Clouds and Credences:

A Probabilistic Approach to Underspecification

Justin D'Ambrosio

University of St. Andrews

jzd1@st-andrews.ac.uk

September 19, 2025

Context Sensitivity

Natural language contains many expressions whose semantic values depend on the context.

Suppose I utter (1):

(1) We had roast duck last night.

In order to know what I've said, you need to know:

1. who is speaking;
2. who the other people are that compose the "we";
3. what day it is; and
4. the fact that in the context, "had" means "ate for dinner".

Following King [2021], call expressions whose values depend on context *supplementives*.

Felicitous Underspecification

But often, the information in the context is insufficient to determine a unique semantic value for a supplementive.

- In such cases, context leaves open a *range* or *cloud* of candidate semantic values for a context-sensitive expression, and so leaves open a cloud of candidate contents for the speaker's utterance.

Felicitous Underspecification

The phenomenon in which context does not determine a unique content for an utterance—instead leaving open a cloud of candidate contents—and yet that utterance is felicitous.

Felicitous Underspecification: Adjectives

Playing Field We're sitting in the stands at a high school soccer match. Our friend's son Jimmy is the star of one of the teams, and has scored three goals. I utter (2):

(2) Jimmy is really athletic.

1. C_1: Jimmy is really athletic *for a member of the soccer team.*
2. C_2: Jimmy is really athletic *for a high schooler.*
3. C_3: Jimmy is really athletic *for someone his age.*

$$R = \{C_1, C_2, C_3\}$$

Felicitous Underspecification: Modals

Peer Pressure: We're out after school and we're trying to decide what to do. You suggest a movie: "do you want to go see Moana 2?" I reply:

(3) I have to do my homework.

1. C_1: I have to do my homework *if I want to pass.*
2. C_2: I have to do my homework *if my mom isn't going to get mad at me.*
3. C_3: I have to do my homework *if I want to get into Stanford.*

$$R = \{C_1, C_2, C_3\}$$

Felicitous Underspecification: Incomplete Expressions

Bad Dinner Party: We're at a dinner party that has turned out to be awful. You notice that I'm getting exasperated. We escape for a minute and you ask: "what's wrong?" I reply:

(4) I've had enough of this.

1. C_1: I've had enough of this *pretentious conversation*.
2. C_2: I've had enough of this *bad food*.
3. C_3: I've had enough of this *awful music*.

$$R = \{C_1, C_2, C_3\}$$

The Ubiquity of Felicitous Underspecification

Felicitous underspecification is ubiquitous in communication. Each of the following kinds of expression give rise to it:

- modals
- conditionals
- quantifiers
- gradable adjectives
- possessives
- incomplete expressions
- demonstratives and indexicals
- polysemous and contested terms

Communicating with Underspecification

A starting point:

Content Uncertainty

When context leaves open a cloud of candidate contents for a felicitous indicative utterance, the audience is *uncertain* about the content of the speaker's assertion.

How does communication work in such cases? Three questions:

Epistemic question: How should the audience update their beliefs and presuppositions?

Semantic question: What is the content of the speaker's assertion?

Success question: What constitutes successful communication?

Goals of the Talk

1. Argue that one leading approach to felicitous underspecification, due to King [2018, 2021], is mistaken.

2. Develop a new, probabilistic model of how interlocutors communicate under uncertainty: the *best guess* model.

3. Show how this theory can, at least in principle, be integrated into the common-ground framework.

The Plan

The Plan

Common Ground

Every conversation involves a set of interlocutors. At any point in the conversation, there is some information that these interlocutors share.

Common Ground The set of propositions commonly accepted by all interlocutors at a particular point in a conversation, denoted $CG_{w,t}$.

Context Set The intersection of the propositions in the common ground at a particular point in the conversation, denoted $CS_{w,t}$.

Presupposition

Following Stalnaker [2014], we can define what it is to presuppose that p in terms of the common ground:

Presupposition An interlocutor i presupposes that p just in case i believes that p is common ground.

Defective Context A context is *defective* if and only if interlocutors have different presuppositions.

Question Under Discussion

The Question Under Discussion (QUD) serves as a model of the topic of conversation. [Roberts, 1996]

- The QUD partitions the context set into exclusive and exhaustive cells.
- Each cell represents a *complete answer* to the question.
- Partial answers to the QUD are disjunctions (unions) of complete answers.

The goal of conversation is to add to the common ground until it entails a complete answer to the QUD.

Two Roles for Common Ground

The common ground plays two distinct roles that we typically associate with context:

Metasemantic Role A body of shared information on which interlocutors rely in planning and interpreting utterances, especially context-sensitive utterances.

Dynamic Role A body of shared information that evolves in response to conversational actions.

Assertion

In the common-ground framework, communication by means of assertion is governed by the following rule:

Assertion Rule If accepted, an assertion changes the context by adding its content to the common ground.

Stalnaker [1974/2002]

If the context set at the time of assertion is $CS_{w,t}$ and the content of my assertion is p, then if my assertion is accepted, the resulting context set is $CS_{w,t'} \cap p$.

Felicitous Underspecification and Common Ground

How do assertions update the common ground in cases of Content Uncertainty?

> *"that the conversational participants left open the candidate semantic values in context that they did **suggests** that conversational purposes at that point in the conversation are adequately served by associating a range of candidate semantic values in context with the supplementive and don't require assigning it a unique semantic value in context. Precisely because of this, conversational participants construct a propositional update that involves **all** the candidate semantic values in context and that adequately serves conversational purposes. These propositional updates include conjoining and disjoining the candidate propositions as well as doesn't-matter updates."*
>
> *King [2021, p. 37]*

From Uncertainty to Indifference

In place of Content Uncertainty, King proposes the following principle:

Content Indifference

When context leaves open a cloud of contents for a felicitous utterance, the audience can be certain, in context, that the speaker did not assert any specific content in the cloud.

Here is the reasoning underlying this principle:

- A speaker would only have underspecified—leaving open a cloud—if doing so were sufficient for conversational purposes. If specifying further mattered to the conversation, they would have done it.

From Indifference to Disjunction (or Conjunction)

Given Content Indifference, King claims that a rational audience member should update their presuppositions with either the *disjunction* or *conjunction* of the propositions in the cloud.

(5) I've had enough of this.

1. C_1: I've had enough of this *conversation*.
2. C_2: I've had enough of this *food*.
3. C_3: I've had enough of this *music*.

Candidate Updates: $C_1 \vee C_2 \vee C_3$ or $C_1 \wedge C_2 \wedge C_3$

Best-Answer Updating

Update with the disjunction unless updating with the conjunction provides a better answer to the QUD.

Uncertainty about Indifference

The problem is that Content Indifference is false!

- **Examples:** in our cases above, not only does context leave open a range of contents, but it also leaves open *whether I intended something specific*.
- **Messaging costs:** It costs more to fully specify, which can make it rational for speakers to underspecify, even when they intend something specific.
- **Strategic Speech:** Sometimes speakers have strategic reasons to underspecify, even when they mean something specific.

Moral: In cases of underspecification, audiences are often uncertain about how specific the speaker's intentions are.

Deeper Motivation: Uniformity and Certainty

Content Indifference is motivated by a principle central to the common ground framework:

Uniformity In order for an assertion to update the common ground, the context must determine a unique propositional content for that assertion. Stalnaker [1974/2002, p. 90]

Content Certainty In order for an assertion to update the common ground, each interlocutor must be certain of its propositional content, given the information in the common ground.

Conversation is Deterministic

Conversational Determinism: Rational communication consists in updating the common ground, and the common ground can only be updated deterministically.

> *"an assertion is a move in conversational game with a rule that says how the context evolves. To allow assertions with content to which the players of the game do not have access is to allow a situation in which the players will be unable to apply the rule. It would be like a card game in which one is dealt a card, face down, and the rule requires that you draw another card if and only if the card you were dealt is red."* *Stalnaker [2009]*

Stalnaker has his own proposal for how audiences update under uncertainty: they update with the *diagonal proposition.*

Why this is a very strange view

Common-ground views eliminate Content Uncertainty via reinterpretation.

- In doing so, they eliminate any role for standard tools that epistemologists use to model uncertainty.
- But Content Uncertainty is plausibly just one more type of uncertainty, on par with many others that we encounter in interactive situations.

General Claim: Communication is best modelled as a game of *incomplete information*—a Bayesian game—in which audiences reason under uncertainty about what the speaker means.

Four Desiderata

What we want is a view that:

- accommodates uncertainty about *both* what the speaker has said and how specific their intentions are.
- makes the right predictions about the contents of assertions.
- employs our best tools for dealing with uncertainty, and doesn't just eliminate uncertainty via reinterpretation.
- preserves the idea that there's body of shared information on which speakers rely in planning their expressions and audiences rely in interpreting them—a common ground.

Probabilistic Contexts

The common ground does not just serve to rule candidate contents in or out—the common ground can make interlocutors *more or less confident* that the speaker has asserted a particular content.

It is a body of *evidence* that can rationalise certain credences.

- In cases of felicitous underspecification, we need to consider the credences that rational interlocutors have about the contents of assertions conditional on the information in the common ground.

Credences over the Cloud

Bad Dinner Party, Redux: We're at a dinner party that has turned out to be awful in lots of ways. It's common ground that I'm particularly annoyed by the conversation I'm having. We escape for a minute and you ask: "what's wrong?" I reply:

(6) I've had enough of this.

In this context, you are not indifferent between the candidate contents left open.

C_1: I've had enough of this conversation.	$P(C_1 \mid CG) = .7$
C_2: I've had enough of this food.	$P(C_2 \mid CG) = .2$
C_3: I've had enough this music.	$P(C_3 \mid CG) = .1$

You have a distribution of *credences over the cloud*.

Rational Probabilistic Communication

Suppose that, guided by your credences, you take me to have asserted C_1, and update your presuppositions with the proposition that I've had enough of this bad conversation.

Suppose further that you were correct: I had a Gricean intention to communicate C_1.

Three claims:

1. this is an instance of successful communication in which we've both acted rationally;
2. the content of my assertion is C_1;
3. and this is a violation of Uniformity.

Against Reinterpretation

Any view that reinterprets my utterance is going to give the wrong verdict here.

- King's view entails that the content of my assertion is $C_1 \vee C_2 \vee C_3$. (wrong)
- Stalnaker's view entails that the content of my utterance is the diagonal proposition. (wrong)
- Both views entail that that we have not rationally, successfully communicated. (wrong)

$\Rightarrow$ What we need is a view that allows interlocutors to update their presuppositions in a way that is guided by the probabilistic information provided by the context.

Latif's Choice

In inquiry, we're often faced with questions without being certain of the answers. Suppose I ask you where you think Latif will go to school.

Here are the statistics for people with the same choice:

Yale	Harvard	Stanford	NYU
38%	30%	20%	12%

Holguin [2022]
Dorst and Mandelkern [2023]

Good Guesses

These guesses seem to be permissible:

(7)
a. Yale
b. Yale or Harvard
c. Yale or Harvard or Stanford
d. Yale or Harvard or Stanford or NYU

Bad Guesses

These guesses seem to not make any sense:

(8)
a. Harvard
b. NYU
c. Yale or Stanford
d. Not Yale.
e. Harvard, Stanford, or NYU.
f. Yale, or he has a birthmark on his left toe.

Five Principles of Good Guessing

Improbable Guessing It's sometimes permissible to answer p even when $P(p) < 0.5$.

Question Sensitivity Whether p is a permissible answer depends not just on the guesser's credence in p but also in what question is being answered.

Optionality Given any question Q, for any $k : 1 \leq k \leq |Q|$, it's permissible for your guess about Q to be the union of exactly k cells of Q.

Filtering A guess about Q is permissible only if it is filtered: if it includes a complete answer p, it must include all complete answers that are more probable than p.

Fit If a guess crosscuts a complete answer, it's impermissible.

Five Principles of Good Guessing

Optionality Given any question Q, for any $k : 1 \leq k \leq |Q|$, it's permissible for your guess about Q to be the union of exactly k cells of Q.

Five Principles of Good Guessing

Filtering A guess about Q is permissible only if it is filtered: if it includes a complete answer p, it must include all complete answers that are more probable than p.

Informativity-Accuracy Tradeoff

According to Dorst and Mandelkern [2023], what explains these principles is that good guesses optimize a tradeoff between accuracy and informativity.

Jamesian Expected Answer-Value :

$$E^J_Q(p) = P(p) \cdot J^{Q_p}$$

- $P(p)$ is the credence you have in guess p.
- Q_p is the proportion of answers to Q ruled out by a guess p.
- J is a parameter that weights informativity.

Good Guesses

Good guesses are either complete answers in which you have the highest credence, or filtered disjunctions of complete answers.

There are four good guesses:

Guess	Probability	Q_p
Yale	.38	3/4
Yale or Harvard	.68	2/4
Yale or Harvard or Stanford	.88	1/4
Yale or Harvard or Stanford or NYU	1	0/4

Notice that the guess "Yale" is highly informative, but its probability is low, while the guess "Yale or Harvard or Stanford of NYU" is totally uninformative, but certain.

Your Best Guess

Your *best guess* at a question Q maximizes Jamesian expected answer value for some choice of J, where $J \geq 1$.

- When $J > 1$, informativity matters, and the larger J is the more it matters.
- When $J = 1$, informativity does not matter, so your best guess will be nonspecific.

Different J-values represent different ways of weighting or trading off accuracy and informativity, which are inversely related.

Violations of Uniformity Shift the QUD

In **Dinner Party**, my utterance violates UNIFORMITY, and so leaves interlocutors uncertain about what I've asserted.

I claim that in such cases, the QUD, whatever it was before, shifts to:

(9) What have I asserted?

You need to answer this question before conversation can proceed. The three complete answers are:

(10) a. I've asserted that I've had enough of this conversation.
b. I've asserted that I've had enough of this food.
c. I've asserted that I've had enough of this music.

Probabilities

Given the information in the common ground, you will have a distribution of credences over (10-a-c), $P(R|CG)$:

Conversation	Food	Music
.7	.2	.1

So the candidate good guesses are as follows:

Guess	Probability	Q_p
Conversation	.7	2/3
Conversation or Food	.9	1/3
Conversation or Food or Music	1	0/3

Best-Guess Updating

Given this distribution of credences, define the *expected communicative value* of a guess p as:

$$P(p|CG) \cdot J^{Q_p}.$$

In communication, your best guess maximises expected communicative value.

Best-Guess Updating

Update your presuppositions with your best guess concerning the content of the speaker's assertion, if you accept it (the guess).

Specificity of Guesses

What is your best guess? It depends on your J-value.

- You may set your J-values high and guess C_1: that I've had enough of this conversation.
- Or you may set your J-values low and guess $C_1 \vee C_2 \vee C_3$: that I've had enough of *something or other*.

How are J-values fixed?

- J-values are fixed partly by your estimate of the specificity of my intention, given what you know of my goals.
- But they are also fixed by your willingness to take a communicative risks and update with something specific, even when I didn't fully specify.

Setting J-values

Here is an intuitive idea: speakers specify only as much as they need to, given their goals.

If their only goal is to answer the (original) QUD, then J-values are set by the QUD:

Minimal J-value In cases of communicative uncertainty, guess with the lowest J-value such that any lower value would yield a worse answer to the QUD.

But often, this is not the only goal that speakers have. So your J-value may take into consideration other goals of mine, such as messaging costs and other non-communicative goals.

QUDs and Bad Conversation

Returning to our example, consider two possible QUDs:

(11) Have you had enough of this? Should we go?

(12) What's wrong?

Relative to (11), your J-value will be low, because the disjunction answers the question perfectly.

Relative to (12), your J-value will be high, because every guess that isn't fully specific fails to give a complete answer to the QUD.

Taking Stock

So far we've accomplished the first two goals of the talk:

1. We've seen that the common-ground model of communication runs into problems dealing with communicative uncertainty.
2. We've developed a new account of updating—Best-Guess Updating—that does better.

Remaining question:

3. Under what conditions do individual best-guess updates yield an update to the common ground?

The Components of Best-Guess Updates

Best-guess updates are determined by two things: interlocutors' distribution of credences and their J-values.

The distribution of credences is conditional on the common ground, and calculated according to Bayes' rule:

$$P(i|CG) = \frac{P(CG|i) \times P(i)}{\sum_{(i' \in I)} P(CG|i') \times P(i')}$$

These credences are then weighted by J^{Q_p}.

Two Assumptions

Two assumptions guarantee that interlocutors guess in the same way:

Common Prior Assumption If interlocutors have common priors, and update their credences on common evidence, then their posteriors must be the same.

Common J-Values If interlocutors weight their identical posterior distributions by the same J-value, then their guesses must be the same.

Transparency Assumptions

But in order to guarantee that best-guess updates are not only the same, but are *common ground*, we need further assumptions:

Transparent Prior The common prior is common ground.

Transparent Bayesian Updating The fact that interlocutors update their credences according to Bayes' rule is be common ground.

Transparent Best-Guess Updating Best-Guess Updating is common ground.

Transparent J-Values The J-value with which interlocutors guess is common ground.

Sufficient, but not Necessary

Together, these conditions are sufficient, but they are not necessary.

Best-guess updates can yield common-ground updates even when these conditions are violated.

1. One such case is when it's common ground that J-values are minimal.
2. Another case is when interlocutors have similar enough credences and J-values to guess in the same way, and their guesses are common ground.

Weakening the Conditions

Here, I think, are the weakest conditions that guarantee that individual best-guess updates yield a common-ground update:

Common Guess Assumption Interlocutors have similar enough priors and J-values to issue in the same guess.

Transparent Guess Assumption Interlocutors common guess is common ground.

These two conditions are both necessary and sufficient for best-guess updates to yield a common-ground update.

Contextual Fragmentation

When interlocutors have different enough priors or J-values to yield different guesses, or when these are not common ground to begin with, best-guess updating will yield a defective context.

- Call the phenomenon in which an underspecified utterance creates a defective context *contextual fragmentation*.

Downstream conversation may *defragment* the context by making it clear, retroactively, what the speaker meant by an earlier assertion.

A More Realistic Picture

The result is what I take to be a more realistic picture of the relationship between assertion and the common ground.

- Given the ubiquity of underspecification, assertions will often fail to update the common ground, at least initially.

- But the content of an assertion can still become common ground retroactively as interlocutors learn more about one another, or as subsequent conversation further specifies the assertion's content.

Moral: The common ground isn't something that is updated *automatically*. The richer the common ground is, the easier it is for speech acts to update it.

Thanks very much!

Contact details:
justin.z.dambrosio@gmail.com
jzd1@st-andrews.ac.uk

Kevin Dorst and Matthew Mandelkern. Good guesses. *Philosophy and Phenomenological Research*, 105(3): 581–618, 2023.

Ben Holguin. Thinking, guessing, believing. *Philosopher's Imprint*, 22(1):1–34, 2022.

Jeffrey C. King. Strong contextual felicity and felicitous underspecification. *Philosophy and Phenomenological Research*, XCVII(3):631–657, November 2018.

Jeffrey C. King. *Felicitous Underspecification: Contextually Sensitive Expressions Lacking Unique Semantic Values in Context*. Oxford University Press, Oxford and New York, 2021.

Craige Roberts. Information structure in discourse: Towards an integrated formal theory of pragmatics. In Jae-Hak Yoon & Andreas Kathol, editor, *OSU Working Papers in Linguistics, Vol 49: Papers in Semantics*. The Ohio State University Department of Linguistics, 1996.

Robert Stalnaker. On hawthorne and magidor on assertion, context, and epistemic accessibility. *Mind*, 118(470):399–409, April 2009.

Robert Stalnaker. *Context*. Oxford University Press, Oxford and New York, 2014.

Robert C. Stalnaker. Assertion. In Paul Portner and Barbara H. Partee, editors, *Formal Semantics: The Essential Readings*, pages 147–162. Blackwell Publishers, 1974/2002.

An alternative: Diagonalization

[Stalnaker, 1974/2002] claims that in cases of Content Uncertainty, we should reinterpret utterances with the *diagonal proposition*:

$\mathfrak{A}$	a	b	c
a	T	F	F
b	F	F	T
c	F	T	F

$\xrightarrow{\text{Diagonalize}}$

$\mathfrak{A}$	a	b	c
a	T	F	F
b	T	F	F
c	T	F	F

Above are what Stalnaker calls *propositional concepts*, where a, b, and c are the three worlds in the context set.

Roughly, the diagonal proposition for (12) is the proposition that I have had enough of *whatever it is that I mean*. On this view, the content of my assertion in (12) is $\{a\}$.

An Argument Against Diagonalization

Question: what should your credence be that the speaker has asserted the diagonal proposition?

Answer: zero. Why?

- Stalnaker reinterprets utterances with the diagonal proposition because an audience is uncertain about the speaker's intentions.
- But unless the diagonal proposition is one of the candidate contents in the cloud, it will differ from each of them.
- Thus, your credence that it is one of the candidate contents asserted should be zero.

Worries/Concerns/Objections

1. "All of this was unnecessary. Common-ground updating is just an idealization anyway."
2. "You say you're interested in modeling non-ideal conversation. But Bayesian models are themselves highly idealized."
3. "You've argued that there can be successful communication when Uniformity is violated. But didn't you just tell us how to get Uniformity back?
4. "You've said that your view doesn't yield implausible verdicts about the contents of assertions, but you haven't given us an account of assertion."