

Zachary Dimovski

The Origins of Classes of Projective Content

“Projection” can be understood as a linguistic phenomenon in which content implied in a sentence persists even when we negate the sentence. Much of the discussion of projective content has centered on presupposition, but focusing only on presupposition would exclude many implications that do project but typically have not been understood as a type of presupposition (Simons et al., 2011). Tonhauser et al. (2013) presents a taxonomy of projective contents that groups the implications into four distinct classes. The taxonomy divides the projective contents along two axes, depending on whether the contents impose a Strong Contextual Felicity Constraint (SCFC) or whether they have an Obligatory Local Effect (OLE). While the authors provide diagnostics to determine which contents impose the SCFC or have an OLE, they do not offer a reason why particular projective contents are associated with particular phenomena. In this paper, I will offer an account of why projective contents belong to the class of contents that they do.

Simons et al. (2011) discusses projection and provides a way of understanding why contents project. The authors define projection as an implication that persists under an entailment-cancelling operator. The sentence in (1a) carries at least two implications; the first is that Archie used to eat meat, and the second is that Archie no longer eats meat.

(1a). Archie stopped eating meat.

The implication that Archie does not eat meat is the entailment of sentence (1a), but sentences (1b) and (1c) do not carry this entailment.

(1b). Archie did not stop eating meat.

(1c). Did Archie stop eating meat?

This means that negating the sentence and questioning the sentence put (1a) under an entailment-cancelling operator. Asserting that Archie has not stopped eating meat clearly does not entail that Archie does not eat meat –it entails the opposite – but it still does imply that Archie used to eat meat. Asking whether Archie has stopped eating meat entails neither that Archie does not eat meat nor the opposite, but it too carries the implication that Archie used to eat meat. This shows that there is at least one implication that persists in both of these cases.

The same authors present a similar understanding of projection in Tonhauser et al. (2013). There, they develop a diagnostic for determining whether an implication projects. If a trigger imposes an SCFC (something that I will cover in detail later), then it is projective. If it does not impose an SCFC, then we can determine if an implication projects based on the “family of sentences” diagnostic. Roughly, a content that does not impose an SCFC is projective when it is implied when a sentence containing a trigger for that projection is asserted, negated, asked as a question, used as the antecedent of a conditional (1d), or when used in an epistemic modal (1e) (Simons et al. 2011).

(1d). If Archie stopped eating meat, then he will probably order the salad.

(1e). It is possible that Archie stopped eating meat.

Simons et al. (2011) gives an account of projection that bases the phenomenon on what is at-issue in discourse. The authors point out that projections are separate from the main point of the utterance; instead, they argue, projections seem to be something that a speaker takes for granted in uttering a sentence. Sentential operators – such as negation, conditionals, and modals – are limited in scope to the main point of the sentence, and therefore, would have no effect on projective content. By “at-issue,” the authors mean that something is at-issue with regard to a “Question Under Discussion” (QUD). The QUD is a question that corresponds to the topic of

discourse, which interlocutors address in different ways with many potential methods throughout a given conversation. The authors' hypothesis about "what projects and why" is that all implications that are not at issue relative to the QUD can project. This is the case because operators target only at issue-content, meaning that they will have no effect on projections.

The categorization of projective contents is the focus of Tonhauser et al. (2013). The authors divide the different sorts of projective contents along two axes. These are the "Strong Contextual Felicity Constraint" and the "Obligatory Local Effect." In the authors' terms, the SCFC involves an utterance's acceptability in an "m-neutral" context, where m-neutral context is a context that implies neither a projective content nor its negation. If a sentence containing a trigger of the projective content, *m*, is acceptable in an m-neutral context, then the projection does not impose an SCFC. If a sentence containing a trigger of *m* is unacceptable in an m-neutral context, but acceptable in an m-positive context (one in which the projection, *m*, is implied), then the projection imposes an SCFC. For example, in a context that does not imply that there is some salient person being discussed, it would be infelicitous to use an utterance containing a pronoun, because the referent of the pronoun would only be clear in the speaker's mind, but not in the context of discourse. This implies that something like a pronoun triggers a projection that imposes an SCFC.

The Obligatory Local Effect concerns projective content embedded under a belief predicate. When a sentence containing a trigger of a projection, *m*, is the complement of a belief predicate, *m* has an OLE if and only if *m* is necessarily a part of the content that is believed. We can understand a projective trigger to have an OLE when belief in the projection is necessarily attributed to the belief-holder. For example, it would be wrong to say that I believe that it has stopped raining, but also that it was not raining. This implies that "stop" triggers some sort of

projection that has an Obligatory Local Effect. The authors provide an in-depth diagnostic for OLE covering five different scenarios, which for the purposes of this paper, it will not be necessary to reproduce.

Although Tonhauser et al. (2013) do not use this terminology, we can also understand the SCFC in terms of common ground. A roughly equivalent way of characterizing the SCFC is to say that something imposes an SCFC when a felicitous utterance of a projective content trigger requires that the projection be a part of the interlocutors' common ground.

Common ground can be understood as interlocutors' mutual recognition of what is commonly believed. In order for some Φ to be common ground, interlocutors must believe Φ and believe that all other members believe Φ and believe that all other members believe that all other members believe Φ and so on. This preliminary understanding, however, is somewhat inadequate. Stalnaker (2002) discusses common ground and through this discussion, shows why a characterization of common ground as being based on common belief must be refined and expanded upon. He criticizes the idea that common ground should be based upon beliefs, and this this comes about through his discussion of the phenomenon known as accommodation.

Accommodation is a fairly intuitive and commonplace phenomenon. If a person were to tell a stranger that she must meet her brother for lunch, the stranger would likely come to believe that the person has a brother. This, however, was not asserted by the person, but was merely presupposed. A speaker can use an assertion to add to the common ground in a clear way – by making some fact known and expecting that his or her interlocutor will accept it.

Presuppositions, however, do not function this way. Since presuppositions seem to involve what a speaker takes for granted in an utterance, felicitous use of a presupposition-containing sentence should require that this presupposition already be in the common ground. This does not match

our intuitions, since it seems perfectly natural that someone could mention to a stranger that she must meet her brother, even when it is clear that the stranger could not know beforehand that the person had a brother. Accommodation is what fixes this problem. With accommodation, the presupposition is added to the common ground at the time of utterance. According to Stalnaker (2002), this can occur because of the common belief that a person would not purposefully be dishonest. It is reasonable, then, for a person to believe that her interlocutor would come to believe that she has a brother. Since the person believes that she has a brother and believes that their interlocutor will come to believe this, and since the interlocutor does indeed come to believe the presupposition, it can be considered a common belief and therefore, part of the common ground.

Accommodation, however, is not simply limited to beliefs. Stalnaker (2002) points to some examples that show that this is the case and use these to expand the notion of common ground beyond simply common belief. He discusses how accommodation and beliefs can diverge in the case of “defective contexts,” or contexts in which not all beliefs about the common ground are correct. A defective context can become clear when a speaker makes an utterance that indicates that he or she thinks that it is a common belief that something is true, but when the listener does not believe that this presupposition is true. A speaker might point out “the man drinking a martini,” and the listener might accept this, even when he knows the liquid to be water. Stalnaker (2002) argues, then, that rather than belief, we should characterize accommodation and common ground in terms of “acceptance.” He characterizes acceptance as an attitude toward a proposition that includes belief, presumption, or assuming something for the sake of argument.

The definition of common ground with which Stalnaker (2002) concludes is that “it is common ground that Φ in a group if all members *accept* (for the purpose of the conversation) that Φ , and all *believe* that all accept that Φ , and all *believe* that all *believe* that all accept that Φ , etc.” (716).

With this, we can understand accommodation as being less about updating and accepting shared beliefs. Instead, it is a way to keep a conversation moving forward. When a presupposition is accommodated, it is because the listener sees no reason to object, whether it be because they have come to have a new belief or because they find a presupposition unobjectionable enough to grant it to the speaker in order to keep the conversation flowing. The common ground, rather than being about mutual alignment of shared beliefs, is the set of assumptions that speakers judge to be relatively unobjectionable.

Tonhauser et al. (2013) does not study pragmatics and common ground in general, but instead focuses on projection specifically in Guarani and English. The research on projective content in the South American language, Guarani, and its comparison to the completely unrelated English language provides one of the most important aspects of the analysis of projection in Tonhauser et al. (2013), since it allows the authors to suggest that the taxonomy of projective content that they present is universal.

The authors presented sentences with projective content in different contexts to native speakers of Guarani. Based on both the responses of the Guarani speakers and the authors’ native English-speaking intuitions, they determined that projective content functions the same way in both Guarani and English. The same contents project, and each projection has the same features in both languages, e.g. the polar implications of *almost* and *aimete* (Guarani for almost) do not impose the SCFC and have an OLE in both languages. The only differences between Guarani

and English with regard to projection arise due to non-pragmatic linguistic differences. For example, there is a gender-attribution projection associated with third-person singular pronouns in English, but the Guarani third-person singular pronoun does not mark gender. The fact that, where applicable, projection functions the same way in the two unrelated languages suggests that projection and features associated with particular projective contents are universal. This could mean that projection is based on some common feature of language, discourse, or human cognition. In this paper, I will attempt to determine what universal rules of discourse might lead to the SCFC and OLE.

One of the classes of projective content in Tonhauser et al. (2013) imposes the SCFC, but does not have the OLE. The authors label this “Class D Projective Content” and claim that this class of projective content has been previously unrecognized. The projections included in Class D are salience conditions for triggers such as “*too*.” This projection implies that there is a salient referent to the word “*too*.” In order for a felicitous utterance of a sentence such as (2), there must be at least one person other than Matthew driving his or her car.

(2). Matthew is driving his car *too*.

This existence projection, however, is not enough. If someone were to ask me what Matthew is doing, and I responded with (2), this would be infelicitous. Surely, at least one other person in the world is driving his or her car, but this is not enough to make the sentence acceptable. There must also be a salient referent of “*too*” that the listener can identify in order for a felicitous utterance of (2).

Other examples of Class D projective content are the salience projection associated with “*that* (Noun),” and although not included in the original taxonomy, the salience projection associated with pronouns. That there is a salience projection associated with pronouns that

belongs to Class D can be demonstrated using the diagnostic presented in Tonhauser et al. (2013). The salience implication of a sentence in (3a) has strong contextual felicity because uttering it is unacceptable in a context where neither the identifiability nor non-identifiability of the referent are implied.

(3a). She talks too fast.

(3b). George believes she talks too fast.

Uttering (3a) unexpectedly, when it is not clear that there is a relevant “she” that is being referenced, would, according to Tonhauser et al. (2013), form an m-neutral context since the speaker fails to make the reference clear. Since (3a) is unacceptable in this context, but would be acceptable in a context where the speaker can easily identify the referent (m-positive), then, according to the diagnostic, it is also projective. If we take (3a) and embed it under a belief-predicate, then we get a sentence like (3b). The sentence in (3b) is acceptable when “she” has an identifiable discourse referent but George is not aware of this. Let us say that in this case, Maria is the referent of “she.” George need not know that Maria is a salient referent in John and Michael’s conversation in order for it to be acceptable that Michael utters (3b). This indicates that the salience projection of pronouns does not have an OLE, and thus belongs to the class of projective contents called Class D in Tonhauser et al. (2013).

It is important to note that the other class of projective contents that impose the Strong Contextual Felicity Constraint, Class A, contains the same triggers as Class D. Rather than projecting a salience condition, Class A triggers project an existence condition. For example, there are at least two projections associated with a pronoun. The first is that the referent is salient and this projection belongs to Class D, while the other is that the referent exists and this belongs to Class A.

At this point, it would be appropriate to discuss the classification of projective contents associated with possessive noun phrases and definite descriptions. The classification of the projections triggered by these phrases will be crucial to my discussion of the SCFC.

The possessive noun phrase projects a possessive relation presupposition. By referring to an object as “my x,” I presuppose my possession of the object. Tonhauser et al. (2013) analyze this projection in both Guarani and English and show that it does not have an SCFC and does not impose an OLE. This places the projection within “Class B.”

The possessive noun phrase should also project an existence presupposition. As discussed earlier, in order to utter felicitously a sentence like, “I have to pick up my sister,” the existence of my sister must already be in the common ground. These presuppositions, however, are readily accommodated. Upon utterance, the existence of my sister will be added to the common ground unless my interlocutor has some strong reason to find this presupposition objectionable.

This is also the case for definite descriptions. In order to assert that “the King of France is bald,” the existence of the King of France must be common ground. This is another example of readily and easily accommodated projective content. If a listener knows very little about the French political system or finds that my false belief is not too objectionable for the purposes of conversation, then the existence of the King of France will become common ground.

In order to determine whether these projections have a Strong Contextual Felicity constraint, we must examine these projections while taking their ease of accommodation into account. The SCFC involves acceptability in a context that neither implies the projection nor its opposite. If something does not have an SCFC, it is appropriate to utter it when the context does not imply it or when it is not in the common ground. If accommodation did not occur, then existence projections triggered by phrases like possessive noun phrases and definite descriptions

would have an SCFC, since they do require their presuppositions be in the common ground. Since, however, the existence presuppositions are so easily accommodated, they *are* appropriate in an m-neutral context. Tonhauser et al. (2013) seem to agree that existence projections of possessive noun phrases and definite descriptions do not have strong contextual felicity. Although the authors do not classify these projections, they do mention that projections that impose an SCFC have strong constraints on context that cannot be satisfied by accommodation and that those projections that do not have an SCFC might either have a weak constraint on context that can be satisfied through accommodation or could have no contextual constraint whatsoever.

These projections, like all other projections concerning existence, do have an Obligatory Local Effect. The fact that sentences like (4) are contradictory demonstrates this.

(4). #Jones believes the King of France is bald, but that he does not exist.

If Jones has a belief about the King of France, then we necessarily attribute to Jones the belief in the existence of the King of France.

Since the existence projections triggered by possessive noun phrases and definite descriptions do not impose an SCFC and do not have an OLE, we can put them in the class of projective content called “Class C” in Tonhauser et al. (2013). If accommodation did not occur, they would impose an SCFC and would be grouped in “Class A” along with the existence projections associated with pronouns or the word “too.”

If we were to take the everyday English sentences that give rise to Class A and D projective contents and change them so that they are no longer indexical (for example, replacing all pronouns with the nouns to which they refer), we would no longer have Class A or D projective contents. A sentence such as 5a would simply become 5b.

(5a). She plays the saxophone.

(5b). My sister plays the saxophone.

Class D contents, which rely on identifiability or salience, would not exist in these situations, because identifiability would no longer be relevant. Whereas felicitous use of “she” involves a demonstrative “pointing-to” an identifiable discourse-referent, uttering “my sister” does not. In these cases, Class A content would not exist simply because they would be the same as Class C contents, e.g. the projection that “she” exists would become the projection that “my sister” exists. What is significant here is that the existence projection associated with “she” imposes a Strong Contextual Felicity Constraint, but the existence projection of “my sister” does not. Since Class D contents go away and Class A contents become Class C, doing away with the need for referents to be identifiable does away with the SCFC. This suggests a link between the Strong Contextual Felicity Constraint and identifiability that I will investigate here.

Since, as discussed earlier, imposing an SCFC corresponds to an inability to be accommodated, a good way to approach an analysis of the SCFC is to work out the relationship between identifiability and accommodation.

Accommodation occurs when a hearer does not take issue with a projection. They accept this projection for the purposes of conversation. Accommodation can occur when our belief toward a projective content is neutral or when we disbelieve the projection and are willing to grant it to the speaker. If, when talking to a close friend, he says to me the sentence in (6), I will not question the existence of his brother.

(6). I am meeting my brother for lunch.

If I know that this friend has a brother, then we are in agreement and there is no need for accommodation. I take this projection for granted just as much as the speaker does. If I am

thoroughly convinced that the speaker does not have a brother, then accommodation likely will not occur. The speaker and I are in disagreement, and I might ask the speaker what exactly he or she means, or I could interpret this as a non-standard use of the word “brother,” or I might worry about the speaker’s mental health. I will not begin simply to believe that the speaker has a brother.

There are, however, cases in which I can accommodate a projection despite believing its opposite. If I think that the speaker was an only child, but do not know him well, I will assume that I am mistaken, and as a result, I will update my belief. Accommodation here will occur based on some subjective measure of how convinced I am of my belief in the opposite of the projection and my trust in the speaker. This is pretty clearly something that varies person-to-person. A more trusting and timid person will likely accommodate information they disbelieve more readily than others will.

Cases in which I have no belief one way or another in the truth or falsehood of a projection are likely cases where accommodation is most likely. This state of mind roughly corresponds to the m-neutral context used in the SCFC in Tonhauser et al. (2013). If I meet a stranger on the street, and he says that he has to meet his brother for lunch, I will easily accommodate this information and come to believe that the speaker has a brother.

Since imposing an SCFC is approximately the same as an inability to be accommodated, projections with an SCFC are projections where we will always be unwilling to grant the projected content to the speaker; however, there must be some reason as to why listeners do not accommodate Class A and D projective contents. My claim is that the triggers that cause Class A and D projective contents never project something that a listener is able to accommodate, since

they necessarily imply something that a speaker simply cannot accept for the purposes of conversation. Saliency illustrates this phenomenon.

One way in which accommodation can occur is through adding common beliefs to the common ground, but the saliency or identifiability projections associated with Class D content imply something about which a listener cannot be neutral. Since a listener necessarily *does* have a definite belief as to whether he or she can identify a salient discourse referent, the listener *must* either agree or disagree with the speaker. The speaker believes that the listener can identify a referent. Either the listener can identify the referent, in which case, they believe that they can identify the referent, and are in agreement with the speaker, or they cannot identify the referent, believe that they cannot identify the referent, and are in disagreement with the speaker. This leaves no room for accommodation through updating beliefs, since there is no neutral state (m-neutral context) in which the listener is not sure that they can identify a referent; there is only total agreement or total disagreement.

Updating beliefs, however, is not the only way that accommodation occurs; other forms of “acceptance” can cause accommodation. In these cases, the speaker appears to accept a presupposition in order to keep the discourse moving forward. In uttering a sentence containing a pronoun, a speaker takes for granted that the listener can identify the referent of this pronoun. Not only is this something that a listener cannot be neutral about, it is also not something that the listener will not object to for the sake of conversation, because failing to object would hinder conversation. If I think that a listener does not have a brother, I might not take issue with the projected existence of the brother so that I do not disrupt conversation, but when the presupposition is that I can identify a referent, conversation will fail if I cannot identify that referent. If two interlocutors are unclear about the referent of their conversation, they are not

clear about the topic of discourse. These projections will not be accommodated, because failing to make an anaphoric referent identifiable is always objectionable. Accommodation occurs so that we can continue moving a conversation forward, but such a presupposition failure on the part of the speaker will cause miscommunication, meaning that accommodating the projection will be more detrimental to discourse than will granting the identifiability to the speaker.

Stalnaker (2002) makes a similar point. He argues that that accommodation is impossible in defective contexts in which a speaker fails to make a reference clear. This is the case because of the expectation that speakers obeying a cooperative principle of communication will assume that all information needed to interpret their utterances is available to the listener. The speaker, in uttering a Class D projective trigger in a situation in which the speaker is unable to identify the anaphoric referent, has both failed at proper communication and has made some improper assumption about my mental state. The speaker, therefore, is being uncooperative. Unlike other cases on uncooperativity, such as lying or deceiving, the listener will not (and likely cannot) mistakenly believe that the speaker is being cooperative. A listener cannot be deceived into identifying an anaphoric referent.

If it is indeed the case that the SCFC occurs because triggers associated with Class A and D contents fail to put the listener in an m-neutral context, we would expect that in a scenario where a listener is neutral as to whether they are able to identify a referent, they would accommodate Class A and D projections. These situations might be unlikely or even impossible, but I think there are certain scenarios where they could occur, or at least, there are certain scenarios that closely approximate such a situation. Take, for example, the story presented in example (7).

(7). Bella is at a gathering and must pretend to be her identical twin sister, Claire. Claire does not wear glasses, but Bella does, and she is nearly blind without them. Bella goes to the gathering without her glasses in order to look more like her sister. A guest at the gathering comes up to Bella and says, "He is very tall." Bella is not able to see the person that the guest is talking about, but in order not to blow her cover, she says, "Indeed, he is."

In (7) Bella puts herself into a situation in which she is not sure whether she should be able to identify something referenced in the visual field, because she is pretending to be Claire and is not sure whether Claire would be able to identify the referent in this situation. In this case, she accommodates that she should be able to identify the referent of the word "he." It is difficult to say whether she accommodates the existence (Class A) projection associated with the pronoun, as believing that a referent is identifiable coincides with believing that a referent exists. In this scenario, it seems as if the primary focus is on the identifiability, rather than the existence. Even if Bella does not necessarily believe that she can identify the referent, at the very least, she plays along, but unless she is very untrusting, she will likely believe that she ought to be able to identify the anaphoric referent.

Accommodation can occur in (7), because of Bella's assumptions about the speaker. It is not common ground that the referent is identifiable in this case. Bella, however, assumes that her interlocutor will not be dishonest on purpose. Bella has no reason to believe that the man would approach her and utter the sentence if the referent were not seemingly identifiable. She therefore finds it unobjectionable that she ought to be able to identify the referent. That the referent of "he" is identifiable becomes accommodated.

These scenarios may not necessarily be as strange as the previous example makes them seem. The example in (8) presents a much more commonplace scenario, and here, it seems as if the Class A existence projection of the pronoun can be accommodated. The previous example tries to construct a case in which the listener is unsure whether she can identify the referent, but in a more natural scenario, a listener might accommodate a Class A projection through simply granting to the speaker that the referent exists, and making some small adjustment in order to better be able to identify it. In this way, we can come up with a scenario in which the listener does seem to accommodate existence conditions of a pronoun.

(8). Archie and Bella are walking down the street and a very tall man is directly ahead of them. Bella has taken off her glasses in order to wear her non-prescription sunglasses.

Without corrective lenses, Bella is unable to see at far distances. Archie makes a gesture pointing forward and says, "He is very tall." Bella switches to her prescription glasses, spots the man, and agrees with Archie.

In (8), Bella is unsure whether there is any actual person to which Archie is referring. Archie may or may not know that Bella cannot identify the referent. If Bella thinks that Archie assumes that she can identify the referent, then this scenario might be somewhat similar to (7). If however, Bella thinks that Archie assumes that she is unable to identify the referent without switching her glasses, then the existence of the referent takes precedent. Although Bella has no way of knowing whether the referent exists (as she cannot even identify him), she will accept Archie's presupposition that he exists. If Bella agreed with Archie that she could not currently identify anything, she would not bother switching her glasses if she did not accept that there was something out there for her to see. In a way, we can understand Bella's switching to her regular

glasses as evidence that she has accommodated the fact that there is a person out there, and she puts on her glasses in order to see him.

It would be completely disruptive and unnatural for Bella to deny that the referent existed. If Bella accused Archie of lying, and said, “There is no one out there, you fool,” Archie would likely be offended and tell her to switch her glasses.

This account of accommodation might pose some problems for the Strong Contextual Felicity Constraint as presented in Tonhauser et al. (2013). It seems that there might be cases in which the projections that should have an SCFC can be found to be acceptable in m-neutral contexts. On top of this, there seems to be a gray area in the existence projections of Class C and Class A existence presuppositions. As shown before, a definite description would belong in Class C. Pronouns, on the other hand, as well as demonstratives, are clearly within Class A. These determiner phrases seem to exist on a continuum from pronoun and demonstrative type phrases to more definite-description-like phrases. It becomes difficult to decide how exactly to classify phrases such as “the man,” “the man who is king,” “the king.” There seems here to be a continuum of contextual felicity. If we rank determiner phrases from those with stronger contextual felicity constraints to those with weaker contextual felicity constraints, it seems difficult to draw a clear line that separates these phrases into those that impose an SCFC and those that do not.

If the Strong Contextual Felicity Constraint is caused by our inability to be neutral toward matters of salience, it is clear why the salience projections of Class D impose the constraint, but I have not yet explained why the existence projections in Class A impose the constraint. In order to account for this, I will focus particularly on one projective trigger. The “*that* (Noun)” trigger, as in the sentence in (9), projects at least three distinct projective contents.

(9). That jerk talked on his phone the whole time.

When a sentence such as (9) is subjected to the Family of Sentences Diagnostic, there are at least three implications that persist. The first is that this referent exists. The second is that the hearer should be able to identify the referent. The third is the property attribution – that the referent is a jerk. These contents are the existence of a referent, its salience, and property attribution, and the three contents belong to Class A, Class D, and Class B, respectively. What is significant here is that the property attribution Class B projection does not impose the Strong Contextual Felicity Constraint. This indexical trigger (which is dependent upon salience) projects at least one content that is appropriate in a neutral context and that can be accommodated. Since it is not necessary that indexical triggers impose the constraint, then in order for my account of this phenomenon to hold water, I need some way of describing why the existence projection of the “*that* (Noun)” trigger imposes the Strong Contextual Felicity Constraint while the property attribution projection does not.

It is worth noting that the “*that* (Noun)” trigger behaves exactly like the pronoun. Both have a Class A existence projection and a Class D salience projection. Tonhauser et al. (2013) showed that the “*that* (Noun)” trigger has a Class B property attribution projection, but pronouns also carry property attributions. The pronoun, “she,” not only points at a thing in the world, it also implies that the referent is female. This property attribution can also be shown to belong to Class B. Consider the scenario in (10).

(10). My brother is waiting for me at some predetermined location. As I approach, I can see him talking with someone, but he or she is so distant that I cannot make out any features of the person. By the time I get there, this person has left, so I ask my brother to whom he was speaking. He responds by saying, “She is an old friend.”

The fact that the referent is a female projects – that the referent is female is not at issue. It also does not impose the SCFC because this is a neutral context – the context does not suggest either that the referent is a female or that the referent is not a female – and the utterance is still appropriate. We can also see that the property attribution projection of pronouns does not have an OLE, because the demonstrative use of “she” in a sentence such as (11) does not attribute contradictory beliefs to anyone.

(11). My uncle believes that she is a man.

We can think of pronouns as in some way being a way of saying “*that* (Noun).” “She” could be equivalent to saying, “that female.” It is conceivable that there could be a language where we must say “that female” when we mean “she,” and it is also conceivable that a language could have a pronoun meaning, “that jerk.” The following discussion of the three projections of the “*that* (Noun)” trigger should generalize to the three projections of pronouns.

If a friend utters (12) without indicating a suitable referent, it is clearly infelicitous. I would be inclined to ask him whom he is referring, but I would never ask him what property he was attributing to this person. I still know that my friend believes that someone is a jerk, even if I do not know who that “someone” is. My understanding of the property attribution projection is in some way separate from my understanding of the salience projection.

(12). That jerk stole my wallet.

If my friend were to utter (12) without there being a salient referent, it would make sense for me to ask whether the referent existed. It is obviously a strange question since it could come off as impolite, but it does make sense to wonder whether a referent exists given that it cannot be identified. If I were a psychiatrist and a patient that suffers from paranoia and hallucinations utters the sentence in (12), it makes sense for me to question the existence of the referent. If I

believe that the patient's utterance stems from his paranoia, then I would question the salience condition – does he believe that the nurse stole his wallet? In this case, the referent would exist, but the patient would have failed to make the referent of “that” identifiable. If I believe it is a result of his hallucinations then I will question the existence condition – is he imagining some person that has stolen his wallet? In this case, the patient might very well believe that he is referencing a salient jerk standing right in front of him, but I cannot see this jerk because he or she does not exist. When the patient utters the sentence, I must ask myself whether he has failed to refer to something because he is referencing something that is not salient, or because he is referencing something that does not exist.

This, then, might be a way to account for the Strong Contextual Felicity Constraint in Class A projective contents. The existence of a referent might be so intimately intertwined with the identifiability of the referent (in a way that property attribution is not) because an unidentifiable referent could be the result of either a referent that is not salient or a referent that does not exist. The property attribution projection, then, does not impose the SCFC because it is not linked to salience in such a way; if the referent does not have the property attributed to it, this does not make it impossible for the listener to identify the referent, although it could lead to confusions.

Another possibility is that existence projections in Class A impose the Strong Contextual Felicity Constraint because they require that the listener first concerns herself with the salience of the referent before judging other features. In order to judge whether a referent exists, a listener must first understand what is being referenced. If “he” in the sentence, “He is bald,” refers to the King of France, the listener must first determine that “he” refers to the “King of France” before she can go out into the world and look for the existence of the King of France or before she can

conjure up her knowledge of French government. In this case, the existence conditions would have an SCFC because they are determined by means of salience-dependent reference. Here, judging the existence of the referent depends on the listener's ability to identify it. Since the ability to identify a referent is something that a listener cannot accommodate, the existence of the referent is something that the listener will not accommodate as well. The property attribution of "*that* (Noun)", then, does not have an SCFC because the property attribution is not determined through the referential use of "that."

In order for this to be correct, there must be some reason as to why definite descriptions do not impose an SCFC. It could be the case that the descriptive content of a definite description is rich enough to make itself identifiable to the listener. Whereas a demonstrative functions somewhat like a gesture, pointing to its referent, the definite description refers because its content can be descriptive enough to make the referent clear to the listener. This would fit well with the "continuum" of contextual felicity mentioned earlier. Phrases with more descriptive content are less likely to impose an SCFC and more likely to belong to Class C, while less descriptive phrases that function more like demonstratives or even gestures are more likely to impose an SCFC and belong to Class A.

A similar understanding of reference can be found in P.F. Strawson's work. In Strawson (1959), he discusses demonstrative identification of particulars. Strawson (1959) describes demonstratives as words that directly locate the particulars being referred to. He argues that noun phrases such as "that car" are cases of demonstratives despite being helped by descriptive words. This demonstrative identification involves the triggers that give rise to Class A and D projective contents and they refer by, in a way, "pointing-to" its referent. Non-demonstrative identification, then, relies on descriptions, names, or both. Strawson (1959) points out that if a listener does not

know of a particular object that fits a description, they may at that instant learn of such a particular. This demonstrates the ability of Class B and C projective contents to be accommodated.

Yet another possible way to account for the SCFC in Class A contents is to suggest that existence projections impose the Strong Contextual Felicity Constraint because they are somehow “downstream” from the salience conditions, because existence is processed only after we process salience. This would imply that all projections downstream from the salience projection (if there are any other than existence) would also impose the Strong Contextual Felicity Constraint, and that any projections “upstream” from salience (i.e. processed before salience) would not impose the constraint. This could imply that the property attribution projection is upstream. First, we understand that the speaker is attributing a particular property attributed to something. Then, we understand that this thing with this property is something that we should be able to identify. Finally, we understand that this thing with this property should exist.

If it is a matter of processing, this explanation should be empirically testable. One potential way of doing this would involve devising a response time test. In this, we could give someone a sentence similar to (12) and measure how quickly participants judge that the referent is a jerk, that the referent is identifiable, and that the referent exists. Another possible study is one in which we measure some sort of physiological response associated with surprise. In this experiment, we could give participants three scenarios in which the different projections are likely to cause surprise. We could have one case with a strange property attribution (“That ostrich-child stole my wallet!”), one in which the referent is not identifiable, and one in which

the referent clearly does not exist. We could assume that a surprise response that takes longer to occur is processed later.

Up to this point, I have only discussed possible origins of the Strong Contextual Felicity Constraint, but this analysis would be incomplete without discussing the Obligatory Local Effect. What is the cause of the split between Class A and D contents and Class B and C contents?

The OLE has to do with how projections behave when embedded in another person's belief predicate. I will examine the causes of the OLE by contrasting Class A contents from Class D contents. Class A has an Obligatory Local Effect while Class D does not. The sentence in (13a) does not attribute contradictory beliefs to my mother, while (13b) does.

(13a). My mother believes that he is tall and does not know that we are talking about him.

(13b). #My mother believes that he is tall and that he does not exist.

I suggest that projections that have an OLE have it because it is necessary for a third person to believe that projection in order to have the belief that the speaker attributes to them. In this way, the conventional meaning of triggers for projections that have an OLE causes the OLE. For the sentence in (13b), my mother could not possibly believe that the referent of "he" is tall if she does not believe that he exists.

This works just as well for Class C contents. Class C existence projections work the same way as Class A existence projections – my mother must believe that the King of France exists in order to believe that he is bald. A similar reasoning can be used to explain projective content triggers such as "stop." In order for my mother to believe that Archie stopped eating meat, she necessarily must believe that Archie used to eat meat; otherwise, she would simply believe that Archie does not eat meat. This necessity derives from the meaning of the word "stop." It implies

that a person used to do something; however, this implication is not at-issue, so it is a projective implication.

Contents that do not have OLE are contents for which it is not necessary for third person to hold in order to have a specific belief. Class B contents largely have to do with different names for an object and for different properties of that object. It is not necessary that a person knows every name for or property of a thing in order to have a belief about it. For example, the possessive relation implication of a possessive noun phrase belongs to Class B. It is not necessary that a person knows that “my bike” in (14) is mine in order to have some belief about it – in fact they could hold the belief that my bike is theirs. Class D contents deal with context-dependent identifiability. A third person does not need to know that I am currently able to identify the object in order to have a belief about it, as can demonstrated by the acceptability of (15).

(14). Abigail believes that my bike is her bike.

(15). Abigail believes that it is hers, but of course, she does not know that it is currently a salient part of our conversation.

Something does not have an OLE when it is not a necessary condition for the third person’s belief because a listener understands it as the speaker’s way of referring to or denoting the state of affairs about which the third person has a belief. If I utter the sentence in (12), it is clear that the property attribution is being used to help the listener identify the object rather than being a part of Abigail’s belief. If something does have an OLE, putting the projection and its opposite under a belief predicate is just a contradiction. It in no way helps the listener identify the state of affairs about which the third person has a belief, so it is not interpreted as being used by the speaker to assist the listener.

In this paper, I have looked at some potential explanations for origin of the classes of projective content found in Tonhauser et al. (2013). The Strong Contextual Felicity Constraint concerns the acceptability of sentences in contexts that cannot be said to suggest the projection. I argued that this occurs because the SCFC deals with the identifiability of referents and that we are usually unwilling to accommodate that we ought to be able to identify something. The Obligatory Local Effect concerns the behavior of projections when sentences carrying these projections are the complements to belief predicates. I have argued that this arises due to the necessity of these projections in having the associated beliefs.

References

- Simons, Mandy; Judith Tonhauser; David Beaver; and Craige Roberts. 2011. What projects and why. *Proceedings of Semantics and Linguistic Theory (SALT)* 22.309–27.
- Stalnaker, R. (2002). Common Ground. *Linguistics and Philosophy*, 25, 701-721.
- Strawson, P. F. (1959). 1. Bodies. In *Individuals: An Essay in Descriptive Metaphysics*. London: Methuen & Co Ltd.
- Tonhauser J., Beaver D., Roberts C., Simons M. (2013) Toward a taxonomy of projective content. *Language* 89(1): 66–109