atom is a formula (called *atomic formula*); and if α and β are formulæ, then so are $\alpha \vee \beta$, $\alpha \wedge \beta$, $\alpha \rightarrow \beta$, and $\sim\alpha$. We might use \rightarrow facing the other way, so that $\alpha \leftarrow \beta$ denotes exactly the same formula as $\beta \rightarrow \alpha$. Furthermore, we use $\alpha \leftrightarrow \beta$ as an abbreviation for the wff $(\alpha \rightarrow \beta) \wedge (\beta \rightarrow \alpha)$, and consider it to have the lowest precedence among these connectives. The *language* \mathcal{L}_0 is the set of all wffs.

The *literals* of \mathcal{L}_0 are its atomic formulæ (*positive literals*) and their \sim -negations (*negative literals*), and we use Lit_0 to denote their set.

2.2 The language \mathcal{L}_1 of first-order, predicate logic

Even though we will use almost exclusively the propositional language \mathcal{L}_0 , we mention here the first-order language \mathcal{L}_1 , since we will need to refer to it a couple of times. We assume given the countable sets of: variables Var , constant symbols Con , and function symbols Fun_n for any arity $n > 0$, which we use to inductively define the set of terms Term : a variable is a term, a constant symbol is a term, and if t_1, \dots, t_n are terms, and $f \in \mathit{Fun}_n$, then $f(t_1, \dots, t_n)$ is also a term. We also assume for every $n \in \mathbb{N}$, a countable set of n -ary predicate symbols Pred_n , which we use, in turn, to define the set of atomic formulæ At_1 of \mathcal{L}_1 : for every $n \in \mathbb{N}$, if $t_1, \dots, t_n \in \mathit{Term}$ and $R \in \mathit{Pred}_n$, then $R(t_1, \dots, t_n) \in \mathit{At}_1$, and At_1 contains no other elements besides these. By induction we finally define the set \mathcal{L}_1 of wffs: an atomic formula is a wff; if α and β are wffs, and $x \in \mathit{Var}$, then $\alpha \vee \beta$, $\alpha \wedge \beta$, $\alpha \rightarrow \beta$, $\sim\alpha$, $(\exists x \alpha)$, and $(\forall x \alpha)$ are also wffs. Again, the literals are the atomic formulæ and their negations, and we denote their set by Lit_1 . We will omit the subscripts and simply use At and Lit if whether we mean the propositional case or the first-order one is either immaterial or well-understood from the context.

2.3 Truth value spaces

The standard semantics of \mathcal{L}_0 is provided by Boolean logic, by mapping each binary connective above to the corresponding boolean operation on $\mathbb{B} = \{\mathbf{F}, \mathbf{T}\}$.

As it turns out, we will need more truth values to handle negation, and therefore we will not tie ourselves to the booleans. Abstracting away the useful properties of \mathbb{B} that we need, we reach the following general definition:

Definition 2.1. A *truth value space* is a completely distributive Heyting algebra with an additional unary operator \sim .¹

The canonical example of a truth value space is \mathbb{B} , in which \sim is defined as the classical negation that flips the two values. This space turns out to be too poor for languages that actually use negation as failure, and so we investigate spaces with more values in the next section.

Definition 2.2 (valuation function). Let \mathcal{V} be a truth value space. A function $val : \mathcal{A}t \rightarrow \mathcal{V}$ is called a (propositional) *valuation function*. We denote the set of valuations ($\mathcal{A}t \rightarrow \mathcal{V}$) by $Val_{\mathcal{V}}$, omitting the subscript if it is obvious from the context. Any valuation function val on $\mathcal{A}t$ is extended to a valuation over all wffs of \mathcal{L}_0 by consulting the operations of \mathcal{V} :

$$\begin{aligned} val(\alpha \wedge \beta) &= val(\alpha) \wedge val(\beta); & val(\sim\alpha) &= \sim val(\alpha); \\ val(\alpha \vee \beta) &= val(\alpha) \vee val(\beta); & val(\alpha \rightarrow \beta) &= val(\alpha) \Rightarrow val(\beta). \end{aligned}$$

Definition 2.3 (\mathcal{V} -tautology). Let \mathcal{V} be a truth value space. The wff α of \mathcal{L}_0 is a \mathcal{V} -*tautology*, iff $v(\alpha) = \max \mathcal{V}$ for any valuation function v .

Definition 2.4 ($\Gamma \models_{\mathcal{V}} \alpha$, $\alpha \equiv_{\mathcal{V}} \beta$). Let Γ be a set of wffs of \mathcal{L}_0 and let \mathcal{V} be a truth value space. We write $\Gamma \models_{\mathcal{V}} \alpha$ and say that the wff α is a *logical consequence of Γ in \mathcal{V}* iff $(\bigwedge \Gamma) \rightarrow \alpha$ is a \mathcal{V} -tautology. Two wffs α and β of \mathcal{L}_0 are called *logically equivalent in \mathcal{V}* iff $\alpha \leftrightarrow \beta$ is a \mathcal{V} -tautology. In symbols, $\alpha \equiv_{\mathcal{V}} \beta$. Whenever \mathcal{V} is clear from the context we will omit the subscripts in these notations.

2.4 The spaces \mathbb{V}_{κ}

Even though three-valued logics have been used for many years in the study of negation in logic programming (e.g., [VGRS91], [Fit85], and [Kun87]) we jump directly to a family of infinite-valued logics on which we will eventually base our semantics of negation-as-failure. We are actually dealing with refinements of the usual three-valued logic that was originally used for the well-founded semantics, enjoying some additional convenient properties. Spaces of this kind were first introduced in [RW05], and further studied in [GRW08] and [Lüd11].

Definition 2.5 (The truth value spaces \mathbb{V}_{κ}). Let $\kappa \geq \omega$ be an ordinal number. The structured set²

$$\mathbb{V}_{\kappa} = (\mathbb{V}_{\kappa}; \vee, \wedge, \Rightarrow, \sim)$$

¹See Appendix A for more information about lattices and Heyting algebras.

²Here we follow the usual practice of abusing the notation by identifying the structured set with its carrier.

consists of an infinite number of distinct elements, which we separate into three disjoint sets:

$$\mathcal{F}_\kappa \triangleq \{\mathbf{F}_\alpha \mid \alpha < \kappa\}; \quad \mathbf{U} \triangleq \{\mathbf{U}\}; \quad \mathcal{T}_\kappa \triangleq \{\mathbf{T}_\alpha \mid \alpha < \kappa\}.$$

We denote their union by $\mathbb{V}_\kappa \triangleq \mathcal{F}_\kappa \cup \mathbf{U} \cup \mathcal{T}_\kappa$, and equip it with the total order

$$\mathbf{F}_0 < \mathbf{F}_1 < \cdots < \mathbf{F}_\alpha < \cdots < \mathbf{U} < \cdots < \mathbf{T}_\alpha < \cdots < \mathbf{T}_1 < \mathbf{T}_0.$$

This turns \mathbb{V}_κ into a complete bounded lattice, thus determining \vee , \wedge , and \Rightarrow :

$$x \vee y = \max\{x, y\}, \quad x \wedge y = \min\{x, y\}, \quad \text{and} \quad x \Rightarrow y = \begin{cases} \mathbf{T}_0 & \text{if } x \leq y, \\ y & \text{otherwise.} \end{cases}$$

But for \mathbb{V}_κ to be a valid candidate for a truth value space, it remains to define the operator \sim :

$$\sim x \triangleq \begin{cases} \mathbf{T}_{\alpha+1} & \text{if } x = \mathbf{F}_\alpha, \\ \mathbf{F}_{\alpha+1} & \text{if } x = \mathbf{T}_\alpha, \\ \mathbf{U} & \text{if } x = \mathbf{U}. \end{cases}$$

Unless explicitly mentioned, we will simply write \mathbb{V} instead of \mathbb{V}_ω in case $\kappa = \omega$.

Definition 2.6. Let κ be an ordinal. The mappings $order : \mathbb{V}_\kappa \rightarrow \mathbf{ON}$ and $collapse : \mathbb{V}_\kappa \rightarrow \mathbb{V}_1$ are then defined by

$$order(v) \triangleq \begin{cases} \alpha & \text{if } v \in \{\mathbf{F}_\alpha, \mathbf{T}_\alpha\} \text{ for some ordinal } \alpha < \kappa \\ \kappa & \text{otherwise, i.e., if } v = \mathbf{U}; \end{cases}$$

$$collapse(v) \triangleq \begin{cases} \mathbf{T}, & \text{if } v \in \mathcal{T}_\kappa, \\ \mathbf{F}, & \text{if } v \in \mathcal{F}_\kappa, \\ \mathbf{U}, & \text{otherwise.} \end{cases}$$

The intuition behind these truth values is easy to explain: we identify \mathbf{F}_0 and \mathbf{T}_0 with the usual boolean values \mathbf{F} and \mathbf{T} , i.e., absolute truth and absolute falsity. The ordinal in the subscript corresponds to a level of doubt that we have, so that \mathbf{F}_1 represents a “false” value but with a little doubt, \mathbf{F}_2 one with a little more, etc., and similarly for the “true” values. In the middle lies \mathbf{U} , which we use in the case that we only have doubts without any bias towards truth or falsity: it is entirely uncertain.

Property 2.1. For any $\kappa \geq \omega$, \mathbb{V}_κ enjoys the following properties:

- (i) it is bounded;
- (ii) it is a chain;
- (iii) it is a distributive lattice;
- (iv) it is a complete lattice;
- (v) it is a Heyting algebra;
- (vi) $\mathbb{V}_\kappa \cong \mathbb{V}_\kappa^\partial$.

Lemma 2.2. For any $\kappa \geq \omega$, \mathbb{V}_κ is algebraic.

Proof. We have just observed that \mathbb{V}_κ is a complete lattice (Property 2.1(iv)). We proceed to compute the set of its compact elements $K(\mathbb{V}_\kappa)$. Let $k \in \mathbb{V}_\kappa$

and $S \subseteq \mathbb{V}_\kappa$ such that $k \leq \bigvee S$. For k to be compact, we need a finite $T \subseteq S$ such that $k \leq \bigvee T$. We have four cases to consider:

CASE 1: $k = \mathbf{F}_\alpha$ for some non-limit ordinal α . Then S must contain at least some element $t \geq k$, so that $\{t\}$ can be the T that we need, and k is compact.

CASE 2: $k = \mathbf{F}_\lambda$ for some limit ordinal λ . Now, taking $S = \{x \mid x < k\}$ we see that $k = \bigvee S$ and therefore k cannot be a compact element.

CASE 3: $k = \mathbf{U}$. Similarly to Case 2, we set $S = \mathfrak{F}$ and verify that k is not compact.

CASE 4: $k = \mathbf{T}_\alpha$ for some ordinal α . Whether α is limit or not, in this case S will always contain at least some element $t \in \mathbb{V}_\kappa \setminus \{x \in \mathbb{V}_\kappa \mid x \leq \mathbf{T}_{\alpha+1}\}$, so that $t \geq k$, and setting $T = \{t\}$ we see that k is compact.

We have reached the conclusion:

$$K(\mathbb{V}_\kappa) = (\mathfrak{F} \setminus \{\mathbf{F}_\lambda \mid \text{Limit}(\lambda)\}) \cup \mathfrak{T}.$$

We are now in position to prove that \mathbb{V}_κ is algebraic, i.e., for every $a \in \mathbb{V}_\kappa$,

$$a = \bigvee \{k \in K(\mathbb{V}_\kappa) \mid k \leq a\}.$$

Indeed, in the nontrivial case that a is not itself compact, it is easy to see that it is the supremum of all the compact elements $k \leq a$. ■

Theorem 2A. *For any $\kappa \geq \omega$, \mathbb{V}_κ is a truth value space.*

Proof. As \mathbb{V}_κ has a unary operation \sim , we only need to verify that it is a completely distributive Heyting algebra. This follows immediately by Theorem A.7, since we already have Lemma 2.2 and Properties 2.1(iii) and (vi). ■

We have defined the general structure of a truth value space, and have seen that \mathbb{B} and \mathbb{V}_κ are examples of such spaces. In the next two chapters we define the syntax and model-theoretic semantics of the four logic programming languages that we will investigate.

The syntax of logic programs

In this chapter we describe four logic programming languages, starting with some informal examples meant to indicate how they are used. Then we give formal mathematical definitions that facilitate their study from our perspective, and introduce the reader to some related terminology.

3.1 An informal look

The simplest language that will concern us here is LP. Heads of rules constitute of a single atom, while bodies are conjunctions of atoms. A simple program, for example, is the following:

$$\mathcal{P}_1 = \left\{ \begin{array}{l} \text{sleeps} \leftarrow \text{tired} \\ \text{works} \leftarrow \text{rested} \\ \text{eats} \leftarrow \text{rested}, \text{hungry} \\ \text{rested} \leftarrow \end{array} \right\}.$$

We might use this simple program to describe the daily life of a rather dull person who is currently rested (this is what the last rule states). We load the program to our implementation and interact with it by asking it questions, or queries, i.e., we ask if a proposition is true or not. The answers have to come from a truth value space \mathcal{V} , and in this case we use \mathbb{B} .¹ The implementation, if correct, will answer positively (**T**) if and only if the query is a logical consequence of the program \mathcal{P}_1 . In this case we say that the query has succeeded; otherwise, it has failed. For example, **works** succeeds here thanks to the second rule, which reduces it to **rested**, which in turn succeeds since it is a fact (the fourth rule). On the other hand, **tired** fails, because there is no evidence to the contrary. This illustrates that by default we consider everything to be false, unless provable otherwise. This assumption, called the *closed world assumption* (CWA), is a very important aspect of logic programming, and it is one of

¹As long as there are no negations present in the programs, using any other truth value space \mathcal{V} in place of \mathbb{B} is pointless: only its \mathbb{B} -isomorphic subset $\{\min \mathcal{V}, \max \mathcal{V}\} \subseteq \mathcal{V}$ will end up being used. This would not have been the case if we had allowed for “facts” of the form $\mathbf{a} \leftarrow v$ for every $v \in \mathcal{V}$, but we are not interested in this direction here. As an example of such an approach, the reader is referred to [vE86].

the things that make it so practical as a tool for computing and dealing with database-like information.

In many applications, it is impractical or even literally impossible to explicitly list negative information. Imagine, for example, that we are implementing a program for an airline. We want to ask whether there is a flight between two airports at a given time. Had we not followed this convention, we would have to list, for every pair of airports (x, y) , and for every possible time t , the specific information of whether there is a flight from x to y at time t . Following the CWA we only need to list the flights that actually take place, and assume that if no flight is listed on a particular time, then there must indeed be no flight at this time from x to y .

Extending LP. Even on unrealistically simple LP programs like \mathcal{P}_1 , we can easily spot some limitations of the language: (1) even though **rested** and **tired** are meant to be related in the obvious manner, they are not; and (2) we have no way to express ambiguous information, such as **hungry** \vee **thirsty**.

The first shortcoming leads to LPN, which introduces the operation \sim in bodies of rules for negation-as-failure. We can now write the program

$$\mathcal{P}_2 = \left\{ \begin{array}{l} \text{sleeps} \leftarrow \text{tired} \\ \text{works} \leftarrow \sim\text{tired} \\ \text{eats} \leftarrow \sim\text{tired}, \text{hungry} \end{array} \right\},$$

which describes the typical day of the same dull person, but in a more concise way. To see how we can still conclude that **works** holds, we observe that according to the second rule, it reduces to the (finite) failure of **tired**. And as we have no information to support that **tired** holds, we indeed conclude that $\sim\text{tired}$ succeeds and therefore so does **works**.

The second shortcoming of LP leads to DLP, which relaxes the restriction that the heads can only contain a single atom, by allowing a disjunction of atoms to appear instead. This gives us the flexibility to reason with indefinite information. Consider, the following example:

$$\mathcal{P}_3 = \left\{ \begin{array}{l} \text{mathematician} \leftarrow \text{topologist} \\ \text{mathematician} \leftarrow \text{algebraist} \\ \text{algebraist} \vee \text{topologist} \leftarrow \end{array} \right\}.$$

In this program, we do not have concrete evidence of whether we are dealing with an algebraist or a topologist, but the third rule states as a matter of fact that (at least) one of the two must hold. Therefore, even though neither the first nor the second rule by themselves are enough to deduce that we are in fact dealing with a mathematician, in this disjunctive setting of DLP programs, we are able to deduce that this is indeed the case.

Naturally, we may allow both extensions simultaneously, and doing so we obtain DLPN which allows both \sim in bodies and \vee in heads, and it is the most general language that will concern us here.

3.2 Terminology

To make it explicit that we refer to the set-theoretic interpretation when we speak of a disjunction or a conjunction, we will use the prefix ‘‘L.P.’’ as shown in

the following definition, which formally introduces this and related terminology:

Definition 3.1. An *L.P. disjunction* is a finite subset $D \subseteq Lit$. An *L.P. conjunction* is a finite sequence \mathcal{D} of L.P. disjunctions. For obvious reasons we omit the “L.P.” prefix whenever no confusion arises. A *clause* is a pair (H, \mathcal{D}) , in which the *head* $H = \{a_1, \dots, a_n\}$ is an L.P. disjunction, and the *body* $\mathcal{D} = \langle D_1, \dots, D_m \rangle$ is an L.P. conjunction. If the head of a clause is non-empty we call it a *rule*, while if it is empty and $m = 1$, a *goal*.² A *fact* is a bodiless clause. In logic programs, rules will be written as

$$\underbrace{a_1 \vee \dots \vee a_n}_{\text{head}} \leftarrow \underbrace{\ell_1^1 \vee \dots \vee \ell_{s_1}^1, \dots, \ell_1^m \vee \dots \vee \ell_{s_m}^m}_{\text{body}}.$$

Such a rule is called *disjunctive* (also *proper*) if $n > 1$; it is *clean*, if $s_j = 1$ for all $1 \leq j \leq m$. Therefore, a clean rule looks like this:

$$a_1 \vee \dots \vee a_n \leftarrow \ell^1, \dots, \ell^m.$$

A *clean program* is a countable set of clean rules; it is *disjunctive* (also *proper*), if at least one of its rules is. Note that we have not specified what the atoms in \mathcal{At} really are. One may consider them to simply be propositional variables without any further structure, just like in propositional calculus. In this case, we have a *propositional program*. Another possibility is to let them be the atomic formulæ of a first-order language, built by its predicates, function symbols, variables, and constants. We then call it a *first-order program*.

► *Example 3.1.* Consider the following sets of rules:

$$\mathcal{P} := \left\{ \begin{array}{l} p \leftarrow a \\ p \leftarrow b \\ a \vee b \leftarrow \end{array} \right\}, \quad \mathcal{Q} := \left\{ \begin{array}{l} e \vee p \leftarrow f \vee g, h \\ p \vee q \leftarrow g, e \vee r \end{array} \right\}, \quad \mathcal{R} := \left\{ \begin{array}{l} d \leftarrow f \vee h \\ p \leftarrow g \vee e \end{array} \right\}.$$

The program \mathcal{P} is disjunctive and clean, \mathcal{Q} is also disjunctive but not clean, and \mathcal{R} is neither of the two. ◀

► *Example 3.2.* Here is a first-order logic program:

$$\left\{ \begin{array}{l} \text{daughter}(X, Y) \leftarrow \text{child}(X, Y), \text{female}(X) \\ \text{spouse}(X, Y) \leftarrow \text{married}(X, Y), \text{female}(X) \\ \text{married}(X, Y) \leftarrow \text{married}(Y, X) \\ \text{child}(\text{eva}, \text{maroui}) \vee \text{married}(\text{sam}, \text{eva}) \leftarrow \\ \text{female}(\text{eva}) \leftarrow \\ \text{male}(\text{sam}) \leftarrow \end{array} \right\}.$$

It is disjunctive and clean. ◀

²We have imposed the restriction $m = 1$ for goals. This will simplify the development without any significant loss: to deal with a goal like $\leftarrow D_1, \dots, D_m$, one can simply add the rule $w \leftarrow D_1, \dots, D_m$ to the program, where w is a suitable fresh atom, and query w instead.

Definition 3.2 (Ground). If a term, an atom, or a formula, ρ or a set of formulæ \mathcal{P} of a first-order language \mathcal{L}_1 contains no variables, it is called *ground*.³ If it does, then by replacing all of its variables with ground terms we obtain a *ground instance* of it. We also define:

$$\begin{aligned}\text{ground}(\rho) &\triangleq \{\rho' \mid \rho' \text{ is a ground instance of } \rho\} \\ \text{ground}(\mathcal{P}) &\triangleq \bigcup \{\text{ground}(\phi) \mid \phi \in \mathcal{P}\}.\end{aligned}$$

► *Example 3.3.* Consider the first-order program \mathcal{E} in the language that contains a constant symbol 0, a unary function symbol S, and the unary predicate symbol even:

$$\mathcal{E} := \left\{ \begin{array}{l} \text{even}(0) \leftarrow \\ \text{even}(S(S(X))) \leftarrow \text{even}(X) \end{array} \right\}.$$

We compute:

$$\text{ground}(\mathcal{E}) := \left\{ \begin{array}{l} \text{even}(0) \leftarrow \\ \text{even}(S^2(0)) \leftarrow \text{even}(0) \\ \text{even}(S^3(0)) \leftarrow \text{even}(S(0)) \\ \text{even}(S^4(0)) \leftarrow \text{even}(S^2(0)) \\ \vdots \end{array} \right\}.$$

◀

‡ **REMARK 3.1** (logic programs and logic formulæ). We have seen that there is an obvious correspondence between logic programs and rules on the one hand and logic formulæ on the other. This allows us to directly use some well-known jargon of mathematical logic: we speak of models, theories, consistency, logical consequences, etc. We should not be too eager to jump to conclusions, however. One of the common pitfalls is to think that the \sim of logic programming corresponds to the \neg of mathematical logic. In fact, the very reason we have chosen to use a different symbol for negation-as-failure, is to keep in mind that this is not the case.

Set-theoretic translations. Given any language of logic \mathcal{L} , we will identify a set of \mathcal{L} -literals with the \mathcal{L} -wff of their disjunction. Note that under this convention, we will identify some actually distinct formulæ, e.g., $a \vee b$ and $b \vee a$; this poses no threat as such formulæ are already equivalent up to commutativity and/or associativity of \vee . Similarly, sequences of sets of literals, when regarded as logic formulæ, are identified with conjunctions of disjunctions of literals. Notice that under this convention they are actually formulæ in *conjunctive normal form* (CNF). It is also convenient to consider rules of the form $\alpha \leftarrow \beta$ as pairs (α, β) . Thus, we have a mapping from the world of logic to the language of set theory, which is a very convenient tool for our development.

► *Example 3.4.* The set $\{a, b, c\}$ is understood to stand for the disjunction $a \vee b \vee c$, the sequence $\langle \{a\}, \{b, c\} \rangle$ for the conjunction $a \wedge (b \vee c)$, and the pair $(\{p, q\}, \langle \{a, b\}, \{b, c\} \rangle)$ for the implication $((a \vee b) \wedge (b \vee c)) \rightarrow p \vee q$. ◀

³See [LMR92, Ch. 2] for more details on first-order languages.

Since a program is itself a set of rules, programs can also be translated in the same manner:

► *Example 3.5.* Consider the program

$$\left\{ \begin{array}{l} p \vee q \leftarrow a, b \vee t \\ r \leftarrow \sim a, t \\ t \leftarrow \end{array} \right\}.$$

Translating it into set-theoretic terms, we end up with the following set of pairs:

$$\{(\{p, q\}, \{a\}, \{b, t\}), (\{r\}, \{\sim a\}, \{t\}), (\{t\}, \langle \rangle)\}. \quad \blacktriangleleft$$

Definition 3.3. A logic programming language L is determined by:

- \mathbf{H}_L , the set of heads of rules of L ;
- \mathbf{B}_L , the set of bodies of rules of L ;
- \mathbf{Q}_L , the set of queries, or goal clauses of L .

We define the set of *rules* of L as $\mathbf{R}_L \triangleq \mathbf{H}_L \times \mathbf{B}_L$. A *program* of L is a set of rules of L . We write \mathbf{P}_L for the set of programs of L . In most logic programming languages, the bodies of rules are required to be conjunctions, in which case we denote by \mathbf{C}_L the set of all possible conjuncts out of which bodies are formed; in symbols,

$$\phi \in \mathbf{B}_L \iff \phi = \bigwedge C, \text{ for some finite sequence } C \in \mathbf{C}_L^*.$$

Definition 3.4. On the set of rules \mathbf{R}_L of a language L we define two operators head and body as the projections

$$\begin{aligned} \text{head}((H, \mathcal{B})) &\triangleq H, \\ \text{body}((H, \mathcal{B})) &\triangleq \mathcal{B}. \end{aligned}$$

3.3 Four logic programming languages

To formally define the languages we are interested in, we need to specify for each one of them its determining sets: its heads, its body-conjuncts, and its queries. Here they are:

$$\begin{array}{ll} \mathbf{H}_{\text{LP}} \triangleq \mathcal{P}_1(\mathcal{A}t) & \mathbf{H}_{\text{DLP}} \triangleq \mathcal{P}_f(\mathcal{A}t) \\ \mathbf{C}_{\text{LP}} \triangleq \mathcal{P}_1(\mathcal{A}t) & \mathbf{C}_{\text{DLP}} \triangleq \mathcal{P}_f(\mathcal{A}t) \\ \mathbf{Q}_{\text{LP}} \triangleq \mathcal{P}_1(\mathcal{A}t) & \mathbf{Q}_{\text{DLP}} \triangleq \mathcal{P}_f(\mathcal{A}t) \\ \\ \mathbf{H}_{\text{LPN}} \triangleq \mathcal{P}_1(\mathcal{A}t) & \mathbf{H}_{\text{DLPN}} \triangleq \mathcal{P}_f(\mathcal{A}t) \\ \mathbf{C}_{\text{LPN}} \triangleq \mathcal{P}_1(\text{Lit}) & \mathbf{C}_{\text{DLPN}} \triangleq \mathcal{P}_f(\text{Lit}) \\ \mathbf{Q}_{\text{LPN}} \triangleq \mathcal{P}_1(\text{Lit}) & \mathbf{Q}_{\text{DLPN}} \triangleq \mathcal{P}_f(\text{Lit}). \end{array}$$

Notice that for all of the languages above, the sets \mathbf{C}_L and \mathbf{Q}_L coincide.

Table 3.1: Summary of representative rules for each language

$p \leftarrow a, b$	LP
$p \vee q \leftarrow a, b$	cDLP
$p \vee q \leftarrow a, b \vee c$	DLP
$p \leftarrow \sim a, b$	LPN
$p \vee q \leftarrow \sim a, b$	cDLPN
$p \vee q \leftarrow \sim a, b \vee c$	DLPN

► *Example 3.6.* Here are some sample programs written in these languages:

$$\mathcal{P}_1 = \left\{ \begin{array}{l} p \leftarrow a \\ p \leftarrow b \\ b \leftarrow \end{array} \right\} \in \mathbf{P}_{\text{LP}} \quad \mathcal{P}_3 = \left\{ \begin{array}{l} a \vee b \leftarrow \\ p \leftarrow a \\ p \leftarrow b \end{array} \right\} \in \mathbf{P}_{\text{DLP}}$$

$$\mathcal{P}_2 = \left\{ \begin{array}{l} p \leftarrow \\ r \leftarrow \sim p \\ s \leftarrow \sim q \end{array} \right\} \in \mathbf{P}_{\text{LPN}} \quad \mathcal{P}_4 = \left\{ \begin{array}{l} p \vee q \vee r \leftarrow \\ p \leftarrow \sim q \\ q \leftarrow \sim r \\ r \leftarrow \sim p \end{array} \right\} \in \mathbf{P}_{\text{DLPN}} \quad \blacktriangleleft$$

3.4 Clean programs

We use cDLP to denote the language of clean DLP programs, which is determined by:

$$\begin{array}{ll} \mathbf{H}_{\text{cDLP}} \triangleq \mathcal{P}_f(\mathcal{A}t) & \mathbf{H}_{\text{cDLPN}} \triangleq \mathcal{P}_f(\mathcal{A}t) \\ \mathbf{C}_{\text{cDLP}} \triangleq \mathcal{P}_1(\mathcal{A}t) & \mathbf{C}_{\text{cDLPN}} \triangleq \mathcal{P}_1(\text{Lit}) \\ \mathbf{Q}_{\text{cDLP}} \triangleq \mathcal{P}_1(\mathcal{A}t) & \mathbf{Q}_{\text{cDLPN}} \triangleq \mathcal{P}_1(\text{Lit}). \end{array}$$

Notice that $\mathbf{P}_{\text{cDLP}} \subseteq \mathbf{P}_{\text{DLP}}$.

Following [LMR92, §2.3], a *disjunctive clause* is the universal closure of a logic formula like

$$L_1 \vee \cdots \vee L_k,$$

where the L_i 's are literals. Separating them into positive and negative, and omitting the quantifiers, this clause can be brought to the form

$$a_1 \vee \cdots \vee a_n \vee \neg b_1 \vee \cdots \vee \neg b_m,$$

or, equivalently (by De Morgan) to

$$a_1 \vee \cdots \vee a_n \vee \neg(b_1 \wedge \cdots \wedge b_m),$$

which, in turn, is logically equivalent to the (reverse) implication

$$a_1 \vee \cdots \vee a_n \leftarrow b_1 \wedge \cdots \wedge b_m.$$

We generally adopt this as the logic programming notation of a clause, writing commas instead of \wedge . Coming this way, it is impossible for a disjunction to

appear in the body of a rule. Here though, we bypass this construction as it is often more natural to express ideas using “unclean” rules.

When we are working with clean programs, we can consider disjunctions in bodies as “syntactic sugar”, thanks to the following transformation:

Definition 3.5 ($\widehat{\mathcal{P}}$ and $\widehat{\phi}$). Let \mathcal{P} be a logic program. Then $\widehat{\mathcal{P}}$ is the program that results if we replace every unclean rule $\phi = (H, \langle D_1, \dots, D_n \rangle)$ of \mathcal{P} by all clean rules in

$$\widehat{\phi} \triangleq \{(H, C) \mid C \in D_1 \times \dots \times D_n\}.$$

We call $\widehat{\mathcal{P}}$ the *clean version* of \mathcal{P} . It follows that if $\mathcal{P} \in \mathbf{P}_{\text{DLP}}$ and $\mathcal{Q} \in \mathbf{P}_{\text{DLPN}}$, then $\widehat{\mathcal{P}} \in \mathbf{P}_{\text{cDLP}}$ and $\widehat{\mathcal{Q}} \in \mathbf{P}_{\text{cDLPN}}$.

This is a simple case of the most general *Lloyd–Topor transformation*, which transforms logic programs containing arbitrary formulæ in their bodies into normal ones. See [LT84], [Llo87, Ch. 4] or [LMR92, pp. 188–189] for more details.

► *Example 3.7.* Here is the desugaring of a program $\mathcal{Q} \in \mathbf{P}_{\text{DLP}}$:

$$\mathcal{Q} := \left\{ \begin{array}{l} \mathbf{e} \vee \mathbf{p} \leftarrow \mathbf{f} \vee \mathbf{g}, \mathbf{h} \\ \mathbf{p} \vee \mathbf{q} \leftarrow \mathbf{g}, \mathbf{e} \vee \mathbf{r} \end{array} \right\} \in \mathbf{P}_{\text{DLP}} \quad \xrightarrow{\widehat{}} \quad \widehat{\mathcal{Q}} := \left\{ \begin{array}{l} \mathbf{e} \vee \mathbf{p} \leftarrow \mathbf{f}, \mathbf{h} \\ \mathbf{e} \vee \mathbf{p} \leftarrow \mathbf{g}, \mathbf{h} \\ \mathbf{p} \vee \mathbf{q} \leftarrow \mathbf{g}, \mathbf{e} \\ \mathbf{p} \vee \mathbf{q} \leftarrow \mathbf{g}, \mathbf{r} \end{array} \right\} \in \mathbf{P}_{\text{cDLP}}. \quad \blacktriangleleft$$

Once we have examined the semantics of logic programs in the next chapter, the following property will become apparent:

Property 3.1. *Let \mathcal{P} be a logic program. $\widehat{\mathcal{P}}$ is clean and equivalent to \mathcal{P} .*

Almost all of the fundamental program constructions that we investigate preserve “cleanliness”; the one time that it really makes a difference is in Section 10.6: there, we prove soundness and completeness for cDLP programs first, and then proceed to consider the general case of DLP.

We have explained what we mean by “logic program”, and we have defined the syntax of four main logic programming languages. A study of their semantics comes next.

Model-theoretic declarative semantics

In this chapter, we define the declarative (or denotational) semantics for the logic programming languages we are interested in:¹ We start with the least Herbrand model semantics for LP programs; then we proceed to present the minimal model semantics of Minker (see [Min82] or [LMR92]) which is the *de facto* denotational semantics for DLP programs. We are by no means thorough in this chapter but the interested reader will find pointers for further studying of the semantics presented in each section.

4.1 Definitions and notation

Definition 4.1. Given a *propositional* logic program \mathcal{P} , we define its *Herbrand base* $\text{HB}(\mathcal{P})$ to be the set of all atoms that appear in \mathcal{P} .

Definition 4.2. Let \mathcal{V} be a truth value space. By a *Herbrand \mathcal{V} -interpretation* I of \mathcal{P} , we mean any assignment of truth values to the elements of the Herbrand base, i.e., any function $I : \text{HB}(\mathcal{P}) \rightarrow \mathcal{V}$. Thanks to the structure of \mathcal{V} , such a function I can be extended to all wffs generated from $\text{HB}(\mathcal{P})$ with the logical connectives in $\{\vee, \wedge, \rightarrow, \sim\}$ in the straightforward, respectful manner, just like we did in Definition 2.2. We will use the same symbol I for the extended function. When the truth value space \mathcal{V} is \mathbb{B} or when it is obvious from the context, we may omit the prefix “ \mathcal{V} -”. We say that I satisfies a rule $H \leftarrow \mathcal{B}$, if $I(\mathcal{B} \rightarrow H) = \max \mathcal{V}$. A Herbrand interpretation that satisfies every rule of \mathcal{P} is called a *Herbrand model* of \mathcal{P} .

Definition 4.3. It is customary in the logic programming literature to identify interpretations with sets of atoms when the truth value space is \mathbb{B} : for any interpretation $I : \text{HB}(\mathcal{P}) \rightarrow \mathbb{B}$, we write

$$a \in I \iff I(a) = \mathbf{T}.$$

¹We use the terms “denotational semantics” and “declarative semantics” interchangeably.

Influenced by this, in any truth value space \mathcal{V} , we may use \emptyset to denote the interpretation $I : \text{HB}(\mathcal{P}) \rightarrow \mathcal{V}$ defined by:

$$I(x) = \min \mathcal{V}, \quad \text{for all } x \in \text{HB}(\mathcal{P}),$$

i.e., the bottom element \perp of the poset $(\text{HB}(\mathcal{P}) \rightarrow \mathcal{V})$, equipped with the pointwise ordering.

4.2 LP—the least Herbrand model

For LP programs, \mathbb{B} will be the truth value space. LP programs enjoy the following very useful property:

Model intersection property. *The intersection of a non-empty family of models is itself a model.*

Observe now that for any program \mathcal{P} at least one satisfying Herbrand interpretation exists; to wit, the Herbrand base itself (i.e., the interpretation that assigns the truth value \mathbf{T} to every element of the Herbrand base). Therefore the family of all Herbrand models $\text{HM}(\mathcal{P})$ is always non-empty, which allows us to define the *least Herbrand model* as the intersection of this family:

$$\text{LHM}(\mathcal{P}) \triangleq \bigcap \text{HM}(\mathcal{P}).$$

Definition 4.4 (Least Herbrand model semantics). Let \mathcal{P} be an LP program. The goal $\leftarrow p$ *succeeds* if $p \in \text{LHM}(\mathcal{P})$.

The following result, due to van Emden and Kowalski, justifies the use of $\text{LHM}(\mathcal{P})$ as the denotational semantics of \mathcal{P} .

Theorem 4A. *Let \mathcal{P} be an LP program. Then*

$$\text{LHM}(\mathcal{P}) = \{p \in \text{HB}(\mathcal{P}) \mid p \text{ is a logical consequence of } \mathcal{P}\}.$$

Proof. See [vEK76, §5] or [Llo87, Theorem 6.2]. ■

This concludes our brief summary of LP semantics; consult [Llo87] for more information.

4.3 DLP—the minimal models

We summarize the minimal model semantics of disjunctive logic programming, using the following DLP program as a driving example:

$$\mathcal{P} := \left\{ \begin{array}{l} p \leftarrow a \\ p \leftarrow b \\ a \vee b \leftarrow \end{array} \right\}.$$

First, we compute its Herbrand models:

$$\{a, p\}, \quad \{b, p\}, \quad \{a, b, p\}.$$

If we try to follow the practice of LP, we will want to select the \subseteq -least Herbrand model to provide semantics for \mathcal{P} . But none of them is least! The model intersection property which LP programs enjoy, fails to hold in the presence of disjunctions:

$$\bigcap \text{HM}(\mathcal{P}) = \bigcap \{\{a, p\}, \{b, p\}, \{a, b, p\}\} = \{p\}.$$

And $\{p\}$ is not a model, since according to the third rule of our program, at least one of the atoms a or b must be true. However, the first two models are \subseteq -minimal. In fact, we can rely on the set $\{\{a, p\}, \{b, p\}\}$ of minimal models to obtain a meaningful semantics for \mathcal{P} .

Let \mathcal{P} be a DLP program. We write $\text{MM}(\mathcal{P})$ for the set of all \subseteq -minimal Herbrand models of \mathcal{P} . By the definition that follows, $\text{MM}(\mathcal{P})$ provides the denotational semantics for the DLP program \mathcal{P} , which we call the *minimal model semantics* à la Minker.

Definition 4.5 (Minimal model semantics). Let \mathcal{P} be a DLP program. The goal $\leftarrow \mathbf{G}$ succeeds if \mathbf{G} is true in every minimal Herbrand model of \mathcal{P} .

The equivalent of Theorem 4A for the disjunctive case is due to Minker:

Theorem 4B. Let \mathcal{P} be a DLP program. A clause C is a logical consequence of \mathcal{P} iff C is true in every minimal Herbrand model of \mathcal{P} .

Proof. See [Min82] or [LMR92, Theorem 3.3]. ■

► *Example 4.1.* Consider the DLP program

$$\mathcal{Q} := \left\{ \begin{array}{l} \mathbf{p} \leftarrow \mathbf{a} \\ \mathbf{p} \leftarrow \mathbf{b} \\ \mathbf{b} \leftarrow \mathbf{c} \\ \mathbf{a} \vee \mathbf{c} \leftarrow \end{array} \right\},$$

compute its Herbrand models, and identify the ones that are minimal:

$$\begin{aligned} \text{HM}(\mathcal{Q}) &= \{\{a, p\}, \{a, b, p\}, \{c, b, p\}, \{a, b, c, p\}\}, \\ \text{MM}(\mathcal{Q}) &= \{\{a, p\}, \{c, b, p\}\}. \end{aligned}$$

Under the minimal model semantics, p and $a \vee c$ are both **T**, as

$$\begin{array}{lcl} \{a, p\} \models p & & \{a, p\} \models a \vee c \\ \{c, b, p\} \models p & \text{and} & \{c, b, p\} \models a \vee c, \end{array}$$

while both a and $b \vee c$ are **F** because of

$$\{c, b, p\} \not\models a \quad \text{and} \quad \{a, p\} \not\models b \vee c. \quad \blacktriangleleft$$

✂ **REMARK 4.1.** In [LMR92], another denotational semantics for DLP is defined, which they call the *least model-state semantics*. However, they prove that the two approaches are equivalent, so we stick with the minimal model semantics here.

We have thus given meaning to disjunctive programs. It is time to do so for ones with negation.

4.4 LPN—the well-founded model

As we have already mentioned, interpreting negation in logic programs is a rather controversial subject. Certainly, *negation-as-failure* (see [Cla78]) is the computational interpretation with the most fruitful applications, and no other interpretation will concern us here. But even if we focus only on negation-as-failure, formally giving semantics to it is neither trivial, nor a problem with a unique satisfying solution. It is fair, however, to state that there have been two dominant and widely accepted, model-theoretic solutions: the *stable model semantics*, and the *well-founded model semantics*.

Negation-as-failure. Given a logic programming system, the main idea behind this approach to negation is that a goal $\leftarrow \sim p$ should succeed, in case the goal $\leftarrow p$ terminates with failure. For example, given the program

$$\mathcal{P} = \left\{ \begin{array}{l} p \leftarrow \\ q \leftarrow \sim p \\ r \leftarrow \sim q \end{array} \right\},$$

the query $\leftarrow q$ terminates with failure (because the goal $\leftarrow p$ terminates with success), which in turn means that the query $\leftarrow r$ succeeds. This is an operationally simple to explain idea, but defining declarative semantics for it proved to be a difficult problem; see [AB94] for a survey of the various attempts.

The stable model semantics. This solution was proposed in [GL88], and assigns to each program \mathcal{P} a certain set of well-behaved models, called *stable models*. A program may have zero, one, or more such models, which is in this approach the price to pay in order to support negation. This semantics is compatible with the least Herbrand model semantics of LP, in the sense that if negations do not appear in \mathcal{P} , then there is only one stable model: the least Herbrand model of \mathcal{P} . This school of programming eventually led to what is now known as *answer set programming* (ASP), which deviated from traditional logic programming practices and its theory and applications lie beyond this thesis; consult [Gel08] for further information.

The well-founded semantics. This semantics was introduced in [VGRS91], and actually assigns to each program \mathcal{P} a *unique* model, called *well-founded*. But there is a different price to pay here: the underlying logic is shifted from two-valued, to many-valued. The original formulation was made using a three-valued logic, which contained an additional truth value, \mathbf{U} , representing the unknown truth value. This approach is also compatible with the least Herbrand model semantics, since lack of negations in \mathcal{P} implies that no element of its well-founded model will have the truth value \mathbf{U} , and this two-valued model is indeed the least Herbrand model of \mathcal{P} . Many years after its appearance, an infinite-valued refinement of the well-founded semantics was presented in [RW05]. There is still one truth value representing the unknown, but the truth values \mathbf{T} and \mathbf{F} representing “true” and “false” respectively are each replaced by infinitely many truth values: we will use the \mathbb{V}_κ spaces for this. From this new semantics, we can obtain the original well-founded model by collapsing all “false” values to \mathbf{F} and all “true” values to \mathbf{T} . It is indeed a refinement

of the well-founded model, and in the sequel, we will refer to it simply as “the well-founded semantics”.

The well-founded semantics of LPN

The truth value space we will use for LPN is \mathbb{V}_κ . For finite programs it will suffice to consider $\kappa := \omega$, while for the general, infinite case, we will set $\kappa := \omega_1$. We introduce the following concise notation:

Definition 4.6. let $f : X \rightarrow Y$ be a function and let $y \in Y$. Then

$$f \parallel y \triangleq f^{-1}(\{y\}).$$

Definition 4.7. Let I and J be two \mathbb{V}_κ -interpretations of a program \mathcal{P} , and let $\alpha \in \kappa$. We write

$$\begin{aligned} I =_\alpha J &\triangleq \text{for all } \lambda \leq \alpha, I \parallel \mathbf{T}_\lambda = J \parallel \mathbf{T}_\lambda \text{ and } I \parallel \mathbf{F}_\lambda = J \parallel \mathbf{F}_\lambda \\ I \sqsubseteq_\alpha J &\triangleq I \parallel \mathbf{T}_\alpha \subseteq J \parallel \mathbf{T}_\alpha, I \parallel \mathbf{F}_\alpha \supseteq J \parallel \mathbf{F}_\alpha, \text{ and for all } \lambda < \alpha, I =_\lambda J \\ I \sqsubset_\alpha J &\triangleq I \sqsubseteq_\alpha J \text{ and } I \neq_\alpha J. \end{aligned}$$

We also define:

$$\begin{aligned} I \sqsubset_\kappa J &\triangleq \text{for some } \alpha \in \kappa, I \sqsubset_\alpha J \\ I \sqsubseteq_\kappa J &\triangleq I = J \text{ or } I \sqsubset_\kappa J. \end{aligned}$$

The following proposition justifies the choice of notation:

Proposition 4.1. *On the set of \mathbb{V}_κ -interpretations*

- (i) *for any $\alpha \in \kappa$, $=_\alpha$ is an equivalence relation.*
- (ii) *for any $\alpha \in \kappa$, \sqsubseteq_α is a preorder (i.e., it is reflexive and transitive);*
- (iii) *\sqsubseteq_κ is a partial order (i.e., reflexive, transitive, and antisymmetric).*

Proof. (i) is trivial. (ii) Reflexivity is immediate from the definition, while transitivity follows directly from the transitivity of \subseteq and $=_\alpha$. (iii) Reflexivity and antisymmetry are also immediate. For transitivity, assume that $I \sqsubset_\kappa J$ and $J \sqsubset_\kappa K$, so that there are ordinals α and β such that $I \sqsubset_\alpha J$ and $J \sqsubset_\beta K$. If $\alpha = \beta$, the result follows from the transitivity of \sqsubset_α . Otherwise assume without loss of generality that $\alpha < \beta$, so that $I \sqsubset_\alpha J =_\alpha K$, which implies that $I \sqsubset_\alpha K$. By the definition, this means that $I \sqsubset_\kappa K$, as we wanted. ■

Next, we define a very useful operator acting on interpretations of programs:

Definition 4.8. Let \mathcal{P} be an LPN program and let I be an interpretation. The operator $T_{\mathcal{P}}$ is defined by:

$$T_{\mathcal{P}}(I)(p) \triangleq \max \{I(\mathcal{B}) \mid (\{p\}, \mathcal{B}) \in \mathcal{P}\}.$$

We call $T_{\mathcal{P}}$ the *immediate consequence operator* of \mathcal{P} .

Given an LPN program \mathcal{P} , its well-founded model $\text{WFM}_\kappa(\mathcal{P})$ is constructed in stages, and it is best described by using $T_{\mathcal{P}}$:² we will approximate $\text{WFM}_\kappa(\mathcal{P})$ by interpretations, starting with the least element \perp of them all, i.e., the interpretation \emptyset which is the constant function with value \mathbf{F}_0 . We repeatedly iterate $T_{\mathcal{P}}$ until the set of atoms with value \mathbf{F}_0 and the set of atoms with value \mathbf{T}_0 have both stabilized, i.e., after, say, m iterations, we have

$$T_{\mathcal{P}}^m(I) \parallel \mathbf{F}_0 = T_{\mathcal{P}}^{m+1}(I) \parallel \mathbf{F}_0 \quad \text{and} \quad T_{\mathcal{P}}^m(I) \parallel \mathbf{T}_0 = T_{\mathcal{P}}^{m+1}(I) \parallel \mathbf{T}_0.$$

Once this happens, we get our next interpretation by keeping the values of these atoms the same, while resetting all other atoms to the “next” false value, in this case \mathbf{F}_1 . We proceed in a similar way until the sets of atoms with values \mathbf{F}_1 and \mathbf{T}_1 have both stabilized as well, and reset the remaining atoms to the value \mathbf{F}_2 . Atoms that do not get a value by following this procedure are set to be \mathbf{U} . Let us see how this procedure works for a concrete example, taken from [RW05]:

► *Example 4.2.* Consider the LPN program

$$\mathcal{P} = \left\{ \begin{array}{l} p \leftarrow \sim q \\ q \leftarrow \sim r \\ s \leftarrow p \\ s \leftarrow \sim s \end{array} \right\}$$

We start the process:

$$\begin{aligned} \emptyset &= \{(p, \mathbf{F}_0), (q, \mathbf{F}_0), (r, \mathbf{F}_0), (s, \mathbf{F}_0)\} \\ T_{\mathcal{P}}(\emptyset) &= \{(p, \mathbf{T}_1), (q, \mathbf{T}_1), (r, \mathbf{F}_0), (s, \mathbf{T}_1)\} \\ T_{\mathcal{P}}^2(\emptyset) &= \{(p, \mathbf{F}_2), (q, \mathbf{T}_1), (r, \mathbf{F}_0), (s, \mathbf{T}_1)\}. \end{aligned}$$

At this point, the values of order 0 have been stabilized. We reset the remaining values to \mathbf{F}_1 and continue in the same manner:

$$\begin{aligned} I_0 &= \{(p, \mathbf{F}_1), (q, \mathbf{F}_1), (r, \mathbf{F}_0), (s, \mathbf{F}_1)\} \\ T_{\mathcal{P}}(I_0) &= \{(p, \mathbf{T}_2), (q, \mathbf{T}_1), (r, \mathbf{F}_0), (s, \mathbf{T}_2)\} \\ T_{\mathcal{P}}^2(I_0) &= \{(p, \mathbf{F}_2), (q, \mathbf{T}_1), (r, \mathbf{F}_0), (s, \mathbf{T}_2)\}. \end{aligned}$$

Now the values of order 1 have been stabilized. Like before, we reset the remaining values, now to \mathbf{F}_2 :

$$\begin{aligned} I_1 &= \{(p, \mathbf{F}_2), (q, \mathbf{T}_1), (r, \mathbf{F}_0), (s, \mathbf{F}_2)\} \\ T_{\mathcal{P}}(I_1) &= \{(p, \mathbf{F}_2), (q, \mathbf{T}_1), (r, \mathbf{F}_0), (s, \mathbf{T}_3)\} \\ T_{\mathcal{P}}^2(I_1) &= \{(p, \mathbf{F}_2), (q, \mathbf{T}_1), (r, \mathbf{F}_0), (s, \mathbf{F}_4)\}, \end{aligned}$$

and values of order 2 have been stabilized. Resetting the remaining values to \mathbf{F}_3 , we compute:

$$\begin{aligned} I_2 &= \{(p, \mathbf{F}_2), (q, \mathbf{T}_1), (r, \mathbf{F}_0), (s, \mathbf{F}_3)\} \\ T_{\mathcal{P}}(I_2) &= \{(p, \mathbf{F}_2), (q, \mathbf{T}_1), (r, \mathbf{F}_0), (s, \mathbf{T}_4)\}. \end{aligned}$$

²Readers familiar with *stratification* will probably recognize the similarities.

Observe now that values of order 3 have disappeared completely, which means that s cannot obtain any value of order less than ω , so we set its value to \mathbf{U} , finally reaching the well-founded model:

$$\text{WFM}_\kappa(\mathcal{P}) = \{(p, \mathbf{F}_2), (q, \mathbf{T}_1), (r, \mathbf{F}_0), (s, \mathbf{U})\}. \quad \blacktriangleleft$$

The \mathbb{V}_κ -valued well-founded model has the following nice property, proved in [RW05]:

Theorem 4C. *Let \mathcal{P} be an LPN program. Then $\text{WFM}_\kappa(\mathcal{P})$ is the \sqsubseteq_κ -least model of \mathcal{P} .*

Given $\text{WFM}_\kappa(\mathcal{P})$, it is immediate to obtain the original, three-valued well-founded model $\text{WFM}(\mathcal{P})$ as defined in [VGRS91]. The following theorem, proven in [RW05], tells us how:

Theorem 4D. *Let \mathcal{P} be an LPN program, and let $p \in \text{HB}(\mathcal{P})$. Then*

$$\text{WFM}(\mathcal{P}) = \text{collapse} \circ \text{WFM}_\kappa(\mathcal{P}).$$

4.5 DLPN—the infinite-valued minimal models

In [CPRW07], the approach studied above was extended to be applicable to *finite* DLPN programs as well. We impose the same restriction as in [CPRW07] and consider only finite programs throughout this section. This also implies that \mathbb{V}_ω is a sufficiently large truth value space.

The shift from definite to disjunctive in the case of negation-free programs resulted in the “compromise” of using a *set* of minimal models, instead of a single least one. The price to pay is the same as we shift from LPN to DLPN programs, although in this case, naturally, these minimal models are many-valued. The ideas involved are rather intuitive, and so we directly jump to an example before stating the results that we are interested in:

► *Example 4.3.* Consider the program

$$\mathcal{P} := \left\{ \begin{array}{l} \mathbf{a} \vee \mathbf{b} \leftarrow \sim \mathbf{p} \\ \mathbf{b} \vee \mathbf{p} \leftarrow \end{array} \right\}.$$

The only fact in this program will force any model M to include (b, \mathbf{T}_0) or (p, \mathbf{T}_0) . In the first case, we immediately obtain the minimal model

$$M_1 := \{(a, \mathbf{F}_0), (b, \mathbf{T}_0), (p, \mathbf{F}_0)\}.$$

In the second case, the fact that p is \mathbf{T}_0 implies that the value of $\sim p$ must be \mathbf{F}_1 , which in turn forces either either a or b to have a value at least as big as \mathbf{F}_1 , so we reach two more minimal models:

$$\begin{aligned} M_2 &:= \{(a, \mathbf{F}_1), (b, \mathbf{F}_0), (p, \mathbf{T}_0)\} \\ \text{and } M_3 &:= \{(a, \mathbf{F}_0), (b, \mathbf{F}_1), (p, \mathbf{T}_0)\}. \end{aligned}$$

Therefore, the set of \mathbb{V}_ω -valued minimal models of \mathcal{P} is

$$\text{MM}_\omega(\mathcal{P}) = \{M_1, M_2, M_3\}. \quad \blacktriangleleft$$

The following are the main results regarding this semantics:

Theorem 4E. *Every DLPN program \mathcal{P} has a non-empty, finite set of \sqsubseteq_ω -minimal, \mathbb{V}_ω -valued models $\text{MM}_\omega(\mathcal{P})$.*

Proof. See [CPRW07]. ■

It is argued in [CPRW07] that this set of minimal models captures the intended meaning of a DLPN program, and it is implied that it indeed provides a semantics in an analogous way to the DLP case: the value of a query G can be obtained as the minimum of the values that G has in each of those minimal models.

Theorem 4F. *Let \mathcal{P} be a DLPN program that contains n propositional variables in its rules, and let M be a minimal model of \mathcal{P} . Then $\text{Range}(M) \subseteq \mathbb{V}_n$.*

Proof. See [CPRW07]. ■

The main reason why we focus on this semantics for DLPN programs is the fact that it is purely model-theoretic and it naturally extends both the well-founded semantics of LPN and the minimal model semantics of DLP, as the following theorem states:

Theorem 4G. *Let \mathcal{P} be a DLPN program. Then*

- *if \mathcal{P} is an LPN program, $\text{MM}_\omega(\mathcal{P}) = \text{WFM}_\omega(\mathcal{P})$;*
- *if \mathcal{P} is a DLP program, $\text{MM}_\omega(\mathcal{P}) = \text{MM}(\mathcal{P})$.*

Proof. See [CPRW07]. ■

4.6 First-order programs vs. infinite propositional ones

Even though we focus completely on propositional logic programs, we summarize below the corresponding notions of the ones we described above, for the first-order case. They are presented in detail in [LMR92, §2.2, §2.4, §3.2]. The reader can safely skip this section on a first reading, coming back to it only before the very last result of Section 10.6, viz. Corollary 10.33.

Definition 4.9. The *Herbrand universe* of \mathcal{L}_1 , HU , is the set of all ground terms which can be formed out of the constant and function symbols of \mathcal{L}_1 . The *Herbrand base* of a logic program \mathcal{P} , $\text{HB}(\mathcal{P})$, is the set of all ground atoms which can be formed using the predicate symbols that appear in \mathcal{P} on the ground terms of HU .

Definition 4.10. A *Herbrand interpretation* of \mathcal{L}_1 assigns: to each n -ary predicate symbol R of Pred_n , an n -ary function from HU^n to \mathbb{B} ; to each n -ary function symbol $f \in \text{Fun}_n$ the function that maps the element $(t_1, \dots, t_n) \in \text{HU}^n$ to the term $f(t_1, \dots, t_n) \in \text{Term}$; and to each constant symbol, itself. A Herbrand interpretation is naturally extended to evaluate all formulæ of \mathcal{L}_1 according to the usual truth tables. If a Herbrand interpretation I is a model of \mathcal{P} we call it a *Herbrand model* of \mathcal{P} .

It is not difficult to see that infinite, propositional programs are as powerful as finite, first-order ones. Hence, for simplicity, we will be dealing with propositional logic programs unless mentioned otherwise. To see how we end up with infinite programs, start from a non-propositional, finite program \mathcal{P} , containing at least one function symbol, and replace each of its rules by all of its ground instances. What you get is a countably infinite program with equivalent denotational semantics. This is exactly what happens in Example 3.3 (p. 16).

► *Example 4.4.* Let us confine ourselves to propositional (albeit infinite) programs. Then, instead of the finite first-order program \mathcal{E} of Example 3.3 (repeated here for convenience together with $\text{ground}(\mathcal{E})$), we could write, for instance, the infinite propositional program \mathcal{E}_0 :

$$\mathcal{E} := \left\{ \begin{array}{l} \text{even}(0) \leftarrow \\ \text{even}(\text{S}(\text{S}(X))) \leftarrow \text{even}(X) \end{array} \right\},$$

$$\text{ground}(\mathcal{E}) := \left\{ \begin{array}{l} \text{even}(0) \leftarrow \\ \text{even}(\text{S}^2(0)) \leftarrow \text{even}(0) \\ \text{even}(\text{S}^3(0)) \leftarrow \text{even}(\text{S}(0)) \\ \text{even}(\text{S}^4(0)) \leftarrow \text{even}(\text{S}^2(0)) \\ \vdots \end{array} \right\}, \quad \mathcal{E}_0 := \left\{ \begin{array}{l} \mathbf{e}_0 \leftarrow \\ \mathbf{e}_2 \leftarrow \mathbf{e}_0 \\ \mathbf{e}_3 \leftarrow \mathbf{e}_1 \\ \mathbf{e}_4 \leftarrow \mathbf{e}_2 \\ \vdots \end{array} \right\},$$

in which we have chosen a propositional variable \mathbf{e}_i for each first-order atomic formula $\text{even}(\text{S}^i(0))$ of $\text{ground}(\mathcal{E})$. ◀

Notice, that since the Herbrand models of a first-order logic program \mathcal{P} are constructed using only ground instances of rules in \mathcal{P} , the following property is immediate:

Property 4.2. *Let \mathcal{P} be a first-order DLP. Then \mathcal{P} and $\text{ground}(\mathcal{P})$ have the same minimal models:*

$$\text{MM}(\mathcal{P}) = \text{MM}(\text{ground}(\mathcal{P})).$$

✂ **REMARK 4.2.** As long as we are studying denotational semantics, the above property allows us to focus on infinite, propositional programs. This is a very common practice in the field of logic programming semantics. In fact, we assume given an implementation of our programming language (based on some operational semantics—which, is immaterial). We then load our finite, first-order program \mathcal{P} , and the implementation computes answers to our queries. Then, a denotational semantics of the equivalent, propositional, and infinite DLP program $\text{ground}(\mathcal{P})$, provides a correctness criterion for those answers. Thus, we avoid variables and function symbols at the cost of finiteness, but this is a fair bargain, as no difficulties arise on the declarative side of semantics (see also [Fit99] for a relevant discussion). In exactly the same sense, when we are giving a fixpoint semantics, we only use the set of ground instances of clauses in \mathcal{P} to define $T_{\mathcal{P}}$ (see [Llo87] and [LMR92]).

We have defined the syntax and the model-theoretic declarative semantics of the logic programming languages that interest us. In the next part, we define an abstract

framework for semantics and develop a toolkit for dealing with disjunctive programs. The goal is to describe an operator that acts on semantics of non-disjunctive languages, and yields a new semantics for a corresponding disjunctive language.

ξ

Part II

From non-disjunctive to disjunctive semantics

An abstract framework for semantics

In order to study logic programming languages and their semantics, we introduce in this chapter a formal framework of semantics and examine the four languages we have met with respect to this framework.

5.1 The framework

Definition 5.1. Let L be a logic programming language, let \mathcal{M} be a set whose elements we will call *meanings*, and let \mathcal{V} be a truth value space. Then, an \mathcal{M} -*semantics* for L is a function

$$\mathbf{m} : \mathbf{P}_L \rightarrow \mathcal{M};$$

a \mathcal{V} -*answer function* for \mathcal{M} is a function

$$\mathbf{a} : \mathcal{M} \rightarrow \mathbf{Q}_L \rightarrow \mathcal{V};$$

and a \mathcal{V} -*system* for L is a function

$$\mathbf{s} : \mathbf{P}_L \rightarrow \mathbf{Q}_L \rightarrow \mathcal{V}.$$

A pair (\mathbf{m}, \mathbf{a}) is simply called a *semantics* for L .

¶ **REMARK 5.1.** Composing a \mathcal{V} -answer function for \mathcal{M} with an \mathcal{M} -semantics for L we obtain a \mathcal{V} -system for L . Therefore, a semantics (\mathbf{m}, \mathbf{a}) naturally gives rise to the \mathcal{V} -system $\mathbf{a} \circ \mathbf{m}$. In this way, we will be able to use (\mathbf{m}, \mathbf{a}) in a context where a \mathcal{V} -system is expected.

Definition 5.2 (\approx). Let L be a logic programming language. We call two semantics of L $(\mathbf{m}_1, \mathbf{a}_1)$ and $(\mathbf{m}_2, \mathbf{a}_2)$ *equivalent* iff the corresponding \mathcal{V} -systems are equal. In symbols,

$$(\mathbf{m}_1, \mathbf{a}_1) \approx (\mathbf{m}_2, \mathbf{a}_2) \iff \mathbf{a}_1 \circ \mathbf{m}_1 = \mathbf{a}_2 \circ \mathbf{m}_2.$$

Notice that \approx is an equivalence relation. When the context clearly hints the intended \mathcal{V} -answer functions under consideration, we might abuse the notation and simply write $\mathbf{m}_1 \approx \mathbf{m}_2$ instead.

Definition 5.3. Let L be a logic programming language. We say that $(\mathbf{m}_1, \mathbf{a}_1)$ *refines* $(\mathbf{m}_2, \mathbf{a}_2)$ with respect to the operator \mathcal{C} iff:

$$(\mathbf{m}_1, \mathbf{a}_1) \triangleleft_{\mathcal{C}} (\mathbf{m}_2, \mathbf{a}_2) \iff \begin{cases} \mathbf{m}_i : \mathbf{P}_L \rightarrow \mathcal{M}_i \\ \mathbf{a}_i : \mathcal{M}_i \rightarrow \mathbf{Q}_L \rightarrow \mathcal{V} \\ \mathcal{C} : \mathcal{M}_1 \rightarrow \mathcal{M}_2 \\ \mathbf{m}_2 = \mathcal{C} \circ \mathbf{m}_1 \\ \mathbf{a}_1 = \mathbf{a}_2 \circ \mathcal{C}. \end{cases}$$

If only \mathbf{m}_1 , \mathbf{m}_2 , \mathbf{a}_2 , and \mathcal{C} are given, we write $\mathbf{m}_1 \triangleleft_{\mathcal{C}}^{\mathbf{a}_2} \mathbf{m}_2$ for $(\mathbf{m}_1, \mathbf{a}_2 \circ \mathcal{C}) \triangleleft_{\mathcal{C}} (\mathbf{m}_2, \mathbf{a}_2)$, and might even omit the superscript \mathbf{a}_2 when it is obvious or implicit by the context.

Lemma 5.1. *The following implication holds:*

$$(\mathbf{m}_1, \mathbf{a}_1) \triangleleft_{\mathcal{C}} (\mathbf{m}_2, \mathbf{a}_2) \implies (\mathbf{m}_1, \mathbf{a}_1) \approx (\mathbf{m}_2, \mathbf{a}_2). \quad (5.1.1)$$

Proof. We have $\mathbf{a}_1 \circ \mathbf{m}_1 = \mathbf{a}_2 \circ \mathcal{C} \circ \mathbf{m}_1 = \mathbf{a}_2 \circ \mathbf{m}_2$, which by the definition of \approx is equivalent to $(\mathbf{m}_1, \mathbf{a}_1) \approx (\mathbf{m}_2, \mathbf{a}_2)$. \blacksquare

5.2 Semantics of LP

LHM: the least Herbrand model semantics

This is the *least Herbrand model semantics* we met in Section 4.2.

- $\mathcal{V}_{\text{LHM}} \triangleq \mathbb{B}$.
- \mathcal{M}_{LHM} is the set of all possible Herbrand interpretations.
- \mathbf{m}_{LHM} maps an LP program to its least Herbrand model.
- $\mathbf{a}_{\text{LHM}}(M)(p) \triangleq \begin{cases} \mathbf{T}, & \text{if } p \in M \\ \mathbf{F}, & \text{otherwise.} \end{cases}$

5.3 Semantics of DLP

MM: the minimal model semantics

This is the *minimal models semantics* of Section 4.3.

- $\mathcal{V}_{\text{MM}} \triangleq \mathbb{B}$.
- \mathcal{M}_{MM} consists of all sets of Herbrand interpretations.
- \mathbf{m}_{MM} maps a DLP program to the set of its minimal models.
- $\mathbf{a}_{\text{MM}}(\mathcal{M})(Q) \triangleq \begin{cases} \mathbf{T}, & \text{if } Q \text{ is true in every model } M \in \mathcal{M} \\ \mathbf{F}, & \text{otherwise.} \end{cases}$

5.4 Semantics of LPN

WF: the well-founded semantics

This is the *well-founded model semantics* for LPN, as described in Section 4.4.

- $\mathcal{V}_{\text{WF}} \triangleq \mathbb{V}_1 = \{\mathbf{F}, \mathbf{U}, \mathbf{T}\}$.
- \mathcal{M}_{WF} consists of all possible Herbrand \mathbb{V}_1 -interpretations.
- \mathbf{m}_{WF} maps every LPN program to its three-valued, well-founded model.
- $\mathbf{a}_{\text{WF}}(M)(p) \triangleq M(p)$.

WF $^\kappa$: the infinite-valued well-founded semantics

This is the *infinite-valued well-founded semantics* for LPN, also described in Section 4.4.

- $\mathcal{V}_{\text{WF}^\kappa} \triangleq \mathbb{V}_\kappa$.
- $\mathcal{M}_{\text{WF}^\kappa}$ consists of all possible Herbrand \mathbb{V}_κ -interpretations of LPN programs.
- $\mathbf{m}_{\text{WF}^\kappa}$ maps every LPN program to its \mathbb{V}_κ -valued, well-founded model.
- $\mathbf{a}_{\text{WF}^\kappa}(M)(p) \triangleq M(p)$.

✎ REMARK 5.2 (The ordinal κ). The ordinal κ that we use in the truth value spaces \mathbb{V}_κ may vary, depending on our needs. The reader should note at this point that if the programs are finite, an ordinal as small as ω suffices to give us satisfying semantics, in the sense that collapsing the obtained \mathbb{V}_ω -valued model to a three-valued one will always yield the desired well-founded model. See [RW05] and [Lüd11] for more information.

5.5 Semantics of DLPN

MM $^\omega$: the infinite-valued minimal model semantics

The infinite-valued minimal model semantics, which we introduced in Section 4.5, and apply for finite, propositional DLPN programs.

- $\mathcal{V}_{\text{MM}^\omega} \triangleq \mathbb{V}_\omega$.
- $\mathcal{M}_{\text{MM}^\omega}$ consists of all possible \mathbb{V}_ω -valued Herbrand interpretations of DLPN programs.
- $\mathbf{m}_{\text{MM}^\omega}$ maps every DLPN program to the set of its minimal, infinite-valued models.
- $\mathbf{a}_{\text{MM}^\omega}(\mathcal{M})(q) \triangleq \bigwedge \{M(q) \mid M \in \mathcal{M}\}$.

Disjunctive operations

In this chapter, we define some operations that can be applied to various disjunctive components of logic programs, for example rules, or even whole programs. We will extensively rely on these operations in the sequel.

6.1 Restricting

Definition 6.1 (Rule restriction). For any given rule ϕ and any set of atoms A , we define the *restriction* of ϕ to A by

$$\phi|_A \triangleq (\text{head}(\phi) \cap A, \text{body}(\phi)).$$

It follows that a rule's body is unaffected by restriction: $\text{body}(\phi) = \text{body}(\phi|_A)$.

Definition 6.2 (Program restriction). Let \mathcal{P} be a DLP or DLPN program and let $\phi \in \mathcal{P}$. Then for any set of atoms A , we can define the *restricted program* $\mathcal{P}|_A^\phi$ by restricting the rule ϕ of \mathcal{P} to A :

$$\mathcal{P}|_A^\phi \triangleq (\mathcal{P} \setminus \{\phi\}) \cup \{\phi|_A\}.$$

In words, $\mathcal{P}|_A^\phi$ is the program which is identical to \mathcal{P} , with the exception that the rule ϕ has been replaced by $\phi|_A$.

► *Example 6.1.* Let \mathcal{P} be the proper DLP program

$$\mathcal{P} := \left\{ \begin{array}{l} p \vee q \vee r \leftarrow f, g \\ p \vee q \vee r \leftarrow a, c \\ f \leftarrow \end{array} \right\},$$

and let $\phi := p \vee q \vee r \leftarrow a, c$. Here are a couple of restrictions of \mathcal{P} with respect to ϕ :

$$\mathcal{P}|_{\{p,q\}}^\phi := \left\{ \begin{array}{l} p \vee q \vee r \leftarrow f, g \\ p \vee q \leftarrow a, c \\ f \leftarrow \end{array} \right\}, \quad \mathcal{P}|_{\{r\}}^\phi := \left\{ \begin{array}{l} p \vee q \vee r \leftarrow f, g \\ r \leftarrow a, c \\ f \leftarrow \end{array} \right\}. \quad \blacktriangleleft$$

We observe the following useful property:

Property 6.1. *The restriction of a rule ϕ is stronger than the original rule, in the sense that any interpretation which satisfies $\phi|_A$ must also satisfy ϕ . In the same sense, restricting a disjunctive program makes it stronger.*

6.2 Splitting

We will frequently fix a disjunction H , and break it into logically stronger, less disjunctive parts. For this we introduce the notion of a proper partition:

Definition 6.3 (Proper partition). Suppose that H is a set of atoms. Then a tuple $\mathcal{H} := (H_1, \dots, H_n)$ is a proper partition of H iff

- (i) $\emptyset \neq H_i \subsetneq H$ for all i ;
- (ii) $H = H_1 \cup \dots \cup H_n$;
- (iii) $H_i \cap H_j = \emptyset$ for all $i \neq j$.

Once we know how to restrict a program, splitting it with respect to some partition \mathcal{H} becomes trivial:

Definition 6.4 (Splitting a program). Let ϕ be a proper disjunctive rule, and let $\mathcal{H} := (H_1, \dots, H_n)$ be a proper partition of its head. Then the *splitting of \mathcal{P} with respect to ϕ over \mathcal{H}* is the tuple

$$\mathcal{P}|_{\mathcal{H}}^{\phi} \triangleq (\mathcal{P}|_{H_1}^{\phi}, \dots, \mathcal{P}|_{H_n}^{\phi}).$$

► *Example 6.2.* The pair $(\mathcal{P}|_{\{p,q\}}^{\phi}, \mathcal{P}|_{\{r\}}^{\phi})$ of Example 6.1 is the splitting of \mathcal{P} with respect to ϕ over $\mathcal{H} := (\{p, q\}, \{r\})$. ◀

To prove the main result of Chapter 10 we need the following lemma, which relates the models of a splitting of a program with those of the original program.

Lemma 6.2 (Inclusions). *Let $(\mathcal{P}_1, \mathcal{P}_2)$ be the splitting of a DLP program \mathcal{P} with respect to ϕ over (H_1, H_2) . Then*

$$\text{MM}(\mathcal{P}) \subseteq \text{MM}(\mathcal{P}_1) \cup \text{MM}(\mathcal{P}_2) \subseteq \text{HM}(\mathcal{P}).$$

Proof. Let $\phi_1 := \phi|_{H_1}$ and $\phi_2 := \phi|_{H_2}$. For the first inclusion, let $S \in \text{MM}(\mathcal{P})$, and suppose $S \notin \text{MM}(\mathcal{P}_1)$. We need $S \in \text{MM}(\mathcal{P}_2)$. There are two ways in which S can fail to be in $\text{MM}(\mathcal{P}_1)$: either it is a model but not a minimal one, or it is not even a model to begin with.

CASE 1: *S is a non-minimal model of \mathcal{P}_1 .* There exists then, a proper sub-model $S_0 \subsetneq S$ of \mathcal{P}_1 , with $S_0 \models \mathcal{P}_1$. By definition and Property 6.1, this would also be a model of \mathcal{P} , and therefore S would not be minimal in \mathcal{P} , which is a contradiction, and so this case may never be.

CASE 2: *S is not a model of \mathcal{P}_1 .*

$$\begin{aligned} S \in \text{MM}(\mathcal{P}) &\implies S \in \text{HM}(\mathcal{P}) \\ &\implies S \models \psi \text{ for all } \psi \in \mathcal{P} \\ &\implies S \models \psi \text{ for all } \psi \in \mathcal{P}_1 \setminus \{\phi_1\} && (\mathcal{P}_1 \setminus \{\phi_1\} \subset \mathcal{P}) \\ &\implies S \not\models \phi_1 && \text{(by case hypothesis)} \end{aligned}$$

Since (ϕ_1, ϕ_2) is the splitting of ϕ over (H_1, H_2) , S is forced to satisfy ϕ_2 , and therefore satisfies every element of \mathcal{P}_2 , so that $S \in \text{HM}(\mathcal{P}_2)$.

It remains to show that it is minimal in \mathcal{P}_2 . But if it was not, we would arrive at the same contradiction as in case 1; therefore, $S \in \text{MM}(\mathcal{P}_2)$. We have proved the inclusion

$$\text{MM}(\mathcal{P}) \subseteq \text{MM}(\mathcal{P}_1) \cup \text{MM}(\mathcal{P}_2). \quad (1)$$

Obviously now, $\text{MM}(\mathcal{P}_1) \subseteq \text{HM}(\mathcal{P}_1)$ and $\text{MM}(\mathcal{P}_2) \subseteq \text{HM}(\mathcal{P}_2)$, so by taking unions on both sides we obtain

$$\text{MM}(\mathcal{P}_1) \cup \text{MM}(\mathcal{P}_2) \subseteq \text{HM}(\mathcal{P}_1) \cup \text{HM}(\mathcal{P}_2). \quad (2)$$

As \mathcal{P} is a weaker program than its restrictions (by Property 6.1), we also have the inclusions $\text{HM}(\mathcal{P}_1) \subseteq \text{HM}(\mathcal{P})$ and $\text{HM}(\mathcal{P}_2) \subseteq \text{HM}(\mathcal{P})$. Hence, by taking unions for one last time,

$$\text{HM}(\mathcal{P}_1) \cup \text{HM}(\mathcal{P}_2) \subseteq \text{HM}(\mathcal{P}). \quad (3)$$

Putting (1)–(3) together:

$$\text{MM}(\mathcal{P}) \subseteq \text{MM}(\mathcal{P}_1) \cup \text{MM}(\mathcal{P}_2) \subseteq \text{HM}(\mathcal{P}_1) \cup \text{HM}(\mathcal{P}_2) \subseteq \text{HM}(\mathcal{P}). \quad \blacksquare$$

¶ **REMARK 6.1.** One might be tempted to believe that some of these inclusions are actually equalities, but none of them holds in general, as the following example demonstrates: consider the DLP program \mathcal{P} and its splitting

$$\mathcal{P} := \left\{ \begin{array}{l} \mathbf{b} \leftarrow \\ \mathbf{a} \vee \mathbf{b} \vee \mathbf{c} \leftarrow \end{array} \right\} \rightsquigarrow \left(\mathcal{P}_1 := \left\{ \begin{array}{l} \mathbf{b} \leftarrow \\ \mathbf{a} \leftarrow \end{array} \right\}, \quad \mathcal{P}_2 := \left\{ \begin{array}{l} \mathbf{b} \leftarrow \\ \mathbf{b} \vee \mathbf{c} \leftarrow \end{array} \right\} \right).$$

Then the corresponding sets of Herbrand models and minimal models are

$$\begin{aligned} \text{HM}(\mathcal{P}) &= \{\{b\}, \{b, a\}, \{b, c\}, \{b, a, c\}\}, & \text{MM}(\mathcal{P}) &= \{\{b\}\}, \\ \text{HM}(\mathcal{P}_1) &= \{\{a, b\}\}, & \text{MM}(\mathcal{P}_1) &= \{\{a, b\}\}, \\ \text{HM}(\mathcal{P}_2) &= \{\{b\}, \{b, c\}\}, & \text{MM}(\mathcal{P}_2) &= \{\{b\}\}, \end{aligned}$$

so that in this example all inclusions are proper:

$$\text{MM}(\mathcal{P}) \subsetneq \text{MM}(\mathcal{P}_1) \cup \text{MM}(\mathcal{P}_2) \subsetneq \text{HM}(\mathcal{P}_1) \cup \text{HM}(\mathcal{P}_2) \subsetneq \text{HM}(\mathcal{P}).$$

6.3 Combinations

When working with a game semantics for disjunctive programs, given a sequence of disjunctions, we will frequently want to *disjunctively combine* them into a single disjunction. In a similar fashion, we wish to combine conjunctions, rules, and later even plays and strategies! Informally speaking, the idea is always the same: we combine two or more “disjunctive things” into a single such thing, by using some kind of logical disjunction. Thus, the resulting combination will be “more disjunctive” than either of them. Even though we use the same notation for all of the different kinds of disjunctive combinations, it will always be clear what type of elements we are dealing with, and so no confusion should arise. This overloaded notation is rather convenient and really pays off in terms of readability.

In the definitions that follow, we introduce combination to deal with disjunctions, conjunctions, and rules.

Definition 6.5 (Combination of disjunctions). The (disjunctive) *combination* of two disjunctions is simply their union:

$$D \vee E \triangleq D \cup E.$$

Definition 6.6 (Combination of conjunctions). Given two conjunctions of disjunctions $\mathcal{D} := \langle D_1, \dots, D_n \rangle$ and $\mathcal{E} := \langle E_1, \dots, E_m \rangle$, their *combination* is another conjunction of disjunctions, logically equivalent to $\mathcal{D} \vee \mathcal{E}$, and defined by the following equation:

$$\mathcal{D} \vee \mathcal{E} \triangleq \langle D_1 \vee E_1, \dots, D_1 \vee E_m, \dots, D_n \vee E_1, \dots, D_n \vee E_m \rangle.$$

Notice that the sequence on the right is empty iff any of \mathcal{D} or \mathcal{E} is empty.

► *Example 6.3.* Combining the following two conjunctions of disjunctions

$$\begin{array}{ll} \mathbf{p, q \vee r, q \vee b} & \text{i.e., } \langle \{p\}, \{q, r\}, \{b, q\} \rangle \\ \mathbf{a \vee b, a \vee r} & \text{i.e., } \langle \{a, b\}, \{a, r\} \rangle \end{array}$$

in the given order, we obtain the conjunction

$$\mathbf{p \vee a \vee b, p \vee a \vee r, q \vee r \vee a \vee b, q \vee r \vee a, q \vee b \vee a, q \vee b \vee a \vee r,}$$

i.e., $\langle \{a, b, p\}, \{a, p, r\}, \{a, b, q, r\}, \{a, q, r\}, \{a, b, q\}, \{a, b, q, r\} \rangle$. ◀

Definition 6.7 (Combination of rules). The *combination* of two disjunctive rules $\phi_1 := (H_1, \mathcal{D}_1)$ and $\phi_2 := (H_2, \mathcal{D}_2)$ is the rule defined by

$$\phi_1 \vee \phi_2 \triangleq (H_1 \vee H_2, \mathcal{D}_1 \vee \mathcal{D}_2).$$

⚡ **REMARK 6.2.** It follows that $\phi_1 \vee \phi_2$ will not be a clean rule in general, unless one of the ϕ_i 's is a fact. This is the only construction that does not preserve cleanliness. However, we will only use it to extract the head of the combined rule (which causes no trouble),¹ and not to create new, potentially unclean rules.

► *Example 6.4.* Combining the following two rules

$$\begin{array}{ll} \mathbf{p \vee q \leftarrow a \vee b, c} & \text{i.e., } (\{p, q\}, \langle \{a, b\}, \{c\} \rangle) \\ \mathbf{p \vee r \leftarrow d, e} & \text{i.e., } (\{p, r\}, \langle \{d\}, \{e\} \rangle) \end{array}$$

we obtain

$$\mathbf{p \vee q \vee r \leftarrow a \vee b \vee d, a \vee b \vee e, c \vee d, c \vee e,}$$

i.e., $(\{p, q, r\}, \langle \{a, b, d\}, \{a, b, e\}, \{c, d\}, \{c, e\} \rangle)$. ◀

Hitherto we have defined combination for pairs of disjunctions, conjunctions, and rules. We can generalize these definitions from pairs to sequences in a straightforward way:

¹Notice that $\text{head}(\phi_1 \vee \phi_2) = \text{head}(\phi_1) \vee \text{head}(\phi_2)$.

Definition 6.8 (Combining sequences). Given a sequence $\langle T_1, \dots, T_n \rangle$ of combinable disjunctions, conjunctions, or rules, we set

$$[\langle T_1, \dots, T_n \rangle] \triangleq T_1 \curlywedge \dots \curlywedge T_n,$$

where \curlywedge is understood to associate to the left, and \emptyset , $\langle \emptyset \rangle$, or $(\emptyset, \langle \emptyset \rangle)$ is its unit, depending on whether we are combining disjunctions, conjunctions, or rules respectively.

An alternative presentation of the same definition uses recursion:

$$[\langle \rangle] \triangleq \begin{cases} \emptyset & \text{(disjunctions)} \\ \langle \emptyset \rangle & \text{(conjunctions of disjunctions)} \\ (\emptyset, \langle \emptyset \rangle) & \text{(disjunctive rules)} \end{cases}$$

$$[\langle T_1, \dots, T_{n+1} \rangle] \triangleq [\langle T_1, \dots, T_n \rangle] \curlywedge T_{n+1}.$$

6.4 Definite instantiations and D-sections

First we need to define the definite instantiations of a disjunctive program \mathcal{D} . Informally, a definite instantiation of \mathcal{D} is what we get by replacing each head of \mathcal{D} by one of its elements. Formally, we define:

Definition 6.9. Let $\phi = (H, \mathcal{B})$ be a disjunctive rule. If $h \in H$, then the definite rule (h, \mathcal{B}) is a *definite instantiation* of ϕ . $D(\phi)$ is the set of all definite instantiations of ϕ .

► *Example 6.5.* Here are some disjunctive rules and their respective sets of definite instantiations:

$$\begin{aligned} \phi_1 &:= \mathbf{a} \vee \mathbf{b} \leftarrow \mathbf{p}, \sim \mathbf{q} & D(\phi_1) &= \left\{ \begin{array}{l} \mathbf{a} \leftarrow \mathbf{p}, \sim \mathbf{q} \\ \mathbf{b} \leftarrow \mathbf{p}, \sim \mathbf{q} \end{array} \right\}, \\ \phi_2 &:= \mathbf{e} \vee \mathbf{f} \vee \mathbf{g} \leftarrow & D(\phi_2) &= \left\{ \begin{array}{l} \mathbf{e} \leftarrow \\ \mathbf{f} \leftarrow \\ \mathbf{g} \leftarrow \end{array} \right\}, \\ \phi_3 &:= \mathbf{p} \vee \mathbf{q} \leftarrow \mathbf{a}, \mathbf{b} \vee \mathbf{c} & D(\phi_3) &= \left\{ \begin{array}{l} \mathbf{p} \vee \mathbf{q} \leftarrow \mathbf{a}, \mathbf{b} \vee \mathbf{c} \\ \mathbf{p} \vee \mathbf{q} \leftarrow \mathbf{a}, \mathbf{b} \vee \mathbf{c} \end{array} \right\}. \quad \blacktriangleleft \end{aligned}$$

Definition 6.10. Let $\mathcal{D} = \{(H_i, \mathcal{B}_i)\}_{i \in I}$ be a disjunctive program, indexed by some set of indices I . A *\mathcal{D} -section* is any choice function $f \in \prod_{i \in I} H_i$. We write $S(\mathcal{D})$ for the set of all \mathcal{D} -sections. If f is a \mathcal{D} -section, we define the *definite instantiation of \mathcal{D} under f* to be the definite program

$$\mathcal{D}_f \triangleq \{(\{f(i)\}, \mathcal{B}_i)\}_{i \in I}.$$

We call \mathcal{P} a *definite instantiation* of \mathcal{D} , if there is a \mathcal{D} -section f such that $\mathcal{P} = \mathcal{D}_f$. Finally, we write $D(\mathcal{D})$ for the set of all definite instantiations of \mathcal{D} .

► *Example 6.6.* Consider the disjunctive program

$$\mathcal{D} := \left\{ \begin{array}{l} \phi_1 := \mathbf{s} \vee \textcircled{\mathbf{t}} \leftarrow \mathbf{p}, \mathbf{b} \vee \mathbf{c} \\ \phi_2 := \textcircled{\mathbf{a}} \vee \mathbf{b} \leftarrow \\ \phi_3 := \textcircled{\mathbf{p}} \leftarrow \mathbf{a} \\ \phi_4 := \textcircled{\mathbf{p}} \leftarrow \mathbf{b} \\ \phi_5 := \textcircled{\mathbf{b}} \vee \mathbf{c} \leftarrow \end{array} \right\}.$$

There are 8 \mathcal{D} -sections in total, and 8 definite instantiations of \mathcal{D} . Let $f, g \in S(\mathcal{D})$ be the following two of them:

$$\begin{aligned} f &:= \{(1, t), (2, a), (3, p), (4, p), (5, b)\} \\ g &:= \{(1, s), (2, b), (3, p), (4, p), (5, b)\}. \end{aligned}$$

From these two \mathcal{D} -sections we obtain two elements of the $D(\mathcal{D})$ set:

$$\mathcal{D}_f = \left\{ \begin{array}{l} \text{t} \leftarrow \text{p}, \text{b} \vee \text{c} \\ \text{a} \leftarrow \\ \text{p} \leftarrow \text{a} \\ \text{p} \leftarrow \text{b} \\ \text{b} \leftarrow \end{array} \right\} \quad \text{and} \quad \mathcal{D}_g = \left\{ \begin{array}{l} \text{s} \leftarrow \text{p}, \text{b} \vee \text{c} \\ \text{b} \leftarrow \\ \text{p} \leftarrow \text{a} \\ \text{p} \leftarrow \text{b} \\ \text{b} \leftarrow \end{array} \right\}.$$

Notice that f and \mathcal{D}_f correspond to the choices that appear circled on the program \mathcal{D} above. ◀

This completes our toolkit that will aid us in manipulating disjunctive programs. It is now time to put these tools into use: in the next chapter, we show how we can extend any semantics of a definite language into a semantics of a corresponding disjunctive language, and in Part **III** we will also use them extensively to define a novel game semantics for DLP and prove its soundness and correctness.

The abstract transformation $(-)^{\vee}$

In this chapter we define the $(-)^{\vee}$ operator, which transforms any given semantics of a non-disjunctive logic programming language into a semantics for the “corresponding” disjunctive one. For reasons of simplicity, *we will assume that all programs in this chapter are clean*. This does not really impose any substantial restriction: for an unclean program \mathcal{D} , we simply use the semantics of its equivalent, clean version $\widehat{\mathcal{D}}$ (see Definition 3.5).

7.1 Definitions and theory

Definition 7.1 ($(-)^{\vee}$). The operator $(-)^{\vee}$ is an overloaded operator that can be applied to:

(1) \mathcal{M} -meanings of LP[N] programs:

$$\begin{array}{ll} \text{if} & \mathbf{m} : \mathbf{P}_{\text{LP[N]}} \rightarrow \mathcal{M}, \\ \text{then} & (\mathbf{m})^{\vee} : \mathbf{P}_{\text{DLP[N]}} \rightarrow \mathcal{P}(\mathcal{M}), \\ \text{is defined by} & (\mathbf{m})^{\vee}(\mathcal{D}) \triangleq \mathbf{m}(\mathbf{D}(\mathcal{D})). \end{array}$$

(2) \mathcal{V} -answers of LP[N] programs:

$$\begin{array}{ll} \text{if} & \mathbf{a} : \mathcal{M} \rightarrow \mathbf{Q}_{\text{LP[N]}} \rightarrow \mathcal{V}, \\ \text{then} & (\mathbf{a})^{\vee} : \mathcal{P}(\mathcal{M}) \rightarrow \mathbf{Q}_{\text{DLP[N]}} \rightarrow \mathcal{V}, \\ \text{is defined by} & (\mathbf{a})^{\vee}(\mathcal{S})(Q) \triangleq \bigwedge_{S \in \mathcal{S}} \bigvee_{q \in Q} \mathbf{a}(S)(q). \end{array}$$

(3) \mathcal{V} -systems of LP[N] programs:

$$\begin{array}{ll} \text{if} & \mathbf{s} : \mathbf{P}_{\text{LP[N]}} \rightarrow \mathbf{Q}_{\text{LP[N]}} \rightarrow \mathcal{V}, \\ \text{then} & (\mathbf{s})^{\vee} : \mathbf{P}_{\text{DLP[N]}} \rightarrow \mathbf{Q}_{\text{DLP[N]}} \rightarrow \mathcal{V}, \\ \text{is defined by} & (\mathbf{s})^{\vee}(\mathcal{D})(Q) \triangleq \bigwedge_{\mathcal{P} \in \mathbf{D}(\mathcal{D})} \bigvee_{q \in Q} \mathbf{s}(\mathcal{P})(q). \end{array}$$

The following theorem justifies the definitions above, and is the driving idea behind them.

Theorem 7A. *Let \mathcal{V} be a truth value space, \mathcal{D} a DLP program, G a DLP goal, and I a \mathcal{V} -interpretation for \mathcal{D} . Then*

$$I(\bigwedge \mathcal{D} \rightarrow \bigvee G) = \bigwedge_{\mathcal{P} \in \mathcal{D}(\mathcal{D})} \bigvee_{g \in G} I(\bigwedge \mathcal{P} \rightarrow g).$$

Proof. Pick a set of indices J to index \mathcal{D} , and denote its rules by R_j , each having a head H_j and a body \mathcal{B}_j , so that $\mathcal{D} := \{R_j \mid j \in J\} = \{H_j \leftarrow \mathcal{B}_j \mid j \in J\}$. Now compute:

$$\begin{aligned} I(\bigwedge \mathcal{D} \rightarrow \bigvee G) &= I(\bigwedge_{j \in J} R_j \rightarrow \bigvee_{g \in G} g) & (1) \\ &= I(\bigwedge_{j \in J} R_j) \Rightarrow I(\bigvee_{g \in G} g) & (*) \\ &= \bigwedge_{j \in J} I(R_j) \Rightarrow \bigvee_{g \in G} I(g) & (*) \\ &= \bigwedge_{j \in J} I(H_j \leftarrow \mathcal{B}_j) \Rightarrow \bigvee_{g \in G} I(g) & (2) \\ &= \bigwedge_{j \in J} I(\bigvee_{h \in H_j} h \leftarrow \mathcal{B}_j) \Rightarrow \bigvee_{g \in G} I(g) & (3) \\ &= \bigwedge_{j \in J} [I(\bigvee_{h \in H_j} h) \leftarrow I(\mathcal{B}_j)] \Rightarrow \bigvee_{g \in G} I(g) & (*) \\ &= \bigwedge_{j \in J} [\bigvee_{h \in H_j} I(h) \leftarrow I(\mathcal{B}_j)] \Rightarrow \bigvee_{g \in G} I(g) & (*) \\ &= \bigwedge_{j \in J} \bigvee_{h \in H_j} [I(h) \leftarrow I(\mathcal{B}_j)] \Rightarrow \bigvee_{g \in G} I(g) & (4) \\ &= \bigvee_{f \in \mathcal{S}(\mathcal{D})} \bigwedge_{j \in J} (I(f(j)) \leftarrow I(\mathcal{B}_j)) \Rightarrow \bigvee_{g \in G} I(g) & (5) \\ &= \bigwedge_{f \in \mathcal{S}(\mathcal{D})} [\bigwedge_{j \in J} (I(f(j)) \leftarrow I(\mathcal{B}_j)) \Rightarrow \bigvee_{g \in G} I(g)] & (6) \\ &= \bigwedge_{f \in \mathcal{S}(\mathcal{D})} [\bigwedge_{j \in J} I(f(j) \leftarrow \mathcal{B}_j) \Rightarrow \bigvee_{g \in G} I(g)] & (*) \\ &= \bigwedge_{f \in \mathcal{S}(\mathcal{D})} [I(\bigwedge_{j \in J} (f(j) \leftarrow \mathcal{B}_j)) \Rightarrow \bigvee_{g \in G} I(g)] & (*) \\ &= \bigwedge_{f \in \mathcal{S}(\mathcal{D})} \bigvee_{g \in G} [I(\bigwedge_{j \in J} (f(j) \leftarrow \mathcal{B}_j)) \Rightarrow I(g)] & (7) \\ &= \bigwedge_{f \in \mathcal{S}(\mathcal{D})} \bigvee_{g \in G} [I(\bigwedge \mathcal{D}_f) \Rightarrow I(g)] & (8) \\ &= \bigwedge_{\mathcal{P} \in \mathcal{D}(\mathcal{D})} \bigvee_{g \in G} [I(\bigwedge \mathcal{P}) \Rightarrow I(g)] & (9) \\ &= \bigwedge_{\mathcal{P} \in \mathcal{D}(\mathcal{D})} \bigvee_{g \in G} I(\bigwedge \mathcal{P} \rightarrow g) & (*) \end{aligned}$$

where each step of the computation is justified as follows: (1) by assumption for \mathcal{D} and G ; (2) by assumption for R_j ; (3) by assumption for H_j ; (4) by Property A.9(III); (5) by the fact that \mathcal{V} , as a truth value space, is completely distributive, and by the definition of $\mathcal{S}(\mathcal{D})$; (6) by Property A.9(IV); (7) by Property A.9(III) again; (8) by the definition of \mathcal{D}_f ; (9) by the definitions of $\mathcal{S}(\mathcal{D})$, $\mathcal{D}(\mathcal{D})$, and \mathcal{D}_f ; and all steps marked by (*) follow from the fact that I is a \mathcal{V} -interpretation. \blacksquare

∇ **REMARK 7.1.** Note that given a set of wffs A , thanks to the idempotence, commutativity and associativity of \vee and \wedge , the various ways of forming a wff

from the sets $\bigvee A$ and $\bigwedge A$ will lead to equal truth values when an interpretation acts on them. We have just made use of this on the proof above.

Lemma 7.1. *Let \mathbf{L} be LP or LPN. Suppose that \mathcal{M} is some set of meanings for \mathbf{L} and \mathcal{V} a truth value space. Let \mathbf{m} and \mathbf{a} be an \mathcal{M} -semantics and a \mathcal{V} -answer function for \mathbf{L} respectively. Then*

$$(\mathbf{a} \circ \mathbf{m})^\vee = (\mathbf{a})^\vee \circ (\mathbf{m})^\vee;$$

or, following Remark 5.1, $(\mathbf{m}, \mathbf{a})^\vee = ((\mathbf{m})^\vee, (\mathbf{a})^\vee)$. It follows that if $(\mathbf{m}_1, \mathbf{a}_1)$ and $(\mathbf{m}_2, \mathbf{a}_2)$ are two semantics for \mathbf{L} , then

$$(\mathbf{m}_1, \mathbf{a}_1) \approx (\mathbf{m}_2, \mathbf{a}_2) \implies (\mathbf{m}_1, \mathbf{a}_1)^\vee \approx (\mathbf{m}_2, \mathbf{a}_2)^\vee. \quad (7.1.1)$$

Proof. We compute:

$$\begin{aligned} ((\mathbf{a})^\vee \circ (\mathbf{m})^\vee)(\mathcal{D})(Q) &= ((\mathbf{a})^\vee((\mathbf{m})^\vee(\mathcal{D}))) (Q) \\ &= \bigwedge_{M \in (\mathbf{m})^\vee(\mathcal{D})} \bigvee_{q \in Q} \mathbf{a}(M)(q) \quad (\text{def. of } (\mathbf{a})^\vee) \\ &= \bigwedge_{M \in \mathbf{m}(\mathcal{D}(\mathcal{D}))} \bigvee_{q \in Q} \mathbf{a}(M)(q) \quad (\text{def. of } (\mathbf{m})^\vee) \\ &= \bigwedge_{\mathcal{P} \in \mathcal{D}(\mathcal{D})} \bigvee_{q \in Q} \mathbf{a}(\mathbf{m}(\mathcal{P}))(q) \\ &= \bigwedge_{\mathcal{P} \in \mathcal{D}(\mathcal{D})} \bigvee_{q \in Q} (\mathbf{a} \circ \mathbf{m})(\mathcal{P})(q) \\ &= (\mathbf{a} \circ \mathbf{m})^\vee(\mathcal{D})(Q). \quad (\text{def. of } (\mathbf{s})^\vee) \quad \blacksquare \end{aligned}$$

Lemma 7.2. *Let \mathcal{V} be a totally ordered, truth value space, and let \mathcal{D} be a clean DLP (or DLPN) program. If M is a model of \mathcal{D} , then there is an LP (or LPN) program $\mathcal{P} \in \mathcal{D}(\mathcal{D})$ such that M is a model of \mathcal{P} . In symbols,*

$$\{M \mid M \text{ is a model of } \mathcal{D}\} \subseteq \{M \mid M \text{ is a model of } \mathcal{P} \text{ for some } \mathcal{P} \in \mathcal{D}(\mathcal{D})\}.$$

Proof. Let us index the rules of \mathcal{D} by some index set J , so that we have $\mathcal{D} = \{R_j \mid j \in J\}$ where for each j , $R_j := H_j \leftarrow \mathcal{B}_j$. Now let M be a model of \mathcal{D} . Therefore, M satisfies every rule R_j of \mathcal{D} , i.e.,

$$\text{for every } j \in J, \quad M(H_j) \geq_{\mathcal{V}} M(\mathcal{B}_j).$$

Since H_j is a finite set of atoms, and since \mathcal{V} is totally ordered, we have

$$M(H_j) = \bigvee \{M(h) \mid h \in H_j\} = \max \{M(h) \mid h \in H_j\} = M(h_j),$$

where h_j is an element of H_j for which the above equality holds. Picking for each $j \in J$ such an h_j , we obtain a \mathcal{D} -section and correspondingly the definite instantiation $\mathcal{P} = \{h_j \leftarrow \mathcal{B}_j \mid j \in J\} \in \mathcal{D}(\mathcal{D})$. We observe that since $M(h_j) = M(H_j) \geq_{\mathcal{V}} M(\mathcal{B}_j)$, M satisfies every rule of \mathcal{P} ; in other words, M is a model of \mathcal{P} , which is what we wanted to show. \blacksquare

As the following counterexample confirms, the above lemma fails to hold in general if we drop the condition that the truth value space has to be totally ordered:

► *Counterexample 7.1.* Consider the truth value space $\mathcal{V} := \mathcal{P}(\{0, 1\})$ ordered under \subseteq , whose elements we will denote by writing \top , \mathbf{L} , \mathbf{R} , and \perp , for $\{0, 1\}$, $\{0\}$, $\{1\}$, and \emptyset respectively. For the DLP program $\mathcal{D} := \{\mathbf{a} \vee \mathbf{b} \leftarrow\}$, we compute

$$D(\mathcal{D}) = \{\mathcal{P}_a, \mathcal{P}_b\}, \quad \text{where} \quad \begin{aligned} \mathcal{P}_a &:= \{\mathbf{a} \leftarrow\} \\ \mathcal{P}_b &:= \{\mathbf{b} \leftarrow\}. \end{aligned}$$

Here is a model M of \mathcal{D} that is neither a model of \mathcal{P}_a , nor of \mathcal{P}_b :

$$M := \{(a, \mathbf{L}), (b, \mathbf{R})\}.$$

In fact, $M(a \vee b) = M(a) \vee M(b) = \mathbf{L} \vee \mathbf{R} = \top$, but $M(a) = \mathbf{L} < \top$ and similarly $M(b) = \mathbf{R} < \top$. ◀

7.2 Applications and examples

As promised, we investigate the application of the $(-)^{\vee}$ operator on the semantics of the non-disjunctive languages that interest us and investigate the equivalences of the resulting semantics.

From LP to DLP

Let us start with the simplest case of LP programs and their least Herbrand model semantics, LHM. As a reminder, we are dealing with $\mathcal{V}_{\text{LHM}} = \mathbb{B}$, \mathcal{M}_{LHM} is the set of all possible Herbrand interpretations, and \mathbf{m}_{LHM} maps an LP program to its least Herbrand model. Finally, $\mathbf{a}_{\text{LHM}}(M)(p)$ is \mathbf{T} exactly when $p \in M$.

We first notice that using $(-)^{\vee}$ on LHM we obtain a semantics for DLP, which we will denote by LHM_{\vee} . We have:

$$\begin{aligned} \mathcal{V}_{\text{LHM}_{\vee}} &= \mathcal{V}_{\text{LHM}} = \mathbb{B} \\ \mathcal{M}_{\text{LHM}_{\vee}} &= \mathcal{P}(\mathcal{M}_{\text{LHM}}). \end{aligned}$$

We proceed following the definitions:

$$\begin{aligned} \mathbf{m}_{\text{LHM}_{\vee}}(\mathcal{P}) &= (\mathbf{m}_{\text{LHM}})^{\vee}(\mathcal{P}) = \mathbf{m}_{\text{LHM}}(D(\mathcal{P})), \\ \mathbf{a}_{\text{LHM}_{\vee}}(\mathcal{S})(Q) &= (\mathbf{a}_{\text{LHM}})^{\vee}(\mathcal{S})(Q) = \bigwedge_{S \in \mathcal{S}} \bigvee_{q \in Q} \mathbf{a}_{\text{LHM}}(S)(q), \\ \mathbf{s}_{\text{LHM}_{\vee}}(\mathcal{D})(Q) &= (\mathbf{s}_{\text{LHM}})^{\vee}(\mathcal{D})(Q) = \bigwedge_{\mathcal{P} \in D(\mathcal{D})} \bigvee_{q \in Q} \mathbf{s}_{\text{LHM}}(\mathcal{P})(q). \end{aligned}$$

Next we show the equivalence of this new, obtained semantics, with the traditional, minimal model semantics MM.

Theorem 7B. *The LHM_{\vee} and the MM semantics are equivalent.*

Proof. To exhibit the equivalence between the minimal model semantics MM and the obtained semantics LHM_{\vee} , we appeal to Lemma 5.1: we define a collector operator $\mathcal{C} : \mathcal{M}_{\text{LHM}_{\vee}} \rightarrow \mathcal{M}_{\text{MM}}$ by

$$\mathcal{C}(\mathcal{M}) \triangleq \{M \in \mathcal{M} \mid M \text{ is } \subseteq\text{-minimal in } \mathcal{M}\},$$

and verify that $(\mathbf{m}_{\text{LHM}_\vee}, \mathbf{a}_{\text{LHM}_\vee}) \triangleleft_{\mathcal{C}} (\mathbf{m}_{\text{MM}}, \mathbf{a}_{\text{MM}})$. Indeed, according to Definition 5.3, this amounts to two things: (1) $\mathbf{m}_{\text{MM}} = \mathcal{C} \circ \mathbf{m}_{\text{LHM}_\vee}$, and (2) $\mathbf{a}_{\text{LHM}_\vee} = \mathbf{a}_{\text{MM}} \circ \mathcal{C}$. The latter is immediate from the definitions of the three objects involved. For the first one, observe first that \mathcal{C} is monotone. Next, suppose that $\mathcal{D} \in \mathbf{P}_{\text{DLP}}$. Using the monotonicity of \mathcal{C} , and Lemma 7.2 (as \mathbb{B} is totally ordered) we compute:

$$\begin{aligned} \mathbf{m}_{\text{MM}}(\mathcal{D}) &= \mathcal{C}(\{M \mid M \text{ is a model of } \mathcal{D}\}) \\ &\subseteq \mathcal{C}(\{M \mid M \text{ is a model of } \mathcal{P} \text{ for some } \mathcal{P} \in \text{D}(\mathcal{D})\}) \\ &= \mathcal{C}(\mathbf{m}_{\text{LHM}_\vee}(\mathcal{D})) = (\mathcal{C} \circ \mathbf{m}_{\text{LHM}_\vee})(\mathcal{D}). \end{aligned}$$

For the other direction,

$$\mathbf{m}_{\text{LHM}_\vee}(\mathcal{D}) \subseteq \{M \mid M \text{ is a model of } \mathcal{D}\},$$

on which we apply the monotone \mathcal{C} on both sides to obtain

$$\begin{aligned} \mathcal{C}(\mathbf{m}_{\text{LHM}_\vee}(\mathcal{D})) &\subseteq \mathcal{C}(\{M \mid M \text{ is a model of } \mathcal{D}\}) \\ &= \mathbf{m}_{\text{MM}}(\mathcal{D}). \end{aligned}$$

Therefore, since \mathcal{D} was arbitrary, we have $\mathbf{m}_{\text{MM}} = \mathcal{C} \circ \mathbf{m}_{\text{LHM}_\vee}$. ■

From LPN to DLPN

Similarly to the LP case, this time we describe the shift from the WF^κ semantics of LPN and obtain a new semantics for DLPN, which we denote by WF_\vee^κ .

Remember, $\mathcal{V}_{\text{WF}^\kappa} = \mathbb{V}_\kappa$, $\mathcal{M}_{\text{WF}^\kappa}$ is the set of all possible \mathbb{V}_κ -interpretations, and $\mathbf{m}_{\text{WF}^\kappa}$ maps an LPN program to its least, \mathbb{V}_κ -valued model. Finally, $\mathbf{a}_{\text{WF}^\kappa}(M)(p)$ is the value of p in M , in other words, $M(p)$.

We denote by WF_\vee^κ the semantics we obtain by using $(-)^{\vee}$ on WF^κ . This semantics has:

$$\begin{aligned} \mathcal{V}_{\text{WF}_\vee^\kappa} &= \mathcal{V}_{\text{WF}^\kappa} = \mathbb{V}_\kappa \\ \mathcal{M}_{\text{WF}_\vee^\kappa} &= \mathcal{P}(\mathcal{M}_{\text{WF}^\kappa}). \end{aligned}$$

Just like in the case of LHM, we follow the definitions and obtain

$$\begin{aligned} \mathbf{m}_{\text{WF}_\vee^\kappa}(\mathcal{P}) &= (\mathbf{m}_{\text{WF}^\kappa})^{\vee}(\mathcal{P}) = \mathbf{m}_{\text{WF}^\kappa}(\text{D}(\mathcal{P})), \\ \mathbf{a}_{\text{WF}_\vee^\kappa}(\mathcal{S})(Q) &= (\mathbf{a}_{\text{WF}^\kappa})^{\vee}(\mathcal{S})(Q) = \bigwedge_{S \in \mathcal{S}} \bigvee_{q \in Q} \mathbf{a}_{\text{WF}^\kappa}(S)(q), \\ \mathbf{s}_{\text{WF}_\vee^\kappa}(\mathcal{D})(Q) &= (\mathbf{s}_{\text{WF}^\kappa})^{\vee}(\mathcal{D})(Q) = \bigwedge_{\mathcal{P} \in \text{D}(\mathcal{D})} \bigvee_{q \in Q} \mathbf{s}_{\text{WF}^\kappa}(\mathcal{P})(q). \end{aligned}$$

Remember that MM^κ is defined only for finite programs, for which ω is a big enough ordinal. Therefore the obtained semantics WF_\vee^κ is in fact more general than MM^κ as it appears in the literature, since WF_\vee^κ gives meaning to any DLPN program, finite or not. Yet, as long as we restrict ourselves to *finite* programs, we have the following theorem:

Theorem 7C. *The WF_\vee^ω and the MM^ω semantics on finite DLPN programs, are equivalent.*

Proof. We define the collector operator $\mathcal{C} : \mathcal{M}_{WF_{\vee}^{\omega}} \rightarrow \mathcal{M}_{MM^{\omega}}$ by

$$\mathcal{C}(\mathcal{M}) \triangleq \{M \in \mathcal{M} \mid M \text{ is } \sqsubseteq_{\omega}\text{-minimal in } \mathcal{M}\},$$

and verify that $(\mathbf{m}_{WF_{\vee}^{\omega}}, \mathbf{a}_{WF_{\vee}^{\omega}}) \triangleleft_{\mathcal{C}} (\mathbf{m}_{MM^{\omega}}, \mathbf{a}_{MM^{\omega}})$, so that the result will again be a direct consequence of Lemma 5.1. The remaining of the proof is similar to the one of Theorem 7B, except that this time we use the fact that \mathbb{V}_{ω} is totally ordered. \blacksquare

This concludes the exposition of our “disjunctive” toolkit. In the next part we turn to games, yet in its last chapter we will reuse the $(-)^{\vee}$ operator, this time on game semantics.

‡

Part III

Game semantics

The LPG semantics

In this chapter we introduce the simplest of games that we will use, and which will form the basis for all the remaining games. As an introduction to the subject of game semantics for logic programs, some notions are only described informally. On the next chapter, they will be strictly formalized.

8.1 Introduction to game semantics

Games made their debut in the logic programming scene with [vE86], where we find the first informal description of a game in the logic programming literature. Similar games can also be found as early as [Acz77]. But it was not until [DCLN98], that a game semantics was systematically studied for the case of LP. It was there shown, that it is in fact sound and complete with respect to SLD resolution. This is a rather involved game, which stays close to the procedural semantics, and therefore directly handles first-order programs, taking into account variables, function symbols, substitutions, etc. Approximately a decade later, this game—or, to be fair, its propositional version—was extended in [GRW08], to cover negation for finite, propositional LPN programs. The LPN game semantics is proven to be equivalent to an infinite-valued refinement of the well-founded model semantics (as defined in [RW05]); it is a denotational game semantics. A couple of years after that, two games to deal with DLP and DLPN (again from a denotational point of view) were described informally in [Tso10].

The history of denotational and game semantics for these four versions of logic programming languages is summarized in Table 8.1.

Lang.	Denotational semantics	Game semantics
LP	least Herbrand model, [vEK76]	LPG, [DCLN98]
DLP	minimal models, [Min82]	DLPG, [Tso13]
LPN	well-founded model, [VGRS91]	LPNG, [GRW08]
DLPN	\mathbb{V}_κ -valued minimal models, [CPRW07]	(LPNG) ^v , § 11.3

Table 8.1: Development of game semantics for logic programming.

Why games? There are various benefits of defining a correct game semantics for each of these programming languages. On the operational side, the LP game helps us prune down the space of SLD derivations, by grouping them together using the much smaller space of strategies. In fact, the *alpha-beta algorithm* was used in [LC00] to speed-up the resolution strategies even for the case of *constraint logic programming*. On the denotational side, these games impart elegant characterizations of these main versions of logic programming. As it turns out, starting from LP, one only needs to add a couple of simple game-rules to its game to arrive at DLP; the addition of a different one brings you to LPN. These rules are fairly modular, so that it even makes sense to consider adding all of them simultaneously to deal with DLPN—although we will choose a different route to obtain a game semantics for DLPN in this text. Contrast the simplicity of this approach to the difficulty of treating disjunction and negation relying solely on model-theoretic tools. In addition, this kind of games is also applicable to *intensional logic programming* (as explained in [NR12]), and even outside of the logic programming world, e.g., to *boolean grammars* (see [KNR11]).

8.2 The LPG game

The general picture, shared by all of the games that we consider here, is based on Lorenzen dialogue games (see [Lor61]). There are two players (Player I vs. Player II, or Opponent vs. Player) and two player rôles: *doubter* vs. *believer*.¹ At each round of the game, one player will be the doubter and the other one will be the believer. In the beginning of the game, Player I is the doubter and Player II the believer. When there is no negation present, the players will never change rôles, and so we may also refer to them as Doubter and Believer respectively.²

Given a program \mathcal{P} and a goal $\leftarrow p$, the game is about the two players arguing over whether the given goal should succeed or not. The player who doubts it (Doubter) begins the game by saying:

Doubter: “Why p ?”.

The defending player (Believer) must give a convincing argument of why he thinks that p is true. He must select a program rule from \mathcal{P} that has p as its head and play it. For instance, selecting $p \leftarrow a, b, c$, he replies:

Believer: “ p because a, b , and c .”,

to which Doubter must respond by doubting a specific conjunct from the body, viz. a, b , or c . Selecting the second one, for example, she replies:

Doubter: “Why b ?”.

¹We agree to refer to Player I as a *she*, and to Player II as a *he*. There is no particular reason for the specific choice, but keeping their genders separate makes talking about them easier. I first encountered this convention in [Vää11], and I adopted it.

²This is not the case for the LPN game: to capture the “unknown” truth value of the well-founded model, we need to allow ties in the game, as well as *rôle-switching* moves. This is explained in Chapter 9.

The game continues in this manner, until a player cannot argue anymore, in which case they lose the game. This means that either Believer played a rule with an empty body (i.e., a fact) or Doubter played an atom which is not the head of any program rule. Doubter has the *benefit of the doubt*, which means that if she can keep doubting forever (in case of an infinite play), she wins. See Section 8.4 for a discussion about the reasoning behind this decision.

This conveys the essence of the games we use for logic programming; all games we define in this text are based on the same principles. We will denote believer moves by β and doubter moves by δ , and refer to the game just described as the LPG game, and denote it by $\Gamma_{\mathcal{P}}^{\text{LPG}}(\leftarrow p)$, for a given program \mathcal{P} and a goal clause $\leftarrow p$. When the game is obvious from the context we may omit the superscript.

8.3 Example plays

► *Example 8.1.* Consider the program

$$\mathcal{P} := \left\{ \begin{array}{l} p \leftarrow q, r \\ q \leftarrow s \\ s \leftarrow \\ r \leftarrow t \\ r \leftarrow \end{array} \right\}.$$

With the goal $\leftarrow p$, three maximal plays in $\Gamma_{\mathcal{P}}^{\text{LPG}}(\leftarrow p)$ might be the following:

$$\pi_1 := \left| \begin{array}{l} \text{goal} : \leftarrow \underline{p} \\ D_0 : p \\ B_0 : p \leftarrow \underline{q}, r \\ D_1 : q \\ B_1 : q \leftarrow \underline{s} \\ D_2 : s \\ B_2 : s \leftarrow \end{array} \right|, \pi_2 := \left| \begin{array}{l} \text{goal} : \leftarrow \underline{p} \\ D_0 : p \\ B_0 : p \leftarrow q, \underline{r} \\ D_1 : r \\ B_1 : r \leftarrow \underline{t} \\ D_2 : t \end{array} \right|, \pi_3 := \left| \begin{array}{l} \text{goal} : \leftarrow \underline{p} \\ D_0 : p \\ B_0 : p \leftarrow q, \underline{r} \\ D_1 : r \\ B_1 : r \leftarrow \end{array} \right|.$$

Conveniently enough, if we underline Doubter's selections—as we have done in the example above—we can condense each play by entirely omitting the lines that correspond to her moves. For the sake of laziness, we will follow this practice from now on. The three plays above are therefore written as:

$$\pi_1 := \left| \begin{array}{l} \text{goal} : \leftarrow \underline{p} \\ B_0 : p \leftarrow \underline{q}, r \\ B_1 : q \leftarrow \underline{s} \\ B_2 : s \leftarrow \end{array} \right|, \pi_2 := \left| \begin{array}{l} \text{goal} : \leftarrow \underline{p} \\ B_0 : p \leftarrow q, \underline{r} \\ B_1 : r \leftarrow \underline{t} \end{array} \right|, \pi_3 := \left| \begin{array}{l} \text{goal} : \leftarrow \underline{p} \\ B_0 : p \leftarrow q, \underline{r} \\ B_1 : r \leftarrow \end{array} \right|.$$

Plays π_1 and π_3 are won by Believer, who has managed to corner the doubter by playing a fact: there are no conjuncts for her to choose from. On the other hand, π_2 is won by Doubter, who managed to doubt an atom which is not the head of any of the rules of \mathcal{P} , and thus Believer cannot make any move and loses the game. ◀

8.4 Benefit of the doubt

Why do we give the benefit of the doubt to the doubter? The following example illustrates how things can go wrong, in case we do not.

► *Example 8.2.* Consider the following program

$$\mathcal{P} := \left\{ \begin{array}{l} \mathbf{a} \leftarrow \mathbf{a}, \mathbf{b} \\ \mathbf{b} \leftarrow \end{array} \right\}$$

and the goal $\leftarrow \mathbf{a}$. A play in $\Gamma_{\mathcal{P}}^{\text{LP}}(\leftarrow \mathbf{a})$ has to begin with Doubter doubting \mathbf{a} , to which Believer responds with $\mathbf{a} \leftarrow \mathbf{a}, \mathbf{b}$. For as long as Doubter doubts \mathbf{a} , Believer always plays the same rule, and therefore plays will never end with a winner. Had we given the benefit of the doubt to Believer, Doubter would eventually have no choice but to switch and doubt \mathbf{b} , at which point Believer would win by playing the fact $\mathbf{b} \leftarrow$. This describes a winning strategy for Believer, but—this being a logic program—we most definitely do not want the goal $\leftarrow \mathbf{a}$ to succeed, because of the denotational semantics:

$$a \notin \text{LHM}(\mathcal{P}) = \{b\}. \quad \blacktriangleleft$$

Who has the benefit of the doubt is a decision similar to the closed/open world assumptions (CWA/OWA) in knowledge representation languages. Giving it to the doubter resembles CWA, while giving it to no player resembles OWA,³ it might be desired in some cases (e.g., description logics used for the semantic web), but it is certainly not the stance we take in the world of logic programming. For more information check [Rei78] and [Min82].

8.5 Semantics from games

Remember—our objective is to use games to provide denotational semantics for logic programs; in other words, to decide which goals should succeed, i.e., which answers are correct.

Notice now, that the three plays π_1 , π_2 , and π_3 of Example 8.1 above are played on the same game, of the same program \mathcal{P} , with the same goal $\leftarrow \mathbf{p}$. Yet Believer wins two of them, and loses another. This shows that *merely winning or losing a particular play in a game, is not enough information to decide if a goal should succeed or not*: maybe we managed to win a play because our opponent played badly enough, or we lost the play because we didn't play smartly enough. This is exactly what happened in π_2 , where Believer chose the rule $\mathbf{r} \leftarrow \mathbf{t}$ to justify his belief on \mathbf{r} , thus enabling the doubter to doubt \mathbf{t} for which there is no rule that supports it. In π_3 he makes the right move and chooses to play the *fact* $\mathbf{r} \leftarrow$, which immediately grants him victory.

It therefore makes no sense to determine the success of a goal by looking at particular plays. Instead, we focus on *strategies*. Informally, a strategy for a game determines what Player will play in each possible position of the game.⁴ If following a strategy σ Player is guaranteed to win against any possible move of Opponent, we call σ a *winning strategy* for that goal.

This allows us to define the LPG semantics:

³This would introduce strategies that are neither winning nor losing: they lead to *ties*.

⁴Formal definitions are given on the next chapter.

Definition 8.1 (LPG semantics). Let \mathcal{P} be an LP program. A goal $\leftarrow \mathbf{p}$ succeeds, if Believer has a winning strategy in $\Gamma_{\mathcal{P}}^{\text{LP}}(\leftarrow \mathbf{p})$.

The LPG semantics under the abstract semantic framework

- $\mathcal{V}_{\text{LPG}} \triangleq \mathbb{B}$.
- \mathcal{M}_{LPG} is the set of strategies based on LP programs.
- \mathbf{m}_{LPG} maps every LP program \mathcal{P} to the set of strategies for the LPG game based on \mathcal{P} .
- $\mathbf{a}_{\text{LPG}}(\Sigma)(q) \triangleq \begin{cases} \mathbf{T}, & \text{if there is a winning strategy } \sigma \in \Sigma \text{ for } q \\ \mathbf{F}, & \text{otherwise.} \end{cases}$

8.6 Soundness and completeness

We state in this section some known results about the LP game.

Theorem 8A (Di Cosmo–Loddo–Nicolet). *The LPG semantics is sound and complete with respect to SLD resolution.*

Theorem 8B (Clark). *SLD resolution is sound and complete with respect to the least Herbrand model semantics.*

The first of these is proven in [DCLN98], while the second one is due to [Cla79] (and can also be found in [Llo87, Theorems 7.1 and 8.6]). Putting the above two theorems together, we arrive at the correctness of the LPG semantics:

Corollary 8.1 (Soundness and completeness of the LP game semantics). *Let \mathcal{P} be an LP program, and $\leftarrow \mathbf{p}$ a goal. Then, there is a winning strategy in the associated game $\Gamma_{\mathcal{P}}^{\text{LP}}(\leftarrow \mathbf{p})$ iff \mathbf{p} belongs to the least Herbrand model of \mathcal{P} .*

Having understood the basic ideas behind the LPG game, we proceed to formally define and study the LPNG game, which will turn out to reduce to the LPG one when no negations are present.

The LPNG semantics

As we have mentioned previously, to deal with negation-as-failure in the way of the well-founded semantics, it is essential to have some unknown truth value \mathbf{U} in the truth value space. In the world of games, we introduce *ties* between the two players. The LPNG game is formally defined for *finite* LPN programs, and thus all programs in this chapter are assumed to be finite. This also has the implication that we can limit ourselves to values in the \mathbb{V}_ω space (see Remark 5.2).

9.1 The LPNG game

The LPNG game is based on the LPG one, and indeed it reduces to it for negation-free programs. Whereas in the LPG game, the rôles of the players stay the same throughout the course of a play, in LPNG, a believer's move can sometimes be *rôle-switching*. For this reason, we cannot refer to the players as “Doubter” and “Believer”, so we speak of “Opponent” and “Player” instead.

The LPNG game is played just like the LPG one, until the moment that the doubter doubts a negated conjunct, say $\sim p$. Essentially we read this as:

Opponent: “Why $\sim p$?”

to which Player (who is the believer) is forced to confirm that he indeed doubts p , which turns him into the doubter, and his opponent into the believer:

Player: “Because I doubt p . Why p ?”

and the game continues, but with the rôles of the players swapped.

It is important to stress that *the doubter retains the benefit of the doubt*, but there is a game-changing difference: it is no longer true that every infinite play will be won by one of the players. In case a player has managed to “secure” their rôle as a doubter from some certain point on, then indeed they will win any infinite play. Otherwise, there must be an infinite number of rôle-switching moves, and *the play results in a tie*.

9.2 Example plays

► *Example 9.1.* Consider the program

$$\mathcal{P} := \left\{ \begin{array}{l} \mathbf{p} \leftarrow \\ \mathbf{q} \leftarrow \sim \mathbf{p} \\ \mathbf{r} \leftarrow \sim \mathbf{q} \end{array} \right\},$$

We show two plays for this program; the first is played on $\Gamma_p^{\text{LPN}}(\leftarrow \mathbf{q})$, the second on $\Gamma_p^{\text{LPN}}(\leftarrow \mathbf{r})$. We mark every rôle-switch by drawing a straight line and we keep the convention of omitting the lines that correspond to doubter moves (by simply underlying the doubts)

$$\pi_1 := \left| \begin{array}{l} \text{goal} : \leftarrow \underline{\mathbf{q}} \\ \mathbf{P}_0 : \mathbf{q} \leftarrow \underline{\sim \mathbf{p}} \\ \hline \mathbf{O}_2 : \mathbf{p} \leftarrow \end{array} \right|, \quad \pi_2 := \left| \begin{array}{l} \text{goal} : \leftarrow \underline{\mathbf{r}} \\ \mathbf{P}_0 : \mathbf{r} \leftarrow \underline{\sim \mathbf{q}} \\ \hline \mathbf{O}_2 : \mathbf{q} \leftarrow \underline{\sim \mathbf{p}} \\ \hline \mathbf{P}_3 : \mathbf{p} \leftarrow \end{array} \right|.$$

Both plays are won by believers, but in π_1 the believer is Opponent, while in π_2 , where there are two rôle-switching moves, Player is the believer. ◀

► *Example 9.2.* Let us now see an infamous program of LPN:

$$\mathcal{Q} := \left\{ \begin{array}{l} \mathbf{p} \leftarrow \sim \mathbf{q} \\ \mathbf{q} \leftarrow \sim \mathbf{p} \end{array} \right\}.$$

Consider the infinite play

$$\pi_3 := \left| \begin{array}{l} \text{goal} : \leftarrow \underline{\mathbf{p}} \\ \mathbf{P}_0 : \mathbf{p} \leftarrow \underline{\sim \mathbf{q}} \\ \hline \mathbf{O}_2 : \mathbf{q} \leftarrow \underline{\sim \mathbf{p}} \\ \hline \mathbf{P}_3 : \mathbf{p} \leftarrow \underline{\sim \mathbf{q}} \\ \hline \mathbf{O}_5 : \mathbf{q} \leftarrow \underline{\sim \mathbf{p}} \\ \hline \mathbf{P}_6 : \mathbf{p} \leftarrow \underline{\sim \mathbf{q}} \\ \hline \vdots \end{array} \right|.$$

Here no player is able to secure the believer rôle for themselves, so this play is won by neither of them: the outcome is a tie. ◀

9.3 Game semantics

To obtain the refined, infinite-valued model $\text{WFM}_\omega(\mathcal{P})$, we need to take into account the rôle-switching moves played. For this, we will use a payoff function:

Definition 9.1 (Payoff). Let π be a play in some game $\Gamma_{\mathcal{P}}^{\text{LPN}}(\leftarrow p)$. Then the *payoff* functions Φ_{ω} and Φ are defined by:

$$\Phi_{\omega}(\pi) \triangleq \begin{cases} \mathbf{T}_n, & \text{if Player wins in } \pi, \\ \mathbf{F}_n, & \text{if Player loses in } \pi, \\ \mathbf{U}, & \text{otherwise,} \end{cases}$$

where n is the number of rôle-switching moves played in π ; and

$$\Phi \triangleq \text{collapse} \circ \Phi_{\omega},$$

where *collapse* is the “subscript-removing” function of Definition 2.6.

► *Example 9.3.* Consider the plays π_1 , π_2 , and π_3 from the examples 9.1 and 9.2 above. Their corresponding payoffs are:

$$\Phi_{\omega}(\pi_1) = \mathbf{F}_1, \quad \Phi_{\omega}(\pi_2) = \mathbf{T}_2, \quad \Phi_{\omega}(\pi_3) = \mathbf{U}. \quad \blacktriangleleft$$

Definition 9.2 (Non-losing strategy). A strategy σ is *non-losing* iff

$$\bigwedge \{ \Phi_{\omega}(\pi) \mid \pi \in \sigma \} \geq \mathbf{U}.$$

‡ **REMARK 9.1.** At this point, we are really in possession of two different games: one in which there are only three possible outcomes for each play (win, lose, or tie); and one in which by consulting the payoff function Φ_{ω} , we actually have various degrees of winning and losing, and a single level of tie. Notice, however, that the two games only differ in the possible outcomes; they share the same moves, plays, and strategies.

We now proceed to directly define the LPNG and LPNG^{ω} semantics under the framework of Chapter 5:

The LPNG semantics

- $\mathcal{V}_{\text{LPNG}} \triangleq \mathbb{V}_1$.
- $\mathcal{M}_{\text{LPNG}}$ is the set of strategies based on LPN programs.
- \mathbf{m}_{LPNG} maps every LPN program \mathcal{P} to the set of strategies for the LPNG game based on \mathcal{P} .
- Finally,

$$\begin{aligned} \mathbf{a}_{\text{LPNG}}(\Sigma)(q) &\triangleq \begin{cases} \mathbf{T}, & \text{if there is a winning strategy in } \Sigma \text{ for } q \\ \mathbf{U}, & \text{else, if there is a non-losing strategy in } \Sigma \text{ for } q \\ \mathbf{F}, & \text{otherwise.} \end{cases} \\ &= \bigvee \{ \bigwedge \{ \Phi(\pi) \mid \pi \in \sigma \} \mid \sigma \text{ is a strategy in } \Sigma \text{ for } q \} \end{aligned}$$

The LPNG^ω semantics

- $\mathcal{V}_{\text{LPNG}^\omega} \triangleq \mathbb{V}_\omega$.
- $\mathcal{M}_{\text{LPNG}^\omega}$ is the set of strategies based on LPN programs.
- $\mathbf{m}_{\text{LPNG}^\omega}$ maps every LPN program \mathcal{P} to the set of strategies for the LPNG game based on \mathcal{P} .
- $\mathbf{a}_{\text{LPNG}^\omega}(\Sigma)(q) \triangleq \bigvee \{ \bigwedge \{ \Phi_\omega(\pi) \mid \pi \in \sigma \} \mid \sigma \text{ is a strategy in } \Sigma \text{ for } q \}$

9.4 Soundness and completeness

Soundness and completeness results for both games with respect to the corresponding well-founded semantics are proven in [GRW08]. The following theorem summarizes these results:

Theorem 9A (Soundness and completeness of the LPNG semantics). *Let \mathcal{P} be an LPN program, and let Σ be the set of all strategies of LPNG games played on \mathcal{P} . Then,*

$$\begin{aligned} \text{WFM}(\mathcal{P}) &= \mathbf{a}_{\text{LPNG}}(\Sigma), \\ \text{WFM}_\omega(\mathcal{P}) &= \mathbf{a}_{\text{LPNG}^\omega}(\Sigma). \end{aligned}$$

We have thus presented a game semantics for LPN programs, which is sound and complete with respect to the well-founded semantics. Next we turn our attention to the other extension of LP programs: DLP.

Chapter 10

The DLPG semantics

In this chapter a game semantics for DLP is formally defined, studied, and proven correct:

Soundness and completeness of the DLPG semantics (Theorem 10G). *The game semantics DLPG of DLP is equivalent to the minimal model semantics, i.e., given any DLP program \mathcal{P} , and any disjunction D ,*

$$D \text{ is true wrt the DLPG semantics} \iff D \text{ is true wrt the minimal model semantics.}$$

Two key ideas of Chapter 6 are further developed to prove the two directions of the above result: *combination* (for completeness) and *splitting* (for both). In short, we begin with a finite disjunctive logic program \mathcal{P} , and split it in two new DLP programs \mathcal{P}_1 and \mathcal{P}_2 , such that they are in a sense, “less disjunctive”. Now, strategies for games in \mathcal{P} can themselves be split to strategies for games in \mathcal{P}_1 and \mathcal{P}_2 , and vice versa: strategies for such games can be combined to form new strategies for games in \mathcal{P} . By repeated splitting, we eventually arrive at programs that are not disjunctive at all (LP programs), which we know how to deal with since [DCLN98]. Finally, compactness will allow us to extend this result to the general case of infinite DLP programs.

Be aware, that even though the DLPG game developed here can be thought of as another extension of the LPG game, its formalization is drastically different from those of the games of either LP or LPN, and appears to be novel in the field of logic programming. It has been influenced instead by game semantics in the style of Abramsky–Jagadeesan–Malacaria and Hyland–Ong–Nickau, used for PCF and functional programming in general (see [AJM00], [AM99], [HO00] and [Nic94]). Although we assume no such prior knowledge of this field, the initiated reader should hopefully feel at home. Let us play.

10.1 The simplified DLP game

This game extends the LPG game to deal with disjunctions. First of all, the goal here is assumed to consist of a *disjunction* of atoms. The key difference from the LPG (and LPNG) games is that *in the DLPG game the believer can play combo moves: he can use more than one rule to support his belief*. What is more, he can use *not only* rules from the given program, but also *implicit rules*, i.e., rules of the form $a \leftarrow a$. Thanks to rule combination, he can disjunctively

combine his selection of rules into one, say $G \leftarrow D_1, \dots, D_n$, and play it against his opponent:

Believer: “ G because D_1, D_2, \dots , and D_n .”;

and it is now Doubter’s turn to choose which disjunction D_i to doubt:

Doubter: “Why D_i ?”.

And the game goes on.

To sum up, Believer is constantly challenged by Doubter to justify why he believes some *disjunction* δ . He does that by Υ -combining a sequence of DLP rules to form a single rule β , such that the head of β is a *subset* of δ . Doubter must then select which disjunction from the body of β she doubts, and so on. Informally, this can be further summarized thus:

DLP game = LP game + implicit rules + combo moves.

The decision to allow the believer to include implicit rules not from the program is backed up by the following example:

► *Example 10.1.* Consider the DLP program

$$\mathcal{P} := \left\{ \begin{array}{l} p \leftarrow a \\ p \leftarrow b \\ b \leftarrow c \\ a \vee c \leftarrow \end{array} \right\}.$$

For the goal $\leftarrow p$, two plays in this game could look like this:

$$\pi_1 := \left| \begin{array}{l} \text{goal} : \leftarrow \underline{p} \\ B_0 : p \leftarrow \underline{a \vee b} \\ B_1 : a \vee b \leftarrow \underline{a \vee c} \\ B_2 : a \vee c \leftarrow \end{array} \right|, \quad \pi_2 := \left| \begin{array}{l} \text{goal} : \leftarrow \underline{p} \\ B_0 : p \leftarrow \underline{a \vee b} \\ B_1 : b \leftarrow \underline{c} \end{array} \right|.$$

In both plays, Believer justifies the move B_0 by combining the first two program rules. Then, in π_1 , he combines the *program* rule $b \leftarrow c$ with the *implicit* rule $a \leftarrow a$, while in π_2 , he only uses program rules. You can easily verify that without implicit rules it is impossible for him to win this game. ◀

Note that in both plays of Example 10.1, since every believer move has a body with only one element in it, Doubter does not really have any choice to make; she is simply following the lead of Believer. Here is an example where she can actually enjoy the game as well:

► *Example 10.2.* The program now is

$$\mathcal{Q} := \left\{ \begin{array}{l} p \leftarrow a, b \\ q \leftarrow a, b \\ a \leftarrow d, c \\ b \leftarrow \\ c \vee d \leftarrow b \end{array} \right\},$$

and the goal is $\leftarrow p \vee q$. Here are two valid plays for this game:

$$\pi_1 := \left| \begin{array}{l} \text{goal :} \quad \leftarrow \underline{p \vee q} \\ \text{B}_0 : \quad p \vee q \leftarrow \underline{a \vee a}, \underline{a \vee b}, a \vee b, b \vee b \\ \text{B}_1 : \quad b \leftarrow \end{array} \right|,$$

won by Believer, and

$$\pi_2 := \left| \begin{array}{l} \text{goal :} \quad \leftarrow \underline{p \vee q} \\ \text{B}_0 : \quad p \vee q \leftarrow \underline{a \vee a}, a \vee b, a \vee b, b \vee b \\ \text{B}_1 : \quad a \leftarrow \underline{d}, \underline{c} \end{array} \right|,$$

by Doubter. ◀

Having seen the basic idea of the simplified game, it is time to formalize things: we define the actual DLPG game.

10.2 The DLPG game

To study this game and prove it correct, we need to refine it substantially. To be specific, believer moves will not be played as a single (combined) DLP rule as was done in the simplified version. Instead, Believer will now specify the exact sequence of rules β that he combined to create his move. Likewise, Doubter will not select a single disjunction from the combined move of Believer; instead she will select a sequence of occurrences: *one from each body* of the rules in β —which, we stress, is now a *sequence* of rules. Things become clearer and formal after the definitions and the examples that follow.

Given a DLP program \mathcal{P} and a goal clause $\leftarrow \mathbf{G}$, we will define the associated DLPG game, and denote it by $\Gamma_{\mathcal{P}}^{\text{DLP}}(\leftarrow \mathbf{G})$ or simply by $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$ when no confusion may arise.

Just like every game that we investigate in this thesis, this is a two-player game, with the two same player rôles: the *doubter* and the *believer*. As was the case for the LPG game, the rôles of the players never change throughout the game and so we will simply call the players Doubter and Believer again.

Doubter starts by doubting \mathbf{G} , the body of the goal clause, and Believer tries to defend it. A player who cannot play a valid move loses the game. Now let us be precise.

Definition 10.1 (Extended program). Given any set of DLP rules \mathcal{P} , we can extend it to \mathcal{P}_+ , which includes all meaningful implicit rules. In detail,

$$\mathcal{P}_+ \triangleq \mathcal{P} \cup \{a \leftarrow a \mid a \in \text{HB}(\mathcal{P})\}.$$

Definition 10.2 (Moves). A *doubter move* δ is a sequence of occurrences of disjunctions in bodies of DLP clauses; we refer to these occurrences as the *doubts* of δ . A *believer move* β from \mathcal{P} is a finite sequence of DLP rules from the extended set \mathcal{P}_+ . Given a move m , we call $[m]$ the *statement* of the move and refer to the sequence m as the *justification* of the statement.¹ We say that the elements of a believer move β that belong to \mathcal{P} constitute the *proper part* of the justification; the rest, the *implicit*. If $|\beta| > 1$ we call β a *combo move*.

The statement	is read as...
D	“Why D?” or “I doubt D.”
$E \leftarrow D_1, \dots, D_n$	“E because D_1, \dots , and D_n .”
$E \leftarrow$	“E because it is a fact.”

Table 10.1: How to read believer and doubter statements.

Table 10.1 suggests how we can read moves aloud. Notice that they resemble an actual dialogue. This motivates the following definition:

Definition 10.3 (Dialogue). A *quasidialogue* from \mathcal{P} is a finite sequence

$$\pi := \langle \delta_0, \beta_0, \delta_1, \beta_1, \dots \rangle,$$

such that:

- for all i , δ_i is a doubter move and β_i a believer move (if they exist);
- for all i , if β_i exists then $\text{head}([\beta_i]) \subseteq [\delta_i]$;
- for all $i > 0$, if δ_i is $\langle D_1, \dots, D_k \rangle$, and β_{i-1} is $\langle \psi_1, \dots, \psi_{k'} \rangle$ then $k = k'$ and $D_j \in \text{body}(\psi_j)$ for all $1 \leq j \leq k$.

It is a *dialogue* if it also satisfies:

- for all i , if β_i exists then $\beta_i \cap \mathcal{P} \neq \emptyset$,

i.e., the believer always selects at least one rule from \mathcal{P} .

✂ **REMARK 10.1.** Regarding the first restriction imposed on believer moves, one could make the seemingly stronger demand that $\text{head}([\beta_i]) = [\delta_i]$. It is easy to check, however, that this changes nothing in the game, since the believer can bring up implicit rules to combine them with his move in case the subset was proper. To see that he can still win exactly the same arguments, remember that by definition, if there is at least one fact in a combination, the resulting rule is also a fact.

Definition 10.4 (Play). A *(quasi)play* of $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$ is a (quasi)dialogue π from \mathcal{P} which, if non-empty, satisfies the additional property that $[\delta_0] = \mathbf{G}$. We denote the empty play by $\varepsilon \triangleq \langle \rangle$.

Note that dialogues (and plays) are sequences and thus inherit the partial orderings from \sqsubseteq and \sqsubseteq_e .

► **Example 10.3.** Consider the same program \mathcal{Q} as in Example 10.2, repeated here for convenience:

$$\mathcal{Q} := \left(\begin{array}{l} p \leftarrow a, b \\ q \leftarrow a, b \\ a \leftarrow d, c \\ b \leftarrow \\ c \vee d \leftarrow b \end{array} \right)$$

¹ $[-]$ was introduced in Definition 6.8.

(the goal is still $\leftarrow \mathbf{p} \vee \mathbf{q}$). The two plays that follow correspond to the ones we considered previously. Here, the statements of the believer moves that use more than one rule are explicitly shown in parentheses merely for the convenience of the reader; they are not part of the actual plays.

$$\pi_1 := \left(\begin{array}{l} \text{goal :} \quad \leftarrow \underline{\mathbf{p} \vee \mathbf{q}} \\ \beta_0 : \quad \mathbf{p} \leftarrow \underline{\mathbf{a}}, \mathbf{b} \\ \quad \mathbf{q} \leftarrow \mathbf{a}, \underline{\mathbf{b}} \\ ([\beta_0] : \quad \mathbf{p} \vee \mathbf{q} \leftarrow \underline{\mathbf{a} \vee \mathbf{a}}, \underline{\mathbf{a} \vee \mathbf{b}}, \mathbf{b} \vee \mathbf{a}, \mathbf{b} \vee \mathbf{b}) \\ \beta_1 : \quad \mathbf{b} \leftarrow \end{array} \right)$$

is (still) won by Believer, and

$$\pi_2 := \left(\begin{array}{l} \text{goal :} \quad \leftarrow \underline{\mathbf{p} \vee \mathbf{q}} \\ \beta_0 : \quad \mathbf{p} \leftarrow \underline{\mathbf{a}}, \mathbf{b} \\ \quad \mathbf{q} \leftarrow \underline{\mathbf{a}}, \mathbf{b} \\ ([\beta_0] : \quad \mathbf{p} \vee \mathbf{q} \leftarrow \underline{\mathbf{a} \vee \mathbf{a}}, \mathbf{a} \vee \mathbf{b}, \mathbf{b} \vee \mathbf{a}, \mathbf{b} \vee \mathbf{b}) \\ \beta_1 : \quad \mathbf{a} \leftarrow \underline{\mathbf{d}}, \underline{\mathbf{c}} \end{array} \right)$$

by Doubter. ◀

Property 10.1. Any dialogue $\pi := \langle \delta_0, \beta_0, \dots \rangle$ from \mathcal{P} can be considered as a play of $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$, where $\mathbf{G} := [\delta_0]$.

✂ **REMARK 10.2** (Disallowing stalling). We have forced the believer to always include at least one rule from the actual program \mathcal{P} , thus banning what we are about to call “stalling” from our games. Still, this concept will be crucial for our exposition: the introduction of plays in which stalling is allowed (i.e., quasiplays) will make a lot of the statements that follow easier to prove.

Definition 10.5 (Follow). We say that Doubter *follows*, if she has only one possible valid move that she can play in response to a believer move β , i.e., if for every rule $\psi \in \beta$, $|\text{body}(\psi)| = 1$. We denote such a follow move by $\bar{\beta}$. In symbols,

$$\bar{\beta} \triangleq \langle \text{the unique } \mathbf{D} \in \text{body}(\psi) \mid \psi \in \beta \rangle.$$

Definition 10.6 (Stalling). We say that a Believer is *stalling*, if he responds to a doubter move δ by playing the sequence of implicit rules

$$\bar{\delta} \triangleq \langle \mathbf{a} \leftarrow \mathbf{a} \mid \mathbf{a} \in [\delta] \rangle.$$

In this way he is forcing the doubter to follow with $\bar{\bar{\delta}}$, which has the same doubts as δ again. Notice that despite the fact that the occurrences will be different, the actual doubts will be the same, which is what really matters here. Easily, stalling is a valid response to δ since $\text{head}([\bar{\delta}]) = [\delta]$.

✂ **REMARK 10.3** (Notation). Once more we have overloaded a symbol here, but once more the end justifies the means: we increase readability with no possibility of ambiguity, since the follow-bar can only be applied to believer moves, while the stalling-bar only covers doubter moves. This also brings the

handy double-bar notation for the only possible reply to a stalling move, in which case the follow is always defined. In addition, it is easy to verify the following cute, bar-cancellation properties:

$$\overline{\overline{\delta}} = \delta \quad \text{and} \quad \overline{\overline{\beta}} = \beta.$$

✂ **REMARK 10.4.** For any doubter move δ , the move $\overline{\overline{\delta}}$ is equal to δ modulo a change in occurrences: it contains exactly the same doubts. Therefore, it also shares the same statement

$$[\delta] = [\overline{\overline{\delta}}].$$

If we remove stallings and their corresponding follow moves from a quasidialogue, we obtain an actual dialogue; and similar for plays. This is exactly what the function `rmstall` (defined below) does; colloquially speaking, it removes the “quasi-”. Note that if the original quasidialogue ended with a stall move, the resulting dialogue will be of odd length.

Definition 10.7 (`rmstall`). Let τ be a quasidialogue from \mathcal{P} . We define the function `rmstall` recursively by cases on $\ell := |\tau|$:

CASE 1: $\ell \leq 1$. Then either $\tau = \langle \rangle$ or $\tau = \langle \delta \rangle$ for some doubter move δ . Both alternatives contain no believer moves, and hence zero stallings; we set

$$\text{rmstall}(\tau) \triangleq \tau.$$

CASE 2: $\ell = 2$. Then $\tau := \langle \delta, \beta \rangle$, so we set

$$\text{rmstall}(\tau) \triangleq \begin{cases} \langle \delta \rangle & \text{if } \beta = \overline{\delta}, \\ \langle \delta, \beta \rangle & \text{otherwise.} \end{cases}$$

CASE 3: $\ell \geq 3$. Then $\tau = \langle \delta, \beta, \delta' \rangle \uplus \tau'$, and we need to recurse:

$$\text{rmstall}(\tau) \triangleq \begin{cases} \text{rmstall}(\delta :: \tau') & \text{if } \beta = \overline{\delta}, \\ \langle \delta, \beta \rangle \uplus \text{rmstall}(\delta' :: \tau') & \text{otherwise.} \end{cases}$$

Property 10.2 (Validity of `rmstall`). *If τ is a quasidialogue from a program, then `rmstall`(τ) is a dialogue from the same program; and if π is a quasiplay in some game, then `rmstall`(π) is a play in the same game. In addition, `rmstall` leaves dialogues and plays intact:*

$$\pi \text{ dialogue or play} \implies \text{rmstall}(\pi) = \pi$$

Property 10.3 (`rmstall` is \sqsubseteq -monotone).

$$\tau \sqsubseteq \tau' \implies \text{rmstall}(\tau) \sqsubseteq \text{rmstall}(\tau').$$

We now define the important notion of a strategy. Roughly speaking, we want a strategy to dictate how a believer should play against a doubter. If it covers all potential doubter moves, we will call it total, and if it always leads to victory, winning. Formally, we define:

Definition 10.8 (Strategy). A strategy σ in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$ is a set of plays in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$, such that:

- (i) $\sigma \neq \emptyset$;
- (ii) every play in σ has even length;
- (iii) σ is closed under even prefixes (and therefore always contains ε):

$$\pi' \sqsubseteq_e \pi \in \sigma \implies \pi' \in \sigma;$$

- (iv) σ is deterministic: if $\pi \in \sigma$ and $\pi \# \langle \delta \rangle$ is a play, then there exists *at most one* move β such that $\pi \# \langle \delta, \beta \rangle \in \sigma$.

We will call a strategy *combo-free* if none of its plays contains combo moves; in symbols, $\beta \in \pi \in \sigma \implies |\beta| = 1$.

✂ **REMARK 10.5** (Strategies as posets). A strategy σ can be seen as a *poset* (σ, \sqsubseteq_e) of even-length plays with a bottom element $\perp_\sigma = \varepsilon$.

Definition 10.9 (Total strategy). A strategy σ is called *total*, if for every play $\pi \in \sigma$ and every doubter response to it δ (i.e., a δ such that $\pi \# \delta$ is a play), there is *at least one* move β such that $\pi \# \langle \delta, \beta \rangle$ is also in σ .

Siding with Believer, we give the following definition:

Definition 10.10 (Winning strategy). A strategy σ is called *winning*, if it is total and finite.

✂ **REMARK 10.6** (Infinite plays and limits). Let us pretend awhile that dialogues (and plays) might be infinite. We should then change (among other things) the definition of strategy to demand that it is also closed under limits—or, should we? If we do so, nothing essential will change, as a correspondence between such strategies and the ones we are using here is obvious:

- starting from a strategy with infinite plays, simply remove them—all their finite, even prefixes are already included;
- starting from a strategy with only finite plays, add all limits (i.e., lubs of \sqsubseteq_e -chains).

On the other hand, if we do not close under limits, we will end up distinguishing between strategies like σ and σ_∞ , where:

$$\begin{aligned} \sigma &:= \{ \langle \delta_0, \beta_0, \dots, \delta_n, \beta_n \rangle \mid n \in \omega \}, \\ \sigma_\infty &:= \sigma \cup \{ \bigvee \sigma \}. \end{aligned}$$

Such a distinction could potentially be useful for some kind of ordinal extension of games. But for reasons of elegance and simplicity we have chosen not to deal with infinite dialogues altogether, as they are not needed for the development of this game semantics of disjunctive logic programs.

Definition 10.11 (answer_σ). Given a play π and a strategy σ in some game $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$, the play $\text{answer}_\sigma(\pi) \in \sigma$ will be:

$$\text{answer}_\sigma(\pi) \triangleq \begin{cases} \pi & \text{if } \pi \in \sigma, \\ \pi \# \langle \beta \rangle & \text{if there exists a } \beta \text{ such that } \pi \# \langle \beta \rangle \in \sigma, \\ \text{undefined} & \text{otherwise.} \end{cases}$$

This is a well-defined partial function because σ is deterministic, and so in the second branch, if such a β exists, it is necessarily unique. Furthermore, if σ is total, $\text{answer}_\sigma(\pi)$ is undefined iff $\pi^- \notin \sigma$. Hence, if we stick from the beginning to a total strategy, it will always provide us with a next move, an answer to any doubter move.

Property 10.4. *The function answer_σ behaves like a closure operator; i.e., whenever it is defined throughout the following statements, it satisfies them:*

$$\begin{aligned} \pi \sqsubseteq \text{answer}_\sigma(\pi) & \quad (\text{extensive}), \\ \pi \sqsubseteq \pi' \implies \text{answer}_\sigma(\pi) \sqsubseteq \text{answer}_\sigma(\pi') & \quad (\text{monotone}), \\ \text{answer}_\sigma(\text{answer}_\sigma(\pi)) = \text{answer}_\sigma(\pi) & \quad (\text{idempotent}). \end{aligned}$$

The DLPG game we have investigated is in a sense equivalent to the simplified one that we described previously. This is a consequence of the logical equivalence $\mathcal{D} \vee \mathcal{E} \equiv \mathcal{D} \vee \mathcal{E}$ and of the following proposition:

Proposition 10.5. *Let \mathcal{D} , \mathcal{E} , and \mathcal{F} be L.P. conjunctions such that $\mathcal{F} \equiv \mathcal{D} \vee \mathcal{E}$. Then for every disjunction $F \in \mathcal{F}$ there is some disjunction $D_F \in \mathcal{D}$ and some disjunction $E_F \in \mathcal{E}$ such that $D_F \cup E_F \subseteq F$.²*

Proof. Pick any disjunction F of \mathcal{F} . Consider the interpretation $\alpha = \mathcal{A}t \setminus F$. By definition, an assignment that makes F false, must falsify \mathcal{F} as well. But \mathcal{F} is logically equivalent to $\mathcal{D} \vee \mathcal{E}$, which means that both \mathcal{D} and \mathcal{E} are false under α . Since they are both conjunctions, there is at least one element $D_F \in \mathcal{D}$ that is false, and similarly for an element $E_F \in \mathcal{E}$. \blacksquare

We have exposed every little detail that we will need in order to develop the DLPG semantics in a formal mathematical setting. Let us do so: in the next two sections, we define the combination and splitting of both plays and strategies. The main point is that these constructions preserve all the properties that we need.

10.3 Plays: combining, restricting, and splitting

Here we show how we can combine arbitrary plays from games of the splitting of a program to create a valid play in the original program's game. The construction we use is slightly technical but hopefully the intuition behind it will be apparent to the reader, who will then have no problem appreciating and accepting the details. We will also see how to work in the opposite direction: starting from a play of a game of the combined program we can split it into two plays, valid in the games of the less disjunctive, restricted programs.

Further notational conventions. To avoid tedious repetitions in what follows, we hereby agree that \mathcal{P} will always be a proper DLP program, $\leftarrow \mathbf{G}$ a goal clause, and ϕ a proper DLP rule of \mathcal{P} ; $\mathcal{H} := (H_1, H_2)$ will be a proper partition of $H := \text{head}(\phi)$, and $(\mathcal{P}_1, \mathcal{P}_2)$ the corresponding splitting of \mathcal{P} with respect to ϕ over \mathcal{H} . Naturally we will write just ϕ_1 and ϕ_2 for the rules $\phi|_{H_1}$

²L.P. disjunctions and conjunctions were introduced in Definition 3.1.

and $\phi|_{H_2}$ respectively. The variable q ranges over the size of the partition: $q = 1, 2$.³ We use π and π_q for dialogues or plays of $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$ and $\Gamma_{\mathcal{P}_q}(\leftarrow \mathbf{G})$ respectively; τ and τ_q for “quasi-”. We remind the reader that we solely use β and δ for believer and doubter moves respectively, and that strategies are usually denoted by σ . In this setting, with this notation, we proceed to define combination, restriction, and splitting of both plays and strategies.

Combining plays

As we are about to witness, while combining plays special care must be taken because a believer move of a restricted play \mathcal{P}_q might not be valid in \mathcal{P} :

Definition 10.12 (Forbidden move). A believer move β of $\Gamma_{\mathcal{P}_q}(\leftarrow \mathbf{G})$ is called *forbidden* in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$ iff it includes ϕ_q . In symbols,

$$\text{Forbidden}_q(\beta) \stackrel{\Delta}{\iff} \phi_q \in \beta.$$

Definition 10.13 (Release). Given a believer move β_i^q of $\Gamma_{\mathcal{P}_q}(\leftarrow \mathbf{G})$, we define β_i^{q*} to be β_i^q after replacing instances of the forbidden rule ϕ_q by ϕ , so that it becomes valid in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$. We call β_i^{q*} the *release* of β_i^q from H_q to H .

Property 10.6. *The release of a move satisfies the following properties:*

- (i) $\beta_i^{q*} = \beta_i^q \iff \phi_q \notin \beta_i^q$;
- (ii) $\text{body}(\beta_i^{q*}) = \text{body}(\beta_i^q)$.

In order to justify believer moves in the combined play, we will need the following proposition:

Proposition 10.7. *Let β_1 and β_2 be two believer moves from \mathcal{P}_1 and \mathcal{P}_2 respectively, and let $\beta := \beta_1^* \uplus \beta_2^*$. Then*

$$\text{head}([\beta]) = \text{head}([\beta_1]) \cup \text{head}([\beta_2]).$$

Proof. We compute

$$\begin{aligned} \text{head}([\beta_1]) \cup \text{head}([\beta_2]) &= \left(\bigcup_{\psi \in \beta_1} \text{head}(\psi) \right) \cup \left(\bigcup_{\psi \in \beta_2} \text{head}(\psi) \right) \\ &= \bigcup_{\psi \in \beta_1 \uplus \beta_2} \text{head}(\psi) \\ &= \text{head}([\beta_1 \uplus \beta_2]) \\ &\stackrel{*}{=} \text{head}([\beta_1^* \uplus \beta_2^*]) \\ &= \text{head}([\beta]), \end{aligned}$$

where the starred equality holds since

$$\text{head}(\phi_1) \cup \text{head}(\phi_2) = H_1 \cup H_2 = \text{head}(\phi). \quad \blacksquare$$

First we define combination for what we call synchronous plays, as it is a lot simpler. Once we have seen how to combine such plays, we proceed to give the most general definition covering arbitrary plays.

³There is nothing special about the number 2 here, and everything that we develop can be stated *mutatis mutandis* for the case in which $|\mathcal{H}|$ is any larger number. However, 2 is as large as we need.

Combining synchronous plays

The three games $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$, $\Gamma_{\mathcal{P}_1}(\leftarrow \mathbf{G})$, and $\Gamma_{\mathcal{P}_2}(\leftarrow \mathbf{G})$ would be identical except that ϕ can only be part of believer moves of $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$, ϕ_1 of those of $\Gamma_{\mathcal{P}_1}(\leftarrow \mathbf{G})$, and ϕ_2 of those of $\Gamma_{\mathcal{P}_2}(\leftarrow \mathbf{G})$. Since moves are sequences, it would be delightful to simply concatenate the moves of π_1 with those of π_2 to obtain a play in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$. Doubter moves are no obstacles to this plan, but believer moves can be troublesome: they may contain the rules ϕ_q , which are not allowed in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$. This problem is easy to solve if *both* moves include their forbidden rules: we simply replace each of them by ϕ —a reasonable action, since $\phi \equiv \phi_1 \dot{\vee} \phi_2$. To deal with this case, we begin with a key definition.

Definition 10.14 (Synchronous). Two quasilays τ_1 and τ_2 in $\Gamma_{\mathcal{P}_1}(\leftarrow \mathbf{G})$ and $\Gamma_{\mathcal{P}_2}(\leftarrow \mathbf{G})$ respectively are called *synchronous* (or *in sync with each other*) if both Believers use their forbidden rules in the exact same turns. In symbols,

$$\text{for all } i \leq \min\{|\tau_1|, |\tau_2|\}, \quad \phi_1 \in \beta_i^1 \iff \phi_2 \in \beta_i^2.$$

Property 10.8. *If τ_1 is in sync with τ_2 , then so is any prefix of τ_1 with any prefix of τ_2 .*

Let π_1 and π_2 be plays of even length in the games $\Gamma_{\mathcal{P}_1}(\leftarrow \mathbf{G})$ and $\Gamma_{\mathcal{P}_2}(\leftarrow \mathbf{G})$ respectively, and suppose that the two plays are synchronous. We will describe a new play $\pi_1 \dot{\vee} \pi_2$: the *synchronous combination of π_1 and π_2 with respect to ϕ over \mathcal{H}* . To better understand this construction, we first present the idea informally, building the combined play turn by turn. We do so as follows: let

$$\begin{aligned} \pi_1 &:= \langle \delta_0^1, \beta_0^1, \delta_1^1, \beta_1^1, \dots \rangle \\ \pi_2 &:= \langle \delta_0^2, \beta_0^2, \delta_1^2, \beta_1^2, \dots \rangle \end{aligned}$$

be the two given plays. We set

$$\pi_1 \dot{\vee} \pi_2 \triangleq \langle \delta_0, \beta_0, \delta_1, \beta_1, \dots \rangle,$$

where the symbols are defined as follows.

The first move is essentially determined by the game: Doubter starts with

$$\delta_0 := \delta_0^1 + \delta_0^2;$$

Assuming that $\phi_i \notin \beta_0^i$, Believer replies with

$$\beta_0 := \beta_0^1 + \beta_0^2 = \beta_0^{1*} + \beta_0^{2*},$$

a valid move since $\text{head}([\beta_0]) \subseteq [\delta_0]$ (Property 10.6, Proposition 10.7). Now Doubter must select one occurrence from each rule-body in β_0 , and the moves δ_1^1 and δ_1^2 provide just that:

$$\delta_1 := \delta_1^1 + \delta_1^2.$$

The game goes on until the turn i in which the believers of π_1 and π_2 both play their forbidden rules ϕ_1 and ϕ_2 in their justifications. At this point, it is of course Believer's turn in the combined play $\pi_1 \dot{\vee} \pi_2$. We set his move to be

$$\beta_i := \beta_i^{1*} + \beta_i^{2*},$$

which is valid in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$. Doubter's turn: δ_{i+1}^1 and δ_{i+1}^2 consist of occurrences from the rule-bodies of β_i^1 and β_i^2 respectively. Since releasing only affects heads (Property 10.6(ii)), all occurrences in the rule-bodies of β_i^q are occurrences in the rule-bodies of β_i^{q*} as well. Therefore, she can copy the selections of δ_{i+1}^1 and δ_{i+1}^2 by playing

$$\delta_{i+1} := \delta_{i+1}^1 \# \delta_{i+1}^2.$$

She does so, and the game goes on until we run out of moves to combine.

✂ **REMARK 10.7.** We have cheated a bit, since we took for granted that the doubter of $\pi_1 \check{\vee} \pi_2$ has in her possession *two* doubter moves δ_k^1 and δ_k^2 from π_1 and π_2 respectively. This will not hold if the plays have unequal lengths; in this case, the combined play ends as soon as the shortest play ends.

We arrive at the following definition, general enough to handle quasiplays:

Definition 10.15 (Synchronous combination). Given two synchronous quasiplays

$$\begin{aligned} \tau_1 &:= \langle \delta_0^1, \beta_0^1, \delta_1^1, \beta_1^1, \dots \rangle \text{ of } \Gamma_{\mathcal{P}_1}(\leftarrow \mathbf{G}), \\ \tau_2 &:= \langle \delta_0^2, \beta_0^2, \delta_1^2, \beta_1^2, \dots \rangle \text{ of } \Gamma_{\mathcal{P}_2}(\leftarrow \mathbf{G}), \end{aligned}$$

their *synchronous combination* $\tau_1 \check{\vee} \tau_2$ over ϕ with respect to \mathcal{H} is the sequence

$$\tau_1 \check{\vee} \tau_2 \triangleq \langle \delta_0, \beta_0, \delta_1, \beta_1, \dots \rangle$$

of length $|\tau_1 \check{\vee} \tau_2| = \min\{|\tau_1|, |\tau_2|\}$, where the symbols involved are defined by

$$\begin{aligned} \delta_i &:= \delta_i^1 \# \delta_i^2, \\ \beta_i &:= \beta_i^{1*} \# \beta_i^{2*}. \end{aligned}$$

Proposition 10.9 (Validity of $\check{\vee}$). *Given two synchronous quasiplays τ_1 and τ_2 of $\Gamma_{\mathcal{P}_1}(\leftarrow \mathbf{G})$ and $\Gamma_{\mathcal{P}_2}(\leftarrow \mathbf{G})$ respectively, $\tau := \tau_1 \check{\vee} \tau_2$ is a quasiplay in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$. It follows that if τ_1 and τ_2 do not stall simultaneously, τ will be a play of the game $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$.*

Proof. First, by the definitions of δ_i and β_i , it is immediate that they are doubter and believer moves respectively. Since each quasiplay τ_q is valid in $\Gamma_{\mathcal{P}_q}(\leftarrow \mathbf{G})$, we know that

$$\begin{aligned} \text{head}([\beta_i^1]) &\subseteq [\delta_i^1], \\ \text{head}([\beta_i^2]) &\subseteq [\delta_i^2], \end{aligned}$$

and so by taking unions on both sides and by using Proposition 10.7 we obtain

$$\text{head}([\beta_i]) = \text{head}([\beta_i^1]) \cup \text{head}([\beta_i^2]) \subseteq [\delta_i^1] \cup [\delta_i^2] = [\delta_i^1 \# \delta_i^2] = [\delta_i],$$

which validates every believer move. Doubter moves are justified by the definition of release (which leaves bodies intact) as explained earlier in the sketchy description of play combination (p. 70). To verify that τ is indeed a quasiplay, observe that both τ_1 and τ_2 share the same goal with τ , and so $\delta_0 = \delta_0^1 \# \delta_0^2$ is a correct first move for a quasidialogue from \mathcal{P} to be a quasiplay of $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$.

For the second claim, observe that a stalling move in $\tau_1 \check{\vee} \tau_2$ would imply the existence of two simultaneously played stalling moves, one in τ_1 and one in τ_2 , against the hypothesis. \blacksquare

As plays never stall, we get the following property as a corollary:

Corollary 10.10 (Preservation of plays). *The synchronous combination of two synchronous plays is a play.*

The definition of $\ddot{\gamma}$ immediately yields a couple of more preservation properties:

Property 10.11 (Preservation of parity). *Let τ_1 and τ_2 be two synchronous quasiplays, both of even length. Then $\tau_1 \ddot{\gamma} \tau_2$ will also have even length.*

Property 10.12 ($\ddot{\gamma}$ is monotone). *Let τ_1 and τ_2 be two synchronous quasiplays. Then for any $\tau'_1 \sqsubseteq \tau_1$ and any $\tau'_2 \sqsubseteq \tau_2$ we have $\tau'_1 \ddot{\gamma} \tau'_2 \sqsubseteq \tau_1 \ddot{\gamma} \tau_2$.*

So far, so good. Not every pair of plays is synchronous though, which brings us to the next topic.

Combining arbitrary plays

Starting with two arbitrary quasiplays τ_1 and τ_2 , we build two new, synchronized quasiplays $\hat{\tau}_1$ and $\hat{\tau}_2$ by inserting pairs of stall–follow moves on the turns in which one believer uses his forbidden rule but the other does not, leaving the rest of the moves intact. Now we can simply combine $\hat{\tau}_1$ with $\hat{\tau}_2$. It is exactly this idea that we exploit to extend the definition of $\tau_1 \ddot{\gamma} \tau_2$ to cover asynchronous plays: *synchronize first, then combine.*

Definition 10.16 (Synchronization). Let τ_1 and τ_2 be two quasialogues from \mathcal{P}_1 and \mathcal{P}_2 respectively, and let ℓ_1 and ℓ_2 be their lengths. We recursively define their *synchronization* $\text{sync}(\tau_1, \tau_2)$ by cases depending on $\ell := \min\{\ell_1, \ell_2\}$.

CASE 1: $\ell < 2$. In this case, there are no believer moves at all (in the shortest play); and they are the only ones that can cause asynchronicity. Hence, any such pair is trivially synchronous:

$$\text{sync}(\tau_1, \tau_2) \triangleq (\tau_1, \tau_2).$$

CASE 2: $\ell \geq 2$. Then $\tau_1 := \langle \delta_1, \beta_1 \rangle \uplus \tau'_1$ and $\tau_2 := \langle \delta_2, \beta_2 \rangle \uplus \tau'_2$, and we set:

$$\text{sync}(\tau_1, \tau_2) \triangleq \begin{cases} (\langle \delta_1, \bar{\delta}_1 \rangle \uplus \text{rec}_1, \langle \delta_2, \beta_2 \rangle \uplus \text{rec}_2) & \text{if (a),} \\ (\langle \delta_1, \beta_1 \rangle \uplus \text{rec}_1, \langle \delta_2, \bar{\delta}_2 \rangle \uplus \text{rec}_2) & \text{if (b),} \\ (\langle \delta_1, \beta_1 \rangle \uplus \text{rec}_1, \langle \delta_2, \beta_2 \rangle \uplus \text{rec}_2) & \text{if (c),} \end{cases}$$

where rec_1 and rec_2 wrap the recursive calls

$$(\text{rec}_1, \text{rec}_2) := \begin{cases} \text{sync}(\bar{\delta}_1 \text{ :: } \beta_1 \text{ :: } \tau'_1, \tau'_2) & \text{if (a),} \\ \text{sync}(\tau'_1, \bar{\delta}_2 \text{ :: } \beta_2 \text{ :: } \tau'_2) & \text{if (b),} \\ \text{sync}(\tau'_1, \tau'_2) & \text{if (c),} \end{cases}$$

all according to the subcases: (a) $\phi_1 \in \beta_1$ and $\phi_2 \notin \beta_2$; (b) $\phi_1 \notin \beta_1$ and $\phi_2 \in \beta_2$; (c) otherwise.

✂ **REMARK 10.8** (Dependencies). Even though we have not incorporated ϕ_1 and ϕ_2 into the symbol **sync** of synchronization, we stress that it does, in fact, depend on both of those rules: it uses them to determine the (a)–(c). Since we have fixed ϕ , ϕ_1 , and ϕ_2 into our notation, however, we allow ourselves to simply use **sync** instead of an excessively precise name like $\text{sync}_{\phi_1, \phi_2}$. The same is true for various symbols that we use, such as Υ and $*$.

Proposition 10.13. *Synchronization is a well-defined, total operation.*

Proof. Observe that on every recursive call of **sync** both of its arguments are themselves quasidialogues from \mathcal{P}_1 and \mathcal{P}_2 ; and so it makes sense to recurse on them. Since **sync** is defined by recursion, we must be careful to ensure that it terminates on every input. Note that on every recursive call of case 2 ($\ell \geq 2$), the sum of the lengths of the arguments decreases: by 2 in subcases (a)–(b), and by 4 in (c). Eventually, at least one of them will become short enough and **sync** will reach case 1 ($\ell < 2$), which contains no recursive calls. ■

One can easily verify that synchronization behaves as expected:

Property 10.14 (Validity of synchronization). *Let $(\dot{\tau}_1, \dot{\tau}_2) := \text{sync}(\tau_1, \tau_2)$. Then*

- (i) *the pair $(\dot{\tau}_1, \dot{\tau}_2)$ is synchronous;*
- (ii) *if (τ_1, τ_2) happens to be synchronous, then $\text{sync}(\tau_1, \tau_2) = (\tau_1, \tau_2)$;*
- (iii) $\text{rmstall}(\dot{\tau}_q) = \text{rmstall}(\tau_q)$;
- (iv) *sync never produces two fresh stalling moves in the same turn.*

Proposition 10.15 (Preservation of parity). *If τ_1 and τ_2 have even lengths, then so do the elements of $\text{sync}(\tau_1, \tau_2)$.*

Proof. This is immediate by the very definition of **sync**, which generates its output $\text{sync}(\tau_1, \tau_2)$ incrementally by pairs of moves, while consuming pairs of moves from its inputs τ_1 and τ_2 . Thus, if both of its inputs are of even length, the same will be true for its outputs. ■

Proposition 10.16 (**sync** is monotone). *Let τ'_1 and τ'_2 be two quasidialogues, and let $\tau_1 \sqsubseteq \tau'_1$ and $\tau_2 \sqsubseteq \tau'_2$. Then $\text{sync}(\tau_1, \tau_2) \sqsubseteq \text{sync}(\tau'_1, \tau'_2)$. Moreover, if $\tau_1 \sqsubset \tau'_1$ and $\tau_2 \sqsubset \tau'_2$, then $\text{sync}(\tau_1, \tau_2) \sqsubset \text{sync}(\tau'_1, \tau'_2)$.*

Proof. Just like in the previous proof, we only need to observe how synchronization really works. In fact, the construction of $\text{sync}(\tau_1, \tau_2)$ must end with a call to case 1 of its definition, with either $\ell = 0$ or $\ell = 1$. In either case, extending any of τ_1 or τ_2 results in an extension of the corresponding synchronized quasidialogue. ■

We know how to combine synchronous plays; and we know how to synchronize asynchronous plays. Compose the two and behold: a method to combine arbitrary plays. Just as in the synchronous case, we state the definition in its most general form, which is able to handle quasidialogues.

Definition 10.17 (Combination of quasialogues). Let τ_1 and τ_2 be two quasialogues from \mathcal{P}_1 and \mathcal{P}_2 respectively. We define their *combination* $\tau_1 \Upsilon \tau_2$ by composition:

$$\Upsilon \triangleq \check{\Upsilon} \circ \text{sync}.$$

In other words,

$$\tau_1 \Upsilon \tau_2 \triangleq \check{\tau}_1 \check{\Upsilon} \check{\tau}_2,$$

where $(\check{\tau}_1, \check{\tau}_2) := \text{sync}(\tau_1, \tau_2)$.

✂ **REMARK 10.9.** By Property 10.14(ii), Υ is compatible with $\check{\Upsilon}$, in the sense that it yields the same output in case its input is synchronous.

✂ **REMARK 10.10** (End of $\tau_1 \Upsilon \tau_2$). According to the definitions of synchronization and synchronous combination, the combined quasialogue $\tau_1 \Upsilon \tau_2$ ends exactly when we reach the final move of either τ_1 or τ_2 (whichever comes first).

By defining the general combination as a composition of sync and $\check{\Upsilon}$, we readily get their composable properties as corollaries:

Corollary 10.17 (Validity of combination). *Given two (quasi)dialogues τ_1 and τ_2 from \mathcal{P}_1 and \mathcal{P}_2 respectively, their disjunctive combination $\tau_1 \Upsilon \tau_2$ is a (quasi)dialogue from \mathcal{P} .*

Corollary 10.18 (Preservation of plays). *If τ_1 and τ_2 are two (quasi)plays in $\Gamma_{\mathcal{P}_1}(\leftarrow \mathbf{G})$ and $\Gamma_{\mathcal{P}_2}(\leftarrow \mathbf{G})$ respectively, then $\tau_1 \Upsilon \tau_2$ is a (quasi)play in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$.*

Corollary 10.19 (Preservation of parity). *If π_1 and π_2 have even lengths, then so does $\pi_1 \Upsilon \pi_2$.*

Corollary 10.20 (Υ is monotone). *Let π'_1 and π'_2 be two dialogues, and let $\pi_1 \sqsubseteq \pi'_1$ and $\pi_2 \sqsubseteq \pi'_2$. Then $\pi_1 \Upsilon \pi_2 \sqsubseteq \pi'_1 \Upsilon \pi'_2$. Moreover, if $\pi_1 \sqsubset \pi'_1$ and $\pi_2 \sqsubset \pi'_2$, then $\pi_1 \Upsilon \pi_2 \sqsubset \pi'_1 \Upsilon \pi'_2$.*

✂ **REMARK 10.11** (Losing is not preserved). Even though Doubter may be victorious in two plays, in the combined version she might not be so. As a counterexample, consider the following program and splitting, in which Doubter wins both π_1 and π_2 (i.e., Believer cannot play a valid response to either plays) and yet she is going to lose in the combined play:

$$\mathcal{P} := \left\{ \begin{array}{l} \text{p} \leftarrow \text{a} \\ \text{p} \leftarrow \text{b} \\ \text{a} \vee \text{b} \leftarrow \\ \text{c} \vee \text{d} \leftarrow \end{array} \right\} \rightsquigarrow \left(\mathcal{P}_1 := \left\{ \begin{array}{l} \text{p} \leftarrow \text{a} \\ \text{p} \leftarrow \text{b} \\ \text{a} \vee \text{b} \leftarrow \\ \text{c} \leftarrow \end{array} \right\}, \mathcal{P}_2 := \left\{ \begin{array}{l} \text{p} \leftarrow \text{a} \\ \text{p} \leftarrow \text{b} \\ \text{a} \vee \text{b} \leftarrow \\ \text{d} \leftarrow \end{array} \right\} \right)$$

and the plays

$$\underbrace{\left[\begin{array}{l} \text{goal :} \quad \leftarrow \underline{\text{p}} \\ \beta_0^1 : \quad \text{p} \leftarrow \underline{\text{a}} \end{array} \right]}_{\pi_1} \Upsilon \underbrace{\left[\begin{array}{l} \text{goal :} \quad \leftarrow \underline{\text{p}} \\ \beta_0^2 : \quad \text{p} \leftarrow \underline{\text{b}} \end{array} \right]}_{\pi_2} = \underbrace{\left[\begin{array}{l} \text{goal :} \quad \leftarrow \underline{\text{p}} \\ \beta_0 : \quad \text{p} \leftarrow \underline{\text{a}} \\ \quad \quad \text{p} \leftarrow \underline{\text{b}} \\ [\beta_0] : \quad \text{p} \leftarrow \underline{\text{a} \vee \text{b}} \end{array} \right]}_{\pi}.$$

Notice that Believer cannot move in neither of the starting plays, but he can certainly move in their combination by playing the rule $\mathbf{a} \vee \mathbf{b} \leftarrow$. However, as we will later see, such misbehaviors are evaded if the plays π_q come from a strategy splitting; the impatient can read Corollary 10.28 (p. 81).

✂ **REMARK 10.12** (Combining goals). When combining plays, their common goal $\leftarrow \mathbf{G}$ does not really have to be all that common. We can combine a play in $\Gamma_{\mathcal{P}_1}(\leftarrow \mathbf{G}_1)$ with a play in $\Gamma_{\mathcal{P}_2}(\leftarrow \mathbf{G}_2)$ to get a new play in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$, where $\mathbf{G} := \mathbf{G}_1 \vee \mathbf{G}_2$. To do so, we first extract plays in $\Gamma_{\mathcal{P}_1}(\leftarrow \mathbf{G})$ and $\Gamma_{\mathcal{P}_2}(\leftarrow \mathbf{G})$ by altering only the first move in each play so that it is restricted to \mathbf{G}_q ; then we combine. But we will not need to do such a thing in this work.

We have seen how to combine plays; now it is time to restrict and split them.

Restricting and splitting plays

Definition 10.18 (Play restriction). Suppose that π is a play and consider the sequence obtained by restricting every believer move in π that contains ϕ to an identical move in which ϕ has been restricted to H_q . We denote this sequence by $\pi|_{H_q}^\phi$ and call it the *restriction* of the play π to H_q with respect to ϕ .

Theorem 10A. *For any play π of $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$, its restriction $\pi_q := \pi|_{H_q}^\phi$ is a valid play in $\Gamma_{\mathcal{P}_q}(\leftarrow \mathbf{G})$.*

Proof. Note that the only alteration performed leaves all bodies intact, so that every doubter move remains valid. The only case where a head is altered is when a believer move β of π includes the forbidden rule ϕ . Denoting by β^q the corresponding believer move of π_q , it is evident that $\text{head}([\beta^q]) \subseteq \text{head}([\beta])$ so that in both plays, every believer move is valid as well. ■

Naturally we define splitting as a pair of restrictions:

Definition 10.19 (Play splitting). Given a play π of $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$, the *play splitting* of π with respect to ϕ over \mathcal{H} is the pair

$$\pi|_{\mathcal{H}}^\phi \triangleq (\pi|_{H_1}^\phi, \pi|_{H_2}^\phi).$$

Thanks to Theorem 10A, these plays are indeed valid in the corresponding games.

After all this work on plays, strategies are next; but we have done most of the hard work in this section, so that the following one will be a breeze.

10.4 Strategies: combining, restricting, and splitting

In the previous section we defined combination, restriction, and splitting for plays in a given game, and proved the correctness of these definitions. Now we will do the same for strategies, keeping the same notation that we agreed upon. The combination and the splitting of strategies lie at the hearts of our proofs of completeness and soundness respectively.

Combining strategies

Combination of plays can be extended to strategies in a straightforward way:

Definition 10.20 (Combination of strategies). Given two strategies σ_1 and σ_2 in $\Gamma_{\mathcal{P}_1}(\leftarrow \mathbf{G})$ and $\Gamma_{\mathcal{P}_2}(\leftarrow \mathbf{G})$ respectively, we define their *combination*

$$\sigma_1 \curlywedge \sigma_2 \triangleq \{\pi_1 \curlywedge \pi_2 \mid \pi_1 \in \sigma_1 \text{ and } \pi_2 \in \sigma_2\}.$$

Definition 10.21. Given a play $\pi \in \sigma := \sigma_1 \curlywedge \sigma_2$, we call the elements of the set

$$\mathbf{C}(\pi) \triangleq \{(\pi_1, \pi_2) \in \sigma_1 \times \sigma_2 \mid \pi_1 \curlywedge \pi_2 = \pi\}$$

the *creators* of π from σ_1 and σ_2 . Equipped with the product order induced by either \sqsubseteq or \sqsubseteq_e , the set $\mathbf{C}(\pi)$ becomes a poset; and as we are about to see, it always has a least element: a pair which we naturally call the *shortest creators* of π from σ_1 and σ_2 . Note that the two orderings \sqsubseteq and \sqsubseteq_e coincide in this poset since all plays involved come from strategies, and therefore are of even length.

We will now prove an important “decomposition” property, which essentially allows us to reverse play-combination in case the two plays come from appropriate strategies.

Lemma 10.21 (Reversibility of combination). *Given a play $\pi \in \sigma := \sigma_1 \curlywedge \sigma_2$, we can extract in a unique way, two synchronized quasiplays τ_1 and τ_2 , each of length $|\pi|$, such that $\tau_1 \curlywedge \tau_2 = \pi$ and both τ_q agree with σ_q , i.e., $\text{rmstall}(\tau_q) \in \sigma_q$.*

Proof. Corollary 10.19 guarantees that $|\pi|$ is even, which allows us to prove the lemma by induction on $\ell := |\pi| / 2$:

BASE: ($\ell = 0$). Trivially, $(\langle \rangle, \langle \rangle)$ is the pair that we seek.

INDUCTION STEP: ($\ell = n + 1$). Let

$$\pi := \langle \delta_0, \beta_0, \dots, \delta_n, \beta_n \rangle.$$

By the induction hypothesis, we know that there are two unique, synchronized quasiplays

$$\begin{aligned} \tau'_1 &:= \langle \delta_0^1, \beta_0^1, \dots, \delta_{n-1}^1, \beta_{n-1}^1 \rangle \\ \tau'_2 &:= \langle \delta_0^2, \beta_0^2, \dots, \delta_{n-1}^2, \beta_{n-1}^2 \rangle \end{aligned}$$

which combine into π^- and agree with the corresponding strategies. Notice that if (τ_1, τ_2) satisfies the requested properties for π , then (τ_1^-, τ_2^-) satisfies them for π^- , so that $\tau_1^- = \tau'_1$ and $\tau_2^- = \tau'_2$. This guarantees that the following construction is the only one possible. We need to determine δ_n^q and β_n^q . Since π is a play, the doubter move δ_n is either the first move ($n = 0$), or it consists of doubts from the bodies of the last believer moves in τ'_1 and τ'_2 ($n > 0$). In the first case, we set $\delta_n^q := \delta_n$, while in the second one, we split it in a unique way in two parts,

$$\delta_n := \delta_n^1 \# \delta_n^2,$$

such that δ_n^q is a valid doubter response to τ'_q . We now use these δ_n^q to obtain the corresponding believer moves. For this we appeal to the fact that σ_q (being

$$\begin{array}{ccccc}
T_1 & \xleftarrow{\text{proj}_{\sigma_1}} & \sigma_1 \curlywedge \sigma_2 & \xrightarrow{\text{proj}_{\sigma_2}} & T_2 \\
\downarrow - & & \downarrow - & & \downarrow - \\
T_1 & \xleftarrow{\text{proj}_{\sigma_1}} & \sigma_1 \curlywedge \sigma_2 & \xrightarrow{\text{proj}_{\sigma_2}} & T_2
\end{array}$$

Figure 10.1: Commutative diagram for Property 10.22.

deterministic) can have at most one next-move for the play $\text{rmstall}(\tau'_q \# \langle \delta_n^q \rangle)$, and we know that it has at least one since $\pi \in \sigma_1 \curlywedge \sigma_2$. We first define

$$b_n^q := (\text{answer}_{\sigma_q} \circ \text{rmstall})(\tau'_q \# \langle \delta_n^q \rangle).$$

Note that as this is an answer from σ_q , the quasiplay $\tau'_q \# \langle \delta_n^q, b_n^q \rangle$ remains in agreement with it. Now we finally obtain the believer moves by setting

$$(\beta_n^1, \beta_n^2) := \begin{cases} (\bar{\delta}_n^1, b_n^2) & \text{if } \phi_1 \in b_n^1 \text{ and } \phi_2 \notin b_n^2 \\ (b_n^1, \bar{\delta}_n^2) & \text{if } \phi_1 \notin b_n^1 \text{ and } \phi_2 \in b_n^2 \\ (b_n^1, b_n^2) & \text{otherwise.} \end{cases} \quad \blacksquare$$

Definition 10.22 (Projections of plays). We call the quasiplay τ_q of the above lemma the *projection of π on σ_q* and denote it by $\text{proj}_{\sigma_q}(\pi)$.

By their construction, projections satisfy the following property:

Property 10.22. *Given any $\pi \in \sigma := \sigma_1 \curlywedge \sigma_2$,*

$$\text{proj}_{\sigma_q}(\pi^-) = (\text{proj}_{\sigma_q}(\pi))^-.$$

In other words, denoting by T_q the sets of quasiplays in $\Gamma_{\mathcal{P}_q}(\leftarrow \mathbf{G})$, the diagram in Figure 10.1 commutes.

A link between projections and shortest creators is revealed: starting with π as above, first use Lemma 10.21 to obtain the quasiplay projections of π . Then, using rmstall , remove all existing stall-follow moves; and finally, in case the quasiplay ended in a stall move, appeal to the strategy's totality to append the appropriate answer from σ_q . Following these steps, one actually obtains the shortest creators of π from σ_1 and σ_2 . This will be clarified shortly in a lemma—but first, a definition that we will need:

Definition 10.23. Let $\pi \in \sigma := \sigma_1 \curlywedge \sigma_2$. We define the function $\text{cr}_{\sigma_q} : \sigma \rightarrow \sigma_q$ as the composition

$$\text{cr}_{\sigma_q} \triangleq \text{answer}_{\sigma_q} \circ \text{rmstall} \circ \text{proj}_{\sigma_q},$$

and simply write

$$\text{cr}(\pi) \triangleq (\text{cr}_{\sigma_1}(\pi), \text{cr}_{\sigma_2}(\pi))$$

for the function $\text{cr} : \sigma \rightarrow \sigma_1 \times \sigma_2$. Both functions will turn out to be total.

Lemma 10.23. *Let $\pi \in \sigma := \sigma_1 \dot{\vee} \sigma_2$. Then*

$$\text{cr}(\pi) = \min \mathbf{C}(\pi).$$

Proof. Let $\pi := \langle \delta_0, \beta_0, \dots, \delta_n, \beta_n \rangle \in \sigma_1 \dot{\vee} \sigma_2$ and pick any pair of creators $(\pi_1, \pi_2) \in \mathbf{C}(\pi)$, so that $\pi_1 \dot{\vee} \pi_2 = \pi$. We will show that $\text{cr}(\pi) \sqsubseteq_e (\pi_1, \pi_2)$. Set $(\tau_1, \tau_2) := \text{sync}(\pi_1, \pi_2)$, and write

$$\begin{aligned} \tau_1 &:= \langle \delta_0^1, \beta_0^1, \delta_1^1, \beta_1^1, \dots, \delta_{n_1}^1, \beta_{n_1}^1 \rangle \\ \tau_2 &:= \langle \delta_0^2, \beta_0^2, \delta_1^2, \beta_1^2, \dots, \delta_{n_2}^2, \beta_{n_2}^2 \rangle. \end{aligned}$$

By definition, $\pi = \tau_1 \dot{\vee} \tau_2$, so that

$$\pi = \left\langle \underbrace{\delta_0^1 \dot{+} \delta_0^2}_{\delta_0}, \underbrace{\beta_0^{1*} \dot{+} \beta_0^{2*}}_{\beta_0}, \underbrace{\delta_1^1 \dot{+} \delta_1^2}_{\delta_1}, \underbrace{\beta_1^{1*} \dot{+} \beta_1^{2*}}_{\beta_1}, \dots, \underbrace{\delta_n^1 \dot{+} \delta_n^2}_{\delta_n}, \underbrace{\beta_n^{1*} \dot{+} \beta_n^{2*}}_{\beta_n} \right\rangle,$$

where $n := \min \{n_1, n_2\}$. Now, using the monotonicity of rmstall and answer_{σ_q} (Properties 10.3 and 10.4), we reason:

$$\begin{aligned} \text{proj}_{\sigma_q}(\pi) \sqsubseteq_e \tau_q &\implies \text{rmstall}(\text{proj}_{\sigma_q}(\pi)) \sqsubseteq_e \text{rmstall}(\tau_q) \\ &\implies \text{answer}_{\sigma_q}(\text{rmstall}(\text{proj}_{\sigma_q}(\pi))) \sqsubseteq_e \text{answer}_{\sigma_q}(\text{rmstall}(\tau_q)) \\ &\implies \text{cr}(\pi) \sqsubseteq_e \text{answer}_{\sigma_q}(\text{rmstall}(\tau_q)). \end{aligned}$$

Notice that τ_q cannot end in a stalling move, and so $\text{rmstall}(\tau_q)$ will have even length. According to Definition 10.11, we can then expect that

$$\text{answer}_{\sigma_q}(\text{rmstall}(\tau_q)) = \text{rmstall}(\tau_q),$$

so that by using Property 10.14(iii) we derive

$$\text{cr}(\pi) \sqsubseteq_e \text{rmstall}(\tau_q) = \pi_q,$$

which is what we wanted to prove. ▀

We now relate the shortest creators of a play with those of its even prefixes.

Proposition 10.24. *Let $\varepsilon \neq \pi \in \sigma := \sigma_1 \dot{\vee} \sigma_2$, and let $\pi_q := \text{cr}_{\sigma_q}(\pi)$. Then*

$$\text{cr}(\pi^-) = \begin{cases} (\pi_1^-, \pi_2) & \text{if (a),} \\ (\pi_1, \pi_2^-) & \text{if (b),} \\ (\pi_1^-, \pi_2^-) & \text{if (c),} \end{cases}$$

depending on whether (a) the penultimate believer move of $\text{proj}_{\sigma_2}(\pi)$ is a stalling, (b) the penultimate believer move of $\text{proj}_{\sigma_1}(\pi)$ is a stalling, or (c) otherwise.

Proof. This is a consequence of Property 10.22. Perhaps the weird-looking conditions need further explanation. The last believer move of π_q is needed to form the last move of π^- iff the penultimate move of $\text{proj}_{\sigma_q}(\pi)$ is a stalling. This holds because once we project a play to σ_q , stalling moves might appear to the quasisplay projections. If the penultimate move of $\text{proj}_{\sigma_q}(\pi)$ is such a stalling,

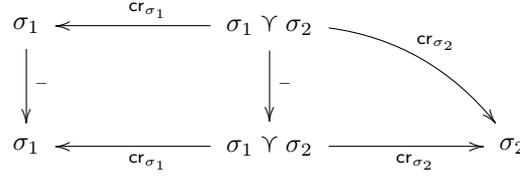


Figure 10.2a: The penultimate believer move of $\text{proj}_{\sigma_2}(\pi)$ is a stalling.

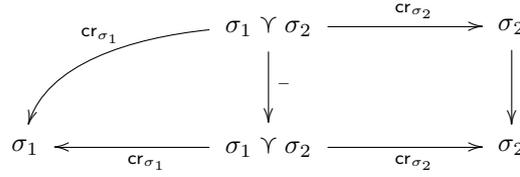


Figure 10.2b: The penultimate believer move of $\text{proj}_{\sigma_1}(\pi)$ is a stalling.

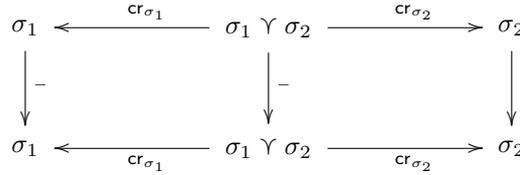


Figure 10.2c: Otherwise.

once we delete the last move from π_q to obtain π_q^- , and use rmstall to remove all stall–follow moves, the resulting play will be of odd length. $\text{answer}_{\sigma_q}(\tau)$ will then reproduce the move we deleted, thus bringing us back to π_q .

Since \curlywedge never concatenates two stalling moves, $\text{proj}_{\sigma_1}(\pi)$ and $\text{proj}_{\sigma_2}(\pi)$ will never stall simultaneously; this proves that the conditions (a)–(c) are mutually exclusive. \blacksquare

This proposition might become clearer after inspecting the three commutative diagrams of Figure 10.2. There, the stated equality is represented by the commutativity of an appropriate diagram, one for each case (a)–(c).

Corollary 10.25. *Let $\varepsilon \neq \pi \in \sigma := \sigma_1 \curlywedge \sigma_2$. Then*

- (i) $\min \mathbf{C}(\pi^-) \sqsubseteq_e \min \mathbf{C}(\pi)$;
- (ii) $\min \mathbf{C}(\pi') \sqsubseteq_e \min \mathbf{C}(\pi)$, for any $\pi' \sqsubseteq_e \pi$.

Theorem 10B. *The set $\sigma := \sigma_1 \curlywedge \sigma_2$, is a strategy in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$.*

Proof. Foremost it is indeed a set of plays in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$ thanks to Corollary 10.17.

Non-empty. Since σ_1 and σ_2 are strategies, they are both non-empty, and so there is at least one play in σ (obtained by combination).

Even-length. This is a direct application of Corollary 10.19.

Even-prefix-closed. This is immediate by Corollary 10.25, since both σ_q are strategies and therefore closed under even prefixes.

Deterministic. Towards a contradiction, assume that π and $\tilde{\pi}$ are plays of even length in σ that differ only in the last believer move:

$$\begin{aligned}\pi &:= \langle \delta_0, \beta_0, \dots, \delta_n, \beta_n \rangle \\ \tilde{\pi} &:= \langle \delta_0, \beta_0, \dots, \delta_n, \tilde{\beta}_n \rangle.\end{aligned}$$

For each of them, use Lemma 10.21 to extract its two quasiplay projections on σ_1 and σ_2 ; (τ_1, τ_2) from π and $(\tilde{\tau}_1, \tilde{\tau}_2)$ from $\tilde{\pi}$.

$$\begin{aligned}\tau_1 &:= \langle \delta_0^1, \beta_0^1, \dots, \delta_n^1, \beta_n^1 \rangle & \tilde{\tau}_1 &:= \langle \delta_0^1, \beta_0^1, \dots, \delta_n^1, \tilde{\beta}_n^1 \rangle \\ \tau_2 &:= \langle \delta_0^2, \beta_0^2, \dots, \delta_n^2, \beta_n^2 \rangle & \tilde{\tau}_2 &:= \langle \delta_0^2, \beta_0^2, \dots, \delta_n^2, \tilde{\beta}_n^2 \rangle,\end{aligned}$$

so that

$$\begin{aligned}\pi &= \langle \delta_0^1 \uplus \delta_0^2, \beta_0^{1*} \uplus \beta_0^{2*}, \dots, \delta_n^1 \uplus \delta_n^2, \beta_n^{1*} \uplus \beta_n^{2*} \rangle \\ \tilde{\pi} &= \langle \delta_0^1 \uplus \delta_0^2, \beta_0^{1*} \uplus \beta_0^{2*}, \dots, \delta_n^1 \uplus \delta_n^2, \tilde{\beta}_n^{1*} \uplus \tilde{\beta}_n^{2*} \rangle.\end{aligned}$$

Now, which of the four statements $\phi \in \beta_n^{q*}$ and $\phi \in \tilde{\beta}_n^{q*}$ hold? By a tedious and trivial inspection of all 16 different cases that arise, one can confirm that every case leads to a contradiction, by obtaining *two* plays of the *same* strategy σ_q , differing only in their final (believer) move. This is of course absurd because strategies are deterministic. \blacksquare

Proposition 10.26 (Preservation of totality). *The combined strategy $\sigma_1 \curlyvee \sigma_2$ is total if both σ_1 and σ_2 are total.*

Proof. Suppose that we are given a play $\langle \delta_0, \beta_0, \dots, \delta_n, \beta_n, \delta_{n+1} \rangle$ such that the immediate prefix

$$\pi := \langle \delta_0, \beta_0, \dots, \delta_n, \beta_n \rangle \in \sigma.$$

We seek a believer move β_{n+1} such that $\pi \uplus \langle \delta_{n+1}, \beta_{n+1} \rangle \in \sigma$. Use Lemma 10.21 to extract the quasiplay projections (τ_1, τ_2) of π on σ_1 and σ_2 , so that

$$\pi = \tau_1 \ddot{\vee} \tau_2 := \langle \delta_0^1 \uplus \delta_0^2, \beta_0^{1*} \uplus \beta_0^{2*}, \dots, \delta_n^1 \uplus \delta_n^2, \beta_n^{1*} \uplus \beta_n^{2*} \rangle.$$

Since δ_{n+1} is a valid next-move for π , it consists of doubts from the bodies of β_n^1 and β_n^2 (Property 10.6(ii)), so that it can be unambiguously split into two sequences of such doubts $\delta_{n+1} := \delta_{n+1}^1 \uplus \delta_{n+1}^2$. By the totality of σ_q , there must be at least one believer move β_{n+1}^q satisfying

$$\pi_q^+ := \text{rmstall}(\tau_q) \uplus \langle \delta_{n+1}^q, \beta_{n+1}^q \rangle \in \sigma_q.$$

Easily now, $\pi_1^+ \curlyvee \pi_2^+$ contains the believer move that we need. \blacksquare

Proposition 10.27 (Preservation of finiteness). *The combined strategy $\sigma_1 \curlyvee \sigma_2$ is finite if both σ_1 and σ_2 are finite.*

Proof. By the definition of strategy combination, $|\sigma|$ is bounded by $|\sigma_1 \times \sigma_2|$, which is finite since both σ_q are finite. \blacksquare

Remembering that total + finite = winning, we arrive at the following corollary:

Corollary 10.28 (Preservation of winning). *The combined strategy $\sigma_1 \curlywedge \sigma_2$ is winning if both σ_1 and σ_2 are winning.*

Restricting and splitting strategies

Definition 10.24 (Strategy restriction). Let σ be a strategy in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$. Then the *restriction* of σ with respect to ϕ on H_q , is the set of plays defined by:

$$\sigma|_{H_q}^{\phi} \triangleq \left\{ \pi|_{H_q}^{\phi} \mid \pi \in \sigma \right\}.$$

Theorem 10C. *The set of plays $\sigma_q := \sigma|_{H_q}^{\phi}$ is a valid strategy in $\Gamma_{\mathcal{P}_q}(\leftarrow \mathbf{G})$.*

Proof. Non-empty. This is trivial, since $\sigma \neq \emptyset$.

Even-length. It is obvious that restriction leaves lengths of plays intact, so that every member of σ_q has even length.

Even-prefix-closed. Let $\pi'_q \sqsubseteq \pi_q \in \sigma_q$, with $\ell := |\pi'_q|$ even. We need $\pi'_q \in \sigma_q$. Since $\pi_q \in \sigma_q$, there is a $\pi \in \sigma$ such that $\pi_q = \pi|_{H_q}^{\phi}$. We now compute:

$$\pi'_q = \pi_q \upharpoonright_{\ell} = \left(\pi|_{H_q}^{\phi} \right) \upharpoonright_{\ell} = \underbrace{\left(\pi \upharpoonright_{\ell} \right)}_{\in \sigma} |_{H_q}^{\phi},$$

which shows that $\pi'_q \in \sigma_q$ as was desired.

Deterministic. Suppose that we have the following plays in σ_q :

$$\begin{aligned} \pi_q &:= \langle \delta_0^q, \beta_0^q, \delta_1^q, \beta_1^q, \dots, \delta_k^q, \beta_k^q \rangle, \\ \widetilde{\pi}_q &:= \langle \delta_0^q, \beta_0^q, \delta_1^q, \beta_1^q, \dots, \delta_k^q, \widetilde{\beta}_k^q \rangle, \end{aligned}$$

so that there are plays

$$\begin{aligned} \pi &:= \langle \delta_0, \beta_0, \delta_1, \beta_1, \dots, \delta_k, \beta_k \rangle, \\ \widetilde{\pi} &:= \langle \widetilde{\delta}_0, \widetilde{\beta}_0, \widetilde{\delta}_1, \widetilde{\beta}_1, \dots, \widetilde{\delta}_k, \widetilde{\beta}_k \rangle, \end{aligned}$$

in σ , such that

$$\pi_q = \pi|_{H_q}^{\phi} \quad \text{and} \quad \widetilde{\pi}_q = \widetilde{\pi}|_{H_q}^{\phi}.$$

Observe that since play-splitting leaves all doubter moves unaltered, we have that $\delta_i = \widetilde{\delta}_i (= \delta_i^q)$ for all $i = 0, \dots, k$. By finite induction on i we show the corresponding equalities for the believer moves. Assume (the inductive hypothesis) that $\beta_j = \widetilde{\beta}_j$ holds for $0 \leq j < i \leq k$. Now look at the plays $\pi \upharpoonright_{2i+2}$ and $\widetilde{\pi} \upharpoonright_{2i+2}$ which both belong in σ since it is prefix-closed. Then $\beta_i = \widetilde{\beta}_i$ since both plays belong in the same (deterministic) strategy σ . This essentially yields $\pi = \widetilde{\pi}$ and consequently $\pi_q = \widetilde{\pi}_q$, which establishes the determinacy of σ_q . \blacksquare

Proposition 10.29 (Preservation of totality). *If σ is total, then so is $\sigma|_{H_q}^{\phi}$.*

Proof. We are given a play $\pi_q \in \sigma_q$ and a doubter move δ such that $\pi_q \# \delta$ is valid in $\Gamma_{\mathcal{P}_q}(\leftarrow \mathbf{G})$, and we seek a believer next-move for it. By the hypothesis, there exists some $\pi \in \sigma$, such that $\pi_q = \pi|_{H_q}^\phi$. Since bodies are left intact by restriction, $\pi \# \delta$ is valid in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$. Hence, by the totality of σ , there is a believer move β such that $\pi \# \langle \delta, \beta \rangle \in \sigma$. This means that $(\pi \# \langle \delta, \beta \rangle)|_{H_q}^\phi \in \sigma_q$, and its last move is the believer move that we sought. \blacksquare

Proposition 10.30 (Preservation of finiteness). *If σ is finite, then so is $\sigma|_{H_q}^\phi$.*

Proof. By definition, every play in $\sigma|_{H_q}^\phi$ is created by restricting a play of σ . This grants finiteness, as $\sigma|_{H_q}^\phi$ is nothing more than the image of a function $(\pi \mapsto \pi|_{H_q}^\phi)$ whose domain is the finite set σ . \blacksquare

Corollary 10.31 (Preservation of winning). *If σ is winning, then so is $\sigma|_{H_q}^\phi$.*

For one last time, we use restriction to obtain splitting:

Definition 10.25 (Strategy splitting). Let σ be a strategy in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$. The *splitting of σ with respect to ϕ over \mathcal{H}* is the pair

$$\sigma|_{\mathcal{H}}^\phi \triangleq (\sigma|_{H_1}^\phi, \sigma|_{H_2}^\phi).$$

This completes our game-theoretic weaponry for DLP programs. We have finally reached the point that we can put all those pieces together to prove that the DLP game semantics is sound and complete with respect to the minimal model semantics.

10.5 Game semantics

This is essentially the same definition as in the LPG semantics:

Definition 10.26 (DLPG semantics). Let \mathcal{P} be a DLP program. A goal $\leftarrow \mathbf{G}$ *succeeds*, if Believer has a winning strategy in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$.

The DLPG semantics

- $\mathcal{V}_{\text{DLPG}} \triangleq \mathbb{B}$.
- $\mathcal{M}_{\text{DLPG}}$ is the set of strategies based on DLP programs.
- \mathbf{m}_{DLPG} maps every DLP program to the set of corresponding winning strategies in the DLPG game.
- $\mathbf{a}_{\text{DLPG}}(\Sigma)(Q) \triangleq \begin{cases} \mathbf{T}, & \text{if there is a winning strategy } \sigma \in \Sigma \text{ for } Q \\ \mathbf{F}, & \text{otherwise.} \end{cases}$

✂ **REMARK 10.13.** Notice at once that if stalling was allowed, Believer would have a winning strategy of “forever stalling”, for any goal. This would make every disjunction derivable! This is one more argument in favor of giving the benefit of the doubt to the Doubter.

✂ **REMARK 10.14** (Backwards compatibility). In the model-theoretic side of semantics, extensions of the language are backwards compatible: for DLP, a program without disjunctions has a unique minimal model, viz. the least Herbrand model; for LPN, the well-founded model of a program without negations is two-valued and coincides with the least Herbrand model as well. Similar compatibilities are desired and achieved in the game-theoretic side. We highlight that strictly speaking, the DLP game is not *directly* compatible with the LP game in the sense that it does not reduce to it when it is played on an LP program. The reason behind this is that in the DLP game, believers can play what we have called combo moves.⁴ Nevertheless, we will prove in Lemma 10.32 that whenever there is a winning strategy in the DLP game of an LP program, then there is a winning strategy in which the believer never plays more than one rule, and is thus compatible with the LP game. In other words, the extra “combo-rule” of the DLP game is unnecessary in the absence of disjunctions.

10.6 Soundness and completeness

Theorem 10D (Soundness of the finite, clean DLP game semantics). *Let \mathcal{P} be a finite, clean DLP program, and $\leftarrow \mathbf{G}$ a goal. If there is a winning strategy in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$, then G is true in every minimal model of \mathcal{P} .*

PROOF is by induction on the number $N(\mathcal{P})$ of disjunction symbols (\vee) that appear in the heads of \mathcal{P} .

BASE. Let σ be a winning strategy in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$. We claim that *there exists a combo-free winning strategy σ' in the same game*. If so, observe that the head of the first believer move of every play of σ' must be one disjunct \mathbf{g} of \mathbf{G} . Hence by altering the first (doubter) moves of its plays to correctly doubt \mathbf{g} , we end up with a combo-free, winning strategy in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{g})$. This is now compatible with the LP game and we can use the soundness of the LP game semantics (Corollary 8.1) to obtain $g \in \text{LHM}(\mathcal{P})$. But this means that G is true in $\text{LHM}(\mathcal{P})$, the only minimal model of \mathcal{P} .

Proof of the claim. If σ contains no combo moves, we are done. Otherwise, pick a maximal play from σ such that it contains at least one combo move and look at the last one, β . We show that we can safely replace β by a non-combo move that contains exactly one of its rules, and still win the game. Towards a contradiction, suppose that no such rule exists. Then every rule in β contains a “bad” atom such that when doubted by the doubter, we cannot win. But this contradicts the fact that σ is winning, as it contains no answer for a doubter move that doubts exactly one bad atom from each rule of β .

INDUCTION STEP. Since \mathcal{P} is a proper DLP program, we can pick a proper DLP rule $\phi \in \mathcal{P}$, and split \mathcal{P} with respect to it over some proper partition (H_1, H_2) of $\text{head}(\phi)$ to get $(\mathcal{P}_1, \mathcal{P}_2)$. Notice that $\max\{N(\mathcal{P}_1), N(\mathcal{P}_2)\} < N(\mathcal{P})$, which allows us to use the induction hypothesis for both programs \mathcal{P}_q .

Let σ be a winning strategy in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$, and split it likewise to derive two strategies σ_1 and σ_2 for $\Gamma_{\mathcal{P}_1}(\leftarrow \mathbf{G})$ and $\Gamma_{\mathcal{P}_2}(\leftarrow \mathbf{G})$ respectively (Theorem 10C). Since σ is winning, σ_1 and σ_2 are also winning (Corollary 10.31), and so by the

⁴The reader may contrast this with the LPN game, which is directly compatible with the LP game when no negations are present: the additional rule of LPN (rôle-switch) can only be used by the believer, and only as an answer to a negative doubt. This cannot happen in an LP program.

induction hypothesis we know that G must be true in every minimal model of \mathcal{P}_1 and in every minimal model of \mathcal{P}_2 . In other words, G is true in the union $\text{MM}(\mathcal{P}_1) \cup \text{MM}(\mathcal{P}_2)$; so by Lemma 6.2, it is true in every element of $\text{MM}(\mathcal{P})$. ■

Theorem 10E (Completeness of the finite, clean DLPG game semantics). *Let \mathcal{P} be a finite, clean DLP program, and $\leftarrow \mathbf{G}$ a goal. If G is true in every minimal model of \mathcal{P} , then there is a winning strategy in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$.*

PROOF is again by induction on $N(\mathcal{P})$:

BASE. Since $\text{LHM}(\mathcal{P}) \models G$, there is at least one $g \in G$ with $g \in \text{LHM}(\mathcal{P})$. Using the completeness of the LP game semantics (Corollary 8.1), we obtain a winning strategy σ in the LP game for \mathcal{P} with the goal $\leftarrow \mathbf{g}$. Without any modification, we can consider this as a winning strategy in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{g})$. It remains to correct all first (doubter) moves by including all extra doubts from \mathbf{G} in them. In this way, we have a winning strategy σ' in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$.

INDUCTION STEP. Again, pick a proper DLP rule $\phi \in \mathcal{P}$, and split \mathcal{P} to get $(\mathcal{P}_1, \mathcal{P}_2)$. We know that G is true in every minimal model of \mathcal{P} . Therefore, it is also true in every model of \mathcal{P} (because a non-minimal model can only make more formulæ true than a minimal one, not less). Using Lemma 6.2 again, $\text{MM}(\mathcal{P}_1) \cup \text{MM}(\mathcal{P}_2) \subseteq \text{HM}(\mathcal{P})$, so that G is true in every minimal model of \mathcal{P}_1 and in every minimal model of \mathcal{P}_2 . By the induction hypothesis, there are two winning strategies σ_1 and σ_2 in the games $\Gamma_{\mathcal{P}_1}(\leftarrow \mathbf{G})$ and $\Gamma_{\mathcal{P}_2}(\leftarrow \mathbf{G})$ respectively. Using Theorem 10B we can combine them to get a new strategy $\sigma_1 \vee \sigma_2$ for $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$, which is winning by Corollary 10.28. ■

In order to generalize these results to general DLP programs we need the following key lemma which connects games on a general DLP program \mathcal{P} with its clean version, $\widehat{\mathcal{P}}$.

Lemma 10.32. *Let \mathcal{P} be a general DLP program, and $\leftarrow \mathbf{G}$ a DLP goal. Then, there exists a winning strategy in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$ iff there is a winning strategy in $\Gamma_{\widehat{\mathcal{P}}}(\leftarrow \mathbf{G})$.*

Proof. “ \Rightarrow ”: We are given a winning strategy σ for $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$. To win in $\Gamma_{\widehat{\mathcal{P}}}(\leftarrow \mathbf{G})$, we move as follows: suppose that the believer following σ plays

$$\beta := \langle \psi_1, \dots, \psi_n \rangle.$$

Now, we might not be able to play the same move, because some of the ψ_i 's may be unclean, and therefore not available in $\widehat{\mathcal{P}}$. But, since each $\rho \in \mathcal{P}$ gives rise to a sequence of clean rules $\widehat{\rho}$ (actually a set, but we can fix an ordering and get a sequence out of it), we play:

$$\widehat{\beta} := \widehat{\psi}_1 \# \dots \# \widehat{\psi}_n.$$

This describes a winning strategy $\widehat{\sigma}$ in $\Gamma_{\widehat{\mathcal{P}}}(\leftarrow \mathbf{G})$. We must verify two things:

(i) that $\widehat{\beta}$ is indeed a valid move, and (ii) that we know how to win from any next doubter move $\widehat{\delta}$. (i) is trivial: it is only the heads that affect the validity of believer moves. Regarding (ii), we make the following claim: *for every $\psi_i \in \beta$ and for any doubts $\widehat{\delta}_i$ on $\widehat{\psi}_i$, there is a disjunction $D^i \in \text{body}(\psi_i)$ such that*

$D^i \subseteq [\widehat{\delta}_i]$. Then, our move in $\Gamma_{\widehat{\mathcal{P}}}(\leftarrow \mathbf{G})$ against a doubter move consisting of such doubts

$$\widehat{\delta} := \widehat{\delta}_1 \# \cdots \# \widehat{\delta}_n$$

is σ 's move for the doubts $\delta := \langle D^1, \dots, D^n \rangle$, valid thanks to the claim and the trivial observation that

$$[\widehat{\delta}] = [\widehat{\delta}_1 \# \cdots \# \widehat{\delta}_n] = [\widehat{\delta}_1] \cup \cdots \cup [\widehat{\delta}_n].$$

Proof of the claim. Assume otherwise: there is a rule $\psi_i \in \beta$ of the form

$$\psi_i := \mathbf{E}_i \leftarrow D_1^i, \dots, D_{k_i}^i,$$

and a $\widehat{\delta}_i$ on $\widehat{\psi}_i$ such that for all $D_j^i \in \mathbf{body}(\psi_i)$, $D_j^i \not\subseteq [\widehat{\delta}_i]$, i.e., there is a $d_j^i \in D_j^i$ with $d_j^i \notin [\widehat{\delta}_i]$. But this implies that $\widehat{\delta}$ did not doubt anything from the body of the rule

$$\mathbf{E}_i \leftarrow \mathbf{d}_1^i, \dots, \mathbf{d}_{k_i}^i \in \widehat{\psi}_i,$$

which is impossible.

“ \Leftarrow ”: In this direction we are given a strategy σ in $\Gamma_{\widehat{\mathcal{P}}}(\leftarrow \mathbf{G})$. Again, to win in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$, we follow σ for as long as it does not instruct us to include transformed rules, i.e., rules that do not exist in \mathcal{P} . Suppose now that $\beta := \langle \psi_1, \dots, \psi_n \rangle$ is the move that σ would play, where some of the ψ_i 's are transformed. We then play

$$\beta^\vee := \langle \psi_1^\vee, \dots, \psi_n^\vee \rangle,$$

where $\psi_i^\vee :=$ the first $\rho \in \mathcal{P}$ such that $\psi_i \in \widehat{\rho}$.

And this describes a winning strategy σ^\vee in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$. We must verify the same claims (i)–(ii) as above. (i) is trivial for the same reason. For (ii), let $\delta^\vee := \langle D_1, \dots, D_n \rangle$ be such a doubter response to β^\vee so that $D_i \in \mathbf{body}(\psi_i^\vee)$. This translates easily to a doubt δ on β , namely

$$\delta := \langle \{d_1\}, \dots, \{d_n\} \rangle,$$

where $d_i \in D_i$ and $\{d_i\} \in \mathbf{body}(\psi_i)$, so that

$$[\delta] = \{d_1, \dots, d_n\} \subseteq D_1 \cup \cdots \cup D_n = [\delta^\vee].$$

This allows us to copycat the move that σ would play in $\Gamma_{\widehat{\mathcal{P}}}(\leftarrow \mathbf{G})$ against δ . We go on in this manner until the winning strategy σ plays a fact ψ among its rules; then ψ^\vee will also be a fact. \blacksquare

Theorem 10F (Soundness and completeness for finite DLP). *Let \mathcal{P} be a finite, general DLP program, and $\leftarrow \mathbf{G}$ a DLP goal. Then, there exists a winning strategy in $\Gamma_{\mathcal{P}}(\leftarrow \mathbf{G})$ iff G is true in every minimal model of \mathcal{P} .*

Proof. We have proved everything we need:

$$\begin{array}{lll} \text{there is a winning} & \iff & \text{there is a winning} \\ \text{strategy in } \Gamma_{\mathcal{P}}(\leftarrow \mathbf{G}) & & \text{strategy in } \Gamma_{\widehat{\mathcal{P}}}(\leftarrow \mathbf{G}) \quad (\text{Lemma 10.32}) \\ & \iff & G \text{ is true on every} \\ & & \text{minimal model of } \widehat{\mathcal{P}} \quad (\text{Theorem 10D } \Rightarrow) \\ & & \quad (\text{Theorem 10E } \Leftarrow) \\ & \iff & G \text{ is true on every} \\ & & \text{minimal model of } \mathcal{P}. \quad (\text{Property 3.1}) \quad \blacksquare \end{array}$$

We now drop the heavy restriction of finiteness. This has the remarkable consequence that we can handle $\mathbf{ground}(\mathcal{P})$ for any first-order DLP program \mathcal{P} (see also Remark 4.2).

Theorem 10G (Soundness and completeness for general DLP). *Let \mathcal{P} be a general DLP program, and $\leftarrow G$ a disjunctive goal. Then, there exists a winning strategy in $\Gamma_{\mathcal{P}}(\leftarrow G)$ iff G is true in every minimal model of \mathcal{P} .*

Proof. Completeness. Assume G is true in every minimal model of \mathcal{P} , and therefore also in every model of \mathcal{P} , in symbols $\mathcal{P} \models G$. By the compactness theorem of propositional calculus, there exists a finite subset $\mathcal{P}_G \subseteq \mathcal{P}$ such that $\mathcal{P}_G \models G$. Specifically, G is true in every minimal model of \mathcal{P}_G , which allows us to use Theorem 10F to obtain a winning strategy σ_G in $\Gamma_{\mathcal{P}_G}(\leftarrow G)$. Observe now that this very strategy is still winning in the “bigger” game $\Gamma_{\mathcal{P}}(\leftarrow G)$, since the believer will never pick any of the additional rules, and therefore the course of the game cannot change.

Soundness. Suppose we have a winning strategy σ for $\Gamma_{\mathcal{P}}(\leftarrow G)$. This can only use a finite number of rules from \mathcal{P} , which form a finite subset $\mathcal{P}_\sigma \subseteq \mathcal{P}$. Obviously, σ is still winning in $\Gamma_{\mathcal{P}_\sigma}(\leftarrow G)$. This means that G is true in every minimal model of \mathcal{P}_σ (by Theorem 10F again), and therefore in every model of \mathcal{P}_σ . Since there are no negations involved, G remains true in every model of the superset \mathcal{P} ; specifically in every minimal one. \blacksquare

Corollary 10.33 (Soundness and completeness for first-order DLP). *Let \mathcal{P} be a first-order DLP program, and C a positive ground clause. Then $\mathcal{P} \models C$ iff there exists a winning strategy in $\Gamma_{\mathbf{ground}(\mathcal{P})}(\leftarrow C)$.*

Proof.

$$\begin{aligned} \mathcal{P} \models C &\iff C \text{ is true in every m.m. of } \mathcal{P} && \text{(Theorem 4B)} \\ &\iff C \text{ is true in every m.m. of } \mathbf{ground}(\mathcal{P}) && \text{(Property 4.2)} \\ &\iff \text{there is a winning } \sigma \text{ in } \Gamma_{\mathbf{ground}(\mathcal{P})}(\leftarrow C) && \text{(Theorem 10G)} \quad \blacksquare \end{aligned}$$

So far, we have studied model-theoretic semantics for all four languages, while game semantics for all but DLPN. In the next chapter we will apply the tools we have developed so far to fill this gap, and also obtain a new, different game for DLP, which still provides equivalent semantics.

Applications of $(-)^{\vee}$ on game semantics

In this chapter, we use the $(-)^{\vee}$ operator on the game semantics of LP and LPN of Chapters 8 and 9 to obtain new, game-theoretic semantics for DLP and DLPN. For DLP, the obtained game semantics is drastically different, albeit equivalent to DLPG, and for DLPN, this appears to be its first game-theoretic semantics. For the same reasons of simplicity that we outlined in the beginning of Chapter 7, *we again assume that all programs are clean in this chapter*. And as there has been no formal definition of a game semantics for *infinite* LPN programs, *programs with negation are also assumed to be finite*.

11.1 A different game semantics for DLP

We pick the simplest of the games we have encountered, LPG, and apply the $(-)^{\vee}$ operator on it. Following the definition, we have:

$$\begin{aligned}\mathcal{V}_{\text{LPG}_{\vee}} &= \mathcal{V}_{\text{LPG}} = \mathbb{B} \\ \mathcal{M}_{\text{LPG}_{\vee}} &= \mathcal{P}(\mathcal{M}_{\text{LPG}}).\end{aligned}$$

We proceed following the definitions:

$$\begin{aligned}\mathbf{m}_{\text{LPG}_{\vee}}(\mathcal{P}) &= (\mathbf{m}_{\text{LPG}})^{\vee}(\mathcal{P}) = \mathbf{m}_{\text{LPG}}(\mathbf{D}(\mathcal{P})), \\ \mathbf{a}_{\text{LPG}_{\vee}}(\mathcal{S})(Q) &= (\mathbf{a}_{\text{LPG}})^{\vee}(\mathcal{S})(Q) = \bigwedge_{S \in \mathcal{S}} \bigvee_{q \in Q} \mathbf{a}_{\text{LPG}}(S)(q), \\ \mathbf{s}_{\text{LPG}_{\vee}}(\mathcal{D})(Q) &= (\mathbf{s}_{\text{LPG}})^{\vee}(\mathcal{D})(Q) = \bigwedge_{\mathcal{P} \in \mathbf{D}(\mathcal{D})} \bigvee_{q \in Q} \mathbf{s}_{\text{LPG}}(\mathcal{P})(q).\end{aligned}$$

Translating the last equation, we see that

$$\begin{aligned}\mathbf{s}_{\text{LPG}_{\vee}}(\mathcal{D})(Q) = \mathbf{T} &\iff \bigwedge_{\mathcal{P} \in \mathbf{D}(\mathcal{D})} \bigvee_{q \in Q} \mathbf{s}_{\text{LPG}}(\mathcal{P})(q) = \mathbf{T} \\ &\iff \text{for every } \mathcal{P} \in \mathbf{D}(\mathcal{D}), \text{ there exists a } q \in Q, \\ &\quad \text{such that } \mathbf{s}_{\text{LPG}}(\mathcal{P})(q) = \mathbf{T} \\ &\iff \text{for every } \mathcal{P} \in \mathbf{D}(\mathcal{D}), \text{ there exists a } q \in Q, \\ &\quad \text{such that there is a winning strategy for } \Gamma_{\mathcal{P}}^{\text{LP}}(\leftarrow q).\end{aligned}$$

Easily, we can interpret these quantifiers as player moves, so that in the LPG_{\vee} game, Doubter begins by playing a definite instantiation $\mathcal{P} \in \text{D}(\mathcal{D})$ (essentially what she chooses is a \mathcal{D} -section). The following move of the Believer is to choose an element of the goal $q \in Q$, and after this point, the players begin playing the game $\Gamma_{\mathcal{P}}^{\text{LP}}(\leftarrow q)$.

Theorem 11A. *The game semantics LPG_{\vee} and the minimal model semantics MM are equivalent.*

Proof. We begin with the equivalence

$$\mathbf{sLPG} = \mathbf{sLHM}, \quad (\text{by Corollary 8.1})$$

on whose both sides we apply the $(-)^{\vee}$ operator and compute:

$$\begin{aligned} (\mathbf{sLPG})^{\vee} &= (\mathbf{sLHM})^{\vee} && (\text{by Lemma 7.1}) \\ &= \mathbf{sMM} && (\text{by Theorem 7B}) \end{aligned} \quad \blacksquare$$

Using Theorem 10G, we obtain the following

Corollary 11.1. *The game semantics LPG_{\vee} and DLPG are equivalent.*

11.2 An encoding of DLP to LP

Suppose now that we are given a *finite* DLP program \mathcal{D} and a DLP goal $G := \{g_1, \dots, g_m\}$. We outline here how we can encode both objects into ones of LP. We will define an operator

$$\text{encode} : \mathbf{P}_{\text{DLP}} \times \mathbf{Q}_{\text{DLP}} \rightarrow \mathbf{P}_{\text{LP}} \times \mathbf{Q}_{\text{LP}},$$

so that if $\text{encode}(\mathcal{D}, G) = (\mathcal{P}, g)$, we will then be able to use the game $\Gamma_{\mathcal{P}}^{\text{LP}}(\leftarrow g)$ to obtain an answer for the initial DLP query G , with respect to the initial program \mathcal{D} .

Before diving into the definition of the encoding, note that since \mathcal{D} is finite, then so is $\text{D}(\mathcal{D}) := \{\mathcal{P}_1, \dots, \mathcal{P}_n\}$. Now define

$$\text{encode}(\mathcal{D}, G) \triangleq (\mathcal{P}, g)$$

where

$$\begin{aligned} \mathcal{P} &:= \mathcal{P}_1 \uplus \dots \uplus \mathcal{P}_n \cup \text{restrictors}(\mathcal{D}, G) \cup \{\text{definitizer}(\mathcal{D}, G)\}, \\ \text{restrictors}(\mathcal{D}, G) &\triangleq \{\mathbf{p}_i \leftarrow \mathbf{g}_j^i \mid 1 \leq i \leq n, 1 \leq j \leq m\}, \\ \text{definitizer}(\mathcal{D}, G) &\triangleq \mathbf{g} \leftarrow \mathbf{p}_1, \dots, \mathbf{p}_n, \end{aligned}$$

and where all atoms p_i and g , are distinct and fresh, i.e., not appearing anywhere in \mathcal{D} , G , or any of the \mathcal{P}_i 's, while every occurrence of the atom g_j in the program \mathcal{P}_i , gives rise to an occurrence of the “tagged” atom g_j^i inside the disjoint union $\uplus \text{D}(\mathcal{D})$.

Some explanations are due. Firstly, we take the disjoint union of the definite instantiations of \mathcal{D} , tagging each atom that appears in \mathcal{P}_i with an i superscript, as we have already seen for the g_j 's. Next, the rôle of the definitizer rule, is to select a specific definite instantiation $\mathcal{P}_i \in \text{D}(\mathcal{D})$. On the other hand, for

each instantiation, there are corresponding restrictors whose job is to restrict the disjunctive goal into one of its disjuncts, a single atom.

Examining the game $\Gamma_{\mathcal{P}}^{\text{L}^{\text{P}}}(\leftarrow \mathbf{g})$ will shed some light to this: The Opponent begins by doubting \mathbf{g} . Next, the Player is forced to play the only rule whose head is g :

$$\mathbf{g} \leftarrow \mathbf{p}_1, \dots, \mathbf{p}_n.$$

Opponent now has to select a \mathbf{p}_i . This corresponds to her choice of a definite instantiation of \mathcal{D} . As we have seen, Player next gets to decide which element $g_j \in G$ he wants to restrict to, and this is exactly what the restrictors are for: Player will choose the rule $\mathbf{p}_i \leftarrow \mathbf{g}_j^i$. Opponent has no choice but to doubt the only conjunct in the body of that rule, \mathbf{g}_j^i , and finally we have arrived in this tagged atom, and so the game is successfully restricted within the rules of the correspondingly tagged \mathcal{P}_i .

Notice that in essence, the game just described is really the one of Section 11.1, but without having to add those artificial rules that we did there: it is the encoding that does their job instead.

11.3 A first game semantics for DLPN

As we have already mentioned, there appears to be no game semantics for DLPN in the literature. In this section, we ameliorate this situation: we obtain two game semantics by using the $(-)^{\vee}$ operator on LPNG and LPNG $^{\omega}$.

According to its definition,

$$\begin{aligned} \mathcal{V}_{\text{LPNG}_{\vee}} &= \mathcal{V}_{\text{LPNG}} = \mathbb{V}_1, & \mathcal{V}_{\text{LPNG}_{\vee}^{\omega}} &= \mathcal{V}_{\text{LPNG}^{\omega}} = \mathbb{V}_{\omega}, \\ \mathcal{M}_{\text{LPNG}_{\vee}} &= \mathcal{P}(\mathcal{M}_{\text{LPNG}}); & \mathcal{M}_{\text{LPNG}_{\vee}^{\omega}} &= \mathcal{P}(\mathcal{M}_{\text{LPNG}^{\omega}}). \end{aligned}$$

Focusing on LPNG $^{\omega}$, we have

$$\begin{aligned} \mathbf{m}_{\text{LPNG}_{\vee}^{\omega}}(\mathcal{P}) &= (\mathbf{m}_{\text{LPNG}^{\omega}})^{\vee}(\mathcal{P}) = \mathbf{m}_{\text{LPNG}^{\omega}}(\mathbf{D}(\mathcal{P})), \\ \mathbf{a}_{\text{LPNG}_{\vee}^{\omega}}(\mathcal{S})(Q) &= (\mathbf{a}_{\text{LPNG}^{\omega}})^{\vee}(\mathcal{S})(Q) = \bigwedge_{S \in \mathcal{S}} \bigvee_{q \in Q} \mathbf{a}_{\text{LPNG}^{\omega}}(S)(q), \\ \mathbf{s}_{\text{LPNG}_{\vee}^{\omega}}(\mathcal{D})(Q) &= (\mathbf{s}_{\text{LPNG}^{\omega}})^{\vee}(\mathcal{D})(Q) = \bigwedge_{\mathcal{P} \in \mathbf{D}(\mathcal{D})} \bigvee_{q \in Q} \mathbf{s}_{\text{LPNG}^{\omega}}(\mathcal{P})(q); \end{aligned}$$

and similarly for LPNG $_{\vee}$.

Again, interpreting these in terms of game rules is straightforward: Opponent begins by playing a definite instantiation $\mathcal{P} \in \mathbf{D}(\mathcal{D})$, Player then chooses an element of the goal $q \in Q$, and after this point, the players begin playing the game $\Gamma_{\mathcal{P}}^{\text{L}^{\text{PN}}}(\leftarrow q)$ normally, and the outcome of their play in it becomes the outcome of the play on the new game.

Theorem 11B. *For finite DLPN programs, the game semantics LPNG $_{\vee}^{\omega}$ and the model-theoretic semantics MM $^{\omega}$ are equivalent.*

Proof. Completely analogously to the proof of Theorem 11A, starting from the equivalence

$$\mathbf{s}_{\text{LPNG}^{\omega}} = \mathbf{s}_{\text{WF}^{\omega}}, \quad (\text{by Theorem 9A})$$

we apply the $(-)^{\vee}$ operator on both sides, and compute:

$$\begin{aligned} (\mathbf{s}_{\text{LPNG}^\omega})^{\vee} &= (\mathbf{s}_{\text{WF}^\omega})^{\vee} && \text{(by Lemma 7.1)} \\ &= \mathbf{s}_{\text{MM}^\omega} && \text{(by Theorem 7C).} \quad \blacksquare \end{aligned}$$

Since we can collapse any infinite-valued space into the three-valued \mathbb{V}_1 , we have chosen to present only the more general, infinite-valued semantics. But if for any reason we want to restrict ourselves to only using the three-valued space \mathbb{V}_1 , we can easily obtain the analogous results.

✂ REMARK 11.1. The encoding we presented in Section 11.2 can be applied *mutatis mutandis* for the case of DLPN programs, encoding them into LPN ones. Note that since negations only occur inside the bodies of the \mathcal{P}_i 's, the rôle-switching moves may only be played once the play has already been restricted to a specific definite instantiation \mathcal{P}_i .

ξ

Part IV
Appendices

Lattices and Heyting algebras

Everything in this appendix can be found in standard texts on such subjects, e.g., [DP02].

Definition A.1. Let L be a non-empty ordered set. L is a *lattice* iff both $x \vee y$ and $x \wedge y$ exist for all $x, y \in L$. L is a *complete lattice* iff $\bigvee S$ and $\bigwedge S$ exist for all $S \subseteq L$. We say that L has a *one* iff there is an element $1 \in L$ such that $x = x \wedge 1$ for all $x \in L$, and a *zero* iff there is an element $0 \in L$ such that $x = x \vee 0$ for all $x \in L$. Other notations for these two elements are \top and \perp respectively. If such two elements exist in a lattice, we call it *bounded*.

Definition A.2. Given a lattice $(L; \leq_L)$, we form its *dual lattice* $(L^\partial; \leq_{L^\partial})$ by defining

$$x \leq_{L^\partial} y \iff y \leq_L x.$$

Definition A.3. A lattice L is *distributive* iff for all $a, b, c \in L$,

$$a \wedge (b \vee c) = (a \wedge b) \vee (a \wedge c).$$

Definition A.4. A complete lattice L is said to be *completely distributive* iff for any doubly indexed family $\{x_{j,k}\}_{j \in J, k \in K_j} \subseteq L$,

$$\bigwedge_{j \in J} \bigvee_{k \in K_j} x_{j,k} = \bigvee_{f \in F} \bigwedge_{j \in J} x_{j,f(j)},$$

where F is the set of all choice functions f that choose for each $j \in J$, an element $f(j) \in K_j$. That is,

$$F = \left\{ f : J \rightarrow \bigcup_{j \in J} K_j \mid f(j) \in K_j \right\} = \prod_{j \in J} K_j.$$

Definition A.5. Let L be a complete lattice. $k \in L$ is *compact*, if for every $S \subseteq L$, there is a finite $T \subseteq S$, such that

$$k \leq \bigvee S \implies k \leq \bigvee T.$$

The set of all compact elements of L is denoted by $K(L)$.

Definition A.6. A complete lattice L is said to be *algebraic* if, for each $a \in L$,

$$a = \bigvee \{k \in K(L) \mid k \leq a\}.$$

Theorem A.7. Let L be a lattice. Then the following are equivalent:

- (i) L is distributive and both L and L^∂ are algebraic;
- (ii) L is completely distributive and algebraic.

PROOF is in [DP02, Theorem 10.29]. ■

Definition A.8. Let H be a (complete) bounded lattice. H is a (complete) *Heyting algebra* iff for every $a, b \in H$ there is an element $a \rightarrow b$ satisfying

$$c \leq a \rightarrow b \iff c \wedge a \leq b$$

for every $c \in H$.

Property A.9 (Some equalities). Let H be a Heyting algebra. Then the following equalities hold between elements of H :

- (i) $a \wedge (b \vee c) = (a \wedge b) \vee (a \wedge c)$
- (ii) $a \vee (b \wedge c) = (a \vee b) \wedge (a \vee c)$
- (iii) $a \Rightarrow (b \vee c) = (a \Rightarrow b) \vee (a \Rightarrow c)$
- (iv) $(a \vee b) \Rightarrow c = (a \Rightarrow c) \wedge (b \Rightarrow c)$.

If in addition H is complete, the following equalities also hold:

- (I) $a \wedge \bigvee_{s \in S} s = \bigvee_{s \in S} (a \wedge s)$
- (II) $a \vee \bigwedge_{s \in S} s = \bigwedge_{s \in S} (a \vee s)$
- (III) $a \Rightarrow \bigvee_{s \in S} s = \bigvee_{s \in S} (a \Rightarrow s)$
- (IV) $(\bigvee_{s \in S} s) \Rightarrow c = \bigwedge_{s \in S} (s \Rightarrow c)$.

References

- [AB94] Krzysztof R. Apt and Roland Bol. Logic programming and negation: A survey. *Journal of Logic Programming*, 19:9–71, 1994.
- [Acz77] Peter Aczel. An introduction to inductive definitions. In Jon Barwise, editor, *Handbook of Mathematical Logic*, pages 739–782. North-Holland, Amsterdam, 1977.
- [AJM00] Samson Abramsky, Radha Jagadeesan, and Pasquale Malacaria. Full abstraction for PCF. *Information and Computation*, 163(2):409–470, December 2000.
- [ALM09] Gianluca Amato, James Lipton, and Robert McGrail. On the algebraic structure of declarative programming languages. *Theoretical Computer Science*, 410(46):4626–4671, 2009.
- [AM99] S. Abramsky and G. McCusker. Game semantics. In H. Schwichtenberg and U. Berger, editors, *Computational Logic: Proceedings of the 1997 Marktoberdorf Summer School*, pages 1–56. Springer-Verlag, 1999.
- [And92] Jean-Marc Andreoli. Logic programming with focusing proofs in linear logic. *Journal of Logic and Computation*, 2:297–347, 1992.
- [AP91] Jean-Marc Andreoli and Remo Pareschi. Logic programming with sequent systems. In Peter Schroeder-Heister, editor, *Extensions of Logic Programming*, volume 475 of *Lecture Notes in Computer Science*, pages 1–30. Springer Berlin Heidelberg, 1991.
- [Apt90] Krzysztof R. Apt. Logic programming. In *Handbook of theoretical computer science (vol. B): formal models and semantics*, pages 493–574. MIT Press, Cambridge, MA, USA, 1990.
- [BM09] Filippo Bonchi and Ugo Montanari. Reactive systems, (semi-)saturated semantics and coalgebras on presheaves. *Theoretical Computer Science*, 410(41):4044–4066, 2009.
- [BZ13] Filippo Bonchi and Fabio Zanasi. Saturated semantics for coalgebraic logic programming. In Reiko Heckel and Stefan Milius, editors, *Algebra and Coalgebra in Computer Science*, volume 8089

- of *Lecture Notes in Computer Science*, pages 80–94. Springer Berlin Heidelberg, 2013.
- [Cla78] Keith Clark. Negation as failure. *Logic and Databases*, pages 293–322, 1978.
- [Cla79] Keith Clark. Predicate logic as a computational formalism. 1979.
- [CM92] Andrea Corradini and Ugo Montanari. An algebraic semantics for structured transition systems and its applications to logic programs. *Theoretical Computer Science*, 103(1):51–106, 1992.
- [CPRW07] Pedro Cabalar, David Pearce, Panos Rondogiannis, and William Wadge. A purely model-theoretic semantics for disjunctive logic programs with negation. In Chitta Baral, Gerhard Brewka, and John Schlipf, editors, *Logic Programming and Nonmonotonic Reasoning*, volume 4483 of *Lecture Notes in Computer Science*, pages 44–57. Springer Berlin / Heidelberg, 2007.
- [DCLN98] Roberto Di Cosmo, Jean-Vincent Loddo, and Stephane Nicolet. A game semantics foundation for logic programming. In Catuscia Palamidessi, Hugh Glaser, and Karl Meinke, editors, *Principles of Declarative Programming*, volume 1490 of *Lecture Notes in Computer Science*, pages 355–373. Springer Berlin / Heidelberg, 1998. 10.1007/BFb0056626.
- [DP02] Brian A. Davey and Hilary A. Priestley. *Introduction to Lattices and Order (2. ed.)*. Cambridge University Press, 2002.
- [Fit85] Melvin Fitting. A kripke-kleene semantics for logic programs. *Journal of Logic Programming*, 2(4):295–312, 1985.
- [Fit99] Melvin Fitting. Fixpoint semantics for logic programming—a survey. *Theoretical Computer Science*, 278:25–51, 1999.
- [Gel08] Michael Gelfond. Answer sets. In F. van Harmelen, V. Lifschitz, and B. Porter, editors, *Handbook of Knowledge Representation*, chapter 7. Elsevier, 2008.
- [GL88] Michael Gelfond and Vladimir Lifschitz. The stable model semantics for logic programming. pages 1070–1080. MIT Press, 1988.
- [GRW08] Chrysidia Galanaki, Panos Rondogiannis, and William W. Wadge. An infinite-game semantics for well-founded negation in logic programming. *Annals of Pure and Applied Logic*, 151(2-3):70–88, 2008. First Games for Logic and Programming Languages Workshop.
- [HO00] J. M. E. Hyland and C.-H. Ong. On full abstraction for PCF: I. Models, observables and the full abstraction problem; II. Dialogue games and innocent strategies; III. A fully abstract and universal game model. *Information and Computation*, 163:285–408, 2000.
- [Hod09] Wilfrid Hodges. Logic and games. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Spring 2009 edition, 2009.

- [KMP10] Ekaterina Komendantskaya, Guy McCusker, and John Power. Coalgebraic semantics for parallel derivation strategies in logic programming. In Michael Johnson and Dusko Pavlovic, editors, *AMAST*, volume 6486 of *Lecture Notes in Computer Science*, pages 111–127. Springer, 2010.
- [KNR11] Vassilis Kountouriotis, Christos Nomikos, and Panos Rondogiannis. A game-theoretic characterization of boolean grammars. *Theoretical Computer Science*, 412(12-14):1169–1183, 2011.
- [KP96] Yoshiki Kinoshita and A. John Power. A fibrational semantics for logic programs. In Roy Dyckhoff, Heinrich Herre, and Peter Schroeder-Heister, editors, *ELP*, volume 1050 of *Lecture Notes in Computer Science*, pages 177–191. Springer, 1996.
- [KP11] Ekaterina Komendantskaya and John Power. Coalgebraic semantics for derivations in logic programming. In *Proceedings of the 4th international conference on Algebra and coalgebra in computer science*, CALCO’11, pages 268–282, Berlin, Heidelberg, 2011. Springer-Verlag.
- [KPS13] Ekaterina Komendantskaya, John Power, and Martin Schmidt. Coalgebraic logic programming: from semantics to implementation. *CoRR*, abs/1312.6568, 2013.
- [Kun87] Kenneth Kunen. Negation in logic programming. *Journal of Logic Programming*, 4(4):289–308, 1987.
- [LC00] Jean-Vincent Loddó and Roberto Di Cosmo. Playing logic programs with the alpha-beta algorithm. In Michel Parigot and Andrei Voronkov, editors, *LPAR*, volume 1955 of *Lecture Notes in Computer Science*, pages 207–224. Springer, 2000.
- [Llo87] John Wylie Lloyd. *Foundations of Logic Programming*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1987.
- [LMR92] Jorge Lobo, Jack Minker, and Arcot Rajasekar. *Foundations of disjunctive logic programming*. MIT Press, Cambridge, MA, USA, 1992.
- [Lor61] Paul Lorenzen. Ein dialogisches konstruktivitätskriterium. In *Infinite Methods: Proceedings of the Symposium on Foundations of Mathematics, Warsaw, 2–9 September 1959*, pages 193–200. Pergamon Press, 1961.
- [LT84] J. W. Lloyd and R. W. Topor. Making Prolog more expressive. *Journal of Logic Programming*, 3:225–240, 1984.
- [Lüd11] Rainer Lüdecke. Every formula-based logic program has a least infinite-valued model. In Hans Tompits, Salvador Abreu, Johannes Oetsch, Jörg Pührer, Dietmar Seipel, Masanobu Umeda, and Armin Wolf, editors, *INAP/WLP*, volume 7773 of *Lecture Notes in Computer Science*, pages 155–172. Springer, 2011.

- [Min82] Jack Minker. *On indefinite databases and the closed world assumption*, volume 138 of *Lecture Notes in Computer Science*. Springer-Verlag, 1982.
- [MN12] Dale Miller and Gopalan Nadathur. *Programming with Higher-Order Logic*. Cambridge University Press, June 2012.
- [MNPS91] Dale Miller, Gopalan Nadathur, Frank Pfenning, and Andre Scedrov. Uniform proofs as a foundation for logic programming. *Annals of Pure and Applied Logic*, 51:125–157, 1991.
- [MS06] Dale Miller and Alexis Saurin. A game semantics for proof search: Preliminary results. *Electronic Notes in Theoretical Computer Science*, 155:543–563, 2006.
- [Nic94] Hanno Nickau. Hereditarily sequential functionals. In Anil Nerode and Yuri Matiyasevich, editors, *LFCS*, volume 813 of *Lecture Notes in Computer Science*, pages 253–264. Springer, 1994.
- [NR12] Christos Nomikos and Panos Rondogiannis. A game semantics for intensional logic programming. Presented in Games for Logic and Programming Languages (GaLoP) VII, Dubrovnik, Croatia, 2012.
- [NV07] Jonty Needham and Marina De Vos. A games semantics of ASP. In Véronica Dahl and Ilkka Niemelä, editors, *Logic Programming*, volume 4670 of *Lecture Notes in Computer Science*, pages 460–461. Springer Berlin Heidelberg, 2007.
- [PR05] David Pym and Eike Ritter. A games semantics for reductive logic and proof-search. In Dan R. Ghica and Guy McCusker, editors, *Games for Logic and Programming Languages (GALOP 2005)*, *University of Edinburgh, 2–3 April 2005*, pages 107–123, 2005.
- [Rei78] R Reiter. On closed world data bases. *Logic and Databases*, pages 55–76, 1978.
- [RW05] Panos Rondogiannis and William W. Wadge. Minimum model semantics for logic programs with negation-as-failure. *ACM Trans. Comput. Logic*, 6(2):441–467, 2005.
- [Tso10] Thanos Tsouanas. Semantic approaches to logic programming. Master’s thesis, 2010.
- [Tso13] Thanos Tsouanas. A game semantics for disjunctive logic programming. *Annals of Pure and Applied Logic*, 164(11):1144–1175, 2013.
- [Vää11] Jouko Väänänen. *Models and Games*. Cambridge series in advanced mathematics. Cambridge University Press, 2011.
- [vE86] M. H. van Emden. Quantitative deduction and its fixpoint theory. *Journal of Logic Programming*, 3:37–53, 1986.
- [vEK76] M. H. van Emden and R. A. Kowalski. The semantics of predicate logic as a programming language. *Journal of the ACM*, 23:569–574, 1976.

-
- [VGRS91] Allen Van Gelder, Kenneth A. Ross, and John S. Schlipf. The well-founded semantics for general logic programs. *J. ACM*, 38(3):619–649, 1991.

Index of symbols

<p> $- \Upsilon -$ $D \Upsilon E$, 40 $\mathcal{D} \Upsilon \mathcal{E}$, 40 $\phi_1 \Upsilon \phi_2$, 40 $\sigma_1 \Upsilon \sigma_2$, 76 $\tau_1 \Upsilon \tau_2$, 74 $\tau_1 \ddot{\Upsilon} \tau_2$, 71 </p> <p> $- \sqsubset -$ $\mathcal{P} _A^\phi$, 37 $\mathcal{P} _{\mathcal{H}^c}^\phi$, 38 $\phi _A$, 37 $\pi _{H_q}^\phi$, 75 $\pi _{\mathcal{H}^c}^\phi$, 75 $\sigma _{H_q}^\phi$, 81 $\sigma _{\mathcal{H}^c}^\phi$, 82 </p> <p> $\mathcal{A}t$, 68 $N(\mathcal{P})$, 83 T_P, 4 $\Gamma_{\mathcal{P}}(\leftarrow G)$, 63 \bigvee, 67 \mathbf{F}, 7, 23 \mathbf{T}, 7, 23 $:=$, 7 $[-]$, 41, 63 $\hat{\mathcal{P}}$, 19 $\mathbf{C}(\pi)$, 76 $\#$, 7 $::$, 7 $\sqsubseteq_e, \sqsubset_e$, 7 \sqsubseteq, \sqsubset, 7 $s _n$, 7 β_i^{q*}, 69, 71, 80 $-$, 7 \equiv, 7, 68, 70 </p>	<p> T_q, 77 $\text{HB}(\mathcal{P})$, 21, 63 $\text{HM}(\mathcal{P})$, 22, 38, 39 $\text{LHM}(\mathcal{P})$, 22, 83 $\text{MM}(\mathcal{P})$, 23, 38, 39 $\text{answer}_\sigma(\pi)$, 67, 77–79 $\text{body}(\phi)$, 17 $\text{cr}(\pi)$, 77, 78 cr_{σ_q}, 77 $\text{ground}(\mathcal{P})$, 16 $\text{ground}(\rho)$, 16 $\text{head}(\phi)$, 17 $\text{proj}_{\sigma_q}(\pi)$, 77 $\text{rmstall}(\tau)$, 66 $\text{sync}(\tau_1, \tau_2)$, 72 $\bar{\beta}$, 65 $\bar{\delta}, \overline{\delta}$, 65 π^-, 68, 76–78 τ^-, 7 \triangleq, 7 ε, 64, 67 \mathcal{P}_+, 63 </p>
--	--

General index

- AJM games, 61
- alpha-beta pruning algorithm, 52
- answer
 - computed, **5**
 - correct, **5**
- answer set programming, 6
- ASP, 6

- benefit of the doubt, 54, 82
- body, **3**
- boolean grammar, 52

- chain, 67
- clause
 - Horn, 6
- closed world assumption, **54**
- CNF, 16
- coalgebraic semantics, **5**
- combination, 61
 - of synchronous plays, 70
- combo move, 83
- commutative diagram, 77, 79
- completeness, **61**, 83
- computed answer, **5**
- constraint logic programming, 52
- correct answer, **5**
- CWA, *see* closed world assumption

- declarative programming, 6
- declarative semantics, *see* denotational semantics
- decomposition, 76
- denotational semantics, **21**
- description logic, 54
- desugaring, 19
- diagram
 - commutative, 77, 79

- dialogue games, 5, 52
- disjunctive combination, *see* combination
- DLP, **4**
 - compatibility, 83
 - game
 - simplified, **61**
 - game semantics
 - completeness, 84
 - soundness, 83
- DLPN, **4**

- extensive operator, **68**

- first-order
 - language, 15, 28
 - program, 28
- fixpoint semantics, **4**, 29
- functional programming, 61

- game
 - LP, 61
 - LPN, 61
 - simplified DLP, **61**
- game semantics, **5**, **51**
 - DLPG, 82
- goal
 - success, 22, *see also* semantics
- grammar
 - boolean, 52
- ground
 - atom, **16**
 - formula, **16**
 - term, **16**

- head, **3**
- Herbrand

- least model, *see also* LHM(\mathcal{P})
 - model
 - least, 3, 51
- Horn clause, 6
- idempotent
 - operator, **68**
- inference rule
 - resolution, 4
 - SLD, 4
 - SLI, 4
- knowledge representation, 54
- language
 - first-order, *see* first-order language
- least Herbrand model, *see* Herbrand model, 55
- least upper bound, 67
- limit, 67
- linear logic, 6
- logic program
 - syntax, **13**
- logic programming, **3**
 - constraint, 52
- Lorenzen dialogue games, 5, 52
- LP, **4, 15**, 55
 - denotational semantics, *see also* least Herbrand model game, 55, 61
 - game semantics, **82**
 - completeness, 55
 - soundness, 55
- LPN, **4**
 - game, 61, 83
- lub, *see* least upper bound
- minimal model semantics, 51
- Minker semantics, *see* minimal model semantics
- mip, *see* model intersection property
- model
 - Herbrand, *see* Herbrand model
 - minimal, 51
 - well-founded, 51, 83
- model intersection property, **22**
- model-theoretic semantics, 4
- monotone
 - answer $_{\sigma}(\pi)$, 68, 78
 - rmstall(τ), 66, 78
 - sync(τ_1, τ_2), 73
 - $\tau_1 \Upsilon \tau_2$, 74
 - $\tau_1 \check{\Upsilon} \tau_2$, 72
 - operator, **68**
- move
 - combo, *see* combo move
- notation, 7, 65, 68
- open world assumption, **54**
- operational semantics, *see* procedural semantics
- operator
 - extensive, **68**
 - idempotent, **68**
 - monotone, **68**
- OWA, *see* open world assumption
- PCF, 61
- play, *see also* preservation combination
 - synchronous, **71**
 - in sync with, *see* synchronous
- preservation
 - of finiteness
 - $\sigma|_{H_q}^{\phi}$, 82
 - $\sigma_1 \Upsilon \sigma_2$, 80
 - of parity
 - sync(τ_1, τ_2), 73
 - $\tau_1 \Upsilon \tau_2$, 74, 76, 79
 - $\tau_1 \check{\Upsilon} \tau_2$, 72
 - of plays
 - $\tau_1 \Upsilon \tau_2$, 74
 - $\tau_1 \check{\Upsilon} \tau_2$, 72
 - of totality
 - $\sigma|_{H_q}^{\phi}$, 81
 - $\sigma_1 \Upsilon \sigma_2$, 80
 - of winning
 - $\sigma|_{H_q}^{\phi}$, 82, 83
 - $\sigma_1 \Upsilon \sigma_2$, 81, 84
- procedural semantics, **4**, 52
- programming
 - answer set, 6
 - declarative, 6
- proof search, 6
- quasiplay, **64, 76**
- release, *see also* β_i^{q*}

- resolution, 4
 - SLD, 4, 55
 - SLI, 4
- rule
 - body, **3**
 - head, **3**
- semantic web, 54
- semantics
 - coalgebraic, **5**
 - declarative, *see* denotational
 - denotational, 3, **21**
 - fixpoint, 4, 29
 - game, **5**, **51**
 - minimal model, *see* minimal model semantics
 - model-theoretic, **4**
 - operational, *see* procedural
 - procedural, 4, 52
 - stable model, 6
 - well-founded, *see* well-founded
- sequent calculus, 6
- SLD resolution, 4, 55, *see also* procedural semantics
- SLI resolution, 4
- soundness, **61**, 83
- splitting, 61
- stable model semantics, 6
- strategy, *see also* preservation
 - combo-free, 83
- success of a goal, *see* goal success
- synchronization, *see also* sync
- synchronous
 - combination, *see* combination of synchronous plays
- syntactic sugar, 19

- total strategy, *see also* preservation
- truth value space, **10**
- truth values, *see also* **T**, **F**

- well-founded model, *see* model

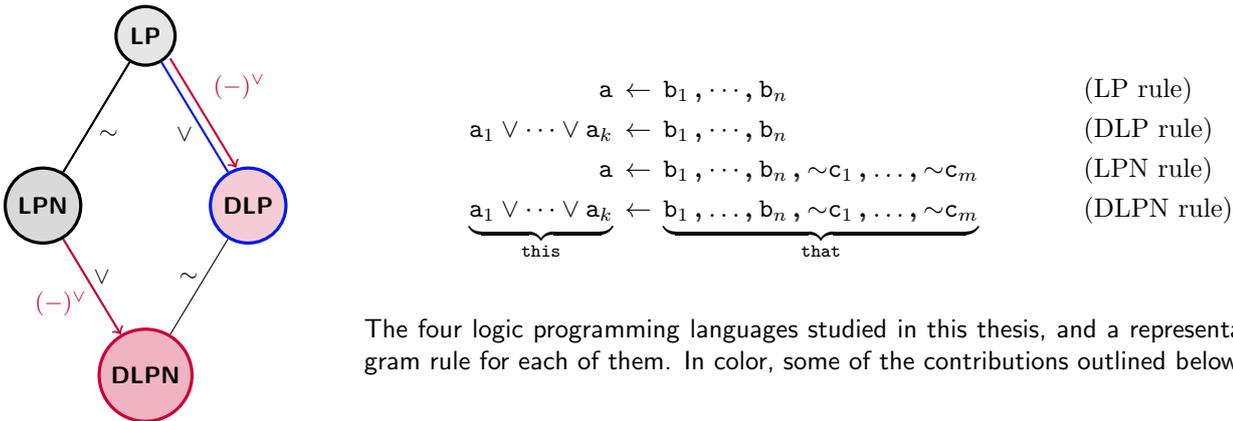
~
~
~
~
~
~

:wq

Programming paradigms. Programming languages come in different forms and flavors. Still, there are some major paradigms of programming, that most languages can be seen to belong to one or another, or to even combine features of more than one of them. Under this perspective, the two main families of languages are *imperative* and *declarative*. Roughly speaking, in imperative languages we program by specifying *how* to reach a solution, in a way that resembles how one would write down a cooking recipe. On the other hand, when programming in declarative languages we define *what* the problem is and *what* is a solution of it, but we leave the step-by-step details (the *how*) to the implementation of the language. Code in declarative languages tends to be more elegant and such languages are also much better suited for reasoning about programs and proving useful properties. The two main paradigms that fall under this family of languages are *functional programming* and *logic programming*. In the first, we program by mathematically defining and evaluating functions; in the latter, a program consists of logic formulæ of the form **this** \leftarrow **that**, meant to be read as “*this* holds if *that* holds”, or “to solve *this* problem, it suffices to solve *that* one”.

Syntax and semantics. A programming language is a two-sided coin, and to understand one and use it well, we need to know both of its sides: its *syntax* and its *semantics*. Syntax dictates what constitutes a well-formed program for a language: how expressions are formed, what symbols or words are allowed and where, etc. But what exactly is *the meaning* of a correctly written program? Answering this question is the job of the language’s semantics, which also come in different kinds, the three main ones being *denotational*, *operational*, and *axiomatic*. In denotational semantics one usually defines and uses appropriate *mathematical objects* to denote meanings of expressions and programs; without such a formalism we cannot hope to prove much about a program. For logic programs, the traditional denotational semantics are *model-theoretic*, but recently *game-theoretic* semantics have been studied as well. An operational semantics describes (either in a step-by-step fashion or in a so-called “big-step” style) the execution of programs and is thus closer to the implementation of a language. Finally, axiomatic semantics specify what effect the execution of statements and programs will have on a set of assertions that are assumed to hold before we perform it.

Four logic programming languages. *In this thesis, we study denotational semantics of four logic programming languages:* (1) **LP** which is the simplest case of logic programs; (2) **DLP** which extends LP by adding disjunctions (\vee); (3) **LPN** which extends LP by adding negations (\sim); and (4) **DLPN** which allows both disjunctions and negations. The following figure represents these extensions schematically:



The four logic programming languages studied in this thesis, and a representative program rule for each of them. In color, some of the contributions outlined below.

Contributions. The three main contributions of this dissertation can be summarized as follows:

- An *abstract framework for logic programming semantics* is defined, and all semantic approaches that we study are placed within this framework. In this process we define the general notion of a *truth value space* as an appropriate algebraic structure that satisfies a set of axioms. The booleans form the canonical example of such a space, but we need to consider much more general ones when dealing with negation. For this we define and study an infinite family of spaces, parametrized by an ordinal number.
- A game semantics for LP has been known for quite a long time. In 2005, it was extended to a new one for the language of LPN. Here a *novel game semantics for DLP programs* is developed in full detail and we prove that it is sound and complete with respect to the standard, minimal model semantics of Minker. What is more, this is done for infinite programs, which has the remarkable consequence that it provides a denotational semantics even for first-order programs. Even though the game itself can be seen as an extension of the LP game, the formalization and notation we have used is influenced by the games used in functional programming instead.
- We define a *semantic operator* $(-)^{\vee}$ which transforms any given semantics of a non-disjunctive logic programming language to a semantics of the “corresponding” disjunctive one. We exhibit the correctness of this transformation by proving that it preserves equivalences of semantics, and that the semantics it yields behave as one would expect them to. We present some *applications of this operator* to obtain some new semantics for disjunctive logic programming languages: notably, applying this method on the game semantics for LPN, we obtain a game semantics for DLPN, filling the gap that remained to complete the game semantics part of the picture sketched above.