A signaling account of contrastive focus[1]

Jon Scott Stevens — Zentrum für Allgemeine Sprachwissenschaft, Berlin

Abstract. This paper outlines a model of contrastive focus placement based on signaling games (Lewis 1969, Franke 2011). First, a simple model of sentence-level focus in question-answer situations is developed. Then, the model is extended to apply to sub-sentential phrases. Finally, an iterative procedure is developed for determining foci at each level of syntactic structure. This extends simple noise-based explanations of focus placement (Schmitz 2008, Bergen and Goodman 2014) to account for more difficult cases such as farmer sentences.

Keywords: contrastive focus, information structure, game theory.

1. Introduction

It seems reasonable, in a sense, that in intonation languages such as English and German, in which the placement of stress is relatively fluid, higher prosodic prominence should be given to elements within a sentence that are crucial for interpretation, as opposed to elements that are redundant or recoverable from the context. This intuition arises from the simple fact that prosodic prominence, though conceived of as an abstract phonological notion, correlates systematically with phonetic prominence, such that phonologically stressed words and phrases are more likely to survive the effects of *noise* (Shannon 1948) on information transmission. The pragmatic relevance of noise, the stochastically determined deletion of parts of a meaningful signal, has been noted in recent years (Schmitz 2008, Benz 2012, Bergen and Goodman 2014), and, when combined with some basic principles of economy, provides a clear explanation of why e.g. the focus structure in (1a) is preferred to those in (1b) and (1c).

(1) Who is teaching phonetics this semester?
  a. **[Bill]$_\mathbf{F}$** is teaching phonetics this semester
  b. # **[Bill is teaching phonetics this semester]$_\mathbf{F}$**
  c. # Bill is **[teaching]$_\mathbf{F}$** phonetics this semester

The question under discussion (QUD) sets up a set of alternatives {'Bill is teaching phonetics this semester', 'Sue is teaching phonetics this semester', etc.}, and everything in the answers in (1) except for the subject is redundant against that set (Roberts 1996). Thus the subject is what Schmitz (2008) calls the 'i-critical' material, i.e., the only material critical to listener interpretation. This is clear from the fact that everything but the subject can be elided in the answer. In fact, the elided version is arguably more natural, given that it takes less effort to produce.

[1]Any useful contributions of this paper were made possible in part by input from Robin Clark, Chris Ahern, numerous conference attendees, and anonymous reviewers.

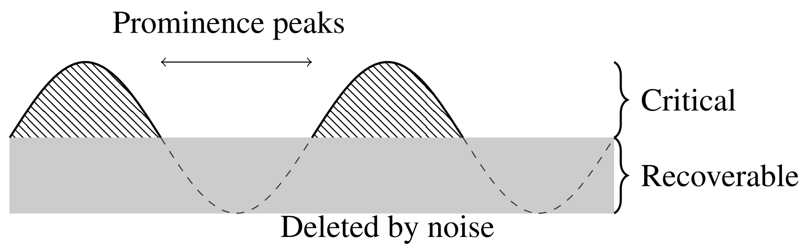

Figure 1: An ideal signal gives more prominence to what is critical, and less prominence to what the hearer could recover from context.

(2) Q: Who is teaching phonetics this semester?
A: Bill.

This principle of least effort is encoded in theories of focus placement as a pressure to focus mark as little as possible (Rooth 1992, Schwarzschild 1999). The linking hypothesis behind this is that focus, being a marker of phonological prominence in languages like English, and thus correlating with phonetic prominence, tends to create greater effort in speech production. What must be focused, then, is that minimal 'i-critical' material whose denotation could not be recovered from context by the hearer in the event that noise prevents the successful transmission of that material. This is visualized in Fig.1: Ideally, the prominence peaks created by prosodic focus marking ensure the successful transmission of the critical material in the presence of noise, thereby ensuring the recoverability of the entire signal. We might expect conventions to develop in intonation languages which approximate this ideal, and thus a noise-based model of focus placement can serve to supplement formal descriptions (e.g., Rooth 1992, Wagner 2012) by explaining *why* such a system should exist in a variety of languages, and not, for example, an inverted system which marks critical information by de-accenting it.

It is not immediately obvious how this noise-based picture of focus placement extends to cases where elision is not possible, namely contrastive focus of the type seen in farmer sentences.

(3) An **[American]**$_{\mathbf{F}}$ farmer punched a **[Canadian]**$_{\mathbf{F}}$ farmer…

Here, it is not the case that there is a QUD of the form {'a P farmer punched a Q farmer'}, and therefore nothing is recoverable from the global discourse context. However, we claim that noise can still play a role in this example, in that the twofold presence of *farmer* reduces the communicative need to elevate either instance of that word to prominence, given that, under certain pragmatic assumptions, if one instance survives the noise, the other will be recoverable.

This paper provides a noise-based account of contrastive focus extending to these cases. Focus placement is modeled as a signaling game between speaker and hearer, where the goal of the game is for the speaker to choose the minimal critical information to send to the hearer, who must recover the speaker's intended meaning from this signal. By hypothesis, this minimal critical information is what determines focus placement. We begin by outlining the information-structural assumptions underlying the analysis (1.1) and giving an informal summary of the approach (1.2), before providing some background on signaling games (2), introducing the game-theoretic model (3) and deriving key examples (4).

1.1. Assumptions

For current purposes we assume a distinction between the marking of contrastive focus and the marking of givenness, where only the former requires a contrast roughly in the sense of Wagner (2012), which is illustrated in (4).

(4) Mary's uncle, who buys and sells high-end convertibles for a living, is coming to Mary's wedding. I wonder what he got her as a present...
   a. He got her a **[cheap]$_{\mathbf{F}}$** convertible
   b. ? He got her a **[cheap convertible]$_{\mathbf{F}}$**
   c. # He got her a **[red]$_{\mathbf{F}}$** convertible
   d. He got her a **[red convertible]$_{\mathbf{F}}$**

There is a genuine semantic contrast between the contextually salient 'high-end convertibles' and 'cheap convertible', insofar as they are mutually exclusive descriptions, whereas no such contrast exists between 'high-end convertibles' and 'red convertible' (a car can be both high-end and red). Only when there is some contrast along these lines can the adjective be focused.[2]

In contrast to (4c), where accent cannot shift away from the modified noun *convertible* unless it is shifting onto a contrastive element, we find that non-modified XPs (i.e., maximal projections) can easily be 'de-accented' in the sense of Ladd (1996) merely in virtue of their being contextually salient (Stevens 2014). This shifts stress leftward onto an element that does not need to contrast with anything in the discourse context.

(5) A: Mary just arrived in her new convertible. What do you think she wants to do tonight?
   B: She wants to paint PICTURES of convertibles

[2]Contextual assumptions play a role, as noted by Katzir (2013). The adjectives 'red' and 'high-end' can create a contrast, e.g., in a context where car collectors are assumed to collect cars with one primary desired attribute, as in (i).

(i) Alice collects **[high-end]$_{\mathbf{F}}$** convertibles, and Bob collects **[red]$_{\mathbf{F}}$** convertibles.

We follow Beaver and Clark (2008) and Stevens (2014) in positing that XPs in discourse whose meanings are highly salient—or *given* (Schwarzschild 1999)—can become 'inactive' for purposes of stress assignment, independently of focus placement. This would give B's utterance in (5) a structure like in (6), where the PP is 'inactive' due its being marked as given ('G-marked'), forcing the main stress within the focused VP to fall on *pictures*.

(6) She wants to **[paint pictures [$_{\mathbf{PP}}$ of convertibles]$_{\mathbf{G}}$]$_{\mathbf{F}}$**

In any case, we are only concerned here with cases where a meaningful contrast is present. This can either take the form of a semantic partition, e.g., over convertibles as in (4), or a pragmatic one, e.g., a well-defined QUD like in (1). The goal of the current analysis will be to define the placement of these foci in terms of strategies in a game of information transmission.

1.2. Informal summary of the proposal

Imagine a game of communication between two players, a *sender* (S) and a *receiver* (R). The rules are simple: S must select one object from an array of options visible to both players, and the goal is to get R to guess which object S has chosen. If this happens, both players win a cash prize. S can send R a written message indicating which object R should select, but only the first word is free. For each subsequent word contained in S's message, the prize money is reduced.

Now, imagine a specific instance of this game where the players are given the context in (4), which sets up the expectation that Mary's uncle bought her an expensive convertible for her wedding. Both players are presented with an array of object descriptions ('cheap convertible', 'cheap sedan', 'expensive convertible', 'expensive sedan'), one of which describes Mary's uncle's actual gift. S is told the identity of the gift, and is instructed to convey this to R. Informal experiments accord with our intuitions that, if S writes the message "cheap" to send to R, that S most likely intends to convey that R should guess that the cheap convertible was the gift, not the cheap sedan.

The proposed signaling analysis is based on the independently motivated principle that in coordination games—games in which two or more players try to select the same option from an array of options (Schelling 1960)—the players' options are labeled with values of salient *attributes* which create *partitions* over a semantic space, such that "when attributes come to mind they come in clusters…it is nearly impossible to notice that 'U' is a vowel without noticing that other objects are consonants." (Bacharach and Bernasconi 1997). The attribute clusters which come to mind are assumed to be conditioned by context, such that (4) evokes, or 'activates', the relevant attributes CATEGORY (convertibles vs. other types of gifts) and PRICE LEVEL (being expensive vs. cheap). Thus, when one is asked in the context of (4) to select a likely gift, one represents the possibility of Mary's uncle having bought her a convertible as a choice between different values of the attribute PRICE LEVEL, as represented in the AVM below.

(7) $${}_{\textit{description}}\begin{bmatrix}\textsc{category} & \lambda x.\ x \text{ is a convertible}\\ \textsc{price level} & \lambda x.\ x \text{ is expensive}\end{bmatrix} \quad {}_{\textit{description}}\begin{bmatrix}\textsc{category} & \lambda x.\ x \text{ is a convertible}\\ \textsc{price level} & \lambda x.\ x \text{ is cheap}\end{bmatrix}$$

If the context creates an expectation or supposition or default belief that the gift is a convertible, then a message like "cheap" which omits the category information will be taken to specify 'cheap convertible' in opposition to 'expensive convertible', because if the category were different, that important information would have been included in the message. Crucially, this reasoning breaks down if the message makes reference to attributes which are not active in the context. In the context of (4), the game is conceived of as a choice between AVM structures representing CATEGORY and PRICE LEVEL, such that the message "red", which refers to the unrepresented attribute COLOR, is interpreted as a non-sequitur, casting doubt on the rationality of the sender. By hypothesis, it is for this reason that contrastive focus cannot be licensed in (4c).

Representing entity descriptions involves partitioning the space of predicates into meaningful attributes like category, color, price, etc., where each attribute is represented as a set of non-overlapping groups (for example, a car cannot be both expensive and cheap, assuming a fixed comparison class). Analogously, we can represent assertions as a pair of pragmatically motivated partitions: QUDs and answers. There is a similar mutual exclusivity, in that only one question can be the question under discussion currently being addressed, and only one answer is intended. This follows the structured meaning approach to QUDs (Krifka 2001, 2007) in partitioning the type-appropriate semantic space into possible arguments of the QUD, represented below as the value of the attribute ANSWER.

(8) $${}_{\textit{assertion}}\begin{bmatrix}\textsc{qud} & \lambda P.\ \text{Mary's uncle bought her a gift } x \text{ such that } P \text{ is true of } x\\ \textsc{answer} & \lambda x.\ x \text{ is expensive \& } x \text{ is a convertible}\end{bmatrix}$$

Such structures allow us to sufficiently constrain the space of possible interpretations in order to construct a well-defined game-theoretic model, and in doing so, correctly predict the restrictions on contrastive focus placement noticed by Wagner (2012) and others. The mechanics of our game of communication are now rather simple. The sender has incentive to select messages that fully specify which option the receiver should choose. But the sender has an opposing incentive to make her message as short as possible. In a simple case like (1), given the two options below, it is obvious that "Bill" will suffice to coordinate on the first assertion.

(9) $${}_{\textit{assertion}}\begin{bmatrix}\textsc{qud} & \lambda x.\ x \text{ is teaching phonetics}\\ \textsc{answer} & \text{Bill}\end{bmatrix} \quad {}_{\textit{assertion}}\begin{bmatrix}\textsc{qud} & \lambda x.\ x \text{ is teaching phonetics}\\ \textsc{answer} & \text{Sue}\end{bmatrix}$$

The case of (4) is more complex, because, given the message "cheap", the intended meaning could in principle be 'cheap sedan'. However, signaling games allow for probabilistic reasoning. Given a principled way of determining the receiver's *prior beliefs* about what the sender wants her to

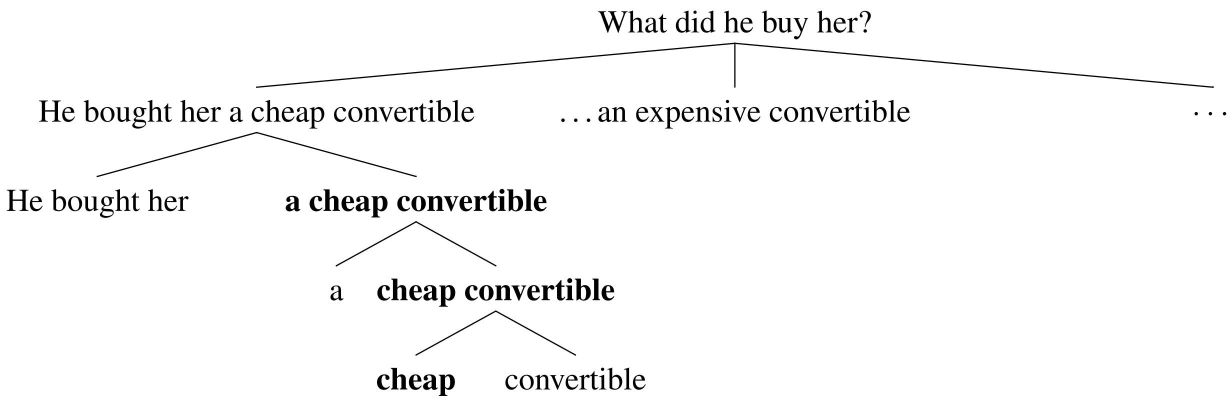

Figure 2: Information structure and partial syntactic structure of an utterance, with critical information marked in bold

choose, the game-theoretic model predicts that the receiver will employ pragmatic reasoning along the following lines.

1. I have received the message "cheap", therefore I assume the value of PRICE LEVEL is $[\![\text{cheap}]\!] = \lambda x.\ x$ is cheap.
2. I assume by default that the most probable value for CATEGORY is $\lambda x.\ x$ is a convertible.
3. I know that the sender knows that I have this belief.
4. Therefore, if the sender had intended anything other than $\lambda x.\ x$ is a convertible, she would have specified the value of CATEGORY.
5. Therefore, the speaker must intend to convey that the gift is a cheap convertible.

The link between this idealized game and the actual facts of intonation is that a formalized version of this game supplies a model of how to calculate, at any level of linguistic structure, what the 'i-critical' information is, and that this critical information is marked as such, which feeds focus assignment in a particular way. To illustrate how this works, consider (4) once more. When we think of a version of our game where the choices before the players are assertions, represented with QUD and ANSWER attributes, the game predicts that only the ANSWER ('a cheap convertible') is critical for successful interpretation. When we zoom in to the content of the NP in the answer ('cheap convertible'), represented with CATEGORY and PRICE LEVEL attributes, taking into account the salience of 'high-end convertible' in the prior context, the critical information is the PRICE LEVEL. If we mark off the critical information at every level of structure, we can obtain a tree structure like in Fig.2. The link between this structure and accent placement is clear: The leaf node of the 'i-critical' sub-tree must be given focus. Moreover, we may tentatively posit that only nodes which are not dominated by any i-critical node can be elided. This predicts (10) below, while at the same time allowing for a principled analysis of farmer sentences, where no elision is possible.

(10) Mary's uncle, who buys and sells expensive convertibles, is coming to her wedding.

Q: What did he buy her as a gift?

A: A **[cheap]$_{\mathbf{F}}$** convertible

Farmer sentences will be handled by assuming that when the signaling game is played at lower levels of structure (e.g., selecting from among possible NP meanings within the larger sentence) the surrounding utterance (e.g., everything outside the NP in the utterance) can be used as context to determine salience and prior probability. The intuition behind this is that, for example, the presence of 'American farmer' elsewhere in the sentence will prime NP meanings with NATIONALITY and PROFESSION attributes, with the antecedent itself making the 'American' and 'farmer' the default values. This in turn cashes out the intuition that doubly representing 'farmer' makes it a safer bet not to accent either instance, because if one gets through, the hearer can guess the identity of the other, assuming that if the value were different from the salient default, the corresponding word would have been focused.

With these intuitions in mind, we now introduce the formal game-theoretic analysis, beginning with some background on signaling games.

2. Signaling games

Signaling games are games of coordination (Schelling 1960) with two players, a *sender* (S) and a *receiver* (R), where the goal is for R to correctly determine a piece of information which is private to S, known as S's *type*. For purposes of pragmatic modeling, the type is typically taken to be a meaning that is to be conveyed to R. Because R cannot get inside S's head, this meaning cannot be observed directly, but rather must be inferred based on a *message* that S sends to R. The meaning that R interprets from S's message is known as the *action*, where actions and types are drawn from a single set of meaningful symbols. As a notational shorthand, a sender type who wants to convey meaning ϕ can be written t_ϕ, and similarly, an action where R selects meaning ϕ can be written a_ϕ. The *utility*—the quantity that is to be maximized by each player— is greater than zero for a type t_i sender and a receiver who takes action a_j if and only if $i = j$.

Formally, the game is a tuple $\langle \{S, R\}, \Phi, M, [\![\cdot]\!], T, \delta, A, U_S, U_R, \mathcal{C} \rangle$, where S and R are the sender and receiver, respectively, Φ is a set of semantic formulae, M is a *language* consisting of a set of possible messages (here, utterances of natural language), $[\![\cdot]\!]$ is a denotation function from M to Φ, T is the set of possible sender types, δ is a prior probability distribution over types in T, A is the set of possible receiver actions (for our purposes identical to T), U_S is the sender's utility function, a function from $T \times M \times A$ to $\mathbb{R}$, U_R is the receiver's utility function, a function from $T \times A$ to $\mathbb{R}$, and finally, $\mathcal{C}$ is a *cost* function, a function from $T \times M$ to $\mathbb{R}$, which is used to subtract a small amount of sender utility for lengthier, more effortful messages.

We assume that $T = A$, both a subset of $\Phi \cup \{\#\}$, where '$\#$' indicates a special type/action known as the *babbling* type/action. The goal of a type $t_\#$ sender is to "babble", i.e., to utter messages without conveying any meaning. We assume that the prior probability $\delta(t_\#)$ of a babbling sender is very close to zero. We further assume that, while a babbler will in principle send any message without regard for semantics or message length, a babbling sender nonetheless tries to signal her type. The inclusion of this low-probability babbling sender allows us to derive proper equilibria which straightforwardly specify optimal responses to otherwise off-equilibrium messages. More specifically, this allows us to include an action $a_\#$ which corresponds to a judgment of infelicity, so that the receiver has a consistent response to messages that seem suboptimal.

$$\begin{aligned} U_R(t,a) &= 1 \text{ iff } a = t \\ &= 0 \text{ otherwise} \end{aligned} \tag{i}$$

$$U_S(t,m,a) = U_R(t,a) - \mathcal{C}(t,m) \tag{ii}$$

The players are tasked with developing utility-maximizing *strategies* which specify how to behave in any possible state of the game. But utility depends on variables which are privately known only by the other player (for S, the action that R will take, and for R, S's type). To find optimal strategies for S and R in this game, we need to find a *perfect Bayesian equilibrium* (Harsanyi 1968, Fudenberg and Tirole 1991) which maximizes each player's *expected utility*, which is the probabilistically weighted average utility for all possible values of any unknown variables. An equilibrium for our purposes is a pair of strategies $\langle \mathcal{S}^*, \mathcal{R}^* \rangle$ such that each strategy maximizes its player's expected utility function given the other, where expected utility is formulated as follows.

$$EU_S(m|t,\mathcal{R}) = \sum_{a \in A} P(a|m,\mathcal{R}) \cdot U_S(t,m,a) \tag{iii}$$

$$EU_R(a|m,\mathcal{S}) = \sum_{t \in T} P(t|m,\mathcal{S}) \cdot U_R(t,a) \tag{iv}$$

As per the standard conception of perfect Bayesian equilibrium, $\mathcal{S}^*$ yields $\arg\max_m EU_S(m|t,\mathcal{R}^*)$ for each $t \in T$, and $\mathcal{R}^*$ yields $\arg\max_a EU_R(a|m,\mathcal{S}^*)$ for each $m \in M$, where crucially, the conditional probabilities $P(a|m,\mathcal{R}^*)$ and $P(t|m,\mathcal{S}^*)$ are rational and consistent beliefs about the private knowledge variables of the game.

One principled way of deriving an equilibrium in such games, given a reasonable set of default beliefs, is the *iterated best response* (IBR) procedure of Franke (2009, 2011), which is based on hierarchical reasoning models of rationality (Camerer et al. 2004, Bardsley et al. 2010) whereby optimal strategies are derived via hierarchical assumptions of the form, 'player 1 believes that player 2 believes that player 1 believes. . . .' The standard form of IBR begins with a receiver R_0

	a_F	a_C	$a_\#$	$m_\emptyset$	m_F	m_C
t_F	1,1	0,0	0,0	✓*	✓*	✗
t_C	0,0	1,1	0,0	✓*	✗	✓*
$t_\#$	0,0	0,0	1,1	✓	✓	✓

Figure 3: A signaling game representation

who plays a *naive* default strategy whereby R assumes that S's message is literally true, i.e., that the sender is of a type t_ϕ such that $\phi \rightarrow [\![m]\!]$. We will also consider the possibility of a 'null message' $m_\emptyset$ of length zero which is compatible with any type. Putting it together: Given a message m, if $[\![m]\!] \in \Phi$, R_0 assumes that S's type entails $[\![m]\!]$, but if $m = m_\emptyset$, R_0 makes no such assumption about S's type. Due its low probability, R_0 never assumes a babbling sender.

We then begin iteration by formulating a sender strategy S_1 which maximizes $EU_S(m|t, R_0)$. Costs for sending lengthier messages are assumed to be very small, such that message length only serves as a tie breaker between possible messages. Therefore, maximizing $EU_S(m|t, R_0)$ can be accomplished by first taking arg max$_m$ $P(t|m, R_0)$, and then, if that set contains more than one message, choosing the shortest. The probability of coordination $P(t|m, R_n)$ is specified as follows.

$$\begin{aligned} P(t|m, R_n) &= \frac{1}{|R_n(m)|} \text{ if } R_n(m) \text{ contains } t \\ &= 0 \text{ otherwise} \end{aligned} \tag{v}$$

Then, given a message m, receiver strategy R_2 takes the set of types that could have produced m, $\{t \in T \mid m \in S_1(t)\}$, and outputs the set of actions equal to the types from that set which maximize the prior δ. We can then keep iterating, constructing S_3 by analogy to S_1 and R_4 by analogy to R_2, etc., until convergence on a stable pair of strategies occurs. This stable fixed point is an equilibrium, and if the strategies entail distinct messages for each type, then it is a *separating equilibrium*. To illustrate with an extremely simple example, consider the following.

(11) Q: Is Mary's uncle a farmer or a car salesman?
A: He is a car salesman.

Let's model this as a signaling game where the two possible non-empty messages are *He is a farmer* and *He is a car salesman* (m_F and m_C, respectively), and the two possible types/actions are the denotations of those messages (t_F/a_F and t_C/a_C). Fig.3 gives the standard representation of this game. Base utilities (not considering the cost term) are given for each player for each type/action combination, and message columns are displayed in increasing order of cost. For each type/message combination, Fig.3 indicates whether that message is a priori *compatible* with that type, i.e., whether R_0 could ever guess that type given that message. Finally, the asterisks in the message columns indicate types with a higher prior probability, i.e., a higher value of $\delta(t)$ (in this case, the non-babbling types). Applying IBR to this game, we converge by S_3 on an equilibrium. The equilibrium is written as a set of tuples of the form $\langle t, m, a\rangle$, where given t, S should send m, and given m, R should select a. Because each message is associated with a unique type, we have achieved a separating equilibrium. Quite simply, if the speaker sends a meaningful message, it should be taken at face value, and if the speaker sends an empty message, the receiver should judge this intentional silence to be infelicitous.

$$\begin{aligned} R_0 &\Rightarrow \{m_F \to \{t_F\}, m_C \to \{t_C\}, m_\emptyset \to \{t_F, t_C\}\} \\ S_1 &\Rightarrow \{t_F \to \{m_F\}, t_C \to \{m_C\}, t_\# \to \{m_F, m_C, m_\emptyset\}\} \\ R_2 &\Rightarrow \{m_F \to \{t_F\}, m_C \to \{t_C\}, m_\emptyset \to \{t_\#\}\} \\ S_3 &\Rightarrow \{t_F \to \{m_F\}, t_C \to \{m_C\}, t_\# \to \{m_\emptyset\}\} \\ \text{Eq.} &\Rightarrow \{\langle t_F, m_F, a_F\rangle, \langle t_C, m_C, a_C\rangle, \langle\, t_\#, m_\emptyset, a_\#\rangle\} \end{aligned} \tag{vi}$$

Having illustrated the basic mechanics of signaling games, we now extend this framework to account for the placement of contrastive foci.

3. Accounting for focus

We now expand on the traditional signaling approach in order to model the use of partial messages (e.g., “Bill”) to convey larger meanings (e.g., ‘Bill is teaching phonetics’). The intuitions outlined in 1.2 about representing meanings as attribute-value structures are cashed out formally by introducing three important notions: *component meaninghood*, *partitioning*, and *contextual availability*, each addressed in turn.

The AVMs in 1.2 have the property that the different attribute values compose together semantically, via function application or predicate modification, to produce the standard semantic denotation of the element in question. For example, the QUD $\lambda x.teach(x, phonetics)$ composes with $bill$ to yield the denotation of “Bill is teaching phonetics.” The QUD and what it composes with are both *component meanings* of the entire assertion. A transitive and reflexive component meaninghood relation $\to_{comp}$ is defined formally as follows.[3]

$$\begin{aligned} A \to_{comp} B \text{ in } M \text{ iff either: } &\text{(i) } \exists\rho \in M : [A([\![\rho]\!]) = B] \vee [[\![\rho]\!](A) = B]\text{, or} \\ &\text{(ii) } \exists C.A \to_{comp} C \;\&\; C \to_{comp} B\text{, or} \\ &\text{(iii) } A = B \end{aligned} \tag{vii}$$

Whereas the standard signaling model outlined in section 2 builds a default hearer strategy R_0 around the assumption that the denotation of S’s message is entailed by S’s type, our model builds R_0 around the assumption that the denotation of S’s message (e.g., $bill$) is a component meaning

[3] $\to_{comp}$ is defined relative to a language M in order to exclude vacuous formulae which are not found in natural language, e.g., $\lambda x.teach(bill, phonetics)$.

of S's type (e.g., $teach(bill, phonetics)$). Where this provides a constraint on the IBR mechanism, a *partition requirement* is needed to constrain the structure of the game itself, namely the space of possible types and actions. Formally, where Φ_u is the set of uncurried set representations of the semantic formulae in Φ, the constraint is formulated as follows.

$$\forall \langle \phi, \phi' \rangle \in \Phi_u \times \Phi_u : \ \phi \neq \phi' \rightarrow \phi \cap \phi' = \{\} \qquad \text{(viii)}$$

Informally, the partition requirement ensures that the space of types and actions are drawn from a set of meanings Φ such that the members of Φ are mutually exclusive semantic descriptions. For example, a valid type space may include American farmers and Canadian car salesmen, but not American farmers and male farmers, since American farmers and male farmers are overlapping sets. This is a formalization of Bacharach's notion of attribute clusters.[4] Finally, as noted in 1.2, attribute clusters must be commonly believed to have 'come to mind' for both players. In other words, sets of alternative meanings, i.e., values of Φ must be *contextually available* in the following sense, borrowing from Schwarzschild (1999) notions of salient common ground (CG_S) and entailment under existential closure (ExClo):

(12) A set of alternatives Φ is contextually available iff:
   a. there is a salient proposition in the shared discourse context that entails that one of the members of Φ has a true existential closure (informal)
   b. $CG_S \Rightarrow [\exists \phi \in \Phi. \text{ ExClo}(\phi)]$ (formal)

This requirement is always met by QUDs because the QUD itself is salient, and the QUD is assumed to have a true answer. This requirement is met by specific attribute clusters if and only if at least one member of the cluster is salient. For example, if $\exists x.\ American(x)$ is a salient fact in the discourse, then an attribute cluster like NATIONALITY containing $\lambda x.\ American(x)$ is available.

These requirements, taken together, allow us to model the placement of contrastive foci as a game where potentially underspecified messages are sent in order to most efficiently guide the receiver toward the correct intended meaning, given that meanings are drawn from a contextually available semantically or pragmatically motivated partition of a meaning space.

Putting it all together, we propose the following IBR procedure for determining the 'critical' information for a given type in a signaling game.

---

[4]For assertions that are structured into QUD and ANSWER attributes, some pragmatic enrichment of ϕ and ϕ' in the formulation of the partition requirement may sometimes be necessary, e.g., an exhaustivity operator to ensure that 'Bill is teaching phonetics' and 'Bill and Sue are both teaching phonetics' form a partition in cases where the speaker is assumed to be giving a maximally informative answer.

1. R_0:
   For all m in M, output the most probable type(s) "compatible" with m:

   - Partially order T from highest to lowest probability according to δ.
   - Output $T^*(m)$, where $T^*(m)$ is the set containing the highest-ranked type(s) t such that $[\![m]\!] \rightarrow_{comp} t$.

2. S_1:
   For all t in T, output the best message(s) to send to R_0:

   - Calculate $M^*(t) \subseteq M$ such that for all m in $M^*(t)$, $P(t|m, R_0) \geq P(t|m', R_0)$ for all $m' \in M$.
   - Partially order $M^*(t)$ by effort (from fewest to greatest number of syllables).
   - Output the set containing the lowest-effort message(s) in $M^*(t)$.

3. R_2:
   For all m in M, output the type(s) that are most likely to send m to S_1:

   - Output the set containing the most probable type(s) t such that $m \in S_1(t)$.

4. For $n \in \{4, 6, 8, \cdots\}$, calculate S_{n-1} and R_n by analogy to S_{n-3} and R_{n-2}, respectively, until convergence occurs.

5. Convergence occurs at a level S_n when for any given type/message pair $\langle t, m\rangle \in T \times M$, $t \in R_{n-1}(m) \leftrightarrow m \in S_n(t)$.

6. If each tuple $\langle t, m, a\rangle$, where $a \in R_{n-1}(m)$ and $m \in S_n(t)$, maps a single type to a distinct (set of) value(s) for m, then the set containing those tuples is a separating equilibrium.

The following section shows how this procedure is applied to concrete examples to derive contrastive foci at different levels of structure.

4. Deriving examples

4.1. Sentence-level focus

We start with a simple example of a question under discussion with only two possible answers.

(13) There are two professors, Bill and Sue, one of which teaches phonetics each semester.
  Q: Who is teaching phonetics?
  A: Bill is teaching phonetics.

	a_{BP}	a_{SP}	$a_{\#}$	$m_{\emptyset}$	m_B	m_S	m_P	m_{BP}	m_{SP}
t_{BP}	1,1	0,0	0,0	✓*	✓*	✗	✓*	✓*	✗
t_{SP}	0,0	1,1	0,0	✓*	✗	✓*	✓*	✗	✓*
$t_{\#}$	0,0	0,0	1,1	✓	✓	✓	✓	✓	✓

Figure 4: Question-answer focus

The QUD in this case is $\{teach(bill, phonetics), teach(sue, phonetics)\}$. Let this be Φ, such that QUD $\cup \{\#\} = T = A$. To pick out the critical information within the sentence "Bill is teaching phonetics" in (13), we derive a separating equilibrium for the corresponding signaling game represented in Fig.4.

Let t_{BP} be a sender who wants to convey 'Bill is teaching phonetics', and t_{SP} a sender who wants to convey 'Sue is teaching phonetics.' R can select either of those two propositions as interpretations of S's message (a_{BP} and a_{SP}, respectively). We consider the following non-null messages.

m_B	"Bill"
m_S	"Sue"
m_P	"is teaching phonetics"
m_{BP}	"Bill is teaching phonetics"
m_{SP}	"Sue is teaching phonetics"

Messages m_B and m_{BP} have denotations which are component meanings of t_{BP}. Messages m_S and m_{SP} have denotations which are component meanings of t_{SP}. Message m_P has a denotation which is a component meaning of both t_{BP} and t_{SP}.

Applying our variant of IBR to this game, we converge on a separating equilibrium at S_3: For type t_{BP} the message "Bill" is best, and for type t_{SP} the message "Sue" is best. The Receiver will assume any other message to have been produced by a babbler, and thus the corresponding focus structures are judged to be infelicitous.

$$\begin{aligned} R_0 &\Rightarrow \{m_{\emptyset}, m_P \rightarrow \{t_{BP}, t_{SP}\}, m_B, m_{BP} \rightarrow \{t_{BP}\}, m_S, m_{SP} \rightarrow \{t_{SP}\}\} \\ S_1 &\Rightarrow \{t_{BP} \rightarrow \{m_B\}, t_{SP} \rightarrow \{m_S\}, t_{\#} \rightarrow M\} \\ R_2 &\Rightarrow \{m_B \rightarrow \{t_{BP}\}, m_S \rightarrow \{t_{SP}\}, m_{\emptyset}, m_P, m_{BP}, m_{SP} \rightarrow \{t_{\#}\}\} \\ S_3 &\Rightarrow \{t_{BP} \rightarrow \{m_B\}, t_{SP} \rightarrow \{m_S\}, t_{\#} \rightarrow \{m_{\emptyset}, m_P, m_{BP}, m_{SP}\}\} \\ \text{Eq.} &\Rightarrow \{\langle t_{BP}, m_B, a_{BP}\rangle, \langle t_{SP}, m_S, a_{SP}\rangle, \langle t_{\#}, \{m_{\emptyset}, m_P, m_{BP}, m_{SP}\}, a_{\#}\rangle\} \end{aligned} \tag{ix}$$

Therefore, when the answer to the QUD is that Bill is teaching phonetics, the focus structure and intonation pattern should be as follows, with contrastive focus only on the portion of the utterance which is most important for interpretation, i.e., that which corresponds to the "winning" message in our signaling game.

(14) a. **[Bill]$_{\mathbf{F}}$** is teaching phonetics
b. BILL is teaching phonetics.

The infelicitous patterns are ruled out on the grounds that they cannot be produced by a rational type t_{BP} Sender. This simple example illustrates how our enrichments to the signaling game come together to model the selection of critical information to be focused. The more interesting case is that of farmer sentences, to which we now turn.

4.2. Farmer sentences

Consider again (3), given again below as (15).

(15) An **[American]$_{\mathbf{F}}$** farmer punched a **[Canadian]$_{\mathbf{F}}$** farmer…

To illustrate how the same game mechanics can derive this example, which is clear under formal accounts like Rooth (1992), but which is *prima facie* problematic for noise-based pragmatic explanations, we start by "zooming in" on the NP *American farmer*. The key claim we make here is that contrastive focus placement within this NP is calculated by treating the NP-external material as context for determining δ. That is, $\delta(t)$ for a given type is proportional to the *salience* of that type, such that the meanings that are salient in the utterance-internal context, *A(n) _____ punched a Canadian farmer*, give a large boost to $\delta(t)$ for types containing that meaning.[5] To illustrate, let's set up a game to model selection of focus at this NP node, considering for simplicity only four meaningful types: t_{aF} for 'American farmer', t_{aW} for 'American watchmaker', t_{cF} for 'Canadian farmer' and t_{cW} for 'Canadian watchmaker'. The game is represented in Fig.5. Crucially, we propose that salience within the NP-external, utterance-internal context determines the prior probability function, such that the following holds.

$$\delta(t_{aW}) < \delta(t_{aF}), \delta(t_{cW}) < \delta(t_{cF}) \quad \text{(x)}$$

We consider the following non-null messages.

[5] It is also possible to consider only context which precedes the target node in the linear order of the sentence. The possibility of both conceptions can account for why the contrastive focus on *American*, but not the contrastive focus on *Canadian*, is optional.

	a_{aF}	a_{aW}	a_{cF}	a_{cW}	$a_{\#}$	$m_{\emptyset}$	m_a	m_c	m_F	m_W	m_{aF}	m_{aW}	m_{cF}	m_{cW}
t_{aF}	1,1	0,0	0,0	0,0	0,0	✓	✓*	✗	✓	✗	✓*	✗	✗	✗
t_{aW}	0,0	1,1	0,0	0,0	0,0	✓	✓	✗	✗	✓	✗	✓*	✗	✗
t_{cF}	0,0	0,0	1,1	0,0	0,0	✓*	✗	✓*	✓*	✗	✗	✗	✓*	✗
t_{cW}	0,0	0,0	0,0	1,1	0,0	✓	✗	✓	✗	✓*	✗	✗	✗	✓*
$t_{\#}$	0,0	0,0	0,0	0,0	1,1	✓	✓	✓	✓	✓	✓	✓	✓	✓

Figure 5: Focus within an NP

m_a	"American"	m_{aF}	"American farmer"
m_c	"Canadian"	m_{aW}	"American watchmaker"
m_F	"farmer"	m_{cF}	"Canadian farmer"
m_W	"watchmaker"	m_{cW}	"Canadian watchmaker"

For any given message, the maximally probable compatible type is indicated with an asterisk. For example, the message "farmer" is compatible with types 'American farmer' and 'Canadian farmer', as well as the babbling type, among which 'Canadian farmer' is maximally probable due to 'Canadian' being salient in the context.

Now applying our solution procedure to this game, we obtain a separating equilibrium which pairs "American" with 'American farmer', "American watchmaker" with 'American watchmaker', the null message with 'Canadian farmer', "watchmaker" with 'Canadian watchmaker', and assumes all other messages to have been generated by $t_{\#}$. Note that the null message in this case corresponds to a lack of contrastive foci within the NP. This does not mean that no prosodic prominence should be assigned at all (see e.g. Selkirk 2007).

$$\begin{aligned}
R_0 &\Rightarrow \{m_a, m_{aF} \to \{t_{aF}\}, m_{\emptyset}, m_c, m_F, m_{cF} \to \{t_{cF}\},\\
&\quad m_W, m_{cW} \to \{t_{cW}\}, m_{aW} \to \{t_{aW}\}\}\\
S_1 &\Rightarrow \{t_{aF} \to \{m_a\}, t_{aW} \to \{m_{aW}\}, t_{cF} \to \{m_{\emptyset}\}, t_{cW} \to \{m_W\}, t_{\#} \to M\}\\
R_2 &\Rightarrow \{m_a \to \{t_{aF}\}, m_{\emptyset} \to \{t_{cF}\}, m_W \to \{t_{cW}\}, m_{aW} \to \{t_{aW}\},\\
&\quad m_{aF}, m_{cF}, m_{cW}, m_c, m_f \to \{t_{\#}\}\}\\
S_3 &\Rightarrow \{t_{aF} \to \{m_a\}, t_{aW} \to \{m_{aW}\}, t_{cF} \to \{m_{\emptyset}\}, t_{cW} \to \{m_W\},\\
&\quad t_{\#} \to \{m_{aF}, m_{cF}, m_{cW}, m_c, m_f\}\}\\
\text{Eq.} &\Rightarrow \{\langle t_{aF}, m_a, a_{aF}\rangle, \langle t_{aW}, m_{aW}, a_{aW}\rangle, \langle t_{cF}, m_{\emptyset}, a_{cF}\rangle, \langle t_{cW}, m_W, a_{cW}\rangle,\\
&\quad \langle t_{\#}, \{m_{aF}, m_{cF}, m_{cW}, m_c, m_f\}, a_{\#}\rangle\}
\end{aligned} \tag{xi}$$

The final step is to formulate a principled procedure for repeating this game at every node and mapping the results to a contrastive focus structure for the whole sentence. We start by specifying a syntactic structure in Fig.6. We iterate the game through this tree as follows.

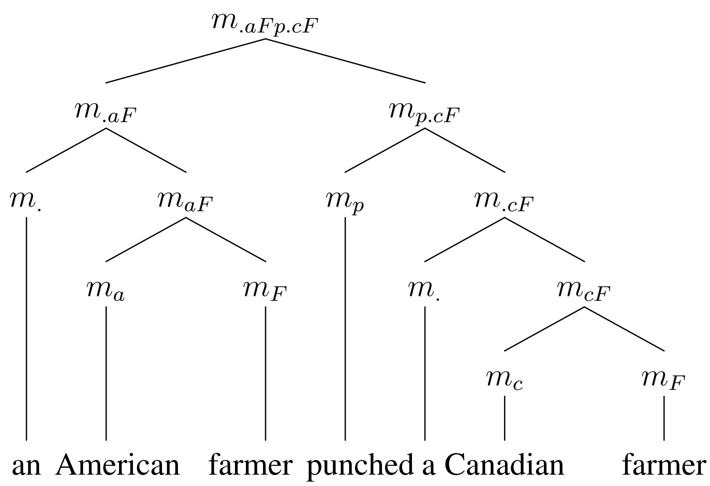

Figure 6: Syntactic structure of a farmer sentence

- Beginning at the root node N of a sentence-level focus as determined by the question under discussion, assuming N has two daughters, A and B:

  1. If there exists a contextually available set Φ_N of mutually exclusive possible meanings of the same semantic type, where Φ_N contains $[\![N]\!]$:
     (a) Let M_N be a set of messages such that for all m in M_N, $[\![m]\!]$ is either a member of Φ_N, or else a component meaning of a member of Φ_N which is of the same semantic type as either $[\![A]\!]$ or $[\![B]\!]$.
     (b) Consider a game where $A = T = \Phi_N \cup \{\#\}$ and $M = M_N \cup \{m_\emptyset\}$; if a separating equilibrium exists, let WINNER(N) be the optimal message for type $t_{[\![N]\!]}$.
     (c) If either $[\![A]\!]$ or $[\![B]\!]$ is a component meaning of $[\![\text{WINNER}(N)]\!]$, then mark the corresponding daughter node as a winning node.
     (d) If either A or B are marked as winning nodes, repeat step 1 at the winning node(s), if they are branching.
  2. If no such contextually available set exists, repeat step 1 at any branching daughter nodes.

- After all iterations, any winning node that does not immediately dominate another winning node is marked as a focus.

Only branching nodes and their immediate daughters are input to the game at a given node. The procedure begins by considering the QUD, as any nodes that are informationally redundant given the QUD do not need to be considered as targets for contrastive focus at lower levels. As in Fig.2 in 1.2, only leaf node winners determine focus. This avoids generating redundant nested foci.

Assume a simplified context for (3) where the possible nationalities are Canadian and American, and the possible professions are farmer and watchmaker, with two possible actions, punching and kicking. If we take the QUD to be something broad like 'what happened?', then prominence should be assigned over the whole utterance. To determine whether further any foci exist further down the tree, we then begin iteration at the root node $m_{.aFp.cF}$, checking whether there is a contextually available set of meanings $\Phi_{aFp.cF}$ which contains the sentence meaning. This would require an antecedent denotation of the form, 'a(n) American/Canadian farmer/watchmaker punched/kicked a(n) American/Canadian farmer/watchmaker.' Insofar as no such antecedent exists, no game can be played at this node, and we simply move on. No game can be played at node $m_{p.cF}$ either. Moving down to $m_{.cF}$, we can construct a valid set of types $T_{.cF} = \{t_{.aF}, t_{.cF}, t_{.aW}, t_{.cW}\}$ due to the salience of 'an American farmer' in the node-external context. The messages considered for this game are $m_{.cF}$ (= 'a Canadian farmer'), $m_{.}$, m_{cF}, m_{aF}, m_{cW} and m_{aW}. IBR yields the following separating equilibrium:

$$\text{Eq.} \Rightarrow \{\langle t_{.aF}, m_{\emptyset}, a_{.aF}\rangle, \langle t_{.aW}, m_{aW}, a_{.aW}\rangle, \langle t_{.cF}, m_{cF}, a_{.cF}\rangle, \langle t_{.cW}, m_{cW}, a_{.cW}\rangle, \\ \langle t_{\#}, \{m_{aF}, m_{.}\}, a_{\#}\rangle\} \tag{xii}$$

The winner at this node is "Canadian farmer". Continuing through the tree, we can derive separating equilibria for $m_{.aF}$ and m_{cF} by analogy to $m_{.cF}$ and m_{aF}, respectively, and then put it all together to obtain the following list of winning messages.

$$\begin{array}{llll} t_{.aF} & \Rightarrow \text{“American farmer”} & t_{.cF} & \Rightarrow \text{“Canadian farmer”} \\ t_{aF} & \Rightarrow \text{“American”} & t_{cF} & \Rightarrow \text{“Canadian”} \end{array} \tag{xiii}$$

Putting it all together, we generate the correct contrastive focus placements.

(16) An **[American]$_{\mathbf{F}}$** farmer punched a **[Canadian]$_{\mathbf{F}}$** farmer

This analysis extends the signaling model of communication, a powerful and flexible formal tool, to account for the role of noise in contrastive focus placement within sentences and sub-sentential phrases. Much work remains to assess the ease with which this approach can be employed to account for the many interesting phenomena related to focus including association with focus, second occurrence focus and more.

References

Bacharach, M. and M. Bernasconi (1997). The variable frame theory of focal points: An experimental study. *Games and Economic Behavior 19*(1), 1–45.

Bardsley, N., J. Mehta, C. Starmer, and R. Sugden (2010). Explaining focal points: Cognitive Hierarchy Theory versus team reasoning. *The Economic Journal 120*(543), 40–79.

Beaver, D. and B. Clark (2008). *Sense and Sensitivity: How Focus Affects Meaning*. Oxford: Blackwell.

Benz, A. (2012). Errors in pragmatics. *Journal of Logic, Language and Information 21*(1), 97–116.

Bergen, L. and N. Goodman (2014). The strategic use of noise in pragmatic reasoning. In P. Bello, M. Guarini, M. McShane, and B. Scassellati (Eds.), *Proceedings of the 36th annual Cognitive Science Conference*, Austin, TX, pp. 182–187. Cognitive Science Society.

Camerer, C. F., T.-H. Ho, and J. K. Chong (2004). A cognitive hierarchy model of games. *The Quarterly Journal of Economics 119*(3), 861–898.

Franke, M. (2009). *Signal to Act: Game Theory in Pragmatics*. Ph. D. thesis, University of Amsterdam.

Franke, M. (2011). Quantity implicatures, exhaustive interpretation, and rational conversation. *Semantic and Pragmatics 4*, 1–82.

Fudenberg, D. and J. Tirole (1991). Perfect Bayesian equilibrium and sequential equilibrium. *Journal of Economic Theory 53*(2), 236–260.

Harsanyi, J. C. (1968). Games of incomplete information played by 'Bayesian' players, part II. *Management Science 14*(5), 320–334.

Katzir, R. (2013). A note on contrast. *Natural language semantics 21*(4), 333–343.

Krifka, M. (2001). For a structured meaning account of questions and answers. In C. Fery and W. Sternefeld (Eds.), *Audiatur Vox Sapientia: A Festschrift for Arnum von Stechow*, pp. 287–319. Berlin: Akademie Verlag.

Krifka, M. (2007). The semantics of questions and the focusation of answers. *Topic and Focus: Studies in Linguistics and Philosophy 82*, 139–150.

Ladd, D. R. (1996). *Intonational Phonology*. New York: Cambridge University Press.

Lewis, D. (1969). *Convention: A Philosophical Study*. Cambridge: Cambridge University Press.

Roberts, C. (1996). Information Structure: Towards an integrated formal theory of pragmatics. In J. Yoon and A. Kathol (Eds.), *OSUWPL Volume 49: Papers in Semantics*.

Rooth, M. (1992). A theory of focus interpretation. *Natural Language Semantics 1*, 75–116.

Schelling, T. (1960). *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.

Schmitz, H.-C. (2008). *Accentuation and Interpretation*. New York: Palgrave MacMillan.

Schwarzschild, R. (1999). Givenness, AvoidF, and other constraints on the placement of accent. *Natural Language Semantics 7*, 141–177.

Selkirk, E. (2007). Contrastive focus, givenness and the unmarked status of "discourse-new". In C. Fery, G. Fanselow, and M. Krifka (Eds.), *Interdisciplinary Studies on Information Structure 6: The Notions of Information Structure*, pp. 125–145.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal 27*(3), 379–423.

Stevens, J. S. (2014). Against a unified analysis of givenness and focus. In R. Santana-LaBarge (Ed.), *Proceedings of the 31st West Coast Conference on Formal Linguistics*, Somerville, MA, pp. 438–446. Cascadilla Press.

Wagner, M. (2012). Focus and givenness: A unified approach. In I. Kučerová and A. Neeleman (Eds.), *Information Structure: Contrasts and Positions*. Cambridge: CUP.