Generalized Multi-Context Systems

Shahab Tasharrofi, Eugenia Ternovska
Simon Fraser University, email: {sta44,ter}@cs.sfu.ca

Abstract

Multi-context systems (MCSs) define a versatile framework for integrating and reasoning about knowledge from different (heterogeneous) sources. In MCSs, different types of non-monotonic reasoning are characterized by different semantics such as equilibrium semantics and grounded equilibrium semantics [Brewka and Eiter, 2007].

We introduce a novel semantics of MCSs, a supported equilibrium semantics. Our semantics is based on a new notion of support. The "strength" of supports determines a spectrum of semantics that, in particular, contains the equilibrium and grounded equilibrium semantics. In this way, our supported equilibrium semantics generalizes these previously defined semantics. Moreover, the "strength" of supports gives us a measure to compare different semantics of MCSs.

Introduction

Plato defined knowledge as justified true beliefs but it took thousands of years before this philosophical standpoint was used to form the intuitionistic view of Brouwer and the S4 provability logic of Godel. Since then, the progress towards more justified belief system have accelerated. Artemov (Artemov 1995) introduced Logic of Proofs (LP) as a formalization that internalizes justifications for statements and several justification logics were defined based on LP by Brezhnev (Brezhnev 2001). Knowledge representation and reasoning (KRR) field has also both benefited from and contributed to this direction. Fitting (Fitting 2005) defined an epistemic semantics for LP and Cabalar (Cabalar 2011) used LP to define causal logic programming under stable model semantics. Inspired by the utility that justification systems bring to true logical statements in general and to non-monotonic logic in particular, this paper uses justifications to remove unintended models of multi-context systems (MCSs) (Brewka and Eiter 2007) in a new semantics that we call *supported*. A MCS is a collection of contexts that are linked using bridge rules. Each context has its own way of representing knowledge (i.e., its own syntax and semantics). Bridge rules define how knowledge can be transferred between contexts. In MCSs, a model has the form of a collection of belief sets (called a *belief state*). The semantics of MCSs that inters us here are as follows.

1. *Equilibrium semantics (ES)* defines intended models as exactly those belief states that, if viewed operationally, remain unchanged after first applying bridge rules and then applying contexts, hence the name of an equilibrium.
2. *Minimal equilibrium semantics (MES)* defines intended models as those equilibriums that are also minimal.
3. *Grounded equilibrium semantics (GES)* defines intended models as the minimal equilibriums of a positive MCS obtained by *reducing* the original MCS. Reducing MCSs is similar (both methodically and intent-wise) to the procedure Gelfond and Lifschitz (Gelfond and Lifschitz 1988) use to define stable model semantics. Unfortunately, this also means that GES is not applicable to MCSs with non-reducible logics, that is logics for which the reduction procedure cannot be applied.

The above semantics are motivated by everyday reasoning about a collection of contexts or agents. In MCSs, some knowledge is shared between different knowledge bases, while some knowledge is kept private/confidential.

Note that that justifications and, in particular, avoiding self-justifications was the main motivation behind the introduction of grounded equilibrium semantics for MCSs (Brewka and Eiter 2007). However, grounded equilibrium semantics (GES) is defined over MCSs in which all contexts are reducible. Thus, even one non-reducible context is enough to render GES non-applicable. Below, example 1 demonstrates this point.

The main goal of this work is to extend MCSs with the notions of justifications and support so that

- a spectrum of semantics for MCSs is obtained by varying one parameter, the support;
- the existing semantics of MCSs are in that spectrum.

We relate variations in support with the level of selectivity of the semantics, and argue that intermediate points in the spectrum are just as useful as the previously known semantics, – ES, MES, GES. We also show the usefulness of our new justification-aware semantics in diagnosis of MCSs.

Recall that normal answer set programs are sets of rules of form $h \leftarrow b_1, \cdots, b_k, \textbf{not}\, b_{k+1}, \cdots, \textbf{not}\, b_n$ such that $h, b_1, \cdots, b_n$ are first order literals (i.e., negated or non-negated atoms with variables). Also, recall that S is a stable model of a normal answer set program P if and only if S is consistent and deductively closed under the positive program obtained by reducing P according to S.

Example 1. *Let $M := (C_1, C_2, C_3)$ be*[1] *a multi-context system with $C_i := (L_i, kb_i, br_i)$ (for $i \in \{1, 2, 3\}$), L_1, L_2 be the logic of normal answer set programs under stable model semantics, L_3 be first-order logic (with $BS_3 := \{\{\}, \{comedian(jk)\}\}$ as possible belief sets), and kb_i, br_i be as follows:*

$$
\begin{aligned}
kb_1 &:= \left\{ \begin{array}{l} forbes400(bg). \\ wealthy(X) \leftarrow forbes400(X). \\ wealthy(X) \leftarrow celebrity(X). \end{array} \right\} \\
br_1 &:= \{celbrity(X) \leftarrow C_2 : famous(X).\} \\
kb_2 &:= \left\{ \begin{array}{l} actor(bp). \\ famous(X) \leftarrow actor(X). \end{array} \right\} \\
br_2 &:= \left\{ \begin{array}{l} famous(X) \leftarrow C_1 : wealthy(X). \\ famous(X) \leftarrow C_3 : comedian(X). \end{array} \right\}, \\
kb_3 &:= \{\} \\
br_3 &:= \{\bot \leftarrow C_3 : comedian(X), \textbf{not}\ C_2 : famous(X).\}
\end{aligned}
$$

with X ranging over the four possibilities of "bg" (stands for Bill Gates), "bp" (stands for Brad Pitt), "jk" (stands for Jimmy Kimmel), and "aj" (stands for Average Joe).

Note that MCS M can only use the equilibrium semantics because context C_3 is not reducible. According to equilibrium semantics, M has equilibriums $S_1, \cdots, S_6$ as follows:

for $i \in \{1, \cdots, 6\} : S_i := (bs_1^i, bs_2^i, bs_3^i)$ where,

$$
\begin{aligned}
bs_1^1 &:= \left\{ \begin{array}{l} forbes400(bg), wealthy(bg), wealthy(bp), \\ celebrity(bg), celebrity(bp) \end{array} \right\}, \\
bs_1^2 &:= bs_1^1 \cup \{wealthy(aj), celebrity(aj)\}, \\
bs_1^3 &:= bs_1^4 := bs_1^1 \cup \{wealthy(jk), celebrity(jk)\}, \\
bs_1^5 &:= bs_1^6 := bs_1^2 \cup bs_1^3, \\[1em]
bs_2^1 &:= \{actor(bp), famous(bg), famous(bp)\}, \\
bs_2^2 &:= bs_2^1 \cup \{famous(aj)\}, \\
bs_2^3 &:= bs_2^4 := bs_2^1 \cup \{famous(jk)\}, \\
bs_2^5 &:= bs_2^6 := bs_2^2 \cup bs_2^3, \\[1em]
bs_3^1 &:= bs_3^2 := bs_3^3 := bs_3^5 := \{\}, \\
bs_3^4 &:= bs_3^6 := \{comedian(jk)\}.
\end{aligned}
$$

If the names of this example are to be taken literally, among all the six equilibrium models above, only S_4 is a reasonable belief state. This is because, first, "Average Joe," by definition, is not famous, wealthy, or a celebrity, and, second, "Jimmy Kimmel" is a famous comedian and a wealthy celebrity.

Note that, in Example 1, grounded equilibrium semantics is not applicable because context C_3 uses the non-reducible first-order logic. However, since kb_3 is empty, one might wonder what would have happened if L_3 was a reducible logic. Example 2 considers that situation.

Example 2. *In Example 1, let M be as before except that L_3 is the reducible logic of normal answer set programs under stable model semantics. Then, equilibrium models of M are as before (i.e., $S_1, \cdots, S_6$) but, now, M also has exactly one grounded equilibrium: S_1.*

[1]We use := to represent "denotes" or "equals by definition".

Note that, as we discussed before, in Example 1, S_4 is the only reasonable equilibrium model. However, as we saw in Examples 1 and 2, neither the equilibrium semantics nor the grounded equilibrium semantics can capture the set of intended models. We saw that equilibrium semantics accepts too many equilibriums, i.e., all of $S_1, \cdots, S_6$, and grounded equilibrium semantics (even if definable) rejects our intended equilibrium model S_4.

Summary of the Paper – First, this paper defines a novel semantics for MCSs that we call supported equilibrium semantics. We show that our semantics generalizes both the equilibrium semantics and the grounded equilibrium semantics under a natural instantiation of justification functions for contexts. For instance, we show that, for our running Example 1, depending on how justification functions are instantiated, our supported equilibrium semantics can either work similar to equilibrium semantics and accept all equilibrium $S_1, \cdots, S_6$, or it can work similar to grounded equilibrium semantics and accept only S_1.

Second, this paper also shows that our supported equilibrium semantics characterizes many interesting cases that cannot be captured by either the equilibrium semantics or by the grounded equilibrium semantics. For instance, we show that under a very natural and easy instantiation of justification functions, our supported equilibrium semantics accepts both equilibrium models S_1 and S_4 of MCS M in Example 1. That is, using such natural justifications, supported equilibrium semantics is neither as relaxed as equilibrium semantics that accepts all $S_1, \cdots, S_6$, nor as restrictive as grounded equilibrium semantics that rejects our intended equilibrium model S_4. In this sense, supported equilibrium semantics properly generalizes its predecessors and better captures our intended equilibrium models.

Finally, the last section of this paper is dedicated to a short discussion on a plausible application of supported equilibrium semantics to detect and repair faulty MCSs. There, we discuss how supported equilibrium semantics helps us to remove equilibrium model S_1 of MCS M from Example 1. This type of debugging and repairing has not previously been possible in MCSs, and its future formal investigation will help MCSs be more suited to modelling practical KRR applications.

Contributions

Our contributions in this paper are as follows:

Supported Equilibrium Semantics – We introduce a novel semantics for multi-context systems by adding the concept of support to all logics and then extending the semantics of equilibrium models so that the concept of support is respected when choosing intended equilibria of a MCS.

Unifying Equilibrium Semantics with Grounded Equilibrium Semantics – We show that our new semantics of supported equilibria naturally extends both the equilibrium semantics and the grounded equilibrium semantics. That is, we show that, by carefully choosing supports of logics, both equilibrium semantics and grounded equilibrium semantics become special cases of supported equilibrium semantics.

Broadening Applications of MCSs – Our supported equilibrium semantics extends the applicability of MCSs to applications in which some/all of justifications cannot be revealed due to privacy, confidentiality or other reasons. This is achieved by allowing contexts to justify their beliefs at varying degrees, and by designing a semantics that works independently of how detailed justifications are. Of course, the more justifications are provided, the better our semantics becomes.

Better Diagnosis and Repair – We show that supported equilibrium semantics paves the way for better diagnoses and repairs in a faulty MCS.

Background

This section briefly reviews multi-context systems using the exposition of (Fink, Ghionna, and Weinzierl 2011) and (Brewka and Eiter 2007). We also use the two following notations here and in other places throughout this paper:

Notation 1 (Negating a Set). *For a set of belief literals X, we use "**not** X" to denote a set that contains the negation of literals in X, i.e.,* **not** $X := \{\textbf{not}\, b \mid b \in X\}$.

Notation 2 (Un-pairing). *For a pair $P := (X, Y)$, we use* $\text{fst}(P)$ *to denote X and* $\text{snd}(P)$ *to denote Y.*

In MCSs (Brewka and Eiter 2007), a *logic* L is a triple $L := \langle KB_L, BS_L, ACC_L\rangle$, where KB_L is a set of knowledge bases (syntactic part of L), BS_L is a set of belief sets (semantic part of L), and $ACC_L : KB_L \mapsto 2^{BS_L}$ maps each knowledge base to a set of acceptable belief sets (the semantics of L). A *multi-context system MCS* $:= (C_1, \cdots, C_n)$ is a collection of contexts $C_i := (L_i, kb_i, br_i)$ with logic L_i, knowledge base $kb_i \in KB_{L_i}$ and bridge rules br_i. In MCSs, bridge rule $r \in br_i$ has the following form:

$$\begin{array}{rcl}(i : s) & \leftarrow & (c_1 : p_1), \cdots, (c_j : p_j),\\ & & \textbf{not}\, (c_{j+1} : p_{j+1}), \cdots, \textbf{not}\, (c_m : p_m).\end{array} \quad (1)$$

where $hd(r) := s$; $body^+(r) := \{(c_k : p_k) \mid 1 \le k \le j\}$; $body^-(r) := \{(c_k : p_k) \mid j+1 \le k \le m\}$; and, $body(r) := body^+(r) \cup (\textbf{not}\, body^-(r))$.

A *belief state* $S := (S_1, ..., S_n)$ is a collection of belief sets, i.e., $S_i \in BS_{L_i}$. Also, we use $X(S)$ for the disjoint union of beliefs in all S_i's, i.e., $X(S) := \{(i : b) \mid 1 \le i \le n \text{ and } b \in S_i\}$. A bridge rule r of form (1) is *applicable* wrt. S, denoted by $S \models body(r)$, iff $p_l \in S_{c_l}$ for $1 \le l \le j$ and $p_l \notin S_{c_l}$ for $j < l \le m$. We define $app_i(S) := \{hd(r) \mid r \in br_i \wedge S \models body(r)\}$ to obtain heads of all applicable bridge rules of context C_i. Belief state S is an *equilibrium* of *MCS* if, for all i, $S_i \in ACC_{L_i}(kb_i \cup app_i(S))$.

Logic L is *monotone* if, for all $kb, kb' \in KB_L$, (1) $ACC(kb)$ is a singleton set $\{S\}$, and, (2) if $ACC(kb) = \{S\}$, $ACC(kb') = \{S'\}$, and $kb \subseteq kb'$ then $S \subseteq S'$. Also, L is *reducible* if (1) subset $KB^*_L \subseteq KB_L$ exists s.t. $L^* := \langle KB^*_L, BS_L, ACC_L\rangle$ is monotone, and, (2) reduction function $red_L : KB_L \times BS_L \mapsto KB^*_L$ exists s.t. a. $red_L(kb, S) = kb$ if $kb \in KB^*_L$, b. $red_L(kb, S') \subseteq red_L(kb, S)$ if $S \subseteq S'$, and, c. $S \in ACC_L(k)$ iff $ACC_L(red_L(k, S)) = \{S\}$.*Context $C := (L, kb, br)$ is reducible* if L is reducible and $red_L(kb \cup H, S) = red_L(kb, S) \cup H$ for all $H \subseteq \{hd(r) \mid r \in br\}$. MCS M is reducible if all of its contexts are reducible.

A reducible MCS $M := (C_1, \ldots, C_n)$ is *definite* if all bridge rules r of M are positive, i.e., $body^-(r) = \emptyset$, and, for all i, $kb_i \in KB^*_{L_i}$. Definite MCSs guarantee monotonic inference and, thus, always have a unique minimal equilibrium (Brewka and Eiter 2007). Also, for reducible MCS $M := (C_1, \cdots, C_n)$, reduction of M under belief state S, denoted by M^S, is a definite MCS $M' := (C^S_1, \cdots, C^S_n)$ where $C^S_i := (L_i, red_i(kb_i, S_i), br^S_i)$ and $br^S_i := \{hd(r) \leftarrow body^+(r) \mid r \in br_i \text{ and } S \models body(r)\}$. Finally, S is a *grounded equilibrium* of reducible MCS M if S is the unique minimal equilibrium of the M^S. For non-reducible MCSs, grounded equilibria are not defined.

Note that other semantics for MCSs and their extensions (such as managed MCSs (Brewka et al. 2011)) exist, but are not considered here because they are not relevant to this paper.

Generalized MCSs

Examples 1 and 2 showed a case where both equilibrium and grounded equilibrium semantics fail to capture our intended models. This section first defines justification functions that justify beliefs of an acceptable belief set with other beliefs of that belief set and formulas in the knowledge base. For example, the belief that a product is unavailable for sale is justified by a combination of other beliefs and formulas as follows. Among the beliefs that contribute towards justifying unavailability of this product, we can mention the belief that all instance of this product in the store are marked as "sold" and that this product is also unavailable in the warehouse. Moreover, a formula also contributes as a justification of unavailability of this product: it says that if a product is available for sale, then it should be either in the warehouse or in the store without being labeled as "sold."

Secondly, in this section, we lift the notion of justification from a particular acceptable belief set in a particular logic to contexts in general, and then to multi-context systems. We show the naturality of our definitions using our running example and leave the formal expressiveness results to the next section.

Support for Logics

In order to define supported equilibrium semantics, we need to introduce the concept of support at the level of logics. To this end, we define *justifications* for acceptable belief sets and, then, extend them to *support for logics*. While we use very natural definitions, we also heavily use our running example to give more intuitions about our definitions.

Consider logic $L := \langle KB, BS, ACC\rangle$, belief set $bs \in BS$, and knowledge base $kb \in KB$ such that bs is an acceptable belief set for kb in L. By a *justification* for bs, we mean a possible explanation of why belief in bs are believed.

Definition 1 (Justification). *Let $L := \langle KB, BS, ACC\rangle$, $kb \in KB$ and $bs \in ACC(kb)$. Then, function $j : bs \mapsto (\mathcal{P}(bs) \times \mathcal{P}(kb))$ is called a* justification *for bs if j is non-circular, i.e., a well-ordering $<_{bs}$ on beliefs in bs exists such that, for all $b, b' \in bs$ with $b \in \text{fst}(j(b'))$, we have $b <_{bs} b'$.*

A justification function as in Definition 1 provides possible explanations for belief sets in a non-circular way. In the following, Example 3 shows a natural justification function in the context of our running example.

Example 3. *Consider multi-context system M from Example 1 and equilibrium models $S_1, \cdots, S_6$. By definition of equilibrium models, we know that, for $i \in \{1, 2, 3\}$ and $j \in \{1, \cdots, 6\}$, belief set bs_i^j is an acceptable belief set for knowledge base $kb_i \cup app_i(S_j)$, i.e., $bs_i^j \in ACC_i(kb_i \cup app_i(S_j))$. In this example, we give two natural justification functions j_1 and j_2 for two belief sets bs_1^1 and bs_1^2 respectively. Note that, by Definition 1, each j_i (non-circularly) maps bs_1^i to a pair consisting of a subset of bs_1^i and a subset of $kb_1 \cup app_1(S_i)$, i.e., $j_i : bs_1^i \mapsto (\mathcal{P}(bs_1^i) \times \mathcal{P}(kb_1 \cup app_1(S_i)))$. Our j_i's are as follows:*

$$
\begin{aligned}
&j_{1,2}(forbes400(bg)) := (\{\}, \{forbes400(bg)\}),\\
&j_{1,2}(wealthy(bg)) := (\{forbes400(bg)\},\\
&\qquad \{wealthy(bg) \leftarrow forbes400(bg)\}),\\
&j_{1,2}(wealthy(bp)) := (\{celebrity(bp)\},\\
&\qquad \{wealthy(bp) \leftarrow celebrity(bp)\}),\\
&j_2(wealthy(aj)) := (\{celebrity(aj)\},\\
&\qquad \{wealthy(aj) \leftarrow celebrity(aj)\}),\\
&j_{1,2}(celebrity(bg)) := (\{\}, \{celebrity(bg)\}),\\
&j_{1,2}(celebrity(bp)) := (\{\}, \{celebrity(bp)\}),\\
&j_2(celebrity(aj)) := (\{\}, \{celebrity(aj)\}).
\end{aligned}
$$

Note that j_1 and j_2 agree on their shared domain, i.e., bs_1^1 and that they both have very intuitive meanings. For instance, $j_1(forbes400(bg)) = (\{\}, \{forbes400(bg)\})$ means that, our belief in $forbes400(bg)$ is independent of all other beliefs and is justified only by a fact from the knowledge base, i.e., $forbes400(bg)$. It is noteworthy that, here, $forbes400(bg)$ appears once as a belief and another time as a formula. Thus, $forbes400(bg)$ does not constitute a self-justification because, for self-justifications to occur, a belief should (directly or indirectly) depend on itself (as a belief).

In Example 3, justification coincides with the consequence relation. That is, if $j_i(b) = (bs, kb)$, $bs \subseteq bs'$, $kb \subseteq kb'$, and $bs' \in ACC_1(kb')$ then $b \in bs'$. However, we want to emphasize that, despite the proximity of the two notions of consequence and justification, the latter is much more flexible than the former. Therefore, as shown later in this paper, the same logic can have many different justification functions (unlike consequence relation which is closely tied to the semantics of a logic). Moreover, since justifications are more flexible, as this paper shows, they can be used to characterize a range of different semantics for the same MCS (which would have been impossible using consequence relation because of their rigidity).

Now, we use justifications to define the support for a logic. Intuitively, $Sup_L(kb, bs)$ denotes a (usually non-exhaustive) set of possible *justifications* for belief set bs.

Definition 2 (Support for a Logic L). *Let $L := \langle KB_L, BS_L, ACC_L\rangle$, $kb \in KB_L$ and $bs \in BS_L$. Then, $Sup_L(kb, bs)$ is a set of justifications for bs according to kb such that if $bs \notin ACC_L(kb)$ then $Sup_L(kb, bs) = \emptyset$.*

Example 4 applies Definition 2 to our running example:

Example 4. *Continuing Example 1, we define Sup_{L_1} to be such that $Sup_{L_1}(kb_1, bs_1^2) = \emptyset$ (because $bs_1^2 \notin ACC_1(kb_1)$) and $Sup_{L_1}(kb_1 \cup app_1(S_2), bs_1^1) = \{j_2, j_3\}$ where j_2 is from Example 3 and $j_3 : bs_1^2 \mapsto (\mathcal{P}(bs_1^2) \times \mathcal{P}(kb_1 \cup app_1(S_2)))$ is as below:*

$$
\begin{aligned}
&j_3(wealthy(bg)) := (\{celebrity(bg)\},\\
&\qquad \{wealthy(bg) \leftarrow celebrity(bg)\}),\\
&j_3(b) := j_2(b) \quad (\text{for } b \in bs_1^2 \setminus \{wealthy(bg)\}).
\end{aligned}
$$

As discussed before, justifications do not uniquely correspond to a logic L of a context. Moreover, as seen in Example 4, different justification functions contribute towards defining different supports for a logic. Among all possible supports for a logic, two trivial but important ones are the *unit support* and the *empty support* (defined below). Unit and empty supports are definable for all logics and, as will be seen later on, they respectively form the minimum and the maximum of a lattice on supports for a logic.

Definition 3 (Unit and Empty Supports). *Consider logic L and two supports uSup_L, eSup_L for L such that, for all kb and bs: (1) $\mathrm{eSup}_L(kb, bs) = \emptyset$, and, (2) $\mathrm{uSup}_L(kb, bs) = \{u^{bs}\}$ where, for all $b \in bs$, $u^{bs}(b) = (\{\}, \{\})$. Also, we call uSup_L and eSup_L respectively as the* unit support for L *and the* empty support for L.

Note that, neither the unit support nor the empty support for a logic does not provide any information about the internal knowledge base. As we show later on, the unit and empty support for a logic define the two ends of an spectrum in which the unit support is the most relaxed support for a logic and empty support is the most rigid support for a logic. In other words, we show later on in this paper that (1) unit support for a logic corresponds to the case where all equilibriums are supported, and, (2) empty support corresponds to the case that no equilibriums is supported.

While empty and unit supports are definable for all logics, they do not give us any insight into why something is (or is not) believed. Thus, one might wonder if more insightful notions of support can be developed for interesting KR logical frameworks. Here, we emphasize that, for many interesting logical frameworks, such as many non-monotonic logics, a more rigid notion of support has already been developed (Lifschitz 2010; Tasharrofi 2013). The following example uses a well-established notion of support for normal logic programs in combination with MCS M from Example 1 to show how supports provide insight into a knowledge base.

Example 5. *Consider context C_1 of Example 1. Since kb_1 is a normal logic program under answer set semantics, one can use a support for logic L of C_1 that is based on the notion of support as defined in (Lifschitz 2010). Using this support function, both justification functions j_1 and j_3 from Examples 3 and 4 belong to $Sup_L(bs_1^1, kb_1 \cup app_1(S_1))$.*

Support for Contexts

Using the definition of supports on the level of logic, we define support on the level of contexts as in Definition 4 that follows. Unlike supports for logics that worked with syntactic (knowledge base formulas) and semantic (beliefs) objects simultaneously, in Definition 4, supports at the level of contexts work solely on semantic objects (beliefs).

Definition 4 (Support for Contexts). *Consider MCS $M := (C_1, \cdots, C_n)$, its context $C_i := \langle L_i, kb_i, br_i\rangle$, and its belief state $S := (S_1, \cdots, S_n)$. Also, let Sup_{L_i} be the support function for logic L_i and recall that $X(S) := \{(i : b) \mid b \in S_i\}$. Support of belief set S_i (from context C_i) under S, denoted by Sup_i^S, is the set of functions $f : S_i \mapsto \mathcal{P}(X(S))$ that are computed by taking a function $g \in Sup_{L_i}(kb_i \cup app_i(S), S_i)$ and tracing the reason for the inclusion of the knowledge that comes from bridge rules. More formally, a function f is included in Sup_i^S if and only if functions $g \in Sup_{L_i}(kb_i \cup app_i(S), S_i)$ and $R : S_i \mapsto \mathcal{P}(br_i)$ exist such that, for all $b \in S_i$, we have:*

$$f(b) := \{i : b' \mid b' \in \text{fst}(g(b))\} \cup \bigcup_{r \in R(b)} body^+(r).$$

and $R(b) \subseteq \{r \mid r \in br_i \text{ and } S \models body(r)\}$ is a minimal subset of applicable bridge rules that justifies the knowledge that comes from bridge rules, i.e., for all formulas $k \in (\text{snd}(g(b)) \setminus kb)$, a rule $r \in R(b)$ exists with $hd(r) = k$.

Note that, in Definition 4, if $\text{snd}(g(b)) \setminus kb = \emptyset$ for some belief b, then $R(b) = \emptyset$. Intuitively, it means that support from other contexts is required only when existing knowledge of a context is not sufficient for supporting a belief. Following example applies the notion of support at the context-level to our running Example 1.

Example 6. *Consider MCS M of Example 1 and equilibrium model S_2 of M. Also, assume that support for logic of C_1 is as in Example 5. Using that support, let us compute one of the functions $f \in Sup_1^{S_2}$. According to Definition 4, such a function f supports beliefs in bs_1^2. Also, by Definition 4, to construct such a function f, we need function $g \in Sup_{L_1}(kb_1 \cup app_1(S_2), bs_1^2)$. By Example 4, we know $Sup_{L_1}(kb_1 \cup app_1(S_2), bs_1^2) = \{j_2, j_3\}$, thus either $g = j_2$ or $g = j_3$. In this example, we take $g = j_2$.*

By Definition 4, to construct f, formulas that are justified (generated) by bridge rules should be justified. In our example, these are formulas of form $celebrity(X)$. Fortunately, each of these formulas appear as the head of has exactly one bridge rule. Therefore, the minimal set $R(b)$ in the case of each of these formulas contains exactly one rule with that belief in its head. For example, when $b = celebrity(aj)$, $R(b)$ only contains rule "$celebrity(aj) \leftarrow C_2 : famous(aj)$". Thus, f is as follows:

$f(celebrity(X)) = \{2 : famous(X)\}$ *(for all X),*
$f(wealthy(X)) = \{1 : celebrity(X)\}$ *(for $X \neq bg$),*
$f(wealthy(bg)) = \{1 : forbes400(bg)\}$,
$f(forbes400(bg)) = \emptyset$.

The numbers 1 and 2 in the function f above refer to contexts C_1 and C_2 respectively.

Example 6 shows how support functions can be extended beyond the boundaries of knowledge bases and belief sets of a logic and into the beliefs from the belief sets of other contexts. According to Definition 4 and as shown in Example 6, this task is achieved by following bridge rules of a context.

Supported Equilibrium Semantics

Now, we use supports at the level of contexts to define supported equilibrium semantics for MCSs. In the following, Definition 5 gives our main notion of a *supported equilibrium*. Informally speaking, a belief state is called a supported equilibrium if all beliefs are well-justified, i.e., they are justified and nothing justifies itself (either directly or indirectly). Similar to Definition 1, self-justifications are avoided by requiring the existence of a well-ordering on the beliefs.

Definition 5 (Supported Equilibrium). *A belief state $S := (S_1, \cdots, S_n)$ of MCS is a* supported equilibrium w.r.t. *$(Sup_{L_1}, \cdots, Sup_{L_n})$ if functions $f_1 \in Sup_1^S, \cdots, f_n \in Sup_n^S$ and well-founded strict partial ordering $<$ on $X(S)$ exist s.t. if $p \in S_i$ and $(j : q) \in f_i(p)$ then $(j : q) < (i : p)$.*

First, note that Definition 5 does not put any special requirement on contexts and works for all contexts and all supports. Therefore, unlike grounded equilibrium semantics of (Brewka and Eiter 2007), introspection in supported equilibrium semantics does not come at the cost of excluding non-reducible contexts.

Second, note that Definition 5 tests a belief state for being supported but not for being an equilibrium. So, one might reasonably suspect that a belief state S might exist such that S is a supported equilibrium (according to Definition 5) but not an equilibrium (according to the original definition of equilibrium semantics (Brewka and Eiter 2007)). However, the following Theorem 1 states that if S is a supported equilibrium then it has to be an equilibrium as well. Therefore, the term "supported equilibrium" is indeed an appropriate and reasonable name for belief states that satisfy the condition of Definition 5.

Theorem 1 (Supported Equilibria $\subseteq$ Equilibria). *Let $M := (C_1, \cdots, C_n)$ and also let $S := (S_1, \cdots, S_n)$ be a supported equilibrium of M w.r.t. $(Sup_{L_1}, \cdots, Sup_{L_n})$. Then, S is also an equilibrium of M.*

Proof. Assume that S is not an equilibrium of M. Then, context C_i should exist such that $S_i \notin ACC(kb_i \cup app_i(S))$. Therefore, by Definition 2, we know that $Sup_{L_i}(kb_i \cup app_i(S), S_i) = \emptyset$. Thus, by Definition 4, $Sup_i^{L_i}$ is also empty. Hence, by Definition 5, S cannot be a supported equilibrium of M which contradicts our assumption. So, S has to be an equilibrium of M. □

Now, let us return to our running example and see if supported equilibrium semantics can be introspective enough to reject equilibrium model S_2 of Example 1 as not supported.

Example 7. *Consider MCS M and equilibrium model S_2 of Example 1. In Example 6, we saw that $Sup_1^{L_1}$ contains two support functions for belief set bs_1^2 of S_2. We also computed support function $f \in Sup_1^{L_1}$. For S_2 to be a supported equilibrium, by Definition 5, functions $f_i \in Sup_i^{L_i}$ (for $i \in \{1, 2, 3\}$) should be found so that beliefs in bs_i^2 ($i \in \{1, 2, 3\}$) are justified and nothing justifies itself. In this example, we show that, indeed, if f_1 is the support function f from Example 6, then a self-justification is inevitable.*

In order to do so, we take f_1 to be function f from Example 6 and show that ordering $<$ cannot exist. We know that $(1 : celebrity(aj)) \in f_1(wealthy(aj))$ and $(2 : famous(aj)) \in f_1(celebrity(aj))$. Thus, ordering $<$ should satisfy $(2 : famous(aj)) < (1 : celebrity(aj)) < (1 : wealthy(aj))$. Now, we look at the possible supports for belief $famous(aj)$ in bs_2^2. According to C_2, belief $famous(X)$ can have three possible supports: $actor(X)$, $comedian(X)$ or $wealthy(X)$. However, for $X = aj$, the only possibility is $wealthy(X)$. Thus, inevitably, we have $(1 : wealthy(aj)) \in f_2(famous(aj))$ and, therefore, ordering $<$ should also satisfy $(1 : wealthy(aj)) < (2 : famous(aj))$ that contradicts the previous constraint on $<$.

Hence, if f_1 is function f from Example 6, ordering $<$ cannot exist and a self-justification inevitably occurs. The reader can check that choosing f_1 differently does not rectify the situation either. So, S_2 is not a supported equilibrium w.r.t. support functions $(Sup_{L_1}, Sup_{L_2}, \mathrm{uSup}_{L_3})$.

Example 7 demonstrates how supported equilibrium semantics avoids undesirable equilibrium models. Moreover, Example 7 shows that supported equilibrium semantics can also single out unfounded beliefs. In Example 7, one can check that all beliefs in S_2 are founded except those about "Average Joe". Even more importantly, Example 7 shows that supported equilibrium semantics allows us to trace back the reason for inclusion of all beliefs in a belief state.

Moreover, recall that grounded equilibrium semantics was not applicable to MCS M from Example 1. Thus, up to now, it was not possible to avoid unintended equilibrium models such as S_2 for MCSs such as M. However, as Example 7 shows, supported equilibrium semantics both (1) applies to M, and, (2) rejects unintended equilibrium S_2 as not supported. Hence, supported equilibrium semantics can indeed express cases that are not expressible using either the equilibrium model semantics or the grounded equilibrium model semantics. It should also be noted that Example 7 concerned a very simple MCS in which all contexts except one were reducible. However, as shown in the next section, the cases that supported equilibrium semantics can express (but equilibrium semantics or grounded equilibrium semantics cannot) are not limited to these easy cases. In the next section, we show that supported equilibrium semantics generalizes and unifies a spectrum of different semantics that can be ordered according to their selectivity.

Degree of Selectivity in Supports

In the previous sections, we defined supported equilibrium semantics and showed that all supported equilibriums are indeed also an equilibrium while vice versa is not necessarily true. That is, as shown in Example 7, there exist equilibriums that are not supported. Moreover, in various parts of this paper, we noted that supported equilibrium semantics is a unifying semantics that can define a spectrum of different semantics ordered according to their flexibility and/or selectivity. In this section, we want to further elaborate on the relationship between different semantics of a MCS.

In this section, we partially order different supports for logics and show that this ordering directly corresponds to the degree of selectivity in supported equilibrium semantics. That is, if support S is less than support S' (according to the ordering in this section), then supported semantics w.r.t. S' is more selective than supported semantics w.r.t. S.

We first define an ordering over justification functions:

Definition 6 (Ordered Justifications). *Let bs be a belief set for a logic L and j_1, j_2 be two justification functions for bs. Then, we say that j_1* is less selective than *j_2, denoted by $j_1 \leq j_2$, if, for all $b \in bs$, we have:*

$$\mathrm{fst}(j_1(b)) \subseteq \mathrm{fst}(j_2(b)) \text{ and } \mathrm{snd}(j_1(b)) \subseteq \mathrm{snd}(j_2(b)).$$

Intuitively, $j_1 \leq j_2$ means that j_1 requires less reasons than j_2 for supporting beliefs b in a belief set. Therefore, since supported equilibriums disallow self-justifying loops, requiring more reasons makes such a loop more probable and, so, shrinks the set of supported equilibriums.

Definition 7 (Ordered Supports). *Let $L := \langle KB, BS, ACC\rangle$ be a logic and Sup_L, Sup'_L be two supports for logic L. Then, we say that Sup_L* is less selective than *Sup'_L, denoted by $Sup_L \leq Sup'_L$ if, for all $kb \in KB$, $bs \in BS$, and $j' \in Sup'_L(kb, bs)$, a justification $j \in Sup_L(kb, bs)$ exists such that $j \leq j'$.*

Informally speaking, Definition 7 says that S_1 is a less selective support than S_2 if and only if, for all justification functions in S_2, a less selective justification function in S_1 can always be found. Hence, again, if $S_1 \leq S_2$ then using S_1 leads to less circularity and thus more supported equilibriums. The following theorem formalizes this reasoning:

Theorem 2. *Let $M := (C_1, \cdots, C_n)$ be a MCS and $S := (S_1, \cdots, S_n)$ be a supported equilibrium of M w.r.t. supports $(Sup_1, \cdots, Sup_n)$. Also, let supports $Sup'_1, \cdots, Sup'_n$ be such that $Sup'_i \leq Sup_i$ (for $1 \leq i \leq n$). Then, S is also a supported equilibrium of M w.r.t. supports $(Sup'_1, \cdots, Sup'_n)$.*

Proof. Let ordering $<$ and functions $f_i \in Sup_i^S$ witness S being a supported equilibrium. Since $f_i \in Sup_i^S$, by Definition 4, functions $g_i \in Sup_i(kb_i \cup app_i(S), S_i)$ and $R_i : S_i \mapsto \mathcal{P}(br_i)$ exist such that, for all $b \in S_i$: (1) $R_i(b)$ is a minimal subset of applicable bridge rules in br_i (with respect to S) that satisfies $\{hd(r) \mid r \in R_i(b)\} \supseteq \mathrm{snd}(g_i(b)) \setminus kb_i$, and, (2) $f_i(b) = \mathrm{fst}(g_i(b)) \cup \bigcup_{r \in R_i(b)} body^+(r)$.

Now, since $Sup'_i \leq Sup_i$, functions $g'_i \in Sup'_i(kb_i \cup app_i(S), S_i)$ exist such that $g'_i \leq g_i$. Also, construct R'_i such that $R'_i(b) = \{r \mid r \in R_i(b)$ and $hd(r) \in (\mathrm{snd}(g'_i(b)) \setminus kb_i)\}$. It is easy to check that $R'_i(b)$ satisfies minimality conditions of Definition 4 and, so, functions $f'_i \in Sup_i'^S$ exist such that $f'_i(b) = \mathrm{fst}(g'_i(b)) \cup \bigcup_{r \in R'_i(b)} body^+(r)$.

Therefore, in order to prove that S is supported w.r.t. $(Sup'_1, \cdots, Sup'_n)$, we only need to show that functions f'_i are non-circular. Note that, since $R'_i(b) \subseteq R_i(b)$ and $\mathrm{fst}(g'_i(b)) \subseteq \mathrm{fst}(g_i(b))$, we have $f'_i(b) \subseteq f_i(b)$ (for all b). Hence, the same ordering $<$ also shows the non-circularity of functions f'_i because:

$$(j : b') \in f'_i(b) \Rightarrow (j : b') \in f_i(b) \Rightarrow (j : b') < (i : b).$$

□

Theorem 2 shows that our ordering on supports is indeed relevant to our supported equilibrium semantics in the sense that greater supports are related to more selective semantics. The following proposition shows that our ordering relation on supports has a minimum and a maximum that respectively coincides with our unit and empty supports.

Proposition 1 (Least and Greatest Supports). *For all logics L and supports Sup_L, we have:* $\mathrm{uSup}_L \leq Sup_L \leq \mathrm{eSup}_L$.

Proof. Showing $Sup_L \leq \mathrm{eSup}_L$ is easy because $\mathrm{eSup}_L(kb, bs) = \emptyset$. Showing $\mathrm{uSup}_L \leq Sup_L$ is also easy because function $u^{bs} \in \mathrm{uSup}_L(kb, bs)$ is less than all $j \in Sup_L(kb, bs)$. This is because $\mathrm{fst}(u^{bs}(b)) = \emptyset \subseteq \mathrm{fst}(j(b))$ and $\mathrm{snd}(u^{bs}(b)) = \emptyset \subseteq \mathrm{snd}(j(b))$. □

According to Proposition 1, unit support is the most flexible support. In the next section, we show that if all logics use the unit support then supported equilibrium semantics coincides with equilibrium semantics. That is, equilibrium semantics is, in fact, a semantics in which beliefs are believed without really knowing why. On the other hand, we will also show in the next section that grounded equilibrium semantics is a particular case of supported equilibrium semantics in which having detailed justifications for beliefs are extremely important.

Generalizing Normal and Grounded Equilibrium Semantics

One of the promises of this paper was that our new supported equilibrium semantics generalizes both of the original semantics for multi-context systems, i.e., the equilibrium semantics and the grounded equilibrium semantics. In this section, we give Theorems 3 and 5 that, respectively, prove that equilibrium semantics and grounded equilibrium semantics are both special cases of supported equilibrium semantics (just using different support functions).

Theorem 3. *Let $M := (C_1, \cdots, C_n)$ be a MCS and $S := (S_1, \cdots, S_n)$ be a belief state of M. Then, S is an equilibrium of MCS if and only if S is a supported equilibrum of MCS w.r.t.* $(\mathrm{uSup}_{L_1}, \cdots, \mathrm{uSup}_{L_n})$.

Proof. ($\Leftarrow$) Directly follows Theorem 1.
($\Rightarrow$) Since S is an equilibrium model of M, we have that $S_i \in ACC(kb_i \cup app_i(S))$ for all $1 \leq i \leq n$. Therefore, by Definition 3, we know that $\mathrm{uSup}_{L_i}(kb_i \cup app_i(S), S_i) = \{u^{S_i}\}$. Now, by Definition 4, $Sup_i^S = \{f^{S_i}\}$ where $f^{S_i}(b) := \emptyset$ for all $b \in S_i$. Hence, we take $(f_1, \cdots, f_n) \in (Sup_1^S \times \cdots \times Sup_n^S)$ such that $f_i := f^{S_i}$. Now, by Definition 5, S is a supported equilibrium because functions $f_1, \cdots, f_n$ do not induce any circular justification. □

Theorem 3 shows that supported equilibrium semantics w.r.t. unit supports naturally extends equilibrium semantics. Previously, we discussed that unit support is associated with a complete black-box view of contexts since it does not provide any information about why something is believed. Hence, by Theorem 3, we know that equilibrium semantics indeed corresponds to the view of contexts as black-boxes.

Next, we show that, if support functions are chosen carefully, grounded equilibrium semantics can also be defined in terms of supported equilibrium semantics. In order to achieve our goal, we first characterize the unique minimal equilibrium of definite MCSs and use this definition to characterize the grounded equilibrium semantics in terms of supported equilibrium semantics.

Definition 8 (Monotonicity-based Support). *Let $L := \langle KB, BS, ACC\rangle$ be a monotone logic, $kb \in KB$ be a knowledge base and $bs \in BS$ be the unique acceptable belief set of kb, i.e., $\{bs\} = ACC(kb)$. Also, let $<$ be a total and well-founded ordering on kb. A function $j_< : bs \mapsto (\mathcal{P}(bs) \times \mathcal{P}(kb))$ is said to be a* monotonicity-based justification according to ordering $<$ *if, for all $b \in bs$, we have $j_<(b) := (bs_1, kb_2)$ where kb_1, kb_2, bs_1 and bs_2 are so that:*

- *either $kb_1 = bs_1 = \emptyset$ or $kb_1 \in KB$, $bs_1 \in BS$ and $\{bs_1\} = ACC(kb_1)$,*
- *$kb_2 \in KB$, $bs_2 \in BS$ and $\{bs_2\} = ACC(kb_2)$,*
- *formulas $k_1, k_2 \in kb$ exist such that $k_1 \leq k_2$, $kb_1 = kb_{<k_1}$, $kb_2 = kb_{\leq k_2}$, and, for all k' with $k_1 < k' \leq k_2$, we have $kb_{<k'} \notin KB$,*
- *$b \notin bs_1$ but $b \in bs_2$.*

We also define the monotonicity-based support of L, *denoted by* mSup_L, *as follows:*

If $\{bs\} \neq ACC(kb)$ then $\mathrm{mSup}_L(kb, bs) = \emptyset$;
Otherwise, $\mathrm{mSup}_L(kb, bs)$ is the set of monotonicity-based justifications $j_<$ according to ordering $<$.

In the following, Theorem 4 shows that, for definite multi-context systems, supported equilibrium semantics and grounded equilibrium semantics coincide.

Theorem 4. *Let $M := (C_1, \ldots, C_n)$ be a definite MCS and $S := (S_1, \cdots, S_n)$ be a belief state of M. Then, S is a supported equilibrium of M w.r.t.* $(\mathrm{mSup}_{L_1}, \cdots, \mathrm{mSup}_{L_n})$ *if and only if S is a grounded equilibrium of M.*

Proof. ($\Leftarrow$) Let S be a grounded equilibrium of M. By Proposition 1 of (Brewka and Eiter 2007), we know that $\{S_i\} = ACC_{L_i}(kb_i^\infty)$ where $kb_i^\infty := \bigcup_\alpha kb_i^\alpha$ and kb_i^α is defined as follows:

$$kb_i^\alpha := \begin{cases} kb_i & \text{if } \alpha = 0, \\ kb_i^\beta \cup app_i(E^\beta) & \text{if } \alpha = \beta + 1, \\ \bigcup_{\beta < \alpha} kb_i^\beta & \text{if } \alpha \text{ is a limit ordinal.} \end{cases}$$

where, for all ordinals α, we have $E^\alpha := (E_1^\alpha, \cdots, E_n^\alpha)$ and $\{E_i^\alpha\} = ACC(kb_i^\alpha)$. So, for ordinals α and β, if $\alpha < \beta$ then $E_i^\alpha \subseteq E_i^\beta$ (for all i). Thus, for all $k \in kb_i^\infty$ (respectively, for all $b \in S_i$), by $rank_i^{kb}(k)$ (respectively, by $rank_i^{bs}(b)$), we denote the minimum ordinal α such that $k \in kb_i^\alpha$ (respectively, $b \in E_i^\alpha$). Now, define orderings $<_i$ (for $i \in \{1, \cdots, n\}$) on kb_i^∞ to be any total ordering that respects the ranks of knowledge in kb_i^α's, i.e., for $k_1, k_2 \in kb_i^\infty$, we have $k_1 <_i k_2$ if $rank_i^{kb}(k_1) < rank_i^{kb}(k_2)$.

Next, we show that functions $f_1 \in \mathrm{mSup}_1^S, \cdots, f_n \in \mathrm{mSup}_n^S$ exist such that $f_1, \cdots, f_n$ avoid self-justification. In order to do that, consider support functions $g_i \in \mathrm{mSup}_{L_i}(kb_i^\infty \cup app_i(S), S_i)$ (for $i \in \{1, \cdots, n\}$) that are generated according to orderings $<_i$ (respectively). Now, construct functions $f_i \in \mathrm{mSup}_i^{L_i}$ using g_i's so that, for $b \in S_i$, we have:

$$f_i(b) := \{(i : b') \mid b' \in \mathrm{fst}(g_i(b))\} \cup \bigcup_{r \in R_b} body^+(r),$$

where R_b is a minimal subset of applicable bridge rules of C_i such that, for each $k \in (\mathrm{snd}(g_i(b)) \setminus kb_i)$, there is a rule $r \in R$ with $hd(r) = k$ and $E_i^{rank_i^{kb}(k)-1} \models body(r)$. Note that $rank_i^{kb}(k) - 1$ is indeed an ordinal because, if $k \notin kb_i$, $rank_i^{kb}(k)$ is always a successor ordinal.

It can be easily checked that, for $i \in \{1, \cdots, n\}$, we have $f_i \in \mathrm{mSup}_i^{L_i}$. We just need to show that ordering $<^*$ exists such that $(i : p) \in f_j(q)$ implies $(i : p) <^* (j : q)$. Define $<^*$ as follows:

$$(i : p) <^* (j : q) \iff \begin{cases} \text{either } i = j \text{ and } p \in g_i(q), \\ \text{or, } rank_i^{bs}(p) < rank_j^{bs}(q). \end{cases}$$

We leave it to the reader to check that $<^*$ is a well-founded ordering. Now, if $(i : p) \in f_j(q)$, we know that either $i = j$ and $p \in \mathrm{fst}(g_j(q))$ or $(i : p) \in body^+(r)$ for a rule $r \in app_i(S)$ with $hd(r) \in \mathrm{snd}(g_j(q))$. In the former case, by construction of $<^*$, we have $(i : p) <^* (j : q)$ as required. In the latter case, let $\alpha := rank_j^{bs}(q)$. By construction of $\mathrm{snd}(g_j(q))$, we know that α is a successor ordinal and that r is applicable under $E^{\alpha-1}$. Thus, $p \in E_i^{\alpha-1}$. Hence, $rank_i^{bs}(p) \leq \alpha - 1 < \alpha = rank_j^{bs}(q)$ and, so, by construction of $<^*$ and as required, $(i : p) <^* (j : q)$. Hence, S is a supported equilibrium of M.
($\Rightarrow$) Since S is a supported equilibrium, functions $f_i \in \mathrm{mSup}_i^{L_i}$ and well-ordering $<_{X(S)}$ exist such that $(i : p) \in f_j(q)$ implies $(i : p) <_{X(S)} (j : q)$. Also, by Definition 4, we know that each f_i is constructed according to some $g_i \in \mathrm{mSup}_{L_i}(kb_i \cup app_i(S), S_i)$. Also, since M is a definite MCS, it has a unique minimal equilibrium. Let $E := (E_1, \cdots, E_n)$ be that minimal equilibrium.

By Theorem 1, we know that S is an equilibrium of M. Thus, we only need to show that S is minimal. Assuming otherwise means that i and p exist such that $p \in S_i$ but $p \notin E_i$. Choose a minimal such i and p w.r.t. ordering $<_{X(S)}$. By Definition 8, $f_i(p) \supseteq \bigcup_{r \in R} body^+(r)$ where R satisfies $\{hd(r) \mid r \in R\} \supseteq (\mathrm{snd}(g_i(p)) \setminus kb_i)$. Now, since $(i : p)$ is minimal w.r.t. $<_{X(S)}$, if $(j : q) \in f_i(p)$ then $q \in E_j$. Thus, for all $r \in R$, we have $E \models body(r)$ and, so, $app_i(E) \supseteq \{hd(r) \mid r \in R\} \supseteq (\mathrm{snd}(g_i(p)) \setminus kb_i)$. Therefore, $(kb_i \cup app_i(E)) \supseteq \mathrm{snd}(g_i(p))$. Now, assume that bs is the unique acceptable belief set of $\mathrm{snd}(g_i(p))$. By Definition 8, $p \in bs$. Moreover, since L_i is monotone and $\{E_i\} = ACC(kb_i \cup app_i(E))$, we have that $E_i \supseteq bs$ and, henceforth, $p \in E_i$. This contradicts our assumption that $p \notin E_i$. Therefore, S is a minimal equilibrium and, since M has only one minimal equilibrium, $S = E$. ☐

Theorem 4 shows that, for definite MCS, grounded equilibrium semantics can be defined in terms of supported equilibrium semantics. In the following, we use monotonicity-based support of M and Theorem 4 to show that, indeed, grounded equilibrium semantics can always be defined in terms of supported equilibrium semantics.

Definition 9 (Reducibility-based Supports). *Let $L := \langle KB, BS, ACC \rangle$ be a reducible logic. Also, let $kb \in KB$ and $bs \in ACC(kb)$. Moreover, let $L^* := \langle KB^*, BS, ACC \rangle$ be the monotone part of logic L. The* reducibility-based support of *L, denoted by* $\mathrm{rSup}_L(kb, bs)$, *is the set of functions $f : bs \mapsto (\mathcal{P}(bs) \times \mathcal{P}(kb))$ such that:*

$$\exists f' \in \mathrm{mSup}_{L^*}(red_L(kb, bs), bs) \text{ s.t.}$$
$$\forall b \in bs : \; f(b) = (\mathrm{fst}(f'(b)), \mathrm{snd}(f'(b)) \cap kb).$$

Needless to say that $\mathrm{rSup}_L(kb, bs) = \emptyset$ *if $bs \notin ACC(kb)$.*

The following lemma suggests that the difference between monotonicity-based supports and reducibility-based supports disappear when we move to the supports to the level of contexts.

Lemma 1. *For reducible MCS $M := (C_1, \cdots, C_n)$ and belief state $S := (S_1, \cdots, S_n)$ of M, let $C'_i := (L_i^*, red_{L_i}(kb_i, S_i), br_i^*)$ where L_i^* is the monotone part of L_i and br_i^* is the reduct of br_i under S. Then, for all i,* $\mathrm{mSup}_i^S = \mathrm{rSup}_i^S$.

Theorem 5. *Let $M := (C_1, \ldots, C_n)$ be a reducible MCS and $S := (S_1, \cdots, S_n)$ be a belief state of M. Then, S is a grounded equilibrium of M if and only if S is a supported equilibrium of M w.r.t.* $(\mathrm{rSup}_{L_1}, \cdots, \mathrm{rSup}_{L_n})$.

Proof. ($\Rightarrow$) Let S be a ground equilibrium of M. Then, S is the unique grounded equilibrium of M^S and, hence, a supported equilibrium of M w.r.t. $(\mathrm{mSup}_{L_1^*}, \cdots, \mathrm{mSup}_{L_n^*})$ (by Theorem 4 and where L_i^*'s are the monotone part of L_i's). Therefore, $f_i^* \in \mathrm{mSup}_i^S$ and strict well-ordering $<$ exist such that $p \in S_i \wedge (j : q) \in f_i(p) \Rightarrow (j : q) < (i : p)$. By Lemma 1, $f_i \in \mathrm{rSup}_i^S$. Thus, S is a supported equilibrium of M w.r.t. $(\mathrm{rSup}_1, \cdots, \mathrm{rSup}_n)$.
($\Leftarrow$) Let S be a supported equilibrium of M. Then, by Lemma 1, S is also a supported equilibrium of M^S (using the same support functions and well-ordering). Therefore, S is the unique grounded equilibrium of M^S (by Theorem 4). Hence, S is a grounded equilibrium of M. ☐

Theorem 5 shows that supported equilibrium semantics naturally extends grounded equilibrium semantics. We also showed previously in Theorem 3 that supported equilibrium semantics also naturally extends equilibrium semantics of MCSs. Thus, our supported equilibrium semantics indeed extends and unifies both the equilibrium semantics and the grounded equilibrium semantics.

Application: Better Diagnoses

In this section, we discuss one of possible applications of supported equilibrium semantics for MCSs: the detailed diagnosis of MCSs in the presence of supports.

Recall the example of shops and warehouses from the beginning of the section on generalized MCSs. There, a product was believed to be unavailable for sale because (1) the product was believed not to be in the warehouse, (2) all the available instances of that product in the store were believed to be "sold," and, (3) our knowledge base asserted that a product can be available only if it is available in the warehouse or if it is available in the store and not marked as "sold." Now, consider an extension of this example in which this product is also believed to be on route from factory to warehouse and our knowledge asserts that a product cannot be both unavailable and on route to warehouse. Here, using supports to trace back the reason for conflict, a more meaningful diagnosis can be obtained: to change our knowledge so that we also believe a product is available if it is on route from factory to warehouse. Note that our diagnosis here proposes a change to the knowledge base of a context; a diagnosis that was impossible without using justifications.

In (Bögl et al. 2010) and (Eiter et al. 2010), authors rightly argue that, due to their distributed behavior and complexity, MCSs are prone to two types of errors: those that originate inside a context and those that originate in the interaction between contexts. Moreover, the authors introduce *inconsistency explanations* and *diagnoses* to repair inconsistent MCSs. For MCS M, an inconsistency-explanation is a pair (E_1, E_2) so that $E_1, E_2 \subseteq br_M$ which, informally speaking, means that if all bridge rules in E_1 are kept and all bridge rules in E_2 are non-applicable then M is inconsistent. Thus, (E_1, E_2) explains why M is inconsistent: because bridge rules in E_1 are all applicable or because bridge rules in E_2 are all non-applicable. Also, authors of (Eiter et al. 2010) study the dual notion of diagnoses (D_1, D_2) which means that excluding bridge rules in D_1 and excluding enough conditions from the body of bridge rules in D_2 makes M consistent. Hence, a diagnosis tells us a way to restore the consistency of M.

The authors of (Eiter et al. 2010) study inconsistency explanations and diagnoses and prove their duality. However, diagnoses and inconsistency explanations only take bridge rules into consideration and cannot deal with inconsistencies that arise from errors internal to a knowledge base. However, we now know that supported equilibrium semantics provides the right means to look inside the contexts of a MCS. Therefore, a mistake can be traced back to its source via justifications. In this section, we consider an example of how better repair is achievable through support functions and leave the full treatment of this subject to a future research.

Moreover, and even more importantly, supported equilibrium semantics enables us to focus on incorrectness of MCSs rather than their inconsistency. It should be clear that inconsistency and incorrectness are different notions. For example, consider MCS M from Example 1. M has two equilibrium models and, thus, is consistent. However, M accepts unintended equilibrium model S_2 and, thus, is incorrect. Similarly, one could argue that some inconsistent multi-context systems are indeed correct. For instance, a tour of American cities that passes each city exactly once might not exist. The non-existence of such a path might make a MCS inconsistent but it does not make it incorrect.

Example 8. *Consider MCS M of Example 1 and its equilibrium model S_1. We know that S_1 is a supported equilibrium model of M w.r.t. the same support functions as in Example 7. According to S_1, "Jimmy Kimmel" is not famous, not wealthy and not a celebrity which is contradictory to the real world. So, we want to guarantee that M should assert $celebrity(jk)$. Under this assumption, M is incorrect because it allows S_1 to be supported. Thus, M is consistent but not correct (because it does not represent what it is intended to represent). Now, using support functions, we understand that one possible repair for our MCS is to guarantee in C_3 that "Jimmy Kimmel" is a comedian. In this case, a non-circular chain of justifications would guarantee that $celebrity(jk)$, $wealthy(jk)$ and $famous(jk)$ will all be true. Thus, we would add knowledge $comedian(jk)$ to the knowledge base of C_3 to repair this example and disallow S_1 from being a supported equilibrium.*

Hence, supported equilibrium semantics enables us to focus on correctness (rather than consistency), and better repair MCSs (by proposing internal changes to contexts).

Conclusion and Future Directions

In this paper, we introduced the notion of justifications and supports to MCSs and invented the novel semantics of supported equilibria for MCSs. Moreover, we showed that supported equilibrium semantics is useful in two important ways: (a) from the theoretical aspect, supported semantics generalizes the two main semantics proposed for multi-context systems, i.e., the equilibrium semantics and the grounded equilibrium semantics, and, (b) from the application point of view, supported equilibrium semantics provides the means to look inside the contexts of MCSs and to better diagnose and repair faulty MCSs. Hence, we believe that supported equilibrium semantics is the most suitable semantics available for MCSs.

Moreover, in this paper, we only argued about two very specific supports, i.e., the trivial unit supports and the reducibility-based supports. We also discussed that supported equilibrium semantics allows the inclusion of many more possible supports that, from the viewpoint of selectivity degree, lie between these two semantics. We believe that many other useful supports exist that have not been investigated in this paper. For example, an interesting and unanswered question about supported equilibrium semantics is to find necessary and/or sufficient conditions under which it is guaranteed that supported equilibria of a MCS will form an anti-chain. Answering such a question gives us natural ways to define a semantics similar to stable model semantics for arbitrary new languages.

References

Artemov, S. 1995. Operational modal logic. Technical report, Cornell University.

Bögl, M.; Eiter, T.; Fink, M.; and Schüller, P. 2010. The MCS-IE system for explaining inconsistency in multi-context systems. In Janhunen, T., and Niemelä, I., eds., *Logics in Artificial Intelligence - 12th European Conference, JELIA 2010, Helsinki, Finland, September 13-15, 2010. Proceedings*, volume 6341 of *Lecture Notes in Computer Science*, 356–359. Springer.

Brewka, G., and Eiter, T. 2007. Equilibria in heterogeneous nonmonotonic multi-context systems. In *Proceedings of the 22nd National Conference on Artificial Intelligence (AAAI'07) - Volume 1*, 385–390. AAAI Press.

Brewka, G.; Eiter, T.; Fink, M.; and Weinzierl, A. 2011. Managed multi-context systems. In Walsh, T., ed., *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, 786–791. IJCAI/AAAI.

Brezhnev, V. 2001. On the logic of proofs. In Striegnitz, K., ed., *Proceedings of the Sixth ESSLLI Student Session, 13th European Summer School in Logic, Language and Information (ESSLLI'01)*, 35 – 46.

Cabalar, P. 2011. Logic programs and causal proofs. In *Logical Formalizations of Commonsense Reasoning, Papers from the 2011 AAAI Spring Symposium, Technical Report SS-11-06, Stanford, California, USA, March 21-23, 2011*. AAAI.

Eiter, T.; Fink, M.; Schüller, P.; and Weinzierl, A. 2010. Finding explanations of inconsistency in multi-context systems. In Lin, F.; Sattler, U.; and Truszczynski, M., eds., *Principles of Knowledge Representation and Reasoning: Proceedings of the Twelfth International Conference, KR 2010, Toronto, Ontario, Canada, May 9-13, 2010*. AAAI Press.

Fink, M.; Ghionna, L.; and Weinzierl, A. 2011. Relational information exchange and aggregation in multi-context systems. In Delgrande, J. P., and Faber, W., eds., *Logic Programming and Nonmonotonic Reasoning - 11th International Conference, LPNMR 2011, Vancouver, Canada, May 16-19, 2011. Proceedings*, volume 6645 of *Lecture Notes in Computer Science*, 120–133. Springer.

Fitting, M. 2005. The logic of proofs, semantically. *Annals of Pure and Applied Logic* 132(1):1 – 25.

Gelfond, M., and Lifschitz, V. 1988. The stable model semantics for logic programming. In *Proceeding of the 5th International Conference and Symposium on Logic Programming (ICLP/SLP'88)*, 1070–1080.

Lifschitz, V. 2010. Thirteen definitions of a stable model. In Blass, A.; Dershowitz, N.; and Reisig, W., eds., *Fields of logic and computation*. Berlin, Heidelberg: Springer-Verlag. 488–503.

Tasharrofi, S. 2013. A rational extension of stable model semantics to the full propositional language. In Rossi, F., ed., *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013*, 1118 – 1124. IJCAI.