Gricean equilibria in a general signaling game

Abstract. The Gricean tradition (Grice, 1957, 1975), and its neo-Gricean extensions (Horn, 2004; Levinson, 1983; Sauerland, 2004; Spector, 2007), formalizes the relationship between literal and pragmatic meaning through conversational maxims such as the maxim of Quality or the maxim of Quantity. Because they know speakers abide by these maxims, hearers draw inferences from the speaker's utterance which go beyond what was literally said. There exist many game-theoretic reconstructions of Gricean reasoning (Asher, Sher, & Williams, 2001; De Jager & Van Rooij, 2007; Jäger, 2007; Rothschild, 2013; Van Rooij, 2003) but implementations typically embed the maxim of Quality as a constraint on possible moves by the speaker. These implementations, where literal meaning is a given *ex ante*, contrast with Lewisian signaling games where any notion of meaning is an emergent property of the equilibrium. The goal of this work is to investigate whether the Quality and Quantity maxims can be emergent properties of an equilibrium in an unconstrained signaling game *à la* Lewis (1969), focusing on simple cases of scalar implicatures. The analysis reveals that under certain assumptions about speakers' and hearers' payoffs and beliefs, there may be no equilibrium in which Gricean maxims are obeyed. This result challenges the notion that these maxims are naturally consequences of principles of cooperation and rationality. Three case studies of scalar implicatures are presented, showing how the existence of Gricean equilibria depends on the existence of priors and payoffs.

1 Introduction

In linguistics, a common assumption has been to distinguish between the content of an utterance and its meaning in context. The content of an utterance is assumed to be determined by the rules of grammar and its meaning in context.

The Gricean tradition, starting from Grice (1957, 1975), and subsequently neo-Gricean tradition (Horn, 2004; Levinson, 1983; Sauerland, 2004; Spector, 2007) has provided ways to formalize the connection between content and meaning in use. According to this tradition, cooperative speakers abide by certain maxims of conversation, some of which are listed below. Knowing that a speaker abides by these maxims entitles a rational hearer to derive *implicatures*, inferences that go beyond what was literally said.

- **Maxim of Quality.** Make your contribution true; so do not convey what you believe false or unjustified.[30]
- **Maxim of Quantity.** Be as informative as required.
- **Maxim of Relation.** Be relevant.

- **Maxim of Manner.** Be perspicuous; so avoid obscurity and ambiguity, and strive for brevity and order.

One example of implicatures are the so-called *scalar implicatures*, which is illustrated by a sentence like (1). While the literal meaning of this sentence is compatible with a situation where all cookies were eaten, hearers tend to infer from (1) that *not all cookies were eaten.*

The maxims help explain the presence of the additional inference: if the speaker believed all cookies were eaten, then, by the maxim of Quantity, they would have uttered *all cookies were eaten*, since it is a more informative statement. But they did not and so, by contraposition, one should conclude that the speaker does not believe that all cookies were eaten.

(1) Some of the cookies were eaten.
$\rightsquigarrow$ *speaker does not believe all cookies were eaten.*

While the maxims are often stipulated, a widespread intuition, since Grice (1975) is that they should follow from *cooperativeness* and *rationality.* Indeed, it seems pre-theoretically intuitive that, in order to be cooperative, one ought to say the truth, say as much as one knows about a particular topic, remain on point, etc. In the words of Fox (2016), the maxims are *virtual truisms.*

Given this background, it is natural then that a rich body of literature (Asher et al., 2001; De Jager & Van Rooij, 2007; Jäger, 2007; Rothschild, 2013; Van Rooij, 2003) has proposed to reinterpret Gricean pragmatics within the framework of game theory, where such notions as cooperativity and rationality can be formally defined. In these works, the act of communication is seen as a signaling game where the speaker picks its utterance among a fixed set of signals; the hearer is tasked to reconstruct the speaker's private information on the basis of the signal received. With some assumptions, it can be shown that, in such sentences as (1), it is indeed strategic for the speaker to pick the message 'all' when the speaker believes all cookies were eaten and to pick the weaker message 'some' to convey that some but not all cookies were eaten.

A critical assumption in many of these works is that speakers are constrained to only produce true messages. Thus, the maxim of Quality is embedded in the rules of the game itself; it is not derived from cooperativity. But, as Franke (2009) puts it:

> It is not reasonable to assume in the context model that the speaker cannot –not even for fun, so to speak– use a signal that is not true. I can very well say whatever I like, whenever I like to whomever I like. I may have to face social or even legal consequences from time to time, but it is not as if the semantics of my language restricts the muscles of my jaw and vocal tract, regulating what I possibly can and what I cannot utter.

These constrained signaling games used in game-theoretic formal pragmatics contrasts with the signaling games introduced to the philosophy literature by Lewis (1969). In these games, there is no notion of meaning *ex ante* and speakers are in principle free to choose any signal they wish. Meaning is an emergent property of certain equilibria: a signal is produced only when the speaker is in a certain epistemic state. Upon receiving a certain signal, a hearer can therefore draw conclusions about said epistemic state and concomitantly the state of the world.

Thus, there remains a question of whether all Gricean maxims - including Quality - can be an emergent property in a general signaling game *à la* Lewis, where utterances may be chosen freely. This would go some way towards showing that the maxims are indeed derivable from more general principle of cooperation and rationality.

This is the question that this work seeks to address. First, I define what it means for a speaker to abide by the Gricean maxims with respect to some literal meaning. Then, we ask whether there is at least *some* equilibrium where the speaker abides by the Gricean maxims. The scope of the investigation is limited to the case of *scalar implicatures*, which is abundantly studied, and the maxims of Quality and Quantity necessary to derive them.

The answer to that question, as I'll show, depends in a critical manner on our assumptions about what the speaker and the hearer hope to achieve (their payoffs) and what the hearer expects about the speaker before any communication occurs (the prior). I show that, under many choices of priors and payoffs, there isn't any equilibrium where the speaker abides by the Gricean maxims. These cases are not pathological; as I'll show, they correspond to relatively natural assumptions. For instance, we'll show that, when the payoff function is continous, then a Gricean equilibrium will only exist if it is impossible for the speaker to be near-certain but not fully certain, which seems to be an unnatural condition.

This conclusion is that, contrary to a widespread intuition, the Gricean maxims are not straightforward consequences, once cooperativity and rationality is assumed, but quite subsantial claims about the prior, payoffs or the structure of the signaling game itself.

This argument will be developed on three case studies, which represent different textbook cases of *scalar implicatures*: simple scalar implicature with *some*, first without ignorance then considering ignorance, and exhaustive interpretations of answers to questions.

Before delving any further, the restricted nature of this work must be emphasized. I will only focus on the maxims of Quality and Quantity, as these are the only one necessary to derive the scalar implicatures we'll consider. I will limit myself to the textbook cases mentioned earlier, and not investigate more intricate cases.

The roadmap is as follows: in section 2, I'll introduce the unconstrained signaling game that will be studied in this piece; in section 3, we define what it means for a speaker to abide by the Gricean maxims; in section 4, I'll present our first case study with sentences like (1), considering only knowledgeable speakers and mainly as an illustration. In the second case study of section 5, I consider the possibility of ignorance and show that, in that case, the existence of Gricean equilibria either places significant and unnatural constraints on the prior or requires a discontinuous payoff function. In case study 3 of section 6, we consider a famous payoff function which is discontinuous, the K-L divergence, and show that, on a different example, it too does not guarantee the existence of Gricean equilibria. I'll conclude in section 7.

2 General signal games

This section gives the general form of the signaling game considered in this piece. It additionally provide notations, assumptions and propositions that will be useful in the rest of this piece.

In this game, a player, the speaker, possesses some information t (its type, among a set $\mathcal{T}$) about the state of the world. This information is provided by nature according a certain probability distribution $\mathbb{P}(T = t)$, which we call *the prior*. The set of types $\mathcal{T}$ may be finite and discrete or infinite and continuous, depending on the case study.

I assume that each type t determines a probability distribution of W, the set of all states the world could be in. In all cases, W will be a finite set. We write $t(w)$ for the probability assigned to w by the probability distribution determined by t. We also write $t(S) = \sum_{w \in S} t(w)$ In what follows, I will assume that types are in correspondence with probability distributions: each type determines a unique probability distribution, each probability distribution is associated with only one type. Thus, I will not always distinguish the two notions in subsequent prose.

On the basis of their type t, the speaker picks a signal s among a set of possible signals S. The second player, the hearer, receives the signal s but is unaware of t; on that basis, they make a guess t' about what the speaker's type might be. Both players receive the same payoff of $\mathcal{U}(t, t')$, making the game fully cooperative. Initially, only minimal assumptions will be made about the payoff function $\mathcal{U}$. We simply guarantee that if the hearer guesses perfectly, they get the maximal possible reward, i.e. property 1.

$$\mathcal{U}(t, t') \leq \mathcal{U}(t, t) \text{ for all } t, t' \tag{1}$$

A (pure) strategy for the speaker is a mapping $\mathcal{S}$ from types to signals ($T \to S$) and a strategy for the hearer, a mapping $\mathcal{H}$ from signals to types ($S \to T$). A combination of a speaker strategy and a hearer strategy $(\mathcal{S}, \mathcal{H})$ is called a strategy profile.

Given this, a speaker and a hearer playing according to $(\mathcal{S}, \mathcal{H})$ can earn an expected payoff of:

$$\text{expected}(\mathcal{S}, \mathcal{H}) := \int \mathcal{U}(t, \mathcal{H}(\mathcal{S}(t))) \cdot \mathbb{P}(\mathrm{d}t) = \mathbb{E}\Big[\mathcal{U}\Big(T, \mathcal{H}(\mathcal{S}(T))\Big)\Big]$$

When the set of types T is finite, this is simply:

$$\text{expected}(\mathcal{S}, \mathcal{H}) = \sum_t \mathbb{P}(T = t) \cdot \mathcal{U}(t, \mathcal{H}(\mathcal{S}(t)))$$

We say a particular strategy profile is *a Nash-Bayes equilibrium* iff

1. For any Hearer strategy $\mathcal{H}'$ different from $\mathcal{H}$, expected$(\mathcal{S}, \mathcal{H}') <$ expected$(\mathcal{S}, \mathcal{H})$ (We say $\mathcal{H}$ and $\mathcal{H}'$ are different iff it's not the case that, almost surely, $\mathcal{H}'(\mathcal{S}(T)) = \mathcal{H}(\mathcal{S}(T))$)
2. For any Speaker strategy $\mathcal{S}'$ different from $\mathcal{S}$, expected$(\mathcal{S}', \mathcal{H}) <$ expected$(\mathcal{S}, \mathcal{H})$ (We say $\mathcal{S}$ and $\mathcal{S}'$ are different iff it's not the case that, almost surely, $\mathcal{S}(T) = \mathcal{S}'(T)$)

When 1 holds, we say that $\mathcal{H}$ is the best response to $\mathcal{S}$. When 2 holds, we say $\mathcal{S}$ is the best response to $\mathcal{H}$.

Many interpretations can be given to the notion of an equilibrium. Here, the intuition we wish to ground the model in is the intuition of a literal "*equilibrium*": in a community, this game is played repeatedly between speakers and hearers, hearers always play according to $\mathcal{H}$, speakers always play according to $\mathcal{S}$ and neither has any strategic reason to deviate from these well-established behaviors. In a nutshell, we think of Nash-Bayes equilibria in the signaling game as Lewisian conventions Lewis (1969). The question of how a particular equilibrium came to be or how it is "*selected*" will not be relevant. We simply seek to answer the following question: is there *some* equilibrium in which the speaker behaves in a way consistent with Gricean maxims?

To find equilibria, it will be convenient to rewrite expected$(\mathcal{S}, \mathcal{H})$ by grouping together terms for which the signal $\mathcal{S}(t)$ is the same. For the finite case:

$$\text{expected}(\mathcal{S}, \mathcal{H}) := \sum_{s} \sum_{\substack{t \\ \mathcal{S}(t)=s}} \mathbb{P}(t) \cdot \mathcal{U}(t, \mathcal{H}(s))$$

In the general case, this amounts to conditionalizing over $\mathcal{S}(T)$. Since T is a random variable with distribution $\mathbb{P}(T = t)$, $\mathcal{S}(T)$ is a random variable[1] and thus, this is coherent:

$$\text{expected}(\mathcal{S}, \mathcal{H}) := \int_{s} \mathbb{E}\Big(\mathcal{U}(T, \mathcal{H}(s)) \Big| \mathcal{S}(T) = s\Big) \mu_{\mathcal{S}(T)}(\mathrm{d}s)$$

Where:

- μ_X is the distribution of the random variable X
- $\mathbb{E}(X \mid Y = y)$ is expected value of X knowing $Y = y$

In this form, we can prove the intuitive fact: if $\mathcal{S}$ is fixed, the maximal value of expected(S, H) is obtained by optimizing independently for every signal s, the response $\mathcal{H}(s)$ to that signal s. More formally, we have the property below:

Proposition 1. If $\mathcal{H}$ maximizes $\mathcal{H} \mapsto$ expected$(\mathcal{S}, \mathcal{H})$, then, for almost all s, $\mathcal{H}(s)$ maximizes $t' \mapsto \mathbb{E}\left(\mathcal{U}(T, t') \mid \mathcal{S}(T) = s\right)$.
The proof, which is straightforward, is relegated to appendix A.1.

3 Gricean equilibria

Our question is: is there any equilibrium where the speaker can be said to follow Gricean maxims? To answer this question, we define a notion of a Gricean speaker strategy. As discussed earlier, Gricean maxims depend on a notion of *literal meaning*; in our set-up, a literal meaning would be a function $[\cdot] : S \to 2^W$ mapping signals to sets of worlds (i.e. propositions).

A speaker strategy $\mathcal{S}$ is said to be Gricean with respect to a literal meaning $[\cdot]$ iff it meets the maxim of Quality and Quantity:

1. **Maxim of Quality:** for all t, $t(w \in [\mathcal{S}(t)]) = 1$ (*the signal is literally true*)

[1] Provided $\mathcal{S}$ is a measurable function, which we will assume throughout.

2. **Maxim of Quality:** for all t, there does not exist s' such that $t(w \in s') = 1$ and $[s'] \subsetneq [\mathcal{S}(t)]$ (*no signal could be sent with a strictly stronger literal meaning*)

One remark: a strategy for the speaker is always Gricean with respect to some literal meaning. It suffices to let $[s]$ be 2^W for all s. With this assumption, every signal is a tautology. In that case, a speaker, no matter which signal they pick, will always be speaking the truth (maxim of Quality) and will always use *a* strongest signal to do so (maxim of Quantity).

In the sequel, we will look at non-trivial literal meanings, derived from fragments of the English language. Thus, we ask whether English-like meanings, regulated by Gricean maxims, could be an Nash-Bayes equilibrium of our unconstrained signaling game. Formally, we say that a Nash-Bayes equilibrium $(\mathcal{S}, \langle\rangle$ is *Gricean* (with respect to $[\cdot]$) iff $\mathcal{S}$ is a Gricean speaker strategy (with respect to $[\cdot]$)

With these formal preliminaries established, we can turn to investigating Gricean equilibria in three case studies: the *some but not all* scalar implicature, first without ignorance, then with ignorance, finally the case of exhaustive answers. In each case, we describe the literal meaning of each signal in English, describe what a Gricean speaker strategy must look like given this literal meaning. Finally, we derive some conditions for this strategy to be part of an equilibrium.

4 Case Study I: *some*

In this first case study, there are only two worlds `sbna`, a world where some but not all cookies were eaten, and `all`, a world where they all were. We assume the speaker knows which world is actual; in other words, there are only two types: t_{sbna}, defined by $t_{sbna}(\texttt{sbna}) = 1$, where the speaker knows that some but not all of the cookies were eaten and t_{all}, defined by $t_{all}(\texttt{all}) = 1$, where the speaker knows all cookies were eaten. We consider two signals 'some' and 'all'.

The English literal meaning for these signals is as described below;

- $[\text{'some'}] = \{\texttt{sbna}, \texttt{all}\}$
- $[\text{'all'}] = \{\texttt{all}\}$

With respect to this literal meaning, there is a unique Gricean speaker strategy. When the speaker is t_{sbna}, the speaker must say 'some', since that is the only true signal (i.e. using 'all' would violate the maxim of Quality). When the speaker is t_{all}, both 'some' and 'all' are true; neither would constitute a violation of the maxim of Quantity. However, since 'all' is stronger than 'some', using 'some' would violate Quantity; thus the only signal that can be used is 'all'. This Gricean strategy $\mathcal{S}$ is defined formally below.

$$\mathcal{S} : \begin{cases} t_{sbna} \mapsto \text{'some'} \\ t_{all} \mapsto \text{'all'} \end{cases}$$

It is clear that if the speaker plays according to the Gricean strategy, a hearer can guess the speaker's type exactly. This is what strategy $\mathcal{H}$ represents.

$$\mathcal{H} : \begin{cases} \text{`some'} \mapsto t_{sbna} \\ \text{`all'} \mapsto t_{all} \end{cases}$$

Since the hearer can perfectly guess the speaker's type, it follows that the strategy profile $(\mathcal{S}, \mathcal{H})$ will be an equilibrium for any payoff function which rewards exact guesses the most (i.e. which meets property 1). On this simple example, it thus does seem that Gricean maxims can be an emergent property of certain equilibria.

But that example is particularly simplistic. The equilibrium $(\mathcal{S}, \mathcal{H})$ can be described as a case where a stronger signal ('all') competes with a weaker one ('some'), but it may equivalently be described as a case where two signals with non-overlapping meanings are used for distinct situations. In other words, this equilibrium is Gricean with respect to (at least) two different and non-trivial literal meanings. First, it is neo-Gricean with respect to the English literal meaning, e.g. $[\text{`some'}] = \{\texttt{sbna}, \texttt{all}\}$ and that $[\text{`all'}] = \{\texttt{all}\}$. But the equilibrium is also neo-Gricean with respect to a different literal meaning: $[\text{`some'}] = \{\texttt{sbna}\}$ and that $[\text{`all'}] = \{\texttt{all}\}$. This corresponds to a language English′ just like English, except that the word 'some' now means *only some*. To rephrase this in in concrete terms, if all one can observe is the behavior of the players on this game, there is no way that an external observer can tell whether the players are speaking English and English′.

One type of evidence used to argue that 'some' behaves as in English and not as in English′ comes from cases where the speaker does not know whether only some or all of the cookies were eaten. In such cases, the use of 'some' *is* warranted; if 'some' meant *only some* (as in English′), it shouldn't.

(2) *Speaker has looked at the content of certain envelopes but has not looked at all of them*
✓ Some of the envelopes have bills in them.

So, for our next case study, we consider a game more faithful to life with additional types for the speaker, corresponding to cases where the speaker might be uncertain as to which of the two worlds is actual. This will allows us to see in greater generality whether Gricean strategies may be part of equilibria.

5 Case study II: *some* in the presence of ignorance

The game is as described in the previous section except that the types are now t_p where p is a real number between 0 and 1 (included) representing the probability that all cookies were eaten, defined by $t_p(\texttt{all}) = p$. So t_{sbna} and t_{all} from the previous game correspond to t_0 and t_1 respectively.

In this game, consider the following strategy for the speaker:

$$\mathcal{S} : t_p \mapsto \begin{cases} \text{`some'} & \text{if } p < 1 \\ \text{`all'} & \text{if } p = 1 \end{cases}$$

It is clear that this strategy is the only Gricean strategy for the English literal meaning, e.g. $[\text{`some'}] = \{\texttt{sbna}, \texttt{all}\}$ and that $[\text{`all'}] = \{\texttt{all}\}$. It is not Gricean for the literal meaning of English′: since, in English′, [‘some’] $= \{\texttt{sbna}\}$, ‘some’ cannot be used whenever the speaker t_p has some uncertainty, i.e. $0 < p < 1$, because the maxim of Quality won't be met: $t_p(w \in [\mathcal{S}(t_p)]) = t_p(w \in [\text{‘some’}]) = 1 - p \neq 1$. (More generally, no strategy can be Gricean with respect to English′ in this game with uncertainty, since every signal of that language has too strong of a meaning for an uncertain speaker)

Could this speaker strategy $\mathcal{S}$ be part of a strategy profile which is a Nash-Bayes equilibrium? The answer here depends in a critical manner on both the prior over types $\mathbb{P}(T = t_p)$ and the payoff function. We can show that for a large class of reasonable payoff functions, if some Gricean strategy profile is an equilibrium, it must be a priori be known that the speaker is either absolutely certain or less confident than a certain threshold α.

Proposition 2. If:

- $(t, t') \mapsto \mathcal{U}(t, t')$ is continuous
- For some $\mathcal{H}$, $(\mathcal{S}, \mathcal{H})$ is a Nash-Bayes equilibrium
- $\mathbb{P}(\mathcal{S}(T) = \text{‘some’}), \mathbb{P}(\mathcal{S}(T) = \text{‘all’}) > 0$,

Then there is some α such that $\mathbb{P}(T \in \{t_\alpha \mid \alpha < p < 1\}) = 0$

As one can see, the assumptions about the payoff function are minimal and *a priori* reasonable. An example of a very natural payoff function meeting these requirements is the payoff function defined below, which is based on expected L^2 distance between the hearer's guess and the actual speaker's type. This payoff function and the L^2 distance it is based on are related to the Brier scoring rule used in scoring forecast (Brier, 1950). Its naturalness derives from the following property: for any given speaker strategy, the hearer maximizes their payoff by selecting the posterior distribution over worlds after hearing the signal s.

$$\mathcal{U}(t, t') = -\sum_w \left(t(w) - t'(w)\right)^2$$

Another assumption needed is the fact that both signals have a non-zero probability of being used: $\mathbb{P}(\mathcal{S}(T) = \text{‘some’}) \neq 0$ and $\mathbb{P}(\mathcal{S}(T) = \text{‘all’}) \neq 0$. Given the definition of $\mathcal{S}$ above, this means that $\mathbb{P}(T = t_1) \neq 0$ and $\mathbb{P}(T \neq t_1) \neq 0$. Note that this assumption imposes that the probability distribution defined by $\mathbb{P}(T = t)$ has an atom in t_1; it for instance rules out for instance a uniform prior over all t_p.

To prove our main result, one can start by considering a strategy profile $(\mathcal{S}, \mathcal{H})$ which is a Nash-Bayes equilibrium. By definition, $\mathcal{H}$ must be a best response to $\mathcal{S}$. So we start by determining best responses to the Gricean strategy $\mathcal{S}$. From proposition 1, we know that it suffices to find, for each signal s with non-zero probability, the best response to that signal. Given the speaker's strategy, the hearer's best response to receiving the signal ‘all’ is to guess t_1, since only a speaker with type t_1 would use such a signal. Formally:

$$\text{for all } t',\ \mathbb{E}(\mathcal{U}(T, t') \mid \mathcal{S}(T) = \text{‘all’}) = \mathbb{E}(\mathcal{U}(t, t') \mid T = t_1) = \mathcal{U}(t_1, t')$$

Which is maximal when $t' = t_1$, following our hypothesis about $\mathcal{U}$. By proposition 1, it follows that $\mathcal{H}(\text{‘all’}) = t_1$ (since the event $T = t_1$ has non-zero probability)

Turning to the case where the hearer receives the message ‘some’, let us simply define $t_\beta := \mathcal{H}(\text{‘some’})$. There are two cases to distinguish, a normal case where $\beta < 1$

and a pathological case where $\beta = 1$. In the pathological case, the hearer always guesses t_1 regardless of the signal received. But then, against $\mathcal{H}$, any $\mathcal{S}$ will achieve the same payoff. And so $\mathcal{S}$ cannot be *the* unique best response to $\mathcal{H}$.

$$\text{expected}(\mathcal{S}, \mathcal{H}) = \int_{t_p} \mathcal{U}(t_p, \mathcal{H}(\mathcal{S}(t_p)))\mathbb{P}(\mathrm{d}t_p) = \int_{t_p} \mathcal{U}(t_p, t_1)\mathbb{P}(\mathrm{d}t_p)$$

So, for $(\mathcal{S}, \mathcal{H})$ to be a Nash-Bayes equilibrium, it must be that $\beta < 1$, i.e. that, upon hearing 'some', the hearer guesses that the speaker is not certain that all cookies were eaten. What the exact value of β is does not matter for the rest of the argument.

Having determined these facts about $\mathcal{H}$, some necessary conditions for $\mathcal{S}$ to be the best response to $\mathcal{H}$ can be determined. The key observation is that if the speaker has a high credence that all cookies were eaten, it becomes better if the hearer guesses t_1 than if they guess t_β. This creates an incentive for the speaker to say 'all' in that case, since that is the signal that prompts the hearer to respond t_1.

Formally, there must be some non-empty interval $(\alpha, 1]$ where $\mathcal{U}(t, 1) > \mathcal{U}(t, \beta)$. This follows from continuity of $\mathcal{U}$ and the fact that $\mathcal{U}(1, 1) > \mathcal{U}(1, \beta)$ (this in turn a consequence of the fact that $\beta \neq 1$). Define $\mathcal{S}'$ such that:

$$\mathcal{S}' : t_p \mapsto \begin{cases} \text{`some'} & \text{if } p \leq \alpha \\ \text{`all'} & \text{if } p > \alpha \end{cases}$$

It can be shown that $\mathcal{S}'$ achieves a better or equal payoff than $\mathcal{S}$ against $\mathcal{H}$:

$$\begin{aligned} \text{expected}(\mathcal{S}', \mathcal{H}) - \text{expected}(\mathcal{S}, H) &= \int_t \mathcal{U}(t, \mathcal{H}(\mathcal{S}'(t))) - \mathcal{U}(t, \mathcal{H}(\mathcal{S}(t)))\mathbb{P}(\mathrm{d}t) \\ &= \int_{\alpha<t<1} \mathcal{U}(t, \mathcal{H}(\mathcal{S}'(t))) - \mathcal{U}(t, \mathcal{H}(\mathcal{S}(t)))\mathbb{P}(\mathrm{d}t) \\ &= \int_{\alpha<t<1} \mathcal{U}(t, \mathcal{H}(\text{`all'})) - \mathcal{U}(t, \mathcal{H}(\text{`some'}))\mathbb{P}(\mathrm{d}t) \\ &= \int_{\alpha<t<1} \underbrace{\mathcal{U}(t, t_1) - \mathcal{U}(t, t_\beta)}_{>0} \mathbb{P}(\mathrm{d}t) \\ &\geq 0 \end{aligned}$$

But since $\mathcal{S}$ is the best response to $\mathcal{H}$, it follows that the difference in expected values between the two strategies is in fact zero. That can only be if $\mathbb{P}(\alpha < T < 1) = 0$ (and concomitantly that $\mathcal{S}(T) = \mathcal{S}'(T)$ almost everywhere). This is the result we wanted to show.

Informally put, this demonstration can be explained as follows: as per our assumptions, the payoff function rewards exact guesses highly and by continuity, close guesses are also highly rewarded. This makes it optimal for a speaker to assert what they have strong credence in, even when they are not absolutely certain of it. A strategy profile in equilibrium therefore does not require full confidence for assertion. This emerging 'norm of assertion' (Williamson, 1996, 2000) is weaker than the Gricean maxim of Quality. Consequently, such a strategy profile will only be Gricean when the prior makes it impossible for a speaker to be very confident ($t(\texttt{all}) \approx 1$) but not absolutely certain ($t(\texttt{all}) = 1$). In this case, the looser confidence-based norm of assertion and the maxim of Quality are indistinguishable.

I take this to be an undesirable conclusion. There is no reason why a prior should exclude near-certainty. Priors that don't exclude it seem plausible descriptions of some situations. As an illustration of one such prior, consider the prior defined by:

- $\mathbb{P}(T = t_0) = \frac{1}{4}$
- $\mathbb{P}(\alpha < T < \beta) = \frac{1}{2}(\beta - \alpha)$
- $\mathbb{P}(T = t_1) = \frac{1}{4}$

The following scenario can ground intuitions about this prior: suppose that there is a phone app that can monitor the quantity of cookies eaten. There is a 50% chance that the speaker has the app. If they do, then they definitely know whether all the cookies were eaten or not. If they don't, then nothing can be said about the speaker's state of knowledge (a uniform prior). Since this prior allows a speaker to be quite confident without being certain, it follows from our result there can be no Gricean equilibrium here.

To preserve the possibility of Gricean strategies in equilibrium without unduly excluding near-certainty, there are two possible lines of response. First, as we noted above, the argument rests on the maxim of Quality being a strong norm of assertion (Williamson, 1996). The maxim of Quality requires absolute certainty on the part of the speaker. One may try to relax the maxim of Quality; perhaps all that is required is that the speaker deems the proposition expressed by the signal more likely than θ, a fixed threshold.

1. **Maxim of Quality:** for all t, $t(w \in [\mathcal{S}(t)]) > \theta$

We leave it to future work to determine whether, with this weaker requirement, Gricean strategies may be obtained. It should be noted, however, that there are arguments from the philosophy literature that strong probabilistic confidence is not enough to warrant an assertion. For instance, it seems incorrect to assert (3) in the context provided even though there is 99% chance that the proposition expressed is true, given speaker's knowledge.

(3) **Context:** *a lottery with 100 tickets. Speaker has not seen the results.*
Your ticket did not win. (Williamson, 2000)

A second line of response, which we address in this work, is that the payoff function is in fact discontinuous. As it turns out, a discontinuous payoff function is commonly used in the game-theoretic pragmatic literature (Frank, Goodman, Lai, & Tenenbaum, 2009; Qing & Franke, 2015; Schreiber & Onea, 2021). It is a payoff function based on the Kullback-Leibler divergence:

$$\mathcal{U}_{KL}(t, t') = -D_{KL}(t, t') = -\sum_w t(w) \log\left(\frac{t(w)}{t'(w)}\right) \tag{2}$$

This payoff function has the property that it is equal to $-\infty$ when the distribution $t'(\cdot)$ is not absolutely continuous with respect to the distribution $t(\cdot)$, i.e. when there is a world w such that $t'(w) = 0$ but $t'(w) = 1$. Put informally, this payoff function gives an infinite penalty whenever a hearer rules out a world that the speaker deems possible.

A corollary of this property is that the payoff function has a discontinuity in its first argument. Using the notation above, $\mathcal{U}_{KL}(t_p, t_1)$ is $-\infty$ whenever $p < 1$ but is equal to zero when $p = 1$. Because it is discontinuous, our results do not apply. In fact, with a Kullback-Leibler payoff, the Gricean speaker strategy $\mathcal{S}$ can be part of a Nash-Bayes equilibrium regardless of the prior.

To see this, first note that one of the results established above carries over to the new payoff function: the best response $\mathcal{H}$ to $\mathcal{S}$ is one that guesses t_1 when the speaker produces 'all' and guessed some t_β where $\beta < 1$ when the speaker produces 'some'. To show that β must be less than 1, we can use the fact that $t \mapsto \mathcal{U}(t, t')$ is strictly convex.

Given such a strategy on the part of the hearer, the speaker's Gricean strategy is also optimal. Indeed, when speaker is type t_1, producing 'some' gets a payoff of $-D_{KL}(t_1, t_\alpha) < 0$ while producing 'all' gets a payoff of $-D_{KL}(t_1, t_1) = 0$. So the Gricean speaker's strategy is optimal. When the speaker is t_β with $\beta < 1$, producing 'some' gets a payoff of $-D_{KL}(t_1, t_\alpha) > -\infty$, while producing 'all' incurs a infinitely negative payoff $-D_{KL}(t_\beta, t_1) = -\infty$. Again, the Gricean speaker, which picks 'some' in the situation, is making an optimal choice.

Yet, the Kullback-Leibler divergence still does not guarantee the existence of Gricean equilibria. To see this, we turn to our last case study where the utility based on Kullback-Leibler divergence is assumed.

6 Case study III: exhaustive answers and discontinuous payoff functions

This case study involves exhaustive answers to questions. When asked (4), a speaker may reply with any of the sentences a. to d. If their answer only mentions a single item (answers a and b), whether coffee or tea, it is inferred that they don't believe the other item is present, even though the literal meaning of their utterance does not convey any information about the other item. When the answer is c, it is inferred that they don't know which of the two items is there.

(4) **Context:** *there are no other items than coffee or tea.*
_ What's in the cupboard?

a. _ Coffee.
$\rightsquigarrow$ *speaker does not believe there's coffee in the cupboard.*

b. _ Tea.
$\rightsquigarrow$ *speaker does not believe there's tea in the cupboard.*

c. _ Coffee or tea.
$\rightsquigarrow$ *speaker doesn't know whether there is tea in the cupboard.*
$\rightsquigarrow$ *speaker doesn't know whether there is tea in the cupboard.*

d. _ Coffee and tea.

All of these inferences are well-explained within the neo-Gricean framework. The maxim of Quantity requires that the speaker's utterance not be (logically) weaker than any of the alternative statements that meets the maxim of Quality. Thus, for an utterance of 'tea' to be valid, it has to be that the speaker does not believe that the stronger statement 'coffee and tea' is true. Similarly, when a speaker utters 'tea or coffee', they must neither believe that 'coffee' is true nor that 'tea' is true. Since they believe that one of the two was put in the cupboard, they must be in the state of ignorance.

6.1 Formalization

In this case study, we now have 3 worlds: $\bar{c}t$, $c\bar{t}$, ct. For simplicity, we ignore the case where there is nothing in the cupboard or other items. There are 4 messages: 'coffee', 'tea', 'coffee or tea', 'coffee and tea'. The English literal meaning for these signals is defined as follows:

- $[\text{'coffee'}] = \{c\bar{t}, ct\}$
- $[\text{'tea'}] = \{\bar{c}t, ct\}$
- $[\text{'coffee or tea'}] = \{c\bar{t}, \bar{c}t, ct\}$
- $[\text{'coffee and tea'}] = \{ct\}$

Speaker types are probability distributions over these three worlds and can thus be parametrized by two parameters. For ease in later demonstration, I choose to describe our speaker types as t_{pq}, in terms of the probability p that there is coffee in the cupboard and the probability q that there is tea in the cupboard.

For example, t_{10} represents a speaker who is certain that there is coffee but no tea in the cupboard, $t_{0.5,0.5}$ represents a speaker who is certain that there is just one item in the cupboard but does not know whether this is tea or coffee, t_{11} a speaker who thinks both coffee and tea are in the cupboard, etc.

From these two parameters, the probability distribution over worlds can be recovered through the following equations:

- $\mathbb{P}(\bar{c}t \mid t_{pq}) = 1 - p$
- $\mathbb{P}(c\bar{t} \mid t_{pq}) = 1 - q$
- $\mathbb{P}(ct \mid t_{pq}) = p + q - 1$

The parameters p and q can take any value between 0 and 1 such that $p + q \geq 1$.

Given these types, Gricean maxims completely determine the behavior of a speaker; there is only one Gricean speaker strategy. Indeed, for any type, there is always a unique strongest message whose literal meaning is compatible with speaker's beliefs. This speaker's strategy is described by:

$$\mathcal{S} : \begin{cases} t_{11} \mapsto \text{'coffee and tea'} & \\ t_{\alpha 1} \mapsto \text{'tea'} & \text{if } \alpha < 1 \\ t_{1\beta} \mapsto \text{'coffee'} & \text{if } \beta < 1 \\ t_{\alpha\beta} \mapsto \text{'coffee or tea'} & \text{if } \alpha, \beta < 1 \end{cases}$$

Can this speaker be part of a Nash-Bayes equilibrium when the payoff function is based on the KL divergence as defined in equation 2? The answer depends once again on the prior. As in case study II in section 5, we assume that every signal has a non-zero probability of being produced. Given $\mathcal{S}$, this constraint translates to the following constraints on the prior:

- $\mathbb{P}(T = t_{11}) > 0$
- $\mathbb{P}(T \in \{t_{\alpha 1} \mid \alpha < 1\}) > 0$
- $\mathbb{P}(T \in \{t_{1\beta} \mid \beta < 1\}) > 0$
- $\mathbb{P}(T \in \{t_{\alpha\beta} \mid \alpha, \beta < 1\}) > 0$

Below, I show an example of an otherwise unobjectionable prior, for which $\mathcal{S}$ cannot be part of an equilibrium. Before we do, let us first start determine the hearer's best response $\mathcal{H}$ to the strategy $\mathcal{S}$, in the general case. A useful property of the KL divergence in that connection is that the expected value of a divergence to a random distribution reaches its minimum at the "*mean*" distribution.

Proposition 3. The function $t' \mapsto \mathbb{E}(D_{KL}(T, t'))$ reaches its minimal value with t' defined by $t'(w) = \mathbb{E}(T(w))$. More conspicuously, we may write $t' = \mathbb{E}(T)$

The proof is given in appendix A.2.

Combining this proposition with proposition 1, we can determine $\mathcal{H}$ to be equal to (almost everywhere):

$$\mathcal{H} : \begin{cases} \text{‘coffee and tea’} \mapsto \mathbb{E}(T \mid \mathcal{S}(T) = \text{‘coffee and tea’}) \\ \text{‘tea’} \mapsto \mathbb{E}(T \mid \mathcal{S}(T) = \text{‘tea’}) \\ \text{‘coffee’} \mapsto \mathbb{E}(T \mid \mathcal{S}(T) = \text{‘coffee’}) \\ \text{‘coffee or tea’} \mapsto \mathbb{E}(T \mid \mathcal{S}(T) = \text{‘coffee or tea’}) \end{cases}$$

Given $\mathcal{S}$, this means:

$$\mathcal{H} : \begin{cases} \text{‘coffee and tea’} \mapsto \mathbb{E}(T \mid T = t_{11}) \\ \text{‘tea’} \mapsto \mathbb{E}(T \mid T = t_{\alpha 1}, \alpha < 1) \\ \text{‘coffee’} \mapsto \mathbb{E}(T \mid T = t_{1\beta}, \beta < 1) \\ \text{‘coffee or tea’} \mapsto \mathbb{E}(T \mid T = t_{\alpha\beta}, \alpha, \beta < 1) \end{cases}$$

Informally put, a hearer who receives the message ‘coffee and tea’ obtains the best outcome by guessing t_{11}, since this is the only type who could produce such a signal. Similarly, a hearer receiving the signal ‘coffee’ ought to guess the mean type of all types $t_{1\beta}$ where $\beta < 1$; this distribution is some t_{1b} with $0 \leq b < 1$. And likewise for ‘tea’, the hearer should guess t_{a1} with some $0 \leq a < 1$. When ‘coffee or tea’ is received, the hearer should guess t_{cd} with both $c, d < 1$.

$$\mathcal{H} : \begin{cases} \text{‘coffee and tea’} \mapsto t_{11} \\ \text{‘tea’} \mapsto t_{a1} \\ \text{‘coffee’} \mapsto t_{1b} \\ \text{‘coffee or tea’} \mapsto t_{cd} \end{cases}$$

For $(\mathcal{S}, \mathcal{H})$ to be an equilibrium, it must be guaranteed that $\mathcal{S}$ is the best response to $\mathcal{H}$. However, as we'll now see, that can't always be. The problem is the following: with certain priors, a speaker of type $t_{1\beta}$ who knows there is coffee but believes there might be tea is better served by saying ‘coffee or tea’ than by saying ‘coffee’, contrary to what the Gricean maxims would command. This is because the signal ‘coffee or tea’ can be used to convey the possibility of tea and a payoff based on Kullback-Leibler divergence rewards such a behavior.

We'll illustrate this point with a simple prior and then make a general statement. Let's consider the following (atomic) prior:

- $\mathbb{P}(T = t_{11}) = 0.3$
- $\mathbb{P}(T = t_{10}) = 0.3$
- $\mathbb{P}(T = t_{01}) = 0.3$
- $\mathbb{P}(T = t_{\frac{2}{3}\frac{2}{3}}) = 0.04$
- $\mathbb{P}(T = t_{1\frac{1}{2}}) = 0.03$
- $\mathbb{P}(T = t_{\frac{1}{2}1}) = 0.03$

The probabilities are chosen carefully so that, overall, there is a 90% chance that the speaker knows precisely the content of the cupboard. This prior is consistent with the so-called 'competence assumption' (Sauerland, 2004; Spector, 2003; van Rooij & Schulz, 2004) used in the neo-Gricean literature. The 'competence assumption' states that, short of any indication to the contrary, the default assumption is that the speaker is knowledgeable. This assumption is used to guarantee strong implicatures (implicatures of the form *the speaker believes it is false that ...*), rather than the weak implicatures generated from the maxim of Quantity (*the speaker does not believe that ...*). Overall then, the prior chosen is a completely reasonable prior from a (neo-)Gricean perspective.

With this prior, the optimal response $\mathcal{H}$ to the Gricean speaker $\mathcal{S}$ is as follows:

$$\mathcal{H} : \begin{cases} \text{`coffee and tea'} \mapsto t_{11} \\ \text{`tea'} \mapsto t_{0.05,1} \\ \text{`coffee'} \mapsto t_{1,0.05} \\ \text{`coffee or tea'} \mapsto t_{\frac{2}{3},\frac{2}{3}} \end{cases}$$

Given the high probability that the speaker is knowledgeable, the optimal hearer strategy tends to guess highly knowledgable types (i.e. distributions of low entropy). This is so except when they hear the signal 'coffee or tea', because this signal can only be produced by an ignorant speaker. Given the hearer's strategy $\mathcal{H}$, consider what happens when the speaker knows there is coffee but deems the presence of tea 50% likely (type $t_{1\frac{1}{2}}$). This type occurs with 3% chance according to the prior. If a speaker chooses to say 'coffee', as the Gricean strategy $\mathcal{S}$ demands, both players get a payoff of $-D_{KL}(t_{1,0.5}, t_{1,0.05}) \approx -1.96$. If they decide to act against the maxim of Quantity and say the weaker 'coffee or tea' instead, then both get a payoff of $-D_{KL}(t_{1,0.5}, t_{\frac{2}{3},\frac{2}{3}}) \approx -0.41$. Thus, the Gricean speaker, which plays according to $\mathcal{S}$, is not giving the best response to $\mathcal{H}$. Since $\mathcal{H}$ is the best response to $\mathcal{S}$ and $\mathcal{S}$ is the only Gricean strategy, it follows that there cannot be Gricean equilibrium with this prior.

We can prove more generally that, if it is more likely for there to be both coffee and tea when the speaker says 'coffee or tea' than it is when they simply say 'coffee', then the prior should make it impossible for a speaker to know that there is coffee and to be near-certain that there is tea. For such a speaker would prefer to say 'coffee or tea' and not 'coffee', the signal required by Gricean maxims. A symmetric result holds for the signal 'tea'.

Proposition 4. The following two statements are true:

- If $t_{cd}(ct) > t_{1b}(ct)$, then for some α, $\mathbb{P}(\{t_{1x} \mid \alpha < x < 1\}) = 0$
- If $t_{cd}(ct) > t_{a1}(ct)$, then for some α', $\mathbb{P}(\{t_{x1} \mid \alpha' < x < 1\}) = 0$

To prove this, let's consider the condition under which it's more advantageous to say 'coffee or tea' over 'coffee' when the speaker knows there is coffee (type t_{1x} for some x):

$$\mathcal{U}(t_{1x}, t_{cd}) - \mathcal{U}(t_{1x}, t_{1b}) > 0$$

Given the definition of the payoff function, this is simply:

$$t_{1x}(c\bar{t}) \log\left(\frac{t_{cd}(c\bar{t})}{t_{1b}(c\bar{t})}\right) + t_{1x}(ct) \log\left(\frac{t_{cd}(ct)}{t_{1b}(ct)}\right) > 0$$

Or:

$$(1-x)\log\left(\frac{t_{cd}(c\bar{t})}{t_{1b}(c\bar{t})}\right) + x \underbrace{\log\left(\frac{t_{cd}(ct)}{t_{1b}(ct)}\right)}_{>0} > 0$$

Since the highlighted quantity is strictly positive, this inequality will hold for all values x in $(\alpha, 1]$ for some α. We can then define $\mathcal{S}'$ a speaker strategy which the same as $\mathcal{S}$ except for values t_{1x} with $\alpha < x < 1$. This strategy, by definition, is better than $\mathcal{S}$: expected$(\mathcal{S}', \mathcal{H}) \geq$ expected$(\mathcal{S}, \mathcal{H})$. Since $(\mathcal{S}, \mathcal{H})$ is by assumption an equilibrium, it must be that $\mathcal{S}$ and $\mathcal{S}'$ coincide almost everywhere and so $\mathbb{P}(\{t_{1x} \mid \alpha < x < 1\}) = 0$, which is what we set out to prove. A symmetric proof is possible considering 'tea' as a signal.

One remark concerning this result is that, the more likely it is that the speaker knows whether there is tea, given that they know there is coffee (i.e. the smaller b is), the easier it is to satisfy the inequality $t_{cd}(ct) > t_{1b}(ct)(= b)$. Thus, a form of the 'competence assumption' will virtually guarantee that either near-uncertainty is impossible or Gricean equilibria don't exist. The prior given in example above is an illustration of this property.

In conclusion, we find that payoffs based on the discontinuous Kullback-Leibler divergence only guarantee Gricean equilibria under some sort of near-certainty prior, just as in case study II of section 5. This is taken to be a similarly problematic conclusion, as it is not clear why these priors could be ruled out on *a priori* grounds.

7 Conclusion

In a general Lewisian signaling game, signals carry no meaning *ex ante* but a certain notion of meaning may be reconstructed within some equilibrium. In this work, we ask whether in an equilibrium, meaning could behave the way it is supposed to according to the (neo-)Gricean tradition. To see that, we considered textbook cases of quantity implicatures and checked whether the behavior encoded in Gricean maxims could be an equilibrium. The main result is that there are conditions under which this is possible but these conditions don't seem *a priori* justifiable.

If these results are taken at face value, they may elicit two different types of response. First, it could be denied that Gricean maxims (especially in their neo-Gricean forms) are an adequate model of actual linguistic behavior, even when full cooperation is assumed. In case study II of section 5, it seemed strategic to use a signal expressing a proposition one did not believe to be true, when that proposition was judged highly likely. In case study III of section 6, it proved strategic to use a signal with a weaker meaning ('coffee or tea') in case that signal would leave open a possibility that a stronger signal would not (i.e. the possibility that tea might be in the cupboard). The second response would be to maintain that neo-Gricean maxims are an appropriate guide of linguistic behavior but deny that linguistic behavior is an equilibrium. This would mean that either naïve speakers or naïve hearers are not offering their best response to the other's strategy or both.

However, we leave open the possibility that these results should not be taken at face value. It is possible that some of these results rest on the particular set-up assumed or the way we translated the maxims into the game-theoretic formalization. In case study II of section 5, we argued that perhaps the maxim of Quality might be weaker (although we pointed out previous arguments that it is as strong as it should be). Likewise, it remains to be investigated whether the hearer's task is adequately modeled as that of guessing a distribution. Despite these limitations, the point remains that reconciling the Gricean behavior with the behavior in the equilibrium of a general signaling game requires substantial assumptions.

A Proofs

A.1 Proof of Proposition 1

Proof. Call E_s the function $t' \mapsto \mathbb{E}\left(\mathcal{U}(T,t') \mid \mathcal{S}(T)=s\right)$. Call SubOptim the set of signals s such that $E_s(\mathcal{H}(s)) \neq \sup_t E(t)$. Our goal is to show that $\mathbb{P}(\mathcal{S}(T) \in \text{SubOptim}) = 0$. By definition, for each signal s in this set, there exists t_s such that $E_s(\mathcal{H}(s)) < E_s(t_s)$. Define $\mathcal{H}'$ such that $\mathcal{H}'(s) = t_s$ when $s \in$ SubOptim and $\mathcal{H}'(s) = \mathcal{H}(s)$ otherwise.

Because $\mathcal{H}$ maximizes $\mathcal{H} \mapsto \text{expected}(\mathcal{S},\mathcal{H})$, we know that:

$$\text{expected}(\mathcal{S},\mathcal{H}) - \text{expected}(\mathcal{S},\mathcal{H}') \geq 0$$

But this is equivalent to:

$$\int_{\text{SubOptim}} \underbrace{\mathbb{E}\Big(\mathcal{U}(T,\mathcal{H}(s))\Big|\mathcal{S}(T)=s\Big) - E\Big(\mathcal{U}(T,\mathcal{H}'(s))\Big|\mathcal{S}(T)=s\Big)}_{<0} \mu_{\mathcal{S}(T)}(\mathrm{d}s) \geq 0$$

For this integral over strictly negative values to be positive, it must be that $\mu_{\mathcal{S}(E)}(\text{SubOptim}) = 0$, which is, by definition, equivalent to $\mathbb{P}(\mathcal{S}(T) \in \text{SubOptim}) = 0$ □

A.2 Proof of Proposition 3

Proof. We aim to minimize the function $E(t') := \mathbb{E}[D_{KL}(T \parallel t')]$.

Step 1: computing gradient

The function we need to minimize is:

$$E(t') = \mathbb{E}[D_{KL}(T \parallel t')] = \mathbb{E}\left[\sum_w t(w) \log\left(\frac{t(w)}{t'(w)}\right)\right].$$

Taking the derivative of this with respect to $t'(w)$, we get:

$$\frac{\partial}{\partial t'(w)} E(t') = -\frac{1}{t'(w)} \mathbb{E}[T(w)].$$

Step 2: normal gradient condition

To minimize $E(t')$, the gradient must be normal to the constraint surface $\sum_w t'(w) = 1$. This occurs when the derivatives with respect to $t'(w)$ are equal for all w, i.e., when:

$$\frac{\mathbb{E}[T(w)]}{t'(w)} = \alpha$$

for some constant α. Solving for $t'(w)$, we obtain:

$$t'(w) = \alpha \mathbb{E}[T(w)].$$

Step 3: enforcing the normalization constraint

To determine α, we use the normalization constraint $\sum_w t'(w) = 1$. Substituting our expression for $t'(w)$:

$$\sum_w t'(w) = \sum_w \alpha \mathbb{E}[T(w)] = \alpha \mathbb{E}\left[\sum_w t(w)\right] = \alpha.$$

Thus, the minimizing condition becomes:

$$t'(w) = \mathbb{E}[T(w)].$$

Step 4: convexity and global minimum

Finally, since $E(t') = \mathbb{E}[D_{KL}(t \parallel t')]$ is a strictly convex function of $t'(w)$, and the KL divergence $D_{KL}(t \parallel t')$ is strictly convex in its second argument, any local minimum must also be a global minimum.

Thus, the function reaches its global minimum when $t'(w) = \mathbb{E}[T(w)]$.

□

References

Asher, N., Sher, I., Williams, M. (2001). Game theoretical foundations for Gricean constraints. *Proceedings of the Thirteenth Amsterdam Colloquium, ILLC, Amsterdam.*

Brier, G.W. (1950). Verification of forecasts expressed in terms of probability. *Monthly weather review*, *78*(1), 1–3,

De Jager, T., & Van Rooij, R. (2007). Explaining quantity implicatures. *Proceedings of the 11th conference on Theoretical aspects of rationality and knowledge - TARK '07* (p. 193). Brussels, Belgium: ACM Press.

Fox, D. (2016). On why ignorance might be part of literal meaning: Commentary on Marie-Christine Meyer. *Handout from the MIT Workshop on Exhaustivity.*

Frank, M., Goodman, N., Lai, P., Tenenbaum, J. (2009). Informative communication in word production and word learning. *Proceedings of the annual meeting of the cognitive science society* (Vol. 31).

Franke, M. (2009). *Signal to act: Game theory in pragmatics.* University of Amsterdam.

Grice, H.P. (1957, May). Meaning. *The Philosophical Review*, *66*(3), 377–388, https://doi.org/10.2307/2182440 2182440

Grice, H.P. (1975). Logic and conversation. *Speech acts* (pp. 41–58). Brill.

Horn, L. (2004). Implicature. *Routledge companion to philosophy of language* (pp. 75–88). Routledge.

Jäger, G. (2007). Game dynamics connects semantics and pragmatics. *Game theory and linguistic meaning* (pp. 103–117). Brill.

Levinson, S.C. (1983). *Pragmatics.* Cambridge university press.

Lewis, D. (1969). *Convention: A philosophical study.* John Wiley & Sons.

Qing, C., & Franke, M. (2015). Variations on a Bayesian Theme: Comparing Bayesian Models of Referential Reasoning. H. Zeevat & H.-C. Schmitz (Eds.), *Bayesian Natural Language Semantics and Pragmatics* (Vol. 2, pp. 201–220). Cham: Springer International Publishing.

Rothschild, D. (2013). Game Theory and Scalar Implicatures. *Philosophical Perspectives*, *27*(1), 438–478, https://doi.org/10.1111/phpe.12024

Sauerland, U. (2004). Scalar implicatures in complex sentences. *Linguistics and philosophy*, *27*(3), 367–391,

Schreiber, A., & Onea, E. (2021). Are narrow focus exhaustivity inferences Bayesian inferences? *Frontiers in Psychology*, *12*, 677223,

Spector, B. (2003). Plural indefinite DPs as Plural-polarity items. *Amsterdam Studies in the Theory and History of Linguistic Science (Series 4)*, 295–314,

Spector, B. (2007). Scalar implicatures: Exhaustivity and gricean reasoning. *Questions in dynamic semantics* (pp. 225–249). Brill.

van Rooij, R., & Schulz, K. (2004, March). Exhaustive Interpretation of Complex Sentences. *Journal of Logic, Language and Information*, *13*(4), 491–519, https://doi.org/10.1007/s10849-004-2118-6

Van Rooij, R. (2003). Quality and quantity of information exchange. *Journal of Logic, Language and Information*, *12*(4), 423–451, https://doi.org/10.1023/A:1025054901745

Williamson, T. (1996). Knowing and asserting. *The Philosophical Review*, *105*(4), 489–523, 2998423

Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press, USA.