Refining the Meaning of Sense Labels in PDTB: “Concession”

Livio Robaldo

University of Turin (Italy)

email: robaldo@di.unito.it

Eleni Miltsakaki

University of Pennsylvania (USA)

email: elenimi@linc.cis.upenn.edu

Jerry R. Hobbs

University of Southern California (USA)

email: hobbs@isi.edu

Abstract

The most recent release of PDTB 2.0 contains annotations of senses of connectives. The PDTB 2.0 manual describes the hierarchical set of senses used in the annotation and offers rough semantic descriptions of each label. In this paper, we refine the semantics of *concession* substantially and offer a formal description of concessive relations and the associated inferences drawn by the reader, utilizing basic notions from Hobbs’s logic, including the distinction between causes and causal complexes (Hobbs, 2005). This work is part of a larger project on the semantics of connectives which aims at developing formal descriptions of discourse relations, useful for processing real data.

1 Introduction

As the demand for more powerful NLP applications increases, there is also an increasing need to develop algorithms for automated processing of discourse relations and models for deriving the inferences drawn by the reader. PDTB 2.0 (Prasad et al., 2008), released in January 2008, contains annotations of discourse connectives and their arguments, attribution, and sense labels giving rough semantic descriptions of the connectives. The availability of such a richly annotated corpus promises to boost our understanding of the structure and meaning of discourse and will facilitate the development of efficient algorithms for identifying discourse connectives and their arguments.

However, in order to be able to derive appropriate inferences associated with discourse relations, we need to develop useful semantic analyses of the meaning of connectives so that they will generate the same range of inferences made by humans. In this paper we take a first step in that direction, offering a simple formal analysis of concessive relations, thus refining the semantics of the concessive sense labels used in PDTB 2.0. Our analysis uses basic notions of causality developed in Hobbs (1998, 2005), capitalizing on the distinction between causes and causal complexes and on the semantics of defeasible causality. Concessive meaning involves the failure of a general defeasible causal relation in this specific instance.

The paper is organized as follows. Section 2 gives an overview of the PDTB 2.0, focusing on the annotation of the senses of connectives, especially “concession”. In Section 3, we present an overview of the framework we are adopting for our formal analysis, namely, Hobbs’s logic of causality, and our basic claims about how the semantics of defeasible causality contributes to the semantics of concession. Section 4 presents the semantic analysis of “concession”. In Section 5, we report briefly on the distribution of concessive labels in PDTB 2.0 and conclude in Section 6.

2 Sense labels in PDTB

The Penn Discourse Treebank provides annotations of the argument structure of discourse connectives, attribution (e.g., ‘ownership’ of the relation by the writer or other individual), and semantic labels for all the annotated connectives (Prasad et al., 2008). This annotation of discourse connectives and their arguments draws on a lexical approach to discourse structure (Webber et al., 2003; Webber and Joshi, 2003), viewing discourse connectives as discourse-level predicates that take two *abstract objects* such as events, states, and propositions (Asher, 1993) as their arguments.

Two major types of discourse connectives are annotated in PDTB: a) explicit connectives including subordinate conjunctions, coordinate conjunctions and adverbials, and b) implicit connectives that are inserted between two adjacent sentences to capture the meaning of the inferred relation when no explicit connective is present. The PDTB 2.0 is, to date, the largest annotation effort at the discourse level, including approximately 40,000 triples in the form (Connective, Arg1, Arg2). Arg2 is the second argument in the text in the case of coordinating conjunctions, and is the complement of subordinating conjunctions. In the case of adverbs, Arg2 is the element which the adverb modifies syntactically. In cases of ambiguity, sense labels indicate the intended sense in the given context. In all other cases, sense labels provide semantic descrip-

tions of the relations conveyed by the connectives, both explicit and implicit.

The tagset of senses is organized hierarchically (Miltsakaki et al., 2008). The top level, or *class level*, has four tags representing four major semantic classes: “TEMPORAL”, “CONTINGENCY”, “COMPARISON” and “EXPANSION”. For each class, a second level of *types* is defined to further refine the semantics of the class levels. For example, “CONTINGENCY” has two types “Cause” (relating two situations via a direct cause-effect relation) and “Condition” (relating a hypothetical scenario with its (possible) consequences). A third level of *subtype* specifies the semantic contribution of each argument. For “CONTINGENCY”, its “Cause” type has two subtypes — “reason” (which applies when the connective indicates that the situation specified in Arg2 is interpreted as the cause of the situation specified in Arg1, as often with the connective *because*) and “result” (which is used when the connective indicates that the situation described in Arg2 is interpreted as the result of the situation presented in Arg1). That is, “reason” occurs when Arg2 causes Arg1; “result” occurs when Arg1 causes Arg2.

Connectives can also be used to relate arguments pragmatically as in *John is in the house because the lights are on* or *If you’re thirsty, there’s beer in the fridge*, where the relation involves the belief in or the telling of the condition rather than the condition itself. For these *rhetorical* or *pragmatic* uses of connectives, a small set of *pragmatic* sense tags has been defined — specifically, “Pragmatic Cause”, “Pragmatic Condition”, “Pragmatic Contrast” and “Pragmatic Concession”.

2.1 “Concession” in PDTB

“Concession” is a type of the class-level category “COMPARISON”. The class tag “COMPARISON” applies when the connective indicates that a discourse relation is established between Arg1 and Arg2 in order to highlight prominent differences between the two situations. Semantically, the truth of both arguments is independent of the connective or the established relation. “COMPARISON” has two types that further specify its semantics. In some cases, Arg1 and Arg2 share a predicate or a property and the difference is highlighted with respect to the values assigned to this property. This interpretation is tagged with the type “Contrast”.

There are also cases in which the highlighted differences are related to expectations raised by one argument which are then denied by the other. This interpretation is tagged with the type “Concession”. According to the description in the PDTB 2.0 manual, the type “Concession” applies when the connective indicates that one of the arguments describes a situation *A* which normally causes *C*, while the other asserts (or implies) $\neg C$. Alternatively, one argument denotes a fact that triggers a set of potential consequences, while the other denies one or more of them.

Two “Concession” subtypes are defined in terms of the argument creating an expectation and the one denying it. Specifically, when Arg2 creates an expectation that Arg1 denies, it is tagged as “expectation”, shown in (1.c-d). When Arg1 creates an expectation that Arg2 denies, it is tagged as “contra-expectation”, shown in (1.e-f). Examples (1.a-b) are made-up sentences we use for explanation and will be discussed here and in the next section. All other examples are taken from PDTB 2.0. Each discourse fragment in (1) distinguishes between a discourse connective (underlined), and two sentence-arguments: Arg1 (*italics*) and Arg2 (**boldface**).

- (1) a. Although John studied hard, he did not pass the exam. (expectation)
- b. Although running is considered healthy, it is not advisable for persons with heart problems. (expectation)
- c. Although they represent only 2% of the population, they control nearly one-third of discretionary income. (expectation)
- d. While acquiring a big brand-name company can be a shortcut to growth, it can also bring a host of unforeseen problems (expectation)
- e. The Texas oilman has acquired a 26.2% stake valued at more than \$1.2 billion in an automotive-lighting company, Koito Manufacturing Co. But he has failed to gain any influence at the company. (contra-expectation)
- f. Mr. Cannell's allegations of cheating "are purely without foundation", and based on unfair inferences. However the state will begin keeping closer track of achievement-test preparation booklets next spring. (contra-expectation)

(1.a) is an example of "expectation": Arg2 (*John studied hard*) creates the expectation that John passed the exam, which is precisely denied by Arg1. The same holds for (1.b-d). Note that (1.b), unlike (1.a, c-d), expresses a general concessive relation, i.e., it does not refer to particular contingent events. (1.e-f) are instances of contra-expectation, where the expectation is created by Arg1. In (1.e), the fact that the Texas oilman acquired the indicated stake value creates the expectation that he gained influence at the company, while, in (1.f), since Mr. Cannell's allegations of cheating are purely without foundation (in the speaker's judgement), we do not expect the state to start tracking the test preparation.

3 Toward a formal definition of "Concession"

Based on our analysis of the range of PDTB tokens tagged with a concessive label, we offer here a more detailed semantic analysis of the meaning of concessive relations. Since the direction of the concessive relation is not relevant, the argument that creates the expectation and the argument that denies it are respectively termed as Arg_{cexp} and Arg_{dexp}. We claim that a concessive relation arises from a contrast between the effects of two causal relations c_c and c_d holding in the domain. c and d stand for "creates" and "denies", respectively. The relation denoted by c_c is the causal relation that creates the expectation, and c_d the one that denies it. The effects of these causal relations, as well as their causes, are taken to be eventualities¹.

In this paper, we use the letter e for most eventualities, possibly with some subscript or superscript.² We make use of the subscripts $x1$ and $x2$, respectively, to distinguish between the causes and the effects in a causal relation c_x . Therefore, the causes in c_c and c_d are indicated by e_{c1} and e_{d1} respectively, and the effects by e_{c2} and e_{d2} , respectively. e_{c2} is the "created expectation"; its cause e_{c1} is conveyed by Arg_{cexp}. e_{d2} is an eventuality that denies e_{c2} , and it is explicitly described in Arg_{dexp}. The cause of

¹The term "eventuality" is borrowed from (Bach, 1981). It covers both standard notions of "state" and "event".

²As we will see, also causal relations are eventualities; so the names c_c and c_d are an exception to this rule.

e_{d2} , i.e., e_{d1} , is usually unknown. Also e_{c2} is, in principle, unknown, but in most cases it can be taken as the negation of e_{d2} .

For instance, in the context of (1.a), the eventuality *John studied hard* (e_{c1}) creates the expectation *John passed the exam* (e_{c2}). Nevertheless, Arg_{dex} says that *John did not pass the exam* actually (e_{d2}). The reason of e_{d2} is unknown and has to be found in the context. In other words, the context, whether explicit or inferred, should include another eventuality that caused John’s failure, despite his studying hard. For example, the next sentence might be *John was very tired during the exam* (e_{d1}).

In order to formalize this account of concession, we need a defeasible notion of causality. Many authors propose such an account of causality, e.g. (Achinstein, 1965; Shoham, 1990; Simon, 1991; Bell, 1999, 2003), and Giunchiglia et al. (2004). The account we use is that of Hobbs (2005). This distinguishes between the monotonic, precise notion of “causal complex” and the nonmonotonic, defeasible notion of “cause”. The former gives us mathematical rigor; the latter is more useful for everyday reasoning and can be characterized in terms of the former. As Hobbs (2005) explains, when we flip a switch to turn on a light, we say that flipping the switch “caused” the light to turn on. But for this to happen, many other factors had to be in place. The bulb had to be intact, the switch had to be connected to the bulb, the power had to be on in the city, and so on. The set of all the states and events that have to hold or happen for an effect e to happen are called the “causal complex” of e . Thus, the flipping of the switch and the normal states of the bulb, the wiring, and the power supply would all be in the causal complex for the turning on of the light. In a causal complex, the majority of participating eventualities are normally true and therefore presumed to hold. In the light bulb case, unless otherwise indicated, it is normally true that the bulb is not burnt out, that the wiring is intact, that the power is on in the city, and so on. But the light switch could be on or off; neither can be presumed. Those eventualities that cannot normally be assumed to be true are identified as causes (cf. Kayser and Nouioua, 2008). They are useful in planning, because they are often the actions that the planner or some other agent must perform. They are useful in explanation and prediction because they frequently constitute the new information. They are less useful in diagnosis, where the whole causal complex has to be considered.

Note that in practice, we can never specify all the eventualities in a causal complex for an event. So while the notion of causal complex gives us a precise way of thinking about causality, it is not adequate for the kind of practical reasoning we do in planning, explaining, and predicting. For this, we need the defeasible notion of “cause”.

3.1 Background on Hobbs’s logic

Hobbs (1998) proposed a wide coverage logical framework for natural language based on the notion of reification. Reification is the action of making states and events first-class individuals in the logic, so they can be referred to by constants and variables. We “reify” eventualities, from the Latin word ‘re(s)’ for ‘thing’: we take them to be *things*. The framework distinguishes two parallel sets of predicates: primed and unprimed. The unprimed predicates are the ordinary predicates we are used to in logical representations of language. For example, (*give a b c*) says that a gives b to c . When we assert this, we are saying that it actually takes place in the real world. The primed predicate is used to talk about the reified eventualities. The expression (*give*’

$e\ a\ b\ c$) says that e is a giving event by a of b to c . Eventualities may be possible or actual. When they are actual, this is simply one of their properties. To say that a state e actually obtains in the real world or that an event e actually occurs in the real world, we write (*Rexist* e). That is, e really exists in the real world. If I want to fly, my wanting really exists, but my flying does not. This is represented as:³

$$(\text{Rexist } e) \wedge (\text{want}' e\ I\ e_1) \wedge (\text{fly}' e_1\ I)$$

Therefore, contrary to $(p\ x)$, $(p'\ e\ x)$ does not say that e actually occurs, only that if it did, it would be a “p” event. The relation between primed and unprimed predicates is then formalized by the following axiom schema:

$$(\text{forall } (x) (\text{iff } (p\ x) (\text{exists}(e) (\text{and}(p'\ e\ x)(\text{Rexist } e))))))$$

Eventualities can be treated as the objects of human thoughts. Reified eventualities are inserted as parameters of such predicates as *believe*, *think*, *want*, etc. These predicates can be applied in a recursive fashion. The fact that *John believes that Jack wants to eat an ice cream* is represented as an eventuality e such that⁴

$$(\text{believe}' e\ \text{John } e_1) \wedge (\text{want}' e_1\ \text{Jack } e_2) \wedge \\ (\text{eat}' e_2\ \text{Jack } Ic) \wedge (\text{iceCream}' e_3\ Ic)$$

In Hobbs’s notation, every relation on eventualities, including logical operators, causal and temporal relations, and even tense and aspect, may be reified into another eventuality. For instance, by asserting $(\text{imply}' e\ e_1\ e_2)$, we reify the implication from e_1 to e_2 into an eventuality e . e has to be thought as ‘the state holding between e_1 and e_2 such that whenever e_1 really exists, e_2 really exists too’. Negation is represented as $(\text{not}' e_1\ e_2)$: e_1 is the eventuality of the e_2 ’s not existing. Some problems arise with negation, in that what is generally negated is an eventuality type rather than an eventuality token or instance. In order to deal with more general cases of concession, we will refer to eventualities that are *inconsistent* with other ones. Two eventualities e_1 and e_2 are said to be inconsistent iff they (respectively) imply two other eventualities e_3 and e_4 such that e_3 is the negation of e_4 . The definition is as follows:

$$(\text{forall } (e_1\ e_3) \\ (\text{iff } (\text{inconsistent } e_1\ e_2) \\ (\text{and } (\text{eventuality } e_1) (\text{eventuality } e_2) \\ (\text{exists } (e_3\ e_4) (\text{and } (\text{imply } e_1\ e_3) \\ (\text{imply } e_2\ e_4)(\text{not}' e_3\ e_4)))))))$$

3.2 Typical elements, eventuality types and tokens

Among the things we can think about are both specific eventualities, like *Fido is barking*, and general or abstract types of eventualities, like *Dogs bark*. We do not want to treat these as radically different kinds of entities. We would like both, at some level, to

³In order to increase readability, we will often make use of the symbol $\wedge$ in place of the unprimed predicate *and*.

⁴The formula expresses the de-re reading of the sentence, where $e_1, e_2, e_3, \text{John}, \text{Jack}, Ic$ are first order constants.

be treated simply as eventualities that can be the content of thoughts. To this end, the logical framework includes the notion of typical element (from Hobbs (1983, 1995, 1998)). The typical element of a set is the reification of the universally quantified variable ranging over the elements of the set (cf. McCarthy (1977)). Typical elements are first-order individuals. The introduction of typical elements arises from the need to move from the standard set-theoretic notation

$$s = \{x \mid p(x)\}$$

or its logical equivalent,

$$(\text{forall } (x) (\text{iff } (\text{member } x \ s) (p \ x)))$$

to a simple statement that p is true of a “typical element” of s by reifying typical elements. The principal property of typical elements is that all properties of typical elements are inherited by the real members of the set.

It is important not to confuse the concept of typical element with the standard concept of “prototype”, which allows defeasibility, i.e., properties that are not inherited by all of the real members of the set. Asserting a predicate on a typical element of a set is logically equivalent to the multiple assertions of that predicate on all elements of the set. Talking about typical elements of sets of eventualities leads to the distinction between eventuality types and eventuality tokens. The logic defines the following concepts, for which we omit formal details⁵: a) *Eventualities types* (aka *abstract eventualities*): eventualities that involve at least one typical element among their arguments or arguments of their arguments (we can call these “parameters”), b) *Partially instantiated eventuality types* (aka *partial instances*): a particular kind of eventuality type resulting from instantiating some of the parameters of the abstract eventuality either with real members of their sets or with typical elements of subsets, and c) *Eventuality tokens* (aka *instances*): a particular kind of partially instantiated eventuality type with no parameters. It is a consequence of universal instantiation that any property that holds of an eventuality type is true of any partial instance of it.

Hobbs’s logical framework is particularly suitable to the study of the semantics of discourse connectives, in that it allows focusing on their meaning while leaving unspecified the details about the eventualities involved. In other words, we can simply assume the existence of two eventualities e_1 and e_2 coming from the two arguments Arg1 and Arg2 respectively. e_1 and e_2 may be either eventuality tokens, on atomic arguments, as in (1.a), or eventuality tokens, on collective arguments, as in (1.c), or (partially instantiated) eventuality types, as in (1.b), or any other kind of eventuality. The semantics of concession proposed below uniformly applies to all these cases.

3.3 Hobbs’s Account of Causality

The account of causality described above in the introduction is represented in terms of two predicates: (*cause*’ $c_x \ e_{x1} \ e_{x2}$) and (*causalComplex*’ $s \ e_{x2}$). *cause*’ says that c_x is the state holding between e_{x1} and e_{x2} such that the former is a non-presumable cause

⁵ Actually, “instance” is slightly more general, since if s is a set, x is its typical element, and y is a member of s , y is an instance of x , even though it is not an eventuality. Nevertheless, in this paper we assume “instances” and “eventuality tokens” to be synonymous.

of the latter. *causalComplex* says that s is the set of all presumable or non-presumable eventualities that are involved in causing e_{x2} . Obviously, e_{x1} belongs to s . Thus, in the light example, the predicate *cause* applies to the flipping of the switch, while the states of the bulb, the wiring, and the power supply would all be in the causal complex s . Several axioms characterize the predicates *cause* and *causalComplex*. Some of them relate causality with time⁶, some relate causality with probability, and so on Hobbs (2005).

It is clear that the theory must *not* include an axiom stating that, whenever a causal relation c_x and its cause e_{x1} really exist, the corresponding effect e_{x2} really exists too. The inclusion of such an axiom would lead to a non-defeasible causality. Rather, we need an axiom stating that an effect really exists just in case all the eventualities in its causal complex really exist:

$$\begin{aligned} &(\text{forall } (s \ e) \\ &\quad (\text{if } (\text{and } (\text{causalComplex } s \ e) \\ &\quad\quad (\text{forall } (e_1)(\text{if } (\text{member } e_1 \ s) (\text{Rexist } e_1)))) \\ &\quad (\text{Rexist } e))) \end{aligned}$$

Nevertheless, as pointed out above, we can never specify all the eventualities in a causal complex. Even in simple sentences like (1.a), the eventualities in the causal complex are not easy to list, and the real causes may not coincide with what we *think* the causes are in that context. For example, recalling our analysis of (1.a) above:

$$\begin{aligned} e_{c1} &= \text{"John studied hard"} \\ e_{c2} &= \text{"John passed the exam"} \\ e_{d1} &= \text{"John was tired during the exam"} \\ e_{d2} &= \text{"John did not pass the exam"} \\ c_c &= \text{"}e_{c1} \text{ causes } e_{c2}\text{"}; c_d = \text{"}e_{d1} \text{ causes } e_{d2}\text{"} \end{aligned}$$

One approach at this point would be to say that both e_{c1} and the negation of e_{d1} belong to the causal complex of e_{c2} , with e_{c1} being the non-presumable cause of e_{c2} . But this would mean that *not being tired during exams* is a kind of “precondition” for *passing exams* by *studying hard*, which is obviously false in many contexts. Note, however, that there is an arbitrary quality to what we designate as being in a causal complex, because causality forms chains and we can start the chain at any point. *John was tired* caused the situation that *he did not manage to concentrate*, which caused the situation that *he made a lot of errors in the exam*, which caused the situation that *the teacher decided to fail him*. One could argue that the last of these eventualities is the real cause of e_{d2} . Similarly, one could argue that e_{c1} is not the real cause of e_{c2} : *John studied hard* causes the situation that *he makes few errors in the exam* . . . and *the teacher decides not to fail him*. The predicate *cause* is defeasibly transitive, however, so these considerations do not affect our account of concession. Furthermore, we do not take the negation of e_{d1} as necessarily belonging to the causal complex for e_{c2} . Rather, we claim that e_{d1} , besides being the cause of e_{d2} , is the cause of another eventuality e_{dp} that is inconsistent with an element e_{cp} in the causal complex for e_{c2} .

⁶As argued also by Giordano and Schwind (2004), the effect caused by an eventuality can take place in the current or in a subsequent instant.

In (1.a), e_{cp} may be simply *John does not have any particular health problem that jeopardizes his passing the exam*. e_{d1} caused both John’s failure and an health status that jeopardizes the passing of his exam. This is what we mean here by “denying of an expectation”.

In our analysis of concession, we distinguish between abstract causalities like *hard studying causes passing exams*, and causality tokens like *John’s tiredness caused John’s failure*. Note that asserting (*Rexist c*) on an abstract causal relation c amounts to asserting (*Rexist c'*) for any (partial) instance c' of c . But recall that *cause* is only defeasible. Both the abstract causal principle and its partial instance are simplified stand-ins for rules that involve entire causal complexes, not all of whose elements may obtain. Thus, just because *hard studying causes passing exams*, we cannot invariably conclude that if John really studied, he really passed the exam.

4 The meaning of concessive relations

Our basic claim is that the meaning of concessive relations is triggered by a contrast between two causal relations c_c and c_d such that one or more eventualities in the causal complex of e_{c2} (the expectation created by c_c), is denied by e_{d2} (the effect of c_d). c_c , c_d , e_{d2} , and e_{c1} (the cause in c_c) really exist in the world, or are at least believed to exist by the speaker/writer. Furthermore, all eventualities in the causal complex for e_{d2} , including the non-presumable cause e_{d1} , which is unknown in many cases, really exist too. Arg_{cexp} conveys e_{c1} , while Arg_{dexp} conveys e_{d2} .

We also claim that in all cases of concession it seems that what really creates the expectation is a causal relation c_c^a that is an abstraction of c_c . c_c really exists in the world precisely because c_c^a really exists and c_c is a partial instance of it. In other words, the real existence of c_c is inherited from c_c^a . On the other hand, there is not necessarily an abstract counterpart c_d^a for c_d that also really exists in the world. For instance, in (1.a), it seems that what creates the expectation is the assumption that the causal relation *studying hard causes passing exams* (c_c^a) really exists in the context. *John’s hard studying causes John’s passing exams* (c_c) is just an instance of c_c^a . This instance really exists in the world too. However, since causality is defeasible, the fact that *John really studied hard* (e_{c1}) does not entail the real existence of *John really passed the exam* (e_{c2}). In fact, this is precisely denied by Arg_{dexp} : *John did not pass the exam* (e_{d2}). The cause of John’s failure, e.g., John’s tiredness (e_{d1}), is (or is the cause of) an eventuality e_{dp} that is) inconsistent with an element e_{cp} of the causal complex for (e_{c2}), namely, *John does not have any particular health problem that jeopardizes the passing of his exam*. Note that we do not necessarily infer that *being tired causes failing an exam*: tiredness was the cause of the failure in this particular scenario only. Therefore, we assert that c_d really exists, but we do not advocate the existence of a more abstract causal relation c_d^a that really exists too.

To summarize, the semantics of concession we propose is formalized in (2). The conjuncts (*Rexist c_c*) and (*Rexist e_{d1}*) have been omitted in (2) because they may be inferred from (*Rexist c_c^a*) and (*Rexist e_{d2}*). s_c is the causal complex associated with c_c . e_{c1} and e_{d2} are given to us in Arg_{cexp} and Arg_{dexp} respectively, while all other eventualities may be inferred by abduction from the contextual knowledge; some hints about how this may be done are provided in Hobbs (2005).

$$\begin{aligned}
(2) \quad & (\text{exist } (c_c \ c_c^a \ e_{c1} \ e_{c2} \ c_d \ e_{d1} \ e_{d2} \ s_c \ e_{cp} \ e_{dp}) \\
& (\text{cause}' \ c_c \ e_{c1} \ e_{c2}) \wedge (\text{cause}' \ c_d \ e_{d1} \ e_{d2}) \wedge (\text{Rexist } c_c^a \ c_c) \wedge \\
& (\text{partialInstance } c_c \ c_c^a \ c_c) \wedge (\text{Rexist } c_d) \wedge (\text{Rexist } e_{c1}) \wedge \\
& (\text{Rexist } e_{d2}) \wedge (\text{cause } e_{d1} \ e_{dp}) \wedge (\text{Rexist } e_{dp}) \wedge \\
& (\text{inconsistent } e_{c2} \ e_{d2}) \wedge (\text{causalComplex } s_c \ e_{c2}) \wedge \\
& (\text{member } e_{cp} \ s_c) \wedge (\text{inconsistent } e_{dp} \ e_{cp}))
\end{aligned}$$

Let us now examine how the semantics given in (2) applies for corpus examples tagged as “expectation” or “contra-expectation”. Let us analyze (1.b) in the light of the semantics proposed in (2). The abstract causality that creates the expectation (c_c^a) is *Something that is considered healthy for humans is advisable for them*⁷. This is partially instantiated in *Since running is considered healthy for persons with heart problems, it is advisable for them* (c_c). Nevertheless, the fact that *running is really considered healthy* in the context (e_{c1}) does not suffice to assert that *running is really advisable for persons with heart problems* (e_{c2}). There is a particular reason why *running is not advisable for persons with heart problems* (e_{d2}), e.g. *their hearts do not tolerate a heartbeat increase* (e_{d1}). Since running causes a heartbeat increase, *the heart can tolerate a heartbeat increase* (e_{cp}) is in the causal complex for e_{c2} and it is inconsistent with e_{d2} .

Similarly, in (1.c), which is taken from the PDTB, it is true that *representing a low percentage of the population causes controlling low percentage of income* (c_c^a). Therefore, *they represent 2% of population* (e_{c1}) causes *they control low percentage of income* (e_{c2}). Nevertheless, e_{c2} does not really exist in the context, in that it is inconsistent with *they control nearly one-third of income* (e_{d2}). There must be another reason for why e_{c2} does not hold. For instance, either *they are very rich*, or *they do not have as many basic expenses as other people*, or a more complex condition. This unknown cause, i.e. e_{d1} , both makes e_{d2} true and e_{c2} false in the context.

The last example highlights the point that finding the eventualities involved in (2) is strongly dependent upon contextual knowledge. 2% is not taken to be a low percentage in *any* context. For instance, 2% mercury in the water may be considered a high percentage of pollution. Analogously, one third may be considered a high percentage in that context, especially if compared with 2% of population, but it may be a low or medium percentage in many other contexts. The analysis of examples (1.d-e) in terms of the definition in (2) is analogous.

5 A survey of concessive relations in PDTB 2.0

PDTB 2.0 contains 1193 tokens of *explicit* connectives which are annotated with one sense tagged as “Concession”, “contra-expectation” and “expectation”. There are also another 20 tokens that have been annotated with double senses, one of which is the concessive type or subtypes. Table (1) shows the distribution of concessive labels for the 1193 tokens. Explicit connectives with a concessive label assigned to less than 10 tokens are grouped under “other”. The rest of the connectives shown in Table (1) amount to 98% of all “contra-expectation” and 95% of all “expectation” tokens. The

⁷This is a paraphrase of *Something being considered healthy for humans causes it to be advisable for humans*.

Table 1: Concessive labels in PDTB 2.0

CONN	“contra-exp.”	“exp.”	“Concession”	Total
although	21	132	1	154 (13%)
but	494	12	2	508 (42.5%)
even if	3	31	1	35 (3%)
even though	15	52	5	72 (6%)
however	70	2	5	77 (6.5%)
nevertheless	19	0	0	19 (1.5%)
nonetheless	17	0	0	17 (1.5%)
still	79	2	1	82 (7%)
though	30	53	1	84 (7%)
while	3	79	1	83 (7%)
yet	32	0	0	32 (2.5%)
other	13	17	0	30 (2.5%)
Total	796	380	17	1193

most common connective annotated with the ‘Concession’ type or one of its two subtypes is “but” with 508 tokens (42% of all concessive labels), followed by “although” with 154 tokens (13% of all concessive labels).

We are currently evaluating the robustness of the proposed refined semantics for concessive labels in PDTB 2.0 starting with the most the most common concessive connectives. While the validation process for the entire corpus is still work in progress, preliminary results on 25% of ‘but’ tokens indicate that the semantics of concession based on defeasible causality applies straightforwardly to more than 60% of the data. In future work, we hope to be able to offer a more comprehensive account of all the concessive labels in PDTB 2.0 including cases of concession in which the created expectation arises from an implication rather than from a causal relation (about 23%), as in (3)

- (3) Although working for U.S. intelligence, Mr. Noriega was hardly helping the U.S. exclusively. (expectation)

In (3), it is strange to say that *working for U.S. intelligence* normally “causes” *helping U.S. exclusively*. Rather, the former seems a kind of necessary condition or job requirement for the latter: *working for U.S. intelligence* implies (among other things) *helping U.S. exclusively*. Suppose that someone discovers that Mr. Noriega is not helping the U.S. exclusively. Mr. Noriega is arguably breaking a rule or flouting an expectation. Therefore, *working for U.S. intelligence* “implies” rather than “causes” *helping U.S. exclusively*.

It is unsurprising that there are cases of concession based on implication rather than causality, because the two concepts are very close to each other. One could think of implication as a kind of abstract, informational, or “denatured” causality. Both obey a kind of (defeasible) modus ponens. When the cause or antecedent happens or holds, so does the effect or consequent. The other key property of causal complexes is that

all the eventualities in it are relevant, in a sense that is made precise in Hobbs (2005). This notoriously does not hold for material implication, but as many have argued, it probably does hold for felicitous uses of our everyday notion of implication. In addition, there are easy conversions between causality and implication. If A causes B , then the fact that A happens (defeasibly) implies that B happens. If P implies Q in the everyday sense, then one's belief in P (defeasibly) causes one's belief in Q . In fact, implicational cases of concession could be viewed as instances of metonymy, where "believe" is the coercion relation, and hence really causal cases of concession.

6 Conclusion

We presented a formal description of the meaning of concession, a substantial refinement of the rough semantics given in the manual of sense annotations of connectives in PDTB 2.0. Our analysis builds on Hobbs's logic of defeasible causality enabled by the crucial distinction between causes and causal complexes. Our basic claim is that concession is triggered by the contrast between two causal relations. The causal relation between the content of one argument of the relation and some implicit eventuality (the expectation created based on the content of the argument) and the content of another causal relation, that between the eventuality described in second argument and *its* implicit cause. This second causal relation picks an element of the causal complex that we would normally assume to hold and challenges it, hence the notion of defeasible causality.

This work illustrates the mutual benefit that corpus annotation and formal analysis can provide to each other. Corpus examples constitute a forcing function on the formal analysis; definitions must accommodate the complexities one finds in the real world. On the other hand, all good annotation rests on solid theory, and formal analysis can help in the adjudication of difficult examples. The particular analysis we give in this paper for the concession relation can clarify issues that arise in annotation, and can also form the basis for recognizing these relations using a knowledge-rich inferencing system.

References

- Achinstein, P. (1965). 'Defeasible' Problems. *The Journal of Philosophy* 62(21), 629–633.
- Asher, N. (1993). *Reference to Abstract Objects*. Kluwer, Dordrecht.
- Bach, E. (1981). On Time, Tense, and Aspect: An Essay in English Metaphysics. In P. Cole (Ed.), *Radical Pragmatics*, pp. 63–81. Academic Press, New York.
- Bell, J. (1999). Primary and secondary events. In M. Thielscher (Ed.), *Proc. of the IJCAI-99 Workshop on Nonmonotonic Reasoning, Action and Change*, pp. 65–72.
- Bell, J. (2003). A common sense theory of causation. In P. Blackburn, C. Ghidini, R. Turner, and F. Giunchiglia (Eds.), *Modeling and Using Context: Fourth International and Interdisciplinary Conference, Context 2003*, Berlin, pp. 40–53. Springer-Verlag.

- Giordano, L. and C. Schwind (2004). Conditional logic of actions and causation. *Artificial Intelligence* 157(1–2), 239–279.
- Giunchiglia, E., J. Lee, V. Lifschitz, N. McCain, and H. Turner (2004). Nonmonotonic causal theories. *Artificial Intelligence* 153(1–2), 49–104.
- Hobbs, J. (1983). An Improper Treatment of Quantification in Ordinary English. In *Proc. of the 21st Annual Meeting of the Association for Computational Linguistics*, Cambridge, Massachusetts, pp. 57–63.
- Hobbs, J. (1995). Monotone Decreasing Quantifiers in a Scope-Free Logical Form. In K. van Deemter and S. Peters (Eds.), *Semantic Ambiguity and Underspecification*, CSLI Lecture Notes, pp. 55–76. CSLI.
- Hobbs, J. (1998). The Logical Notation: Ontological Promiscuity. In *Discourse and Inference*, Chapter 2.
- Hobbs, J. (2005). Towards a Useful Notion of Causality for Lexical Semantics. *Journal of Semantics* 22(2), 181–209.
- Kayser, D. and F. Nouioua (2008). From the Description of an Accident to its Causes. submitted to *Artificial Intelligence*.
- McCarthy, J. (1977). Epistemological Problems of Artificial Intelligence. In *Proc. of International Joint Conference on Artificial Intelligence*, Cambridge, Massachusetts, pp. 1038–1044.
- Miltsakaki, E., L. Robaldo, A. Lee, and A. Joshi (2008). Sense Annotation in the Penn Discourse Treebank. In *Proc. of Computational Linguistics and Intelligent Text Processing*, Volume 4919 of *LNCS*, pp. 275–286. Springer.
- Prasad, R., N. Dinesh, A. Lee, E. Miltsakaki, L. Robaldo, A. Joshi, and B. Webber (2008). The Penn Discourse Treebank 2.0. In *Proc. of the 6th Int. Conf. on Language Resources and Evaluation*.
- Prasad, R., E. Miltsakaki, N. Dinesh, A. Lee, A. Joshi, B. Webber, and L. Robaldo (2008). The Penn Discourse Treebank 2.0. Annotation Manual. Technical Report IRCS-06-01, IRCS Technical Report, Institute of Research in Cognitive Science, University of Pennsylvania.
- Shoham, Y. (1990). Nonmonotonic reasoning and causation. *Cognitive Science* 14, 213–252.
- Simon, H. (1991). Nonmonotonic reasoning and causation: Comment. *Cognitive Science* 49, 517–528.
- Webber, B. and A. Joshi (2003). Anchoring a lexicalized tree-adjoining grammar for discourse. In M. Stede, L. Wanner, and E. Hovy (Eds.), *Discourse Relations and Discourse Markers: Proceedings of the Conference*, pp. 86–92.
- Webber, B., A. Joshi, M. Stone, and A. Knott (2003). Anaphora and discourse structure. *Computational Linguistics* 29(4), 545–587.