

On the Semantics of SPARQL Queries with Optional Matching under Entailment Regimes

Egor V. Kostylev and Bernardo Cuenca Grau

Department of Computer Science, University of Oxford

Abstract. We study the semantics of SPARQL queries with optional matching features under entailment regimes. We argue that the normative semantics may lead to answers that are in conflict with the intuitive meaning of optional matching, where unbound variables naturally represent unknown information. We propose an extension of the SPARQL algebra that addresses these issues and is compatible with any entailment regime satisfying the minimal requirements given in the normative specification. We then study the complexity of query evaluation and show that our extension comes at no cost for regimes with an entailment relation of reasonable complexity. Finally, we show that our semantics preserves the known properties of optional matching that are commonly exploited for static analysis and optimisation.

1 Introduction

SPARQL became the standard language for querying RDF in 2008 [1]. Since then, the theoretical properties of SPARQL have been the subject of intensive research efforts and are by now relatively well-understood [2,3,4,5,6,7]. At the same time, SPARQL has become a core technology in practice, and most RDF-based applications rely on SPARQL endpoints for query formulation and processing.

The functionality of many such applications is enhanced by OWL 2 ontologies [8], which are used to provide background knowledge about the application domain, and to enrich query answers with implicit information. A new version of SPARQL, called SPARQL 1.1, was released in 2013 [9]. This new version captures the capabilities of OWL 2 by means of the so-called *entailment regimes* [10]: a flexible mechanism for extending SPARQL query answering to the W3C standards layered on top of RDF. A regime specifies which RDF graphs and SPARQL queries are legal (i.e., admissible) for the regime, as well as an entailment relation that unambiguously defines query answers for all legal queries and graphs.

The semantics of SPARQL under entailment regimes is specified for the conjunctive fragment, where queries are represented as *basic graph patterns* (i.e., sets of RDF triples with variables) and query answers are directly provided by the entailment relation of the regime. Roughly speaking, to check whether a mapping from variables of the query to nodes in the RDF graph is an answer to the query, one first transforms the query itself into an RDF graph by substituting each variable with the corresponding value, and then checks whether this graph is entailed in the regime by the original data graph [10,11].

When one goes beyond the basic fragment of SPARQL the language becomes considerably more complicated, but the effect of entailment regimes on the query semantics remains circumscribed to basic graph patterns. Thus, to evaluate a query one must first evaluate its component basic patterns using the relevant regime, and then compose the results by means of the SPARQL algebra operations.

Of particular interest from both a theoretical and a practical perspective is the extension of the basic fragment of SPARQL with the *optional matching feature*, which is realised in the language by means of the `OPTIONAL` operator (abbreviated by `OPT` in this paper). This feature allows the optional information to be added to query answers only when the information is available in the RDF data graph: if the optional part of the query does not match the data, then the relevant variables are left unbounded in query answers.

One of the main motivations behind optional matching in SPARQL was to deal with the “lack of regular, complete structures in RDF graphs” (see [9] Section 6) and hence with the inherent incompleteness of information in RDF data sources where only partial information about the relevant Web resources is typically available. In this setting, an unbound variable in an answer mapping is naturally interpreted as a “null” value, meaning that there might exist a binding for this variable if we consider other information elsewhere on the Web, but none is currently available in the RDF graph at hand. An additional (and slightly different) motivation for optional matching was to introduce a mechanism for “not rejecting solutions because some part of the query pattern does not match” [1]; in this sense, one would naturally expect optional matching to either extend solutions with the optional information, or to leave solutions unchanged.

Both readings of optional matching coincide if we focus just on RDF, and they are faithfully captured by the normative semantics. In this paper, however, we argue that they naturally diverge once we consider more sophisticated entailment regimes. Furthermore, the differences that arise, even if subtle, can have a major impact on expected answers.

To make this discussion concrete, let us briefly discuss a simple example of an RDF graph representing the direct train lines between UK cities as well as ferry boat transfers from UK cities to international destinations. Let this graph be exhaustive in its description of rail connections, but much less so in what concerns ferry transfers. We may exploit optional matching to retrieve all direct train connections between cities *X* and *Y*, extended with ferry transfers from *Y* to other cities *Z* whenever possible. Under the normative semantics of SPARQL we may obtain answers $(London, Oxford, -)$ and $(London, Holyhead, -)$ provided the graph has information about direct train lines from *London* to both *Oxford* and *Holyhead*, but no matching can be found in the graph for ferry connections starting from *Oxford* or *Holyhead* to other cities. Suppose next that the data graph is extended to a graph corresponding to an OWL 2 ontology in which it is stated that inland cities do not have ferry connections, and that *Oxford* is an inland city. The ontology establishes a clear distinction between *Oxford* and *Holyhead*: whereas the former is inland and cannot have ferry connections, the latter may still well be (and indeed is) a coastal city offering a number of

transfers to international destinations. The normative OWL 2 direct semantics entailment regime, however, does not distinguish between the case of *Holyhead* (where the information about ferry connections is still unknown) and *Oxford* (where the information is certain), and both answers would be returned. In this way, the normative semantics adopts the reading of optional matching where the optional information is used to complete (but never discard) query answers. In contrast, under the reading of unbounded variables as placeholders for unknown information, one would naturally expect the answer on *Oxford* to be ruled out. Indeed, if our goal were to find rail to ferry transfers starting from *London* and terminating in *Dublin* by first querying this graph and then looking for the missing information elsewhere on the Web, discarding cities like *Oxford* on the first stage would significantly facilitate our task.

In this paper, we propose an alternative semantics for the OPT operator which adopts the aforementioned reading of optional matching as an incomplete “null”. We call our semantics *strict*, which reflects the fact that it rules out those answers in which unbound variables in the optional part cannot be matched to *any* consistent extension of the input graph. Our semantics is given as an extension of the SPARQL algebra and hence satisfies the expected compositionality properties of algebraic query languages. Furthermore, it is backwards-compatible with the normative semantics for regimes in which all legal graphs are consistent, such as the RDF regime [10]. We also study the complexity of query evaluation and show that our extension comes at no cost for regimes in which entailment is not harder than query evaluation under normative semantics for the RDF regime. Finally, we show that our semantics preserves the known properties of optional matching that are commonly exploited for static analysis and optimisation.

2 SPARQL 1.1 under Entailment Regimes

In this section, we formalise the syntax and normative semantics of a core fragment of SPARQL 1.1 with optional matching under entailment regimes. Our formalisation is based on the normative specification documents [9,10,11] and builds on the well-known foundational works on SPARQL [2,3,6].

2.1 Syntax

Let \mathbf{I} , \mathbf{L} , and \mathbf{B} be countably infinite sets of *IRIs*, *literals*, and *blank nodes*, respectively. The set of *RDF terms* \mathbf{T} is $\mathbf{I} \cup \mathbf{L} \cup \mathbf{B}$. An *RDF triple* is a triple $(s\ p\ o)$ from $\mathbf{T} \times \mathbf{I} \times \mathbf{T}$, where s is called *subject*, p *predicate*, and o *object*. An (*RDF*) *graph* is a finite set of RDF triples.

Assume additionally the existence of a countably infinite set \mathbf{V} of variables disjoint from \mathbf{T} . A *triple pattern* is a tuple from $(\mathbf{T} \cup \mathbf{V}) \times (\mathbf{I} \cup \mathbf{V}) \times (\mathbf{T} \cup \mathbf{V})$. A *basic graph pattern (BGP)* is a finite set of triple patterns. *Built-in conditions* are inductively defined as follows:

1. if $?X, ?Y \in \mathbf{V}$ and $c \in \mathbf{T}$ then $\text{bound}(?X)$, $?X = c$, and $?X = ?Y$ are built-in conditions; and

2. if R_1 and R_2 are built-in conditions then $\neg R_1$, $R_1 \wedge R_2$ and $R_1 \vee R_2$ are built-in conditions.

Complex graph patterns are constructed from BGPs using a wide range of available operators that are applicable to graph patterns and built-in conditions. We focus on the AND-OPT-FILTER fragment (i.e., we consider neither union nor projection), which is widely accepted to be the fundamental core of SPARQL [2]. In this setting, *graph patterns* are inductively defined as follows (e.g., see [11]):

1. every BGP is a graph pattern;
2. if P_1 and P_2 are graph patterns that share no blank nodes then $(P_1 \text{ AND } P_2)$ and $(P_1 \text{ OPT } P_2)$ are graph patterns (called AND and OPT patterns); and
3. if P is a graph pattern and R is a built-in condition, then $(P \text{ FILTER } R)$ is a graph pattern (called FILTER pattern).

We denote $\text{vars}(P)$ (resp. with $\text{triples}(P)$) all the variables from \mathbf{V} (resp. all triple patterns) which appear in a graph pattern P .

We conclude with the definition of a special class of graph patterns with intuitive behaviour [2]. A graph pattern is *well-designed* iff (a) for each of its FILTER sub-patterns $(P \text{ FILTER } R)$ all the variables of R are in $\text{vars}(P)$, and (b) for each of its OPT sub-patterns $(P_1 \text{ OPT } P_2)$ the pattern P_1 mentions all the variables of P_2 which appear outside this sub-pattern. Note that all graph patterns in the examples of this paper are well-designed.

2.2 Semantics of BGPs under Entailment Regimes

The semantics of graph patterns is defined in terms of *mappings*; that is, partial functions from variables \mathbf{V} to terms \mathbf{T} . The *domain* $\text{dom}(\mu)$ of a mapping μ is the set of variables on which μ is defined. For a BGP P we denote with $\mu(P)$ the BGP obtained by applying μ to all variables in P from $\text{dom}(\mu)$.

Two mappings μ_1 and μ_2 are *compatible* (written as $\mu_1 \sim \mu_2$) if $\mu_1(?X) = \mu_2(?X)$ for all variables $?X$ which are in both $\text{dom}(\mu_1)$ and $\text{dom}(\mu_2)$. If $\mu_1 \sim \mu_2$, then we write $\mu_1 \cup \mu_2$ for the mapping obtained by extending μ_1 with μ_2 on variables undefined in μ_1 . A mapping μ_1 is *subsumed* by a mapping μ_2 (written $\mu_1 \sqsubseteq \mu_2$) iff $\mu_1 \sim \mu_2$ and $\text{dom}(\mu_1) \subseteq \text{dom}(\mu_2)$. Finally, a set of mappings Ω_1 is *subsumed* by a set of mappings Ω_2 (written $\Omega_1 \sqsubseteq \Omega_2$) iff for each $\mu_1 \in \Omega_1$ there exists $\mu_2 \in \Omega_2$ such that $\mu_1 \sqsubseteq \mu_2$.

Based on [10], an *entailment regime* \mathfrak{R} (or simply a *regime*) can be formalised as a tuple $(\mathbf{R}, \mathcal{G}, \mathcal{P}, \mathcal{C}, \llbracket \cdot \rrbracket)$, where

1. \mathbf{R} is a set of *reserved* IRIs from \mathbf{I} ;
2. \mathcal{G} is the set of *legal* graphs;
3. \mathcal{P} is the set of *legal* BGPs;
4. \mathcal{C} is the set of *consistent* graphs, such that $\mathcal{C} \subseteq \mathcal{G}$; and
5. $\llbracket \cdot \rrbracket$ is the *query answering* function, that takes a graph G from \mathcal{G} and a BGP P from \mathcal{P} and returns either a set $\llbracket P \rrbracket_G$ of mappings μ such that $\text{dom}(\mu) = \text{vars}(P)$, if $P \in \mathcal{C}$; or *Err*, otherwise.

As in most theoretical works on SPARQL [2,3,6,12], we assume that the query answering function returns a set of mappings, rather than a multiset. In the normative specification [10] the value of $\llbracket \cdot \rrbracket$ for inconsistent graphs is explicitly undefined, and the only thing which is guaranteed is that the answer is finite. However, an assumed behaviour is to at least issue a warning for inconsistent graphs. Moreover, in some regimes such as the OWL 2 Direct Semantics Regime issuing an error is mandatory. A regime must also satisfy certain basic additional conditions, which are immaterial to our results (see [10] Section 1.3).

The definitions of query answering and consistency in a regime are based on an entailment relation [10], which is also specified as part of the regime. We do not model the entailment relation explicitly, but assume two conditions that capture the effects of any reasonable entailment relation on legality and consistency. All regimes mentioned in the normative specification satisfy these properties and in this paper we consider only regimes that do so.

- (C1) If graphs G , G_1 and G_2 are legal and there is $h : \mathbf{T} \rightarrow \mathbf{T}$, preserving \mathbf{R} , such that $h(G_1 \cup G_2) \subseteq G$ then $G_1 \cup G_2$ is legal; if, in addition, G is in \mathcal{C} then $G_1 \cup G_2$ is also in \mathcal{C} .
- (C2) If a BGP P is in \mathcal{P} then $\mu(P)$ is in \mathcal{G} for any (total) $\mu : \mathbf{V} \rightarrow (\mathbf{T} \setminus \mathbf{R})$, such that $\mu(P)$ is a graph; if also $\mu(P)$ is in \mathcal{C} then $\mu \in \llbracket P \rrbracket_{\mu(P)}$.

Condition (C1) formalises (a weak form of) the monotonicity of legality and consistency: an illegal graph that is a union of legal ones cannot be made legal by identifying and renaming of non-reserved terms or adding triples to it; moreover, a similar property holds for consistency. Condition (C2) guarantees, that “freezing” variables of a legal BGP to non-reserved terms gives us a legal graph, and, moreover, if such a graph is consistent, then the answer of the BGP on this graph contains the mapping corresponding to the “freezing”.

The notions introduced in the remainder of this paper are parameterised with a regime \mathfrak{R} , which is not mentioned explicitly for brevity.

2.3 Normative Semantics under Entailment Regimes

Following [2], now we show how the query answering function $\llbracket \cdot \rrbracket$ extends to complex graph patterns (we refer to [2] for details). A mapping μ *satisfies* a built-in condition R , denoted $\mu \models R$, if one of the following holds:

1. R is $\text{bound}(?X)$ and $?X \in \text{dom}(\mu)$; or
2. R is $?X = c$, $?X \in \text{dom}(\mu)$ and $\mu(?X) = c$; or
3. R is $?X = ?Y$, $?X \in \text{dom}(\mu)$, $?Y \in \text{dom}(\mu)$ and $\mu(?X) = \mu(?Y)$; or
4. R is an evaluating to true Boolean combination of other built-in conditions.

Given two sets of mappings Ω_1 and Ω_2 , the *join*, *union* and *difference* operations are defined as follows:

$$\begin{aligned} \Omega_1 \bowtie \Omega_2 &= \{\mu_1 \cup \mu_2 \mid \mu_1 \in \Omega_1 \text{ and } \mu_2 \in \Omega_2 \text{ such that } \mu_1 \sim \mu_2\}, \\ \Omega_1 \cup \Omega_2 &= \{\mu \mid \mu \in \Omega_1 \text{ or } \mu \in \Omega_2\}, \\ \Omega_1 \setminus \Omega_2 &= \{\mu_1 \mid \mu_1 \in \Omega_1, \text{ there is no } \mu_2 \in \Omega_2 \text{ such that } \mu_1 \sim \mu_2\}. \end{aligned}$$

Based on these, the *left outer join* operation is defined as follows:

$$\Omega_1 \bowtie \Omega_2 = (\Omega_1 \times \Omega_2) \cup (\Omega_1 \setminus \Omega_2).$$

A graph pattern is *legal* for a regime \mathfrak{R} if all the BGPs it contains are legal. The *normative query answering* function $\llbracket \cdot \rrbracket^n$ is inductively defined for all legal graph patterns P on the base of $\llbracket \cdot \rrbracket$ as follows. For graphs G from \mathcal{C} we have:

1. if P is a BGP then $\llbracket P \rrbracket_G^n = \llbracket P \rrbracket_G$;
2. if P is P_1 AND P_2 then $\llbracket P \rrbracket_G^n = \llbracket P_1 \rrbracket_G^n \times \llbracket P_2 \rrbracket_G^n$;
3. if P is P_1 OPT P_2 then $\llbracket P \rrbracket_G^n = \llbracket P_1 \rrbracket_G^n \bowtie \llbracket P_2 \rrbracket_G^n$; and
4. if P is P' FILTER R then $\llbracket P \rrbracket_G^n = \{\mu \mid \mu \in \llbracket P' \rrbracket_G^n \text{ and } \mu \models R\}$.

If $G \notin \mathcal{C}$ then $\llbracket P \rrbracket_G^n = Err$ for any graph pattern P (which again coincides with $\llbracket P \rrbracket_G$ when P is a BGP). Note, that by these definitions $\mu \in \llbracket P \rrbracket_G^n$ implies that $\text{dom}(\mu) \subseteq \text{vars}(P)$, but this inclusion may be strict if P contains OPT operator.

Having the semantics defined, we say that two legal patterns P_1 and P_2 are *equivalent (under normative semantics)* with respect to a regime \mathfrak{R} , denoted by $P_1 \equiv^n P_2$, if $\llbracket P_1 \rrbracket_G^n = \llbracket P_2 \rrbracket_G^n$ for every RDF graph $G \in \mathcal{G}$.

3 On Optional Matching under the Normative Semantics

One of the main motivations for optional matching in SPARQL was to deal with the “lack of regular, complete structures in RDF graphs” [9]. In contrast to relational databases, RDF data is loosely structured and hence in many applications it is not satisfactory to reject an answer if some relevant information is missing. For example, if we are interested in retrieving the names, emails, and websites of employees, we may not want to discard a partial answer involving the name and email address of a particular employee merely because the information on the employee’s website is not available in the graph.

The normative semantics was designed to deal with such situations: the optional information is included in query answers only when the information is available; otherwise, the relevant variables are left unbounded. An unbound variable in an answer is thus a manifestation of inherent incompleteness of RDF data sources, and the missing information is interpreted as unknown.

This natural interpretation of query results, however, no longer holds if the query is evaluated under certain entailment regimes, as we illustrate next by means of examples. In these and all other examples given later on, we focus on the OWL 2 direct semantics regime. In order for an RDF graph to be legal for this regime, it must correspond to an OWL 2 ontology; similarly, legal BGPs must correspond to an extended ontology in which variables are allowed [10]. Thus, in the examples we express RDF graphs and BGPs in (extended) OWL 2 functional syntax, and use words “ontology” and “graph” interchangeably.¹

¹ Declaration axioms are omitted in ontologies and BGPs. Also, we use shortened names for some constructs, i.e., PropertyDomain instead of ObjectPropertyDomain.

Example 1. Consider the OWL 2 ontology \mathcal{O}_1 consisting of the following axioms:

ClassAssertion(*InlandCity Oxford*), PropertyAssertion(*train London Oxford*),
 ClassAssertion(*CoastalCity Holyhead*), PropertyAssertion(*train London Holyhead*),
 PropertyDomain(*ferry CoastalCity*), DisjointClasses(*CoastalCity InlandCity*).

Consider also the following graph pattern P_1 , which we wish to evaluate over \mathcal{O}_1 :

PropertyAssertion(*train ?X ?Y*) OPT PropertyAssertion(*ferry ?Y ?Z*).

Intuitively, solutions to P_1 provide direct train lines from city X to city Y as well as, optionally, the ferry transfers from Y to other cities Z . Under the normative semantics, the BGPs in P_1 are evaluated separately. In particular, the optional BGP is evaluated to the empty set, and $\llbracket P_1 \rrbracket_{\mathcal{O}_1}^n = \{\mu_1, \mu_2\}$, where

$$\begin{aligned} \mu_1 &= \{?X \mapsto \textit{London}, ?Y \mapsto \textit{Oxford}\}, \text{ and} \\ \mu_2 &= \{?X \mapsto \textit{London}, ?Y \mapsto \textit{Holyhead}\}. \end{aligned}$$

In both answers, variable $?Z$ is unbounded and hence we conclude that \mathcal{O}_1 contains no relevant information about ferry connections starting from *Oxford* or *Holyhead*. However, the nature of the lack of such information is fundamentally different. On the one hand, the connections from *Holyhead* (e.g., to *Dublin*) are missing from \mathcal{O}_1 just by the incompleteness of the information in the graph, which is usual in (and also a feature of) Semantic Web applications. On the other hand, *Oxford* cannot have a ferry connection because it is a landlocked city, and hence the information about its (lack of) ferry connections is certain. Thus, the normative semantics cannot distinguish between unknown and non-existent ferry connections. However, if we adhere to the reading of unbounded variables as incomplete information or “nulls”, then μ_1 should not be returned as an answer.

The issues described in this example become even more apparent in cases where the optional part alone cannot be satisfied, as it is incompatible with the information in the graph, as illustrated by the following example.

Example 2. Consider the ontology \mathcal{O}_2 with the following axioms:

ClassAssertion(*Person Peter*), DisjointProperties(*hasFather hasMother*).

Furthermore, consider the following pattern P_2 :

ClassAssertion(*Person ?X*) OPT ({
 PropertyAssertion(*hasFather ?X ?Y*),
 PropertyAssertion(*hasMother ?X ?Y*)}).

The optional BGP is evaluated to the empty mapping and hence under the normative semantics we have the mapping $\{?X \mapsto \textit{Peter}\}$ forming the answer set $\llbracket P_2 \rrbracket_{\mathcal{O}_2}^n$. Note, however, that the optional BGP in P_2 is in contradiction with the disjointness axiom: under the OWL 2 regime, no solution to P_2 can exist for any ontology containing that axiom.

As these examples suggest, if we interpret unbound variables in answers to queries with optional parts as an indication of unknown information in the data graph, then the normative semantics may yield counter-intuitive answers.

At the core of this issue is the inability of the normative semantics to distinguish between answers in which it is possible to assign values to the missing optional part (a natural reflection of incompleteness in the data), and those where this is impossible (a reflection that the missing information is fundamentally incompatible with the answer). This distinction is immaterial for regimes in which all legal graphs are consistent, but it quickly becomes apparent in more sophisticated regimes, such as those based on OWL 2.

4 Semantics of Strict Optional Matching

In this section, we propose our novel semantics for optional matching under regimes. In a nutshell, our semantics addresses the issues described in Section 3 by ruling out those answer mappings where unbound variables in the optional part cannot be matched to any consistent extension of the input graph. Our semantics is therefore *strict*, in the sense that only answers in which unbound variables are genuine manifestations of incompleteness in the data are returned.

4.1 Definition of Strict Semantics

We start by introducing the notion of a frozen RDF graph for a pattern P and a mapping μ . Roughly speaking, this graph is obtained by taking all the triple patterns in P and transforming them into RDF triples by applying the extension of μ where unbounded variables are “frozen” to arbitrary fresh constants.

Definition 1. *Let $\mathfrak{R} = (\mathbf{R}, \mathcal{G}, \mathcal{P}, \mathcal{C}, \llbracket \cdot \rrbracket)$ be an entailment regime. Let P be a legal graph pattern, and let μ be a mapping from variables \mathbf{V} to RDF terms \mathbf{T} . Then, the freezing G_μ^P of P under μ is the RDF graph $\bar{\mu}(\text{triples}(P))$, where $\bar{\mu}$ is the mapping that extends μ by assigning each variable in $\text{vars}(P)$, which is not in $\text{dom}(\mu)$, to a globally fresh IRI from \mathbf{I} (not belonging to \mathbf{R}).*

The freezing G_μ^P depends only on the candidate mapping μ and the triple patterns occurring in P ; thus, it does not depend either on the specific operators used in P , or on the RDF graph over which the query pattern is to be evaluated.

Example 3. For the pattern P_1 and mappings μ_1 and μ_2 from Example 1 we have the following freezings in functional-style syntax:

$$\begin{aligned} G_{\mu_1}^{P_1} &= \{\text{PropertyAssertion}(\text{train London Oxford}, \\ &\quad \text{PropertyAssertion}(\text{ferry Oxford } w_1)); \\ G_{\mu_2}^{P_1} &= \{\text{PropertyAssertion}(\text{train London Holyhead}, \\ &\quad \text{PropertyAssertion}(\text{ferry Holyhead } w_2)); \end{aligned}$$

where w_1 and w_2 are freshly introduced IRIs.

Intuitively, the freezing represents the simplest, most general, RDF graph over which all the undefined variables in a given solution mapping could be bounded to concrete values. Thus, if G_μ^P together with the input graph G is not a consistent graph for the relevant regime, we can conclude, using condition (C1) of the regime, that the undefined variables in μ will never be matched to concrete values in any consistent extension of G and hence μ should be ruled out as an answer. On the other hand, if $G \cup G_\mu^P$ is consistent, then such an extension exists and, by condition (C2), the undefined variables can be mapped in this extension.

Definition 2. Let $\mathfrak{R} = (\mathbf{R}, \mathcal{G}, \mathcal{P}, \mathcal{C}, [\cdot])$ be an entailment regime. A mapping μ is \mathfrak{R} -admissible for a graph $G \in \mathcal{C}$ and legal graph pattern P if $G \cup G_\mu^P$ is a graph belonging to \mathcal{C} . The set of all \mathfrak{R} -admissible mappings for a consistent graph G and a legal graph pattern P is denoted as $\text{Adm}(G, P)$.

Example 4. Clearly, $\mathcal{O}_1 \cup G_{\mu_1}^{P_1}$ is inconsistent since ferries only depart from coastal cities, but Oxford is an inland city. In contrast, $\mathcal{O}_1 \cup G_{\mu_2}^{P_1}$ is consistent. Thus, we have $\mu_1 \notin \text{Adm}(\mathcal{O}_1, P)$, but $\mu_2 \in \text{Adm}(\mathcal{O}_1, P)$.

We are now ready to formalise our semantics.

Definition 3. Let $\mathfrak{R} = (\mathbf{R}, \mathcal{G}, \mathcal{P}, \mathcal{C}, [\cdot])$ be an entailment regime. The strict query answering function $[\cdot]^s$ is inductively defined for all legal graph patterns P on the base of $[\cdot]$ as follows. For graphs G from \mathcal{C} we have:

1. if P is a BGP then $[[P]]_G^s = [[P]]_G$;
2. if P is P_1 AND P_2 then $[[P]]_G^s = ([[P_1]]_G^s \times [[P_2]]_G^s) \cap \text{Adm}(G, P)$;
3. if P is P_1 OPT P_2 then $[[P]]_G^s = ([[P_1]]_G^s \bowtie [[P_2]]_G^s) \cap \text{Adm}(G, P)$; and
4. if P is P' FILTER R then $[[P]]_G^s = \{\mu \mid \mu \in [[P']]_G^s \text{ and } \mu \models R\}$,

where \cap denotes the standard set-theoretic intersection. If $G \notin \mathcal{C}$ then $[[P]]_G^s = \text{Err}$ for any graph pattern P . Finally, legal patterns P_1 and P_2 are equivalent (under strict semantics), written $P_1 \equiv^s P_2$, if $[[P_1]]_G^s = [[P_2]]_G^s$ for any legal G .

Example 5. The strict semantics behaves as expected for our running examples. For \mathcal{O}_1 and P_1 from Example 1 we have that $[[P_1]]_{\mathcal{O}_1}^s = \{\mu_1\}$, whereas for \mathcal{O}_2 and P_2 from Example 2 we have $[[P_2]]_{\mathcal{O}_2}^s = \emptyset$.

The strict and normative semantics coincide in two limit cases. First, if the entailment regime does not allow for inconsistent graphs (i.e., if $\mathcal{C} = \mathcal{G}$) as is the case for the RDF regime [10], then $[[P]]_G^s = [[P]]_G^n$ for every legal pattern P and graph G . Second, if the relevant pattern P is OPT-free then the freezing for every candidate answer mapping contains no fresh IRIs and is \mathfrak{R} -entailed by G ; thus, we again have $[[P]]_G^s = [[P]]_G^n$ for every legal graph G .

Thus, the difference between the normative semantics $[\cdot]^n$ and strict semantics $[\cdot]^s$ manifests only for regimes that admit inconsistency, and is circumscribed to the presence of OPT in graph patterns, where non-admissible mappings are excluded in the case of the strict semantics. Note, however, that even if a mapping

μ_1 (resp. μ_2) is admissible for a sub-pattern P_1 (resp. P_2) containing OPT, it is possible for $\mu_1 \cup \mu_2$ not to be admissible for the joined pattern $P = P_1$ AND P_2 . Thus, the admissibility restriction is also explicitly reflected in the semantics of AND given in Definition 3. This is illustrated in the example given next.

Example 6. Consider ontology \mathcal{O}_3 , consisting of the following axioms:

```
SubClassOf(
  IntersectionOf(SomeValuesFrom(husband Thing)
    SomeValuesFrom(wife Thing))
  Nothing),
ClassAssertion(Person Mary).
```

The first axiom establishes that a person cannot have both a husband and a wife. Consider also the following well-designed graph pattern P_3 :

```
(ClassAssertion(Person ?X) OPT (PropertyAssertion(husband ?X ?Y))) AND
(ClassAssertion(Person ?X) OPT (PropertyAssertion(wife ?X ?Z))).
```

Clearly, $\mu = \{?X \mapsto \textit{Mary}\}$ belongs to the strict answer to each of the OPT sub-patterns of P_3 since each of them independently can match to a consistent extension of \mathcal{O}_3 . However, μ is not admissible for P_3 since *Mary* has both a husband and a wife in $G_\mu^{P_3}$, and hence $\mathcal{O}_3 \cup G_\mu^{P_3}$ is inconsistent. Thus, $\llbracket P_3 \rrbracket_{\mathcal{O}_3}^s = \emptyset$.

4.2 Comparing the Normative and Strict Semantics

Our previous examples support the expected behaviour of our semantics, namely that its effect is circumscribed to filtering out problematic answers returned under the normative semantics. We next formally show that our semantics behaves as expected *in general*, provided that we restrict ourselves to well-designed patterns and negation-free FILTER expressions (which are rather mild restrictions).

It is known that patterns which are not well-designed easily lead to unexpected answers, even under the normative semantics (we refer to [2] for a detailed discussion). Therefore, it comes at no surprise that the intuitive behaviour of our semantics is only guaranteed under this assumption.

But before stating our main result in this section, we establish first a useful lemma that is applicable to well-designed patterns under both semantics.

Lemma 1. *Let $\mathfrak{R} = (\mathbf{R}, \mathcal{G}, \mathcal{P}, \mathcal{C}, \llbracket \cdot \rrbracket)$ be an entailment regime, and let P be a sub-pattern of a well-designed graph pattern. If a variable appears both inside and outside P , then it is defined in each mapping from $\llbracket P \rrbracket_G^n$ for any consistent G as well as in each mapping from $\llbracket P \rrbracket_G^s$.*

Proof. For the sake of contradiction, let $?X$ be a variable which appears both inside and outside P , but is undefined in some mapping $\mu \in \llbracket P \rrbracket_G^n \cup \llbracket P \rrbracket_G^s$. Since $?X$ is undefined, it appears in the optional (i.e., right) part of an OPT-sub-pattern of P . The overall graph pattern is well-designed, so $?X$ must appear in the mandatory (i.e., left) part of the OPT-sub-pattern of P . If this sub-pattern

is not in the optional part of another higher-level sub-pattern, then $?X$ must be defined in μ , which contradicts the original assumption. Otherwise, we can apply the same argument and show that $?X$ is in the mandatory part of the higher-level sub-pattern. Such reasoning would eventually lead to a contradiction. \square

Theorem 1. *Let $\mathfrak{R} = (\mathbf{R}, \mathcal{G}, \mathcal{P}, \mathcal{C}, \llbracket \cdot \rrbracket)$ be an entailment regime. The inclusion $\llbracket P \rrbracket_G^s \subseteq \llbracket P \rrbracket_G^n$ holds for any graph G from \mathcal{C} and any legal well-designed graph pattern P which does not use negation \neg in FILTER expressions.*

Proof. The proof is by induction on the structure of the graph pattern P .

1. Let P be a BGP. Then $\llbracket P \rrbracket_G^s = \llbracket P \rrbracket_G^n = \llbracket P \rrbracket_G$ by definition of the semantics.
2. Let P be P_1 AND P_2 , $\llbracket P_1 \rrbracket_G^s \subseteq \llbracket P_1 \rrbracket_G^n$, $\llbracket P_2 \rrbracket_G^s \subseteq \llbracket P_2 \rrbracket_G^n$, and μ be a mapping from $\llbracket P \rrbracket_G^s$. By the definition of the semantics of AND patterns we have that $\llbracket P \rrbracket_G^s = (\llbracket P_1 \rrbracket_G^s \bowtie \llbracket P_2 \rrbracket_G^s) \cap \text{Adm}(G, P)$, that is there exist compatible $\mu_1 \in \llbracket P_1 \rrbracket_G^s$ and $\mu_2 \in \llbracket P_2 \rrbracket_G^s$ such that $\mu = \mu_1 \cup \mu_2$. By the inductive assumption, there exist $\mu'_1 \in \llbracket P_1 \rrbracket_G^n$ and $\mu'_2 \in \llbracket P_2 \rrbracket_G^n$ such that $\mu_1 \sqsubseteq \mu'_1$ and $\mu_2 \sqsubseteq \mu'_2$. By Lemma 1, μ'_1 and μ'_2 are compatible. Indeed, if a variable is defined in both of them, it occurs in P_1 and P_2 , so, by the lemma, it is defined in both μ_1 and μ_2 which are compatible; hence the values of this variable in all μ_1, μ_2, μ'_1 and μ'_2 coincide, which implies that the last two are compatible. So, $\mu'_1 \cup \mu'_2$ is in $\llbracket P_1 \rrbracket_G^n \bowtie \llbracket P_2 \rrbracket_G^n = \llbracket P \rrbracket_G^n$ as required.
3. Let P be P_1 OPT P_2 , $\llbracket P_1 \rrbracket_G^s \subseteq \llbracket P_1 \rrbracket_G^n$, $\llbracket P_2 \rrbracket_G^s \subseteq \llbracket P_2 \rrbracket_G^n$, and μ be a mapping from $\llbracket P \rrbracket_G^s$. By the definition of the semantics of OPT patterns we have that $\llbracket P \rrbracket_G^s = (\llbracket P_1 \rrbracket_G^s \bowtie \llbracket P_2 \rrbracket_G^s) \cap \text{Adm}(G, P)$, that is, we have two options.
 - (a) Let $\mu \in \llbracket P_1 \rrbracket_G^s \bowtie \llbracket P_2 \rrbracket_G^s$ and $\mu \in \text{Adm}(G, P)$. In this case the rest of the proof goes the same lines as the case when P is an AND graph pattern.
 - (b) Let $\mu \in \llbracket P_1 \rrbracket_G^s \setminus \llbracket P_2 \rrbracket_G^s$ and $\mu \in \text{Adm}(G, P)$. Hence, in particular, $\mu \in \llbracket P_1 \rrbracket_G^s$. By the inductive assumption, there exists a mapping μ'_1 in $\llbracket P_1 \rrbracket_G^n$ such that $\mu \sqsubseteq \mu'_1$. Hence, as required, there exists a mapping μ' in $\llbracket P \rrbracket_G^n$ (which is either $\mu'_1 \cup \mu'_2$ for some $\mu'_2 \in \llbracket P_2 \rrbracket_G^n$ compatible with μ'_1 , or μ'_1 itself, if such a μ'_2 does not exist), such that $\mu \sqsubseteq \mu'$.
4. Let P be P' FILTER R , $\llbracket P' \rrbracket_G^s \subseteq \llbracket P' \rrbracket_G^n$, R be a built-in condition not using \neg , and $\mu \in \llbracket P \rrbracket_G^s$. By the definition of semantics of FILTER patterns, $\mu \in \llbracket P' \rrbracket_G^s$. Hence, by the inductive assumption, there exists $\mu' \in \llbracket P' \rrbracket_G^n$ such that $\mu \sqsubseteq \mu'$. Since R does not use \neg , $\mu' \models R$. So, $\mu' \in \llbracket P \rrbracket_G^n$, as required. \square

Note that Theorem 1 is formulated in terms of subsumption, instead of set-theoretic containment. The rationale behind this formulation is clarified next.

Example 7. Consider the ontology \mathcal{O}'_1 , which is obtained from \mathcal{O}_1 in Example 1 by removing all axioms involving *Holyhead*, and adding the following axiom:

PropertyAssertion(*bus* *Canterbury* *London*).

Consider also the following graph pattern P'_1 :

PropertyAssertion(*bus* $?U$ $?X$) OPT

(PropertyAssertion(*train* $?X$ $?Y$) OPT

PropertyAssertion(*ferry* $?Y$ $?Z$)).

The following mapping μ is returned by the normative semantics:

$$\{?U \mapsto \textit{Canterbury}, ?X \mapsto \textit{London}, ?Y \mapsto \textit{Oxford}\}.$$

As already discussed, Oxford is an inland city and hence cannot have ferry connections; thus, μ is not returned under strict semantics. However, it may be possible to reach a ferry connection from London (although none is given), and hence the following answer μ' is returned instead of μ under strict semantics.

$$\{?U \mapsto \textit{Canterbury}, ?X \mapsto \textit{London}\}.$$

Clearly, μ' is not returned under normative semantics and hence $\llbracket P'_1 \rrbracket_{\mathcal{O}'_1}^s \not\subseteq \llbracket P'_1 \rrbracket_{\mathcal{O}'_1}^n$; however, $\mu' \sqsubseteq \mu$ and $\llbracket P'_1 \rrbracket_{\mathcal{O}'_1}^s \sqsubseteq \llbracket P'_1 \rrbracket_{\mathcal{O}'_1}^n$.

5 Computational Properties and Static Optimisation

In this section, we first study the computational properties of our semantics. We show that the complexity of graph pattern evaluation under strict and normative semantics coincide, provided that consistency checking is feasible in PSPACE for the regime at hand. Then we focus on static query analysis, and in particular on pattern equivalence. We show that the key equivalence-preserving transformation rules that have been proposed for static optimisation of SPARQL queries continue to hold if we consider equivalence under strict semantics.

5.1 Complexity of Strict Graph Pattern Evaluation

Recall that the graph pattern evaluation is the key reasoning problem in SPARQL. In the context of entailment regimes, it is defined as given next, where x is either n or s , depending of the semantics.

GRAPH PATTERN EVALUATION
<i>Input:</i> Regime \mathfrak{R} , legal graph G , legal graph pattern P , and mapping μ .
<i>Question:</i> Is $\mu \in \llbracket P \rrbracket_G^x$ under the regime \mathfrak{R} ?

Here, when we say that regime \mathfrak{R} is a part of the input, we mean that it includes two oracle functions checking consistency of legal graphs and evaluating legal BGPs over legal graphs, respectively. In what follows, we refer to the problem as **NORMATIVE** if $x = n$, and as **STRICT** if $x = s$.

It is known that the normative graph pattern evaluation problem is in PSPACE for the RDF regime [2]. We next argue that membership in PSPACE holds in general for any regime satisfying the basic properties discussed in Section 2 and for both normative and strict versions of the problem, provided that the complexity of both oracles of the regime is in PSPACE.

Theorem 2. *NORMATIVE and STRICT GRAPH PATTERN EVALUATION problems are in PSPACE, provided the oracles associated to input regimes are in PSPACE.*

Proof. We start with the normative semantics. If the regime \mathfrak{R} allows for inconsistent graphs, then the regime may require to first check the input graph for consistency, which is feasible in PSPACE by the assumptions of the theorem. Then, the same recursive procedure $Eval(\mu, P, G)$ for a mapping μ , graph pattern P and graph G described in [2] for the RDF regime is applicable to evaluate queries under \mathfrak{R} . The procedure is in PSPACE because each BGP in the input pattern can be evaluated by the second oracle in PSPACE by the assumptions.

The situation is analogous for the strict semantics: consistency checking may be required upfront by the regime and the recursive PSPACE procedure $Eval(\mu, P, G)$ can be designed by a straightforward implementation of the definition of strict semantics. The only difference with the normative case is that in each step we need to check for admissibility, which is possible in PSPACE given our assumptions about the oracles of the regime. For example, if P is a pattern of the form $P_1 \text{ OPT } P_2$, then $Eval$ first checks whether $Eval(\mu, P_1 \text{ AND } P_2, G)$ holds, or $Eval(\mu, P_1, G)$ with not $Eval(\mu', P_2, G)$ for any mapping μ' , $\text{dom}(\mu') \subseteq \text{vars}(P_2)$, holds; and then checks, by means of the oracle function, whether $\mu \in \text{Adm}(G, P)$. If both of the checks are positive, then the answer is true, otherwise, it is false. \square

Consequently, the use of our strict semantics does not increase the computational complexity for reasonable regimes. In particular, it follows directly from Theorem 2 that the evaluation problem is in PSPACE under both semantics for the tractable entailment regimes associated to the OWL 2 profiles [13].

It is also well-known that graph pattern evaluation under normative semantics is PSPACE-hard for the RDF regime [2]. To formulate a general hardness result that holds for *any* regime we would need to require additional properties for a regime to qualify as “reasonable”. In order not to unnecessarily complicate the presentation, we simply point out that PSPACE-hardness holds for all the regimes mentioned in the specification under both normative and strict semantics [10]. This is immediate for the normative semantics. The hardness for the strict semantics can be proved by reduction from quantified boolean formula validity as it is done in [2]. This reduction encodes a formula in the query pattern, and relies on a fixed RDF dataset, which corresponds to a set of first-order facts. Thus, the extension of such fixed graph with the freezing of a sub-pattern of the query pattern is always consistent under all the regimes in the specification.

5.2 Static Analysis and Optimisation

Static analysis and optimisation of SPARQL queries has received significant attention in recent years [4,6,14,15,16]. A key ingredient for query optimisation is the availability of a comprehensive catalog of equivalence-preserving transformation rules for SPARQL patterns. Schmidt et al. [4] and Perez et al. [2] provide a rich set of such equivalences, which holds under normative semantics for RDF regime. Some of these equivalences, such as idempotence, commutativity, and associativity of the AND operator, hold without any restrictions (for our core fragment of SPARQL) and are quite easy to prove. However, those that involve OPT are more intricate and hold only for well-designed patterns.

The claim of this section is that these equivalences continue to hold for any entailment regime, under both normative and strict semantics. For brevity, we concentrate only on strict semantics and the equivalences with OPT, leaving the proofs for the rest out of the scope of this paper. But before doing so we state the following auxiliary lemma, which proof is straightforward application of conditions (C1)–(C2) on regimes and omitted for brevity.

Lemma 2. *Let $\mathfrak{R} = (\mathbf{R}, \mathcal{G}, \mathcal{P}, \mathcal{C}, \llbracket \cdot \rrbracket)$ be an entailment regime. If $\mu \in \text{Adm}(G, P)$ for a mapping μ , graph $G \in \mathcal{C}$ and legal graph pattern P , then $\mu' \in \text{Adm}(G, P')$ for any mapping $\mu' \sqsubseteq \mu$ and legal graph pattern P' such that $\text{triples}(P') \subseteq \text{triples}(P)$.*

The following theorem establishes that the known equivalences involving OPT in well-designed patterns hold for any entailment regime under strict semantics.

Theorem 3. *Let \mathfrak{R} be an entailment regime. The following equivalences hold, provided the graph patterns on both sides are legal and well-designed:*

$$(P_1 \text{ OPT } P_2) \text{ FILTER } R \equiv^s (P_1 \text{ FILTER } R) \text{ OPT } P_2, \tag{1}$$

$$P_1 \text{ AND } (P_2 \text{ OPT } P_3) \equiv^s (P_1 \text{ AND } P_2) \text{ OPT } P_3, \tag{2}$$

$$(P_1 \text{ OPT } P_2) \text{ AND } P_3 \equiv^s (P_1 \text{ AND } P_3) \text{ OPT } P_2, \tag{3}$$

$$(P_1 \text{ OPT } P_2) \text{ OPT } P_3 \equiv^s (P_1 \text{ OPT } P_3) \text{ OPT } P_2. \tag{4}$$

Proof. Equivalence (1). For brevity we give a proof only for the fact that for every graph each mapping in the semantics of the pattern on the left is in the semantics of the pattern on the right, and the other direction is similar. Let $\mu \in \llbracket (P_1 \text{ OPT } P_2) \text{ FILTER } R \rrbracket_G^s$ for a consistent graph G , built-in condition R and legal graph patterns P_1 and P_2 , such that the both sides of the equivalence are well-designed. By the definition of strict semantics there are two possibilities.

Case 1: $\mu = \mu_1 \cup \mu_2$ for $\mu_i \in \llbracket P_i \rrbracket_G^s$, such that $\mu \in \text{Adm}(G, P_1 \text{ OPT } P_2)$ and $\mu \models R$. To show that $\mu \in \llbracket (P_1 \text{ FILTER } R) \text{ OPT } P_2 \rrbracket_G^s$ the only thing we need to prove is that $\mu_1 \models R$, and we do it by showing that for every $?X$ which appears in R the values of μ and μ_1 on $?X$ are either the same or both undefined. Indeed, if such $?X$ is in P_2 then, since $(P_1 \text{ OPT } P_2) \text{ FILTER } R$ is well-designed, it appears in P_1 , and hence, by Lemma 1, $?X \in \text{dom}(\mu_1)$, which means that $\mu(?X) = \mu_1(?X)$. On the other hand, if $?X$ is not in P_2 , then $?X \notin \text{dom}(\mu_2)$, that is either the value of μ on $?X$ coincides with the value of μ_1 or both of them are undefined.

Case 2: $\mu \in \llbracket P_1 \rrbracket_G^s$, such that $\mu \in \text{Adm}(G, P_1 \text{ OPT } P_2)$, $\mu \models R$, and there is no $\mu_2 \in \llbracket P_2 \rrbracket_G^s$ such that $\mu \sim \mu_2$. Immediately we have that $\mu \in \llbracket (P_1 \text{ FILTER } R) \text{ OPT } P_2 \rrbracket_G^s$.

Equivalence (2). For brevity, we show the more involved direction that shows containment from right to left. Let $\mu \in \llbracket (P_1 \text{ AND } P_2) \text{ OPT } P_3 \rrbracket_G^s$ for a consistent G , and legal patterns P_1, P_2 and P_3 , such that the both sides of the equivalence are well-designed. We have the following cases by the definition of strict semantics.

Case 1: $\mu = \mu_1 \cup \mu_2 \cup \mu_3$ for $\mu_i \in \llbracket P_i \rrbracket_G^s$, such that $\mu_1 \cup \mu_2 \in \text{Adm}(G, P_1 \text{ AND } P_2)$ and $\mu \in \text{Adm}(G, (P_1 \text{ AND } P_2) \text{ OPT } P_3)$. By Lemma 2 the last inclusion implies

the first. Moreover, it implies that $\mu_2 \cup \mu_3 \in \text{Adm}(G, P_2 \text{ OPT } P_3)$, which, together with the fact that $\text{Adm}(G, (P_1 \text{ AND } P_2) \text{ OPT } P_3) = \text{Adm}(G, P_1 \text{ AND } (P_2 \text{ OPT } P_3))$ guarantees that $\mu \in \llbracket P_1 \text{ AND } (P_2 \text{ OPT } P_3) \rrbracket_G^s$ as required.

Case 2: $\mu = \mu_1 \cup \mu_2$ for $\mu_i \in \llbracket P_i \rrbracket_G^s$, such that $\mu \in \text{Adm}(G, P_1 \text{ AND } P_2)$, $\mu \in \text{Adm}(G, (P_1 \text{ AND } P_2) \text{ OPT } P_3)$, and there is no $\mu_3 \in \llbracket P_3 \rrbracket_G^s$ such that $\mu \sim \mu_3$. By Lemma 2 $\mu_2 \in \text{Adm}(G, P_2 \text{ OPT } P_3)$ holds. Also, since $P_1 \text{ AND } (P_2 \text{ OPT } P_3)$ is well-designed, there is no $\mu'_3 \in \llbracket P_3 \rrbracket_G^s$ such that $\mu_2 \sim \mu'_3$; otherwise there would be a variable $?X \in \text{dom}(\mu_1) \cap \text{dom}(\mu_3)$, but $?X \notin \text{dom}(\mu_2)$, with $\mu_1(?X) \neq \mu_3(?X)$, which by the fact that the pattern is well-designed would imply that $?X$ appears in P_2 and by Lemma 1 would lead to a contradiction. Putting the facts above together with the equality $\text{Adm}(G, (P_1 \text{ AND } P_2) \text{ OPT } P_3) = \text{Adm}(G, P_1 \text{ AND } (P_2 \text{ OPT } P_3))$ we obtain that $\mu \in \llbracket (P_1 \text{ AND } P_2) \text{ OPT } P_3 \rrbracket_G^s$ as required.

Equivalence (3) follows from (2) and commutativity of AND.

Equivalence (4). We concentrate again on forward direction. Let $\mu \in \llbracket (P_1 \text{ OPT } P_2) \text{ OPT } P_3 \rrbracket_G^s$ for a consistent graph G , and legal patterns P_1, P_2 and P_3 such that the both sides of the equivalence are well-designed. We have four possibilities. We prove the first two, omitting the very similar proofs for other two for brevity.

Case 1: $\mu = \mu_1 \cup \mu_2 \cup \mu_3$ for $\mu_i \in \llbracket P_i \rrbracket_G^s$, such that $\mu_1 \cup \mu_2 \in \text{Adm}(G, P_1 \text{ OPT } P_2)$ and $\mu \in \text{Adm}(G, (P_1 \text{ OPT } P_2) \text{ OPT } P_3)$. By Lemma 2 the last inclusion implies $\mu_1 \cup \mu_3 \in \text{Adm}(G, P_1 \text{ OPT } P_3)$, which, together with $\text{Adm}(G, (P_1 \text{ OPT } P_2) \text{ OPT } P_3) = \text{Adm}(G, (P_1 \text{ OPT } P_3) \text{ OPT } P_2)$ guarantees $\mu \in \llbracket (P_1 \text{ OPT } P_3) \text{ OPT } P_2 \rrbracket_G^s$.

Case 2: $\mu = \mu_1 \cup \mu_2$ for $\mu_i \in \llbracket P_i \rrbracket_G^s$, where $\mu \in \text{Adm}(G, P_1 \text{ OPT } P_2)$, $\mu \in \text{Adm}(G, (P_1 \text{ OPT } P_2) \text{ OPT } P_3)$, and there is no $\mu_3 \in \llbracket P_3 \rrbracket_G^s$ with $\mu \sim \mu_3$. By Lemma 2 we have $\mu_1 \in \text{Adm}(G, P_1 \text{ OPT } P_3)$. Similar to Case 2 in the proof of (2), since $(P_1 \text{ OPT } P_2) \text{ OPT } P_3$ is well-designed, by Lemma 1 there is no $\mu'_3 \in \llbracket P_3 \rrbracket_G^s$ such that $\mu_1 \sim \mu'_3$. But then, since $\text{Adm}(G, (P_1 \text{ OPT } P_2) \text{ OPT } P_3) = \text{Adm}(G, (P_1 \text{ OPT } P_3) \text{ OPT } P_2)$ we have $\mu \in \llbracket (P_1 \text{ OPT } P_3) \text{ OPT } P_2 \rrbracket_G^s$ as required.

Case 3: $\mu = \mu_1 \cup \mu_3$ for $\mu_i \in \llbracket P_i \rrbracket_G^s$, such that $\mu_1 \in \text{Adm}(G, P_1 \text{ OPT } P_2)$, $\mu \in \text{Adm}(G, (P_1 \text{ OPT } P_2) \text{ OPT } P_3)$, and there is no $\mu_2 \in \llbracket P_2 \rrbracket_G^s$ such that $\mu_1 \sim \mu_2$.

Case 4: $\mu \in \llbracket P_1 \rrbracket_G^s$, $\mu \in \text{Adm}(G, P_1 \text{ OPT } P_2)$, $\mu \in \text{Adm}(G, (P_1 \text{ OPT } P_2) \text{ OPT } P_3)$, and neither $\mu_2 \in \llbracket P_2 \rrbracket_G^s$ such that $\mu \sim \mu_2$, nor $\mu_3 \in \llbracket P_3 \rrbracket_G^s$ such that $\mu \sim \mu_3$. \square

6 Conclusion

In this paper, we have proposed a novel semantics for optional matching in SPARQL under entailment regimes where unbound variables in answer mappings are naturally interpreted as “null” values. Our *strict* semantics has been designed to deal in a faithful way with the “lack of regular, complete structures in RDF graphs” and hence with the fundamental incompleteness of information on the Semantic Web [1]. We believe that both strict and normative semantics are valid, but one may be more appropriate than the other in certain applications. Both semantics are compatible at a fundamental level and it would be possible to exploit them in the same application by letting users commit to one or the other explicitly when posing queries. Integrating them in a clean way from a syntactic point of view is more tricky, and it is something we leave for future investigation.

Acknowledgements. This work was supported by the Royal Society, and the EPSRC projects Score!, Exoda, and MaSI³.

References

1. Prud'hommeaux, E., Seaborne, A.: SPARQL query language for RDF. W3C Recommendation (2008), <http://www.w3.org/TR/rdf-sparql-query/>
2. Pérez, J., Arenas, M., Gutierrez, C.: Semantics and complexity of SPARQL. *ACM Trans. Database Syst.* 34(3) (2009)
3. Angles, R., Gutierrez, C.: The expressive power of SPARQL. In: Sheth, A.P., Staab, S., Dean, M., Paolucci, M., Maynard, D., Finin, T., Thirunarayan, K. (eds.) *ISWC 2008*. LNCS, vol. 5318, pp. 114–129. Springer, Heidelberg (2008)
4. Schmidt, M., Meier, M., Lausen, G.: Foundations of SPARQL query optimization. In: *ICDT*, pp. 4–33 (2010)
5. Arenas, M., Pérez, J.: Querying semantic web data with SPARQL. In: *PODS*, pp. 305–316 (2011)
6. Letelier, A., Pérez, J., Pichler, R., Skritek, S.: Static analysis and optimization of semantic web queries. *ACM Trans. Database Syst.* 38(4), 25 (2013)
7. Polleres, A.: From SPARQL to rules (and back). In: *WWW*, pp. 787–796 (2007)
8. Motik, B., Patel-Schneider, P.F., Parsia, B.: OWL 2 Web Ontology Language Structural Specification and Functional-style Syntax. W3C Recommendation (2012), <http://www.w3.org/TR/owl2-syntax/>
9. W3C SPARQL Working Group: SPARQL 1.1 Query language. W3C Recommendation (2013), <http://www.w3.org/TR/sparql11-query/>
10. Glimm, B., Ogbuji, C.: SPARQL 1.1 Entailment Regimes. W3C Recommendation (2013), <http://www.w3.org/TR/sparql11-entailment/>
11. Glimm, B., Krötzsch, M.: SPARQL beyond subgraph matching. In: Patel-Schneider, P.F., Pan, Y., Hitzler, P., Mika, P., Zhang, L., Pan, J.Z., Horrocks, I., Glimm, B. (eds.) *ISWC 2010, Part I*. LNCS, vol. 6496, pp. 241–256. Springer, Heidelberg (2010)
12. Pérez, J., Arenas, M., Gutierrez, C.: Semantics and complexity of SPARQL. In: Cruz, I., Decker, S., Allemang, D., Preist, C., Schwabe, D., Mika, P., Uschold, M., Aroyo, L.M. (eds.) *ISWC 2006*. LNCS, vol. 4273, pp. 30–43. Springer, Heidelberg (2006)
13. Motik, B., Cuenca Grau, B., Horrocks, I., Wu, Z., Fokoue, A., Lutz, C.: OWL 2 Web Ontology Language: Profiles (October 27, 2009), <http://www.w3.org/TR/owl2-profiles/>
14. Chekol, M.W., Euzenat, J., Genevès, P., Layaïda, N.: SPARQL query containment under RDFS entailment regime. In: Gramlich, B., Miller, D., Sattler, U. (eds.) *IJCAR 2012*. LNCS, vol. 7364, pp. 134–148. Springer, Heidelberg (2012)
15. Chekol, M.W., Euzenat, J., Genevès, P., Layaïda, N.: SPARQL query containment under *SHI* axioms. In: *AAAI* (2012)
16. Wudage Chekol, M., Euzenat, J., Genevès, P., Layaïda, N.: Evaluating and benchmarking SPARQL query containment solvers. In: Alani, H., et al. (eds.) *ISWC 2013, Part II*. LNCS, vol. 8219, pp. 408–423. Springer, Heidelberg (2013)