

Forgetting auxiliary atoms in forks

Felicidad Aguado¹, Pedro Cabalar¹, Jorge Fandinno¹,
David Pearce², Gilberto Pérez¹ and Concepción Vidal¹

¹ Department of Computer Science
University of Corunna, SPAIN
{aguado,cabalar,jorge.fandino}@udc.es
{gilberto.pvega,concepcion.vidalm}@udc.es

² Universidad Politécnica de Madrid, SPAIN
david.pearce@upm.es

Abstract. In this work we tackle the problem of checking strong equivalence of logic programs that may contain local auxiliary atoms, to be removed from their stable models and to be forbidden in any external context. We call this property *projective strong equivalence* (PSE). It has been recently proved that not any logic program containing auxiliary atoms can be reformulated, under PSE, as another logic program or formula without them – this is known as *strongly persistent forgetting*. In this paper, we introduce a conservative extension of *Equilibrium Logic* and its monotonic basis, the logic of *Here-and-There*, in which we deal with a new connective ‘|’ we call *fork*. We provide a semantic characterisation of PSE for forks and use it to show that, in this extension, it is always possible to forget auxiliary atoms under strong persistence. We further define when the obtained fork is representable as a regular formula.

1 Introduction

Answer Set Programming (ASP [1]) has become a established problem-solving paradigm for Knowledge Representation and Reasoning (KRR). The reasons for this success derive from the practical point of view, with the availability of efficient solvers [2,3] and application domains [4], but also from its solid theoretical foundations, rooted in the *stable models* [5] semantics for normal logic programs that was later generalised to arbitrary propositional [6], first-order [7, 8] and infinitary [9] formulas. An important breakthrough that supported these extensions of ASP has been its logical characterisation in terms of *Equilibrium Logic* [6] and its monotonic basis, the intermediate logic of *Here-and-There* (HT). Despite its expressiveness, a recent result [10] has shown that Equilibrium Logic has limitations in capturing the representational power of auxiliary atoms, which cannot always be forgotten. To illustrate this point, take the following problem.

Example 1. Two individuals, mother and father, both carrying alleles a and b , procreate an offspring. We want to generate all the possible ways in which the offspring may inherit its parents’ genetic information. \square

According to Mendelian laws, we should obtain three possible combinations that, ignoring their frequency, correspond to the sets of alleles $\{a\}$, $\{b\}$ and $\{a, b\}$. The straightforward way to generate these three stable models is to encode an inclusive disjunction with the following 3 rules:

$$a \vee \neg a \quad b \vee \neg b \quad \perp \leftarrow \neg a \wedge \neg b \quad (P_1)$$

A drawback of this representation is that it does not differentiate the information coming from each parent, possibly becoming a problem of elaboration tolerance. For instance, if only the mother's information were available, one would expect to obtain the stable models $\{a\}$ and $\{b\}$ but not $\{a, b\}$, as there is no evidence of that combination without further information about the father. So, the mother alone would be better represented by a regular disjunction $a \vee b$. However, we cannot represent each parent as an independent disjunction like that, since $(a \vee b) \wedge (a \vee b)$ just amounts to $(a \vee b)$ and the combination $\{a, b\}$ is not obtained. A simple way to represent these two disjunctions separately is using auxiliary atoms to keep track of alleles from the mother ($ma \vee mb$) and the father ($fa \vee fb$). This leads to program P_2 :

$$\begin{array}{lll} ma \vee mb & a \leftarrow ma & b \leftarrow mb & (P_m) \\ fa \vee fb & a \leftarrow fa & b \leftarrow fb & (P_f) \end{array}$$

consisting of the mother's contribution P_m and the father's contribution P_f . Four stable models are obtained from P_2 , $\{ma, fa, a\}$, $\{mb, fb, b\}$, $\{ma, fb, a, b\}$ and $\{mb, fa, a, b\}$, but if we *project* them on the original vocabulary $V = \{a, b\}$ (i.e. we remove auxiliary atoms), they collapse to three $\{a\}$, $\{b\}$ and $\{a, b\}$ as expected. Note that, although auxiliary atoms in this example have a meaning in the real world (they represent the effective sources of each inherited allele) they were not part of the original alphabet $V = \{a, b\}$ of Example 1, which does not distinguish between the same effect $\{a, b\}$ but due to different sources $\{ma, fb, a, b\}$ and $\{mb, fa, a, b\}$.

As we have seen, P_1 and P_2 are “ V -equivalent” in the sense that they yield the same stable models when projected to alphabet $V = \{a, b\}$. A natural question is whether this also holds in any context, that is, if $P_1 \cup Q$ and $P_2 \cup Q$ also yield the same V -projected stable models, for any context Q in the target alphabet V (since we want to keep auxiliary atoms local or hidden). This is obviously a kind of *strong equivalence* relation [11] – in fact it is one of the possible generalisations¹ of strong equivalence studied in [12]. In this paper, we will just call it *projective strong equivalence* (PSE) with respect to V , or *V -strong equivalence* for short. The PSE relation has also been used in the literature for comparing a program P and some transformation $tr(P)$ that either extends the vocabulary with new auxiliary atoms [13] (called there *strong faithfulness*) or reduces it for forgetting atoms as in [10] (called there *strong persistence*).

¹ It corresponds to *relativised* strong equivalence (with respect to V) with projection (with respect to V).

As we will see later, programs P_1 and P_2 are indeed V -strongly equivalent, so they express the same combined knowledge obtained from both parents. However, if we want to keep each parent’s contribution separately, i.e., if we try to consider P_m or P_f alone, *there is no possible $\{a, b\}$ -strongly equivalent representation* in Equilibrium Logic. This impossibility follows from a recent result in [10] that shows that forgetting atoms under strong persistence is sometimes impossible. In practice, this means that auxiliary atoms in ASP are more than ‘just’ auxiliary, as they allow one to represent problems that cannot be captured without them. A natural idea is to consider an extension of ASP in which forgetting auxiliary atoms is always possible.

In this paper, we extend logic programs to include a new construct ‘ $|$ ’ we call *fork* whose intuitive meaning is that the stable models of $P | P'$ correspond to the union of stable models from P and P' in any context Q , that is $SM[(P | P') \wedge Q] = SM[P \wedge Q] \cup SM[P' \wedge Q]$. Using this construct, we can represent Example 1 as the conjunction of two forks $(a | b) \wedge (a | b)$, one per each parent. This conjunction of forks is not idempotent but will actually amount to $(a | b | a \wedge b)$. We will show that forgetting is always possible in forks but some of them, such as $(a | b)$, cannot be represented in Equilibrium Logic unless we allow the introduction of auxiliary atoms.

The rest of the paper is organised as follows. Section 2 recalls basic definitions of HT and Equilibrium Logic. Section 3 introduces an alternative characterisation of HT in terms of T -supports. The next section extends the syntax with the fork connective and generalises the semantics to sets of T -supports (so-called T -views). In Section 5 we characterise PSE for forks and relate this property to forgetting. Finally, Section 6 discusses related work and concludes the paper.

2 Preliminaries

We begin by recalling some basic definitions and results related to HT. Let At be a set of atoms called the *propositional signature*. A (*propositional*) *formula* φ is defined using the grammar:

$$\varphi ::= \perp \mid p \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \varphi \rightarrow \varphi$$

where p is an atom. We will use Greek letters φ, ψ, γ and their variants to stand for formulas. We define the derived operators $\neg\varphi \stackrel{\text{def}}{=} (\varphi \rightarrow \perp)$ and $\top \stackrel{\text{def}}{=} \neg\perp$ and $\varphi \leftrightarrow \psi \stackrel{\text{def}}{=} (\varphi \rightarrow \psi) \wedge (\psi \rightarrow \varphi)$. Given a formula φ , by $At(\varphi) \subseteq At$ we denote the set of atoms occurring in φ . For simplicity, we consider finite theories understood as the conjunction of their formulas. The extension to infinite theories is straightforward.

A *classical interpretation* T is a set of atoms $T \subseteq At$. We write $T \models \varphi$ to stand for the usual classical satisfaction of a formula φ . An HT-*interpretation* is a pair $\langle H, T \rangle$ (respectively called “here” and “there”) of sets of atoms $H \subseteq T \subseteq At$; it is said to be *total* when $H = T$. An interpretation $\langle H, T \rangle$ *satisfies* a formula φ , written $\langle H, T \rangle \models \varphi$, if any of the following recursive conditions hold:

- $\langle H, T \rangle \not\models \perp$

- $\langle H, T \rangle \models p$ iff $p \in H$
- $\langle H, T \rangle \models \varphi \wedge \psi$ iff $\langle H, T \rangle \models \varphi$ and $\langle H, T \rangle \models \psi$
- $\langle H, T \rangle \models \varphi \vee \psi$ iff $\langle H, T \rangle \models \varphi$ or $\langle H, T \rangle \models \psi$
- $\langle H, T \rangle \models \varphi \rightarrow \psi$ iff both (i) $T \models \varphi \rightarrow \psi$ and (ii) $\langle H, T \rangle \not\models \varphi$ or $\langle H, T \rangle \models \psi$

By abuse of notation, we use ‘ \models ’ both for classical and for HT-satisfaction: the ambiguity is removed by the form of the left interpretation (a single set T for classical and a pair $\langle H, T \rangle$ for HT). It is not difficult to see that, for total interpretations, $\langle T, T \rangle \models \varphi$ amounts to classical satisfaction $T \models \varphi$.

Proposition 1. *For any interpretation $\langle H, T \rangle$ and any formula φ , $\langle H, T \rangle \models \varphi$ implies $\langle T, T \rangle \models \varphi$ (i.e. $T \models \varphi$ classically).*

A total interpretation $\langle T, T \rangle$ is an *equilibrium model* of a formula φ iff $\langle T, T \rangle \models \varphi$ and there is no $H \subset T$ such that $\langle H, T \rangle \models \varphi$. If so, we say that T is a *stable model* of φ . We write $SM[\varphi]$ to stand for the set of stable models of φ . Moreover, we represent their projection onto some vocabulary V as $SM_V[\varphi] \stackrel{\text{def}}{=} \{T \cap V \mid T \in SM[\varphi]\}$.

Definition 1 (projective strong entailment/equivalence). *Let φ and ψ be formulas and V some vocabulary (set of atoms). We say that φ V -strongly entails ψ , written $\varphi \vdash_V \psi$ if $SM_V[\varphi \wedge \gamma] \subseteq SM_V[\psi \wedge \gamma]$ for any formula γ such that $At(\gamma) \subseteq V$. We further say that φ and ψ are V -strongly equivalent, written $\varphi \cong_V \psi$, if both $\varphi \vdash_V \psi$ and $\psi \vdash_V \varphi$, that is, $SM_V[\varphi \wedge \gamma] = SM_V[\psi \wedge \gamma]$ for any formula γ such that $At(\gamma) \subseteq V$. \square*

When the vocabulary $V \supseteq At(\varphi) \cup At(\psi)$ contains the original language of φ and ψ , the projection has no relevant effect and the previous definitions amount to regular (non-projective) *strong entailment* and *strong equivalence*. In this case, we simply drop the V subindex in the previous notations. The following results, respectively proved in [11] and [14], characterise non-projective strong equivalence and entailment in terms of HT:

Proposition 2. *Let φ, ψ be a pair of formulas. Then*

- i) $\varphi \cong \psi$ iff $\varphi \equiv \psi$ in HT.
- ii) $\varphi \vdash \psi$ iff both $\varphi \models \psi$ classically and any model $\langle H, T \rangle$ of ψ s.t. $T \models \varphi$ is also a model of φ . \square

In the case of projected strong entailment and equivalence, a semantic characterisation was provided in [12], although limited to the case of disjunctive logic programs. We will provide later a characterisation of strong entailment and equivalence for fork formulas that, for the particular case in which the fork operator does not occur, will also constitute an extension of [12] to arbitrary propositional formulas.

3 T -supports

In this section we consider an alternative characterisation of HT-semantics that relies on the idea of *support* for a given classical interpretation T .

Definition 2 (T -support). Given a set of atoms T , a T -support \mathcal{H} is a set of subsets of T , that is $\mathcal{H} \subseteq 2^T$, satisfying $T \in \mathcal{H}$ if $\mathcal{H} \neq \emptyset$. \square

To increase readability of examples, we will just write a support as a sequence of interpretations between square brackets. For instance, possible supports for $T = \{a, b\}$ are $[\{a, b\} \{a\}]$, $[\{a, b\} \{b\} \emptyset]$ or the empty support $[\]$.

Intuitively, \mathcal{H} will be used to capture the set of “here” components H that support the “there” world T as a model of a given formula φ , that is, the set of H 's such that $\langle H, T \rangle \models \varphi$. When \mathcal{H} is empty $[\]$, there is no support for T , so $\langle T, T \rangle \not\models \varphi$ and thus, T is not even a classical model. If \mathcal{H} is not empty, this means we have at least some model $\langle H, T \rangle$ and, by Proposition 1, $\langle T, T \rangle$ must be a model too; this is why we require $T \in \mathcal{H}$ in the set. When not empty, the fewer models in \mathcal{H} , the more supported is T , since it is closer to being stable. Accordingly, the most supported \mathcal{H} would precisely be $\mathcal{H} = [T]$ corresponding to a stable model. This ordering relation is formally defined below.

Definition 3. Given a set of atoms $T \subseteq At$ and two T -supports \mathcal{H} and \mathcal{H}' we write $\mathcal{H} \preceq_T \mathcal{H}'$ iff either $\mathcal{H} = [\]$ or $\mathcal{H} \supseteq \mathcal{H}' \neq [\]$. \square

We usually write $\mathcal{H} \preceq \mathcal{H}'$ instead of $\mathcal{H} \preceq_T \mathcal{H}'$ when clear from the context. As an example, the classical interpretation $T = \{a, b\}$ is more supported in $\mathcal{H}_1 = [\{a, b\} \{a\}]$ than in $\mathcal{H}_2 = [\{a, b\} \{a\} \{b\} \emptyset]$, that is $\mathcal{H}_2 \preceq \mathcal{H}_1$, because \mathcal{H}_2 contains more interpretations and is further from being stable. We write \mathbf{H}_T to stand for the set of all possible T -supports. (\mathbf{H}_T, \preceq) forms a poset whose bottom and top elements are $[\]$ and $[T]$ respectively. Given a T -support \mathcal{H} , we define its complementary support $\overline{\mathcal{H}}$ as:

$$\overline{\mathcal{H}} \stackrel{\text{def}}{=} \begin{cases} [\] & \text{if } \mathcal{H} = 2^T \\ [T] \cup \{H \subseteq T \mid H \notin \mathcal{H}\} & \text{otherwise} \end{cases}$$

We also define \mathcal{H}_V as the projection of every set in \mathcal{H} to the vocabulary V , i.e., $\mathcal{H}_V \stackrel{\text{def}}{=} \{H \cap V \mid H \in \mathcal{H}\}$. The relation between T -supports and formulas is given by the following definition.

Definition 4 (T -denotation). Let At be a propositional signature and $T \subseteq At$. The T -denotation of a formula φ , written $\llbracket \varphi \rrbracket^T$, is a T -support recursively defined as follows:

$$\begin{aligned} \llbracket \perp \rrbracket^T &\stackrel{\text{def}}{=} [\] & \llbracket \varphi \wedge \psi \rrbracket^T &\stackrel{\text{def}}{=} \llbracket \varphi \rrbracket^T \cap \llbracket \psi \rrbracket^T \\ \llbracket p \rrbracket^T &\stackrel{\text{def}}{=} \{H \subseteq T \mid p \in H\} & \llbracket \varphi \vee \psi \rrbracket^T &\stackrel{\text{def}}{=} \llbracket \varphi \rrbracket^T \cup \llbracket \psi \rrbracket^T \\ \llbracket \varphi \rightarrow \psi \rrbracket^T &\stackrel{\text{def}}{=} \begin{cases} \frac{[\]}{\llbracket \varphi \rrbracket^T \cup \llbracket \psi \rrbracket^T} & \text{if } T \not\models \varphi \rightarrow \psi \\ \llbracket \varphi \rrbracket^T \cup \llbracket \psi \rrbracket^T & \text{otherwise} \end{cases} & & \square \end{aligned}$$

The following proposition follows by structural induction and shows that T -denotations can be used as an alternative semantics for the logic HT.

Proposition 3. For any interpretation $\langle H, T \rangle$ and formula φ : $\langle H, T \rangle \models \varphi$ iff $H \in \llbracket \varphi \rrbracket^T$.

Corollary 1. For any set of atoms T and propositional formulas φ, ψ , the following conditions hold:

- i) $T \models \varphi$ iff $\llbracket \varphi \rrbracket^T \neq []$ iff $T \in \llbracket \varphi \rrbracket^T$.
- ii) T is a stable model of φ iff $\llbracket \varphi \rrbracket^T = [T]$. □

As an example, given $At = \{a, b\}$, the T -denotations of the formulas $\neg a \rightarrow b$ and $a \vee b$ are:

T	$\llbracket \neg a \rightarrow b \rrbracket^T$	$\llbracket a \vee b \rrbracket^T$
\emptyset	$[]$	$[]$
$\{a\}$	$[\emptyset \{a\}]$	$[\{a\}]$
$\{b\}$	$[\{b\}]$	$[\{b\}]$
$\{a, b\}$	$[\{a, b\} \{a\} \{b\} \emptyset]$	$[\{a, b\} \{a\} \{b\}]$

Notice that the only stable model of $\neg a \rightarrow b$ is $\{b\}$ because $\llbracket \neg a \rightarrow b \rrbracket^{\{b\}} = [\{b\}]$. Similarly, the stable models of $a \vee b$ are $\{a\}$ and $\{b\}$. Note also that, for all T , we always have $\llbracket \neg a \rightarrow b \rrbracket^T \preceq \llbracket a \vee b \rrbracket^T$, that is, $\neg a \rightarrow b$ is always less supported (further from being stable) than $a \vee b$. In fact, this has an interesting consequence, as stated by the next result:

Proposition 4. For any two propositional formulas φ, ψ the following holds:

- i) $\varphi \sim \psi$ iff $\llbracket \varphi \rrbracket^T \preceq \llbracket \psi \rrbracket^T$ for every set of atoms $T \subseteq At$,
- ii) $\varphi \cong \psi$ iff $\llbracket \varphi \rrbracket^T = \llbracket \psi \rrbracket^T$ for every set of atoms $T \subseteq At$. □

While ii) is an immediate consequence of Proposition 3, item i) states that φ strongly entails ψ iff the former is always less supported than the latter. Note how Proposition 4 is much more readable than Proposition 2, especially regarding strong entailment and its relation to strong equivalence. In our example above, Proposition 4 implies that if we replace $\neg a \rightarrow b$ by $a \vee b$ in any program, we will get the same or perhaps more stable models.

The following result is a rephrasing of Theorem 2 in [15] under the equivalence of Proposition 3 and asserts that any assignment of T -supports for all T 's corresponds to (the denotation of) some formula.

Proposition 5. Let σ be an arbitrary assignment of a support $\sigma(T) \in \mathbf{H}_T$ for each $T \subseteq At$. Then, there exists a formula φ s.t. $\llbracket \varphi \rrbracket^T = \sigma(T)$ for all T . □

4 Forks and T -views

A *fork* is defined using the grammar:

$$F ::= \perp \mid p \mid F \mid F \mid F \wedge F \mid \varphi \vee \varphi \mid \varphi \rightarrow F$$

where φ is a propositional formula and $p \in At$ is an atom. As we can see, the fork operator ‘ \mid ’ cannot occur in the scope of disjunction or negation, since $\neg F$ stands for $F \rightarrow \perp$ and implications do not allow forks in the antecedent. Extending the semantics to arbitrarily nested connectives is left for future work.

As we will see, a fork F will be always reducible to the form $(\varphi_1 | \dots | \varphi_n)$ where each φ_i is a formula. Thus, a natural way to define its semantics is keeping a set of supports $\Delta = \{\mathcal{H}_1, \dots, \mathcal{H}_n\}$ for each classical interpretation T . However, if we have a pair of supports in Δ such that $\mathcal{H}_i \preceq \mathcal{H}_j$, then \mathcal{H}_i is useless since, according to Proposition 4 ii), if it yields a stable model, the latter is always produced by \mathcal{H}_j too. For this reason, we will collect sets of supports that are \preceq -closed, so their *maximal* elements become the representative ones. Formally, given a T -support \mathcal{H} we define the set of \preceq -smaller supports $\downarrow\mathcal{H} = \{\mathcal{H}' \in \mathbf{H}_T \mid \mathcal{H}' \preceq \mathcal{H}\}$. This is usually called the *ideal* of \mathcal{H} . We extend this notation to any set of supports Δ so that $\downarrow\Delta \stackrel{\text{def}}{=} \bigcup_{\mathcal{H} \in \Delta} \downarrow\mathcal{H} = \{\mathcal{H}' \preceq \mathcal{H} \mid \mathcal{H} \in \Delta\}$.

Definition 5 (T -view). A T -view is a non-empty set of T -supports $\Delta \subseteq \mathbf{H}_T$ that is \preceq -closed, i.e., $\downarrow\Delta = \Delta$. \square

We are now ready to extend the concept of T -denotation to forks.

Definition 6 (T -denotation of a fork). Let At be a propositional signature and $T \subseteq At$ a set of atoms. The T -denotation of a fork F , written $\langle\langle F \rangle\rangle^T$, is a T -view recursively defined as follows:

$$\begin{aligned}
\langle\langle \perp \rangle\rangle^T &\stackrel{\text{def}}{=} \{\{\}\} \\
\langle\langle p \rangle\rangle^T &\stackrel{\text{def}}{=} \downarrow\llbracket p \rrbracket^T \quad \text{for any atom } p \\
\langle\langle F \wedge G \rangle\rangle^T &\stackrel{\text{def}}{=} \downarrow\{\mathcal{H} \cap \mathcal{H}' \mid \mathcal{H} \in \langle\langle F \rangle\rangle^T \text{ and } \mathcal{H}' \in \langle\langle G \rangle\rangle^T\} \\
\langle\langle \varphi \vee \psi \rangle\rangle^T &\stackrel{\text{def}}{=} \downarrow\{\mathcal{H} \cup \mathcal{H}' \mid \mathcal{H} \in \langle\langle \varphi \rangle\rangle^T \text{ and } \mathcal{H}' \in \langle\langle \psi \rangle\rangle^T\} \\
\langle\langle \varphi \rightarrow F \rangle\rangle^T &\stackrel{\text{def}}{=} \begin{cases} \{\{\}\} & \text{if } T \models \varphi \text{ and } \langle\langle F \rangle\rangle^T = \{\{\}\} \\ \downarrow\{\llbracket \varphi \rrbracket^T \cup \mathcal{H} \mid \mathcal{H} \in \langle\langle F \rangle\rangle^T\} & \text{otherwise} \end{cases} \\
\langle\langle F \mid G \rangle\rangle^T &\stackrel{\text{def}}{=} \langle\langle F \rangle\rangle^T \cup \langle\langle G \rangle\rangle^T
\end{aligned}$$

It is easy to see that the fork operator ' $|$ ' is commutative, associative and idempotent. We will also see later that conjunction and implication distribute over ' $|$ '. As for the rest of operators, note that the definitions above also cover propositional formulas. The following result shows that this new T -denotation of a propositional formula φ as a T -view, $\langle\langle \varphi \rangle\rangle^T$, is precisely the ideal of its T -denotation as a T -support $\llbracket \varphi \rrbracket^T$.

Proposition 6. Let φ be a propositional formula and $T \subseteq At$ be a set of atoms. Then, $\langle\langle \varphi \rangle\rangle^T = \downarrow\llbracket \varphi \rrbracket^T$. \square

Corollary 2. Given a propositional formula φ , a set of atoms $T \subseteq At$ is a stable model of φ iff $\langle\langle \varphi \rangle\rangle^T = \downarrow\{T\}$. \square

Corollary 2 provides a natural way to define stable models of forks.

Definition 7. Given a fork F , we say that $T \subseteq At$ is a stable model of F iff $\langle\langle F \rangle\rangle^T = \downarrow\{T\}$. $SM[F]$ denotes the set of stable models of F . \square

The intuition behind a fork is that we can collect its stable models independently:

Proposition 7. *Given forks F and G : $SM[F \mid G] = SM[F] \cup SM[G]$. \square*

Once $SM[F]$ is defined, we can immediately extend the definition of V -strong entailment and equivalence to forks in the obvious way, i.e., using forks instead of propositional formulas. We postpone the effect of projecting onto some vocabulary V to the next section and focus on the regular, non-projected versions \vdash and \cong . As in Proposition 4, \vdash and \cong have a simple characterisation in terms of denotations:

Proposition 8. *For any pair of forks F, G the following holds:*

- i) $F \vdash G$ iff $\ll F \gg^T \subseteq \ll G \gg^T$ for every set of atoms $T \subseteq At$,*
- ii) $F \cong G$ iff $\ll F \gg^T = \ll G \gg^T$ for every set of atoms $T \subseteq At$. \square*

This helps us derive the following interesting properties:

Proposition 9. *Let F, G, H be arbitrary forks and φ a formula. Then:*

$$(F \mid G) \wedge H \cong (F \wedge H) \mid (G \wedge H) \quad (1)$$

$$\varphi \rightarrow (F \mid G) \cong (\varphi \rightarrow F) \mid (\varphi \rightarrow G) \quad (2)$$

$$(F \mid G) \cong G \quad \text{if } F \vdash G \quad (3)$$

The first two properties (1) and (2), allow us to reduce any arbitrary fork to a *normal form* $(\varphi_1 \mid \dots \mid \varphi_n)$ where each φ_i is a propositional formula. Moreover, (3) is a kind of *subsumption* property that, without loss of generality, allows us to assume that no pair of formulas $\varphi_i \vdash \varphi_j$ for $i \neq j$ in that expression. As an example of subsumption, take the fork $(\neg a \rightarrow b \mid a \vee b)$. As we saw before, $(\neg a \rightarrow b) \vdash (a \vee b)$ because $\ll \neg a \rightarrow b \gg^T \preceq \ll a \vee b \gg^T$ for all T (Proposition 4). Then, the ideal $\Downarrow \ll \neg a \rightarrow b \gg^T$ is included in $\Downarrow \ll a \vee b \gg^T$ which (by Proposition 6) is the same as saying $\ll \neg a \rightarrow b \gg^T \subseteq \ll a \vee b \gg^T$. But then, $\ll \neg a \rightarrow b \mid a \vee b \gg^T = \ll \neg a \rightarrow b \gg^T \cup \ll a \vee b \gg^T = \ll a \vee b \gg^T$. In other words, $(\neg a \rightarrow b \mid a \vee b) \cong (a \vee b)$.

An important consequence of conjunction-distributivity (1) is that, although \wedge is idempotent on propositional formulas, it ceases to be so when connecting forks. Take, for instance, the formalisation of Example 1 using the expression $(a \mid b) \wedge (a \mid b)$. If we apply distributivity and reduce to a normal form:

$$\begin{aligned} (a \mid b) \wedge (a \mid b) &\cong a \wedge (a \mid b) \mid b \wedge (a \mid b) && \text{distributivity (1)} \\ &\cong (a \mid a \wedge b) \mid (b \wedge a \mid b) && \text{distributivity (1)} \\ &\cong a \mid a \wedge b \mid b \wedge a \mid b && \text{associativity of '}' \\ &\cong a \mid a \wedge b \mid b && \text{commut. of } \wedge \text{ and idempotence of '}' \end{aligned}$$

but then, by Proposition 7, from $a \wedge b$ we get the stable model $\{a, b\} \notin SM[a \mid b]$.

5 Projective strong equivalence/entailment

In this section, we provide a semantic characterisation of projective strong entailment \vdash_V and equivalence \cong_V for some vocabulary $V \subseteq At$. We say that a

T -support \mathcal{H} is² V -vacuous iff there is some $H \subset T$ in \mathcal{H} satisfying $H \cap V = T \cap V$. The reason for the name “vacuous” is that if we take a formula φ with denotation $\llbracket \varphi \rrbracket^T = \mathcal{H}$, then T will never become stable if we can only add a context γ for vocabulary V . To do so, we would need $\llbracket \varphi \wedge \gamma \rrbracket^T = [T]$ but $H \subset T$ should also belong to the support since H and T are indistinguishable for any γ over V .

Definition 8. Let $V \subseteq At$ be a vocabulary and $T \subseteq V$ be a set of atoms. Then, the V - T -denotation of a fork F is a T -view defined as follows:

$$\langle\langle F \rangle\rangle_V^T \stackrel{\text{def}}{=} \downarrow \{ \mathcal{H}_V \mid \mathcal{H} \in \langle\langle F \rangle\rangle^{T'} \text{ s.t. } T' \cap V = T \text{ and } \mathcal{H} \text{ is not } V\text{-vacuous} \} \quad \square$$

In other words, we collect all the non-vacuous supports \mathcal{H} that belong to any T' -denotation $\langle\langle F \rangle\rangle^{T'}$ such that T' coincides with T for atoms in V , and then we project the supports taking \mathcal{H}_V . In doing so, we can just consider maximal \mathcal{H} 's in $\langle\langle F \rangle\rangle^{T'}$ (and add the empty T -support $[\]$ to the result if necessary). As might be expected, projecting the T -denotation of a fork F on a superset $V \supseteq At(F)$ of its atoms produces no effect:

Proposition 10. For any vocabulary $V \subseteq At$, fork F with $At(F) \subseteq V$ and set of atoms $T \subseteq V$, $\langle\langle F \rangle\rangle_V^T = \langle\langle F \rangle\rangle^T$. \square

More interestingly, the V - T -denotation of F can be used to precisely characterise its projected stable models.

Proposition 11. For any vocabulary $V \subseteq At$, fork F and set of atoms $T \subseteq V$, it holds that $T \in SM_V[F]$ iff $\langle\langle F \rangle\rangle_V^T = \downarrow [T]$. \square

As a main result, we naturally extend Proposition 8 to the projective case.

Theorem 1. For any vocabulary $V \subseteq At$, forks F, G , the following holds:

- i) $F \sim_V G$ iff $\langle\langle F \rangle\rangle_V^T \subseteq \langle\langle G \rangle\rangle_V^T$ for every set of atoms $T \subseteq V$, and
- ii) $F \cong_V G$ iff $\langle\langle F \rangle\rangle_V^T = \langle\langle G \rangle\rangle_V^T$ for every set of atoms $T \subseteq V$. \square

Moreover, the following result shows that we can just use a formula as a context instead of an arbitrary fork.

Proposition 12. For any vocabulary $V \subseteq At$, forks F, G , the following holds:

- i) $F \sim_V G$ iff $SM[F \wedge \gamma] \subseteq SM[G \wedge \gamma]$ for any formula γ with $At(\gamma) \subseteq V$
- ii) $F \cong_V G$ iff $SM[F \wedge \gamma] = SM[G \wedge \gamma]$ for any formula γ with $At(\gamma) \subseteq V$.

As an immediate consequence we can extend the characterisation of PSE from disjunctive logic programs in [12] to arbitrary propositional formulas.

Corollary 3. For any vocabulary $V \subseteq At$, formulas φ, ψ , the following holds:

- i) $\varphi \sim_V \psi$ iff $\langle\langle \varphi \rangle\rangle_V^T \subseteq \langle\langle \psi \rangle\rangle_V^T$ for every set of atoms $T \subseteq V$, and
- ii) $\varphi \cong_V \psi$ iff $\langle\langle \varphi \rangle\rangle_V^T = \langle\langle \psi \rangle\rangle_V^T$ for every set of atoms $T \subseteq V$. \square

² This notion is analogous to condition ii) in the definition of V -SE-models that characterises relativised strong equivalence [16].

As we mentioned in the introduction, PSE is closely related to forgetting. Let At be a signature and $V \subseteq At$. Given an expression φ over At , a *forgetting operator* is a partial function $\mathbf{f}(\varphi, V) = \psi$ that assigns a new expression ψ over³ V . Operator \mathbf{f} is said to be *strongly persistent* iff $\varphi \cong_V \mathbf{f}(\varphi, V)$ for every formula φ and set of atoms $V \subseteq At$ for which it is defined. Now, imagine we wish to apply forgetting on a fork F over At keeping atoms $V \subseteq At$. In light of Corollary 3, we can start by obtaining the projected denotations $\ll F \gg_V^T$ for all $T \subseteq V$. This corresponds to a set of T -views that can be precisely captured by another fork over V , as stated by the following result.

Proposition 13. *Given a vocabulary $V \subseteq At$, and assignment σ so that $\sigma(T)$ is some arbitrary T -view for each $T \subseteq V$, there exists a fork G such that $At(G) \subseteq V$ and $\ll G \gg^T = \sigma(T)$ for all $T \subseteq V$. \square*

In other words, we can always define $\mathbf{f}(F, V) \stackrel{\text{def}}{=} G$ by obtaining a fork G over V s.t. $\ll G \gg^T = \ll F \gg_V^T$ and, by Proposition 10 and Theorem 1, $G \cong_V F$.

Theorem 2. *For every fork F and set of atoms $V \subseteq At$, there is a fork G such that $At(G) \subseteq V$ and $F \cong_V G$. Consequently, there is a total, strongly persistent forgetting operator over forks. \square*

Recall from [10], that such a total operator has been shown not to exist for HT. However, since every propositional formula is also a fork, forgetting in HT is now possible if we allow the target language to be extended with the fork ‘|’ operator – “we can always forget as a fork.” Furthermore, the following relation between the T -denotation of forks and of formulas sheds light to the reason why it is not possible to forget inside HT.

Proposition 14. *Given sets of atoms $T \subseteq V \subseteq At$, then:*

- i) any formula φ with $At(\varphi) \subseteq V$ satisfies $\ll \varphi \gg_V^T = \Downarrow \ll \varphi \gg^T$ and, thus, $\ll \varphi \gg_V^T$ has a \preceq -maximum element.*
- ii) for every T -view Δ with a \preceq -maximum element, there is a propositional formula φ with $At(\varphi) \subseteq V$ that satisfies $\ll \varphi \gg_V^T = \Delta$ and $\ll \varphi \gg_V^{T'} = \{[\]\}$ for every $T' \subseteq V$ with $T' \neq T$. \square*

That is, there is a one-to-one correspondence between each assignment of T -views with some \preceq -maximum element and a formula⁴ modulo strong equivalence in the vocabulary V . On the other hand, there are T -views that have more than one \preceq -maximal element and, thus, they *cannot be represented* as formulas. That is, there are more theories modulo \cong_V than theories over V modulo \cong .

As an example, consider program P_m from the introduction interpreted as the conjunction of its rules, and assume we want to forget ma and mb . Let us take all subsets $T' \subseteq \{a, b, ma, mb\}$. Since P_m is a formula, all its T' -views will

³ Note that we are defining the forgetting operator with respect to the projected signature instead of the forgotten atoms $At \setminus V$.

⁴ A way to obtain the formula is collecting the set of HT-countermodels (those not in the T -supports) and applying the method in [17] to get a minimal logic program.

have a \preceq -maximum element, $\llbracket P_m \rrbracket^{T'}$, shown in the left table of Figure 1 where, for brevity, we only show cases of non-empty supports (i.e., when T' is a classical model). Observe that none of these supports is V -vacuous because we never get a strictly smaller $H \subset T$ that coincides with T in $V = \{a, b\}$. Now, according to the definition of $\llbracket P_m \rrbracket_V^T$, for each $T \subseteq V = \{a, b\}$ we must find those T' such that $T' \cap \{a, b\} = T$. For $T = \{a\}$ the only possibility is $T' = \{ma, a\}$ that, after removing ma , yields a maximum support $[\{a\}]$. The case for $T = \{b\}$ is completely analogous, yielding maximum support $[\{b\}]$. But for $T = \{a, b\}$ we get two candidate T' interpretations, $\{ma, a, b\}$ and $\{mb, a, b\}$. After removing ma, mb , their respective supports become $[\{a, b\} \{a\}]$ and $[\{a, b\} \{b\}]$ that are not \preceq -comparable. Therefore, they become the two \preceq -maximal supports in the T -view $\llbracket P_m \rrbracket_V^{\{a, b\}}$. Proposition 14 tells us that this *is not representable* as a propositional formula, although Proposition 13 always guarantees a representation as a fork. In particular, this set of T -views can be captured by the fork $(a \mid b)$ as we expected from the discussion in the introduction. Analogously, forgetting fa, fb in P_f yields a second fork $(a \mid b)$. Thus, the whole program $P_2 = P_m \wedge P_f$ is $\{a, b\}$ -strongly equivalent to $(a \mid b) \wedge (a \mid b)$ that, as we discussed in Section 4, amounts to $(a \mid b \mid a \wedge b)$. It is easy to see that the (non-empty) T -denotations of this fork are:

T	maximal supports in $\llbracket P_2 \rrbracket^T$
$\{a\}$	$[\{a\}]$
$\{b\}$	$[\{b\}]$
$\{a, b\}$	$[\{a, b\} \{a\}]$ $[\{a, b\} \{b\}]$

as the three cases are stable models. But this means that each T -view has a \preceq -maximum T -support, and so, the fork is representable as a formula. Program P_1 has precisely this same T -views, so it is strongly equivalent to P_2 .

T'	max. supports in $\llbracket P_m \rrbracket^{T'}$	T	max. supports in $\llbracket P_m \rrbracket_V^T$
$\{ma, a\}$	$[\{ma, a\}]$	$\{a\}$	$[\{a\}]$
$\{ma, a, b\}$	$[\{ma, a, b\} \{ma, a\}]$	$\{b\}$	$[\{b\}]$
$\{mb, b\}$	$[\{mb, b\}]$	$\{a, b\}$	$[\{a, b\} \{b\}]$ $[\{a, b\} \{a\}]$
$\{mb, a, b\}$	$[\{mb, a, b\} \{mb, b\}]$		

Fig. 1. Forgetting ma and mb in P_m .

To illustrate vacuous supports, take a second example $\alpha = (\neg p \rightarrow q) \wedge (\neg q \rightarrow p)$ and assume we want to forget q . The maximal (non-empty) supports in the views of α are:

T'	maximal supports in $\llbracket \alpha \rrbracket^{T'}$
$\{p\}$	$[\{p\}]$
$\{q\}$	$[\{q\}]$
$\{p, q\}$	$[\{p, q\} \{p\} \{q\} \emptyset]$

Note now that the last support is $\{p\}$ -vacuous because $\{p\} \subset \{p, q\} = T'$ but they coincide in the truth of $V = \{p\}$. Thus, for $T = \{p\}$ we only have the non-vacuous $T' = \{p\}$ and we get $[\{p\}]$ as maximum support. For $T = \emptyset$ we have $T' = \{q\}$ and, after removing atom q , we obtain $[\emptyset]$ as maximum support.

As both views have maximum supports, this is representable as a propositional formula, that in this case corresponds to $p \vee \neg p$.

6 Related work and conclusions

We have extended the syntax and semantics of Here-and-There (HT) to deal with a new type of construct ‘|’ called *fork*. We have studied the property of *projective strong equivalence* (PSE) for forks: two forks satisfy PSE for a vocabulary V iff they yield the same stable models projected on V for any context over V . We also provided a semantic characterisation of PSE that allowed us to prove that it is always possible to forget (under strong persistence) an auxiliary atom in a fork, something recently proved to be false in standard HT [10].

Our work on PSE is strongly related to [12] as, in fact, PSE is one of the cases studied in that paper, which was focused on disjunctive logic programs. Moreover, our semantic characterisation of PSE is in one-to-one relation to the one used there (called *certificates*). In that sense, our Corollary 3 extends the result in [12] to arbitrary formulas in HT and this is still valid for certificates. Previous complexity results can be used to prove that brave and cautious reasoning with forks in normal form are Σ_2^P and Π_2^P -complete, respectively. Similarly, checking PSE of forks in normal form is Π_4^P -complete. These results follow from the fact that every fork in normal form can be replaced (under PSE) by some (log-space constructive) propositional formula using new auxiliary atoms. We also conjecture that these results hold for arbitrary forks, but the formal proof is left for future work.

One natural extension of forks is to consider the addition of probabilities. In that way, for instance, our example about Mendelian laws could reflect the proportion of each possible combination, 1/4 for $\{a\}$ and $\{b\}$ and 1/2 for $\{a, b\}$. Doing so, we conjecture a strong formal connection to *CP-logic* [18], where the use of disjunction behaves as our fork connective. Note that, although forks do not deal with probabilities, they allow a more general syntax than CP-logic programs, which additionally require the well-founded model to be defined on all atoms. Similarly, we also plan a formal comparison with non-deterministic causal laws [19].

References

1. Baral, C.: Knowledge Representation, Reasoning and Declarative Problem Solving. Cambridge University Press (2003)
2. Gebser, M., Kaufmann, B., Schaub, T.: Conflict-driven answer set solving: From theory to practice. *Artificial Intelligence* **187-188** (2012) 52–89
3. Faber, W., Pfeifer, G., Leone, N., Dell’Armi, T., Ielpa, G.: Design and implementation of aggregate functions in the DLV system. *Theory and Practice of Logic Programming* **8**(5-6) (2008) 545–580
4. Erdem, E., Gelfond, M., Leone, N.: Applications of answer set programming. *AI Magazine* **37**(3) (2016) 53–68

5. Gelfond, M., Lifschitz, V.: The stable model semantics for logic programming. In Kowalski, R., Bowen, K., eds.: Proceedings of the Fifth International Conference and Symposium of Logic Programming (ICLP'88), MIT Press (1988) 1070–1080
6. Pearce, D.: A new logical characterisation of stable models and answer sets. In: Non monotonic extensions of logic programming. Proc. NMELP'96. (LNAI 1216). Springer-Verlag (1996)
7. Pearce, D.: Equilibrium logic. *Annals of Mathematics and Artificial Intelligence* **47**(1-2) (2006) 3–41
8. Ferraris, P., Lee, J., Lifschitz, V.: A new perspective on stable models. In: Proc. of the International Joint Conference on Artificial Intelligence (IJCAI'07). (2007) 372–379
9. Harrison, A., Lifschitz, V., Truszczyński, M.: On equivalence of infinitary formulas under the stable model semantics. *Theory and Practice of Logic Programming* **15**(1) (2015) 18–34
10. Gonçalves, R., Knorr, M., Leite, J.: You can't always forget what you want: On the limits of forgetting in answer set programming. In Kaminka, G.A., Fox, M., Bouquet, P., Hüllermeier, E., Dignum, V., Dignum, F., van Harmelen, F., eds.: ECAI 2016 - 22nd European Conference on Artificial Intelligence – Including Prestigious Applications of Artificial Intelligence (PAIS 2016). Volume 285 of *Frontiers in Artificial Intelligence and Applications.*, The Hague, The Netherlands, IOS Press (2016) 957–965
11. Lifschitz, V., Pearce, D., Valverde, A.: Strongly equivalent logic programs. *ACM Transactions on Computational Logic* **2**(4) (2001) 526–541
12. Eiter, T., Tompits, H., Woltran, S.: On solution correspondences in answer-set programming. In Kaelbling, L.P., Saffiotti, A., eds.: Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI'05), Edinburgh, Scotland, UK, Professional Book Center (2005) 97–102
13. Cabalar, P., Pearce, D., Valverde, A.: Reducing propositional theories in equilibrium logic to logic programs. In Bento, C., Cardoso, A., Dias, G., eds.: Progress in Artificial Intelligence, 12th Portuguese Conference on Artificial Intelligence, EPIA 2005, Covilhã, Portugal, December 5-8, 2005, Proceedings. Volume 3808 of *Lecture Notes in Computer Science.*, Springer (2005) 4–17
14. Aguado, F., Cabalar, P., Pearce, D., Pérez, G., Vidal, C.: A denotational semantics for equilibrium logic. *Theory and Practice of Logic Programming* **15**(4-5) (2015) 620–634
15. Cabalar, P., Ferraris, P.: Propositional theories are strongly equivalent to logic programs. *Theory and Practice of Logic Programming* **7**(6) (2007) 745–759
16. Eiter, T., Fink, M., Woltran, S.: Semantical characterizations and complexity of equivalences in answer set programming. *ACM Transactions on Computational Logic* **8**(3) (2007) 17
17. Cabalar, P., Pearce, D., Valverde, A.: Minimal logic programs. In: Proc. of the 23rd Intl. Conf. on Logic Programming, ICLP 2007, Porto, Portugal (2007) 104–118
18. Vennekens, J., Denecker, M., Bruynooghe, M.: CP-logic: A language of causal probabilistic events and its relation to logic programming. *Theory and Practice of Logic Programming* **9**(3) (2009) 245–308
19. Bochman, A.: On disjunctive causal inference and indeterminism. In: Workshop on Non-monotonic Reasoning, Actions and Change (associated to IJCAI'03). (2003) 45–50