

CROATIAN
JOURNAL
OF PHILOSOPHY

Vol. XVI · No. 47 · 2016

Articles

Introduction DUNJA JUTRONIĆ	127
Précis of <i>Imagination and Convention</i> ERNIE LEPORE and MATTHEW STONE	129
Against Lepore and Stone's Sceptic Account of Metaphorical Meaning ESTHER ROMERO and BELÉN SORIA	145
Intentionalism versus The New Conventionalism DANIEL W. HARRIS	173
Cooperation with Multiple Audiences MARILYNN JOHNSON	203
Coordinating with Language JESSICA KEISER	229
Superficially and Deeply Contingent A Priori Truths MARCO RUFFINO	247

Cooperation with Multiple Audiences

MARILYNN JOHNSON

Department of Philosophy,

The Graduate Center, City University of New York, USA

Steven Pinker proposes a game-theoretic framework to help explain the use of veiled speech in contexts where the ultimate aims of the speaker and hearer may diverge—such as cases of bribing a police officer to get out of a ticket and paying a maître d’ to get a table. This is presented as a response to what Pinker sees as the failure in H. P. Grice’s influential theory of meaning to recognize that speakers and hearers are not always cooperating. In this paper I argue that Pinker mischaracterizes Grice’s views on cooperation, and use this to refine a positive picture of what sort of cooperation is demanded by Grice’s Cooperative Principle. This positive picture serves to insulate the Gricean framework from objectors—including Pinker—who overstate the obligations entailed by the adoption of the Cooperative Principle. I then argue that the cases Pinker presents are best treated by recognizing that in each instance the utterance is formulated with two intentions towards two different audiences and detail a resulting revision to Pinker’s game-theoretic framework that reflects this proposal. I conclude by demonstrating how this proposed game-theoretic framework of cooperation with multiple audiences can be used to model the costs and benefits of other types of discourse, including political speech.

Keywords: Grice, cooperation, cooperative principle, multiple audiences, Pinker, game theory.

Introduction

Philosopher of language H. P. Grice is well known for his theory of speaker meaning, which is grounded in his Cooperative Principle and four maxims of conversation. Grice’s work has influenced much current work in philosophy of language, linguistics, philosophy of law, evolutionary psychology, and many other areas. There is one area of Grice’s program, however, that is frequently criticized: his Cooperative Principle. This principle is attacked by theorists across a wide range of

disciplines who point to the numerous ways in which certain communicative interactions are not cooperative¹ (Marmor 2011; Godfrey-Smith and Martínez 2013; Pinker 2007a; Pinker 2007b, 2011; Pinker, Nowak, and Lee 2008; Lee and Pinker 2010). However, as I will argue in this essay, Grice's Cooperative Principle is not as demanding as some have thought. Many interactions where the interlocutors have divergent ultimate aims are cooperative in the sense relevant to Grice's theory of meaning.

Perhaps the most robust and sustained objections to Grice on these grounds have been waged by psychologist Steven Pinker. In recent work, Steven Pinker (2007a, 2007b, 2011); Pinker, Nowak, and Lee 2008; Lee and Pinker 2010) has proposed a game-theoretic framework to help explain the use of implicature in contexts where the ultimate aims of the speaker and hearer may diverge. Pinker proposes that in such contexts implicature can be used as a way to avoid a number of social and financial costs by discussing examples—such as bribing a police officer to get out of a ticket and paying a *maitre d'* to get a table—and provides a game-theoretic framework that is meant to model these costs and benefits. This is presented as a response to what Pinker sees as Grice's failure to recognize that speakers and hearers are not always cooperating.

In this paper I will argue that Pinker seriously mischaracterizes Grice's views on cooperation (see also Terkourafi 2011a; 2011b; Reboul Forthcoming a; Forthcoming b). In the course of doing so, I refine a positive notion of the sort of cooperation that is demanded by Grice's Cooperative Principle. Although I make Pinker my target, this positive notion could be used to dispel similar objections to Grice made by other theorists as well. I then argue that the cases Pinker presents are best treated by recognizing that in each instance the utterance was formulated with two different audiences in mind (See also Grice 1989: 37 and Neale 1999: 29-30 for discussion of multiple audiences): 1) the immediate audience, and 2) a future potential audience. The apparent obscurity arises because there is a clash between the intended communicated content with respect to these different audiences.

H. P. Grice

In works such as “Meaning”, “Logic and Conversation”, “Utterer's Meaning and Intention” (Grice 1989), Grice presents a theory of speak-

¹ Andrei Marmor writes “The standard model in the pragmatics literature focuses on ordinary conversations, in which the parties are presumed to engage in a cooperative exchange of information” (Marmor 2011: 83) but that “the enactment of a law is not a cooperative exchange of information” (Marmor 2011: 96). Peter Godfrey-Smith and Manolo Martínez write “Many theorists have seen communication as a fundamentally cooperative phenomenon. In an evolutionary context, however, cooperation cannot be taken for granted, because of problems of subversion and free-riding” (Godfrey-Smith and Martínez 2013: 1).

er meaning. Speaker meaning captures how a speaker can write something such as “Mr. X’s command of English is excellent, and his attendance at tutorials has been regular” in a letter of recommendation and thereby mean “Mr. X is no good at philosophy” (Grice 1989: 33).

In this discussion Grice presents his Cooperative Principle. The Cooperative Principle states that conversational partners will make their conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange (29). This is a “quazi-contractual matter” (29). Further, Grice writes that “On the assumption that some such general principle as this is acceptable, one may perhaps distinguish four categories under one or another of which will, in general, yield results in accordance with the Cooperative Principle” (26). These are his four conversational maxims of quantity, quality, relation, and manner. The maxim of quantity states

- “1) Make your contribution as informative as is required (for the current purposes of the exchange) and
- 2) Do not make your contribution more informative than is required” (26).

The maxim of quality states

- “Supermaxim: try to make your contribution one that is true.
- 1) Do not say what you believe to be false, and
- 2) Do not say that for which you lack adequate evidence” (27).

The maxim of relation states

- “1) Be relevant” (27).

The maxim of manner states

- “Supermaxim: be perspicuous (clearly expressed or presented, lucid).
- 1) Avoid obscurity of expression,
- 2) Avoid ambiguity,
- 3) Be brief (avoid unnecessary prolixity), and
- 4) Be orderly” (27).

To understand what Grice was up to with his maxims of conversation we must pay attention to an important distinction between metaphysical ‘determination’ of meaning and epistemological ‘determination’ of meaning (Neale Forthcoming). That is, metaphysical determination of meaning is what makes it the case that some speaker’s utterance has some meaning, and epistemological determination of meaning is the inferential processes hearers go through to work out that meaning (Neale Forthcoming). This is a fundamental point about the roles and aims of inquiries into meaning. Grice’s theory of meaning is a theory of what metaphysically determines meaning; other theorists, such as Sperber and Wilson, are engaged in a project of developing a theory of the epistemological determination of meaning (Neale Forthcoming). This distinction allows us to clearly situate Grice’s project and helps to explain certain features of his theory.

Grice's ideas are some of the most enduring in philosophy of language and have been accepted by many linguists and psychologists. His work has also been the subject of much debate and criticism. One source of criticism is of Grice's theory is the Cooperative Principle and the demands this principle places on interlocutors.

Steven Pinker

One such critic is psychologist Steven Pinker. In his 2007 book *The Stuff of Thought* and later papers in 2008, 2010, and 2011 Pinker proposes a number of instances of perceived non-cooperation as problem cases for Grice and his Cooperative Principle. These are sketched out in *The Stuff of Thought* and spelled out in mathematical detail (Pinker, Nowak, and Lee 2008) and experimental results (Lee and Pinker 2010) in later papers.

In setting up his argument against Grice in *The Stuff of Thought*, Pinker writes,

We'll begin with a famous theory from the philosophy of language that tries to ground indirect speech in pure rationality—the demands of efficient communication between two cooperating agents. This Spock-like theory will then be enhanced by a dose of social psychology, which reminds us that people don't just exchange data like modems. (Pinker 2007a: 375)

In a paper published the following year Pinker and his coauthors write,

Existing theories of indirect speech are based on the premise that human conversation partners work together toward a common goal—the efficient exchange of information, in the influential theory of H. P. Grice (5) ... Yet a fundamental insight from evolutionary biology is that most social relationships involve combinations of cooperation and conflict. (Pinker, Nowak, and Lee 2008: 833)

He further presents his proposal in opposition to Grice writing,

Our theory of the Strategic Speaker supplements the traditional approaches with the insight from evolutionary psychology that most social interaction involves mixtures of cooperation and conflict rather than pure cooperation. (Pinker 2011: 2866)

Pinker continues this characterization with,

Grice came to conversation from the bloodless world of logic and said little about why people bother to implicate their meanings rather than just blurting them out. We discover the answer when we remember that people are not just in the business of downloading information into each other's heads... (Pinker 2007a: 379)

Pinker's criticism against Grice has grown into a sustained attack ranging across a wide body of work.

Pinker goes on to propose a number of instances of perceived non-cooperation as problem cases for Grice and his Cooperative Principle. The first is a case of bribing a maître d'. We are to imagine that someone wants a table in a busy restaurant. This person says something like "Is there any way you could speed up my wait?" while handing the maître d' \$20 intending to implicate a bribe. Pinker has us consider the question, "What are the benefits of implicating the bribe rather than stating it literally?" The second case is of bribing a police officer. We are to imagine that someone gets pulled over for speeding. This person says something like "Can't we settle this here?" intending to implicate a bribe. Again, we are to consider, "What are the benefits of implicating the bribe rather than stating it literally?"

Pinker ties these cases to what he sees as the problem for Grice, writing that the police bribe case,

... is inconsistent with the traditional idea that indirect speech is an implementation of pure cooperation: The driver here is using indirect speech not to help the honest officer attain that goal (viz. to enforce the law) but rather to confound that goal. (Pinker, Nowak, and Lee 2008: 834)

It is important to note that with these examples Pinker does not take himself to be presenting an analysis of bribing *per se* but uses these bribing cases as an instance of *perceived non-cooperation*—his true target. As such, my treatment of Pinker's cases here is not meant to be a treatment of these cases *qua* instances of bribing, but *qua* instances of perceived non-cooperation.

Pinker writes that "any scenario like that in which the best course of action depends on the choices of another actor is in the province of game theory" (2007a: 393) and presents game-theoretic frameworks such as the one below to show these costs and benefits:

	Honest Officer	Dishonest Officer
Don't bribe	Ticket	Ticket
Bribe	Arrest for bribe	Pay bribe; go free

	Honest Officer	Dishonest Officer
Don't bribe	Ticket	Ticket
Bribe	Arrest for bribe	Pay bribe; go free
Implicate bribe	Ticket	Pay bribe; go free

The first matrix displays the options that are available to the speaker if implicature is not on the table.² In the first matrix if the speaker has

² Strangely, Pinker characterizes this first matrix as being what an agent acting in a "perfect" Gricean way, Maxim Man, would be faced with. This is a further serious mischaracterization of Grice that evinces Pinker's inadequate grasp of his theory—but which I will not go into in detail on here beyond this footnote. Pinker writes, "Consider a perfect Gricean speaker who says exactly what he means when he says anything at all. Maxim Man is pulled over for running a red light and is

an honest officer the best option is to pay the ticket; if the speaker has a dishonest officer the best option is to pay the bribe and go free. The second matrix shows the options available if implicature is on the table. Note that in the implicature row of the second matrix the speaker has the options of either paying the ticket, the best option if he has an honest officer, or paying the bribe and going free, the best option if he has a dishonest officer. On Pinker's framework, the speaker, in implicating the bribe, has as the possible outcomes the two best results given some officer.

Further, Pinker writes that for the second matrix to obtain, it must be the case that the speaker "knows that the officer can work through the implicature and recognize it as an intended bribe, and he also knows that the officer knows that he couldn't make a bribery charge stick in court because the ambiguous wording would prevent a prosecutor from proving his guilt beyond a reasonable doubt" (2007a: 394). This is a lot of knowledge for someone to have about a police officer he just met, especially considering that at the same time the speaker is not supposed to know whether the officer is honest or dishonest.

In such sections of *The Stuff of Thought* and later work building off these cases (2007a; 2007b; 2011; Pinker, Nowak, and Lee 2008; Lee and Pinker 2010) Pinker advances the following implicit argument.

1. Veiled speech is inefficient.
2. Because veiled speech is inefficient it is uncooperative.
3. Grice only considers communication that occurs in contexts of pure cooperation, as given by his "Cooperative Principle".
4. Therefore, veiled speech cannot be accounted for by Grice, and we need Pinker's account instead.

In what follows I will argue that Pinker's premises 2 and 3 are both false, beginning with premise 3. This has the result that his argument is not sound, and we do not have support for the conclusion.

pondering whether to bribe the officer. Since he obeys the maxims of conversation more assiduously than he obeys the laws of traffic or the laws of bribery, the only way he can bribe the officer is by saying, 'If you let me go without a ticket, I'll pay you fifty dollars'" (Pinker 2007a: 393; See also Pinker, Nowak, and Lee 2008: 834 for this characterization again). The contrasts between this bizarre characterization of Grice and what his theory actually consists of will become apparent when I discuss the ways the maxims can be weighed against the others and in my positive proposal in later sections. To describe Grice's theory as one that does not allow for conversational implicatures demonstrates a misunderstanding of the most fundamental kind.

Third Premise: Pinker and Grice on the Cooperative Principle

Let me begin with my argument that Pinker's third premise is false. That is, I will now argue that it is not true that Grice only considers communication that occurs in contexts of pure cooperation, as given by his Cooperative Principle.

As noted in the previous section, Pinker contrasts his view with Grice's by saying that Grice's theory of conversation is "based on the premise that human conversation partners work together toward a common goal—the efficient exchange of information" (Pinker, Nowak, and Lee 2008: 833). Pinker characterizes Grice's view as relying on "pure cooperation" (Pinker 2011: 2866; Pinker, Nowak, and Lee 2008: 833) and writes that Grice overlooks the fact that most interactions involve both cooperation and conflict. He writes that "in trying to deduce the laws of conversation from a 'Cooperative Principle'" (note the use of scare quotes) Grice is "guilty" of assuming "the speaker and the hearer are working in perfect harmony" (Pinker 2007a: 392).

However, contrary to the way Pinker presents him, Grice does *not* demand full cooperation. What Grice actually says about cooperation is the following,

Our talk exchanges do not normally consist of a succession of disconnected remarks, and would not be rational if they did. They are characteristically, to some degree at least, cooperative efforts; and each participant recognizes in them, to some extent, a common purpose or set of purposes, or at least a mutually accepted direction...at each stage, some possible conversational moves would be excluded as conversationally unsuitable. (26)

We see in this section that there are a number of places where Grice hedges or qualifies his position on cooperation in some way, using words and phrases such as "do not normally", "characteristically", "to some degree", "at least", "to some extent", "or at least", and "some...would be excluded". Grice is not making a sweeping claim that all conversation exchanges are always fully cooperative, but is making the tentative claim that they characteristically involve some degree of cooperation.

Grice goes on to be quite clear about what features a conversation must have to be cooperative³ on his view, paraphrased here.

³ It should be noted that Grice is not making a claim about the sort of cooperation that must have been required for language to have begun in homo sapiens as a species or with respect to particular human populations. Grice's arguments are made within a culture where we already have both the cognitive requirements for language use in general and a system where particular languages exist. There certainly are interesting questions to be asked about the sort of cooperation required for the cognitive capacity for language to evolve and be selected for, and for language use to be a sustainable system in a species (Dawkins and Krebs 1978; Sterelny 2003; Papineau 2005; Dessalles 2007; Tomasello 2010), but we can pull such questions apart from what I consider here.

Characteristic features that jointly distinguish cooperative transactions:⁴

1. Have some common *immediate* aim
2. Contributions of participants are mutually dependent
3. The transaction will continue until it reaches its natural terminus (the interlocutors will not just walk away in the middle of the conversation) (29)

Grice further elaborates on what he means by this requirement that interlocutors have some “common immediate aim”. He writes, “The participants have some common immediate aim, like getting a car mended; their ultimate aims may, of course, be independent and even in conflict—each may want to get the car mended in order to drive off, leaving the other stranded” (Grice 1989: 29; see also 30). We see here that Grice explicitly allows for cases where the ultimate aims of conversational participants are in opposition.

With Grice’s specific claims at hand let us reflect on one of the cases Pinker presented. Recall that Pinker wrote,

[the police bribe case] ... is inconsistent with the traditional idea that indirect speech is an implementation of pure cooperation: The driver here is using indirect speech not to help the honest officer attain that goal (viz. to enforce the law) but rather to confound that goal. (Pinker, Nowak, and Lee 2008: 834)

However, someone who does not want the officer to enforce the law as in Pinker’s police bribe example can still be cooperative with respect to the conversation.

We could say the speaker is being “Communicatively Cooperative” although not “Ultimately Cooperative”,⁵ understanding these two notions in the following way:

Communicative Cooperation

Demanded by Grice’s Cooperative Principle

1. Speakers abide by Grice’s maxims of conversation of quality, quantity, relation, and manner in formulating their utterances.
2. The conversational participants have enough immediate aims in common that they will not abandon the conversation altogether.

⁴ Dan Sperber and Deirdre Wilson argue that the notion of cooperation demanded by Relevance Theory is weaker than what is demanded by Grice. They write “... Grice assumes that communication involves a greater degree of cooperation than we do. For us, the only purpose that a genuine communicator and a willing audience necessarily have in common is to achieve uptake: that is, to have the communicator’s informative intention recognized by the audience” (1989: 161). Although, it is not clear what this difference really amounts to, and how some set of interlocutors could share the desire that the audience recognize the communicator’s informative intention if they did not have even a minimal common purpose.

⁵ I will henceforth capitalize these terms to show I am using them in a technical way according to the definitions I provide here.

Ultimate Cooperation

Not Demanded by Grice's Cooperative Principle

1. Shared desire for some outcome with respect to the conversation, overall interaction, or some further long-term aim.

Having these two notions at hand can help to streamline discussion of cooperation and communication.

It becomes clear, for one thing, that participants often satisfy the requirements of Communicative Cooperation although they have divergent ultimate aims. Arguments, for one—which are certainly uncooperative in some sense—are very often Communicatively Cooperative according to the conditions just provided. Consider the following argument, between two participants, 1 and 2,

- 1: You never put the forks in the right place.
- 2: Well I think it's stupid to separate the salad forks from the dinner forks.
- 1: Just because you don't know the difference between a salad fork and a dinner fork doesn't mean I should have to live like a heathen!
- 2: A heathen? You're the one who leaves towels on the floor!
- 1: When's the last time I left a towel on the floor?

Notice that this conversation is closely in accordance with Grice's Maxims of Conversation. Specifically,

- a) Each reply is directly in response to the previous comment (Maxim of Relation⁶).
- b) They are consistent in tone (Maxim of Manner).
- c) Each participant is expected to only say things that are true and are challenged if not (Maxim of Quality). And,
- d) Each retort is about the same length (Maxim of Quantity).

This means all the maxims of conversation were abided by. At the same time, neither conversational participant abandoned the exchange. Thus, according to the conditions laid out, this is a clear example of Communicative Cooperation.

This example would still be a clear example of Communicative Cooperation even if each person were trying to achieve any of the following ultimate aims: a) make the other feel insecure, b) drive them to madness c) have the silverware drawer sorted so that they have easy access to the knife they plan to use to kill the other, and so on. None of these ultimate aims affects the fact that the dialogue is Communicatively Cooperative.

⁶ As David Lumsden insightfully points out, there may be an important connection between what he calls the "extra-linguistic goals" and the "linguistic goals" insofar as the interlocutors' ideas about what the further aim of the conversation is could shape what counts as relevant to the conversation, as in Grice's petrol case. At the same time, there is not always this interplay between the two, as he also renognizes (Lumsden 2008: 1901).

The silverware drawer argument example just discussed is one that may at first seem to not be cooperative in the sense demanded by Grice, but in fact is, as becomes evident with a clearly defined distinction between Communicative Cooperation and Ultimate Cooperation.

In making his arguments against Grice Pinker conflates Ultimate Cooperation with Communicative Cooperation. The cases Pinker's presents as counterexamples do not, in fact, present a problem for Grice's Cooperative Principle. And neither would other arguments that conflate Ultimate Cooperation with Communicative Cooperation.

Thus, we are now in a position to return to Pinker's implicit argument.

1. Veiled speech is inefficient.
2. Because veiled speech is inefficient it is uncooperative.
3. Grice only considers communication that occurs in contexts of pure cooperation, as given by his "Cooperative Principle".
4. Therefore, veiled speech cannot be accounted for by Grice, and we need Pinker instead.

I have argued here that the third premise—that Grice only considers communication that occurs in contexts of pure cooperation, as given by his Cooperative Principle—is false.

Failing to be Communicatively Cooperative

With the silverware argument example above I presented a case that might at first appear to be not cooperative in the relevant sense for Grice, but which, upon defining Communicative Cooperation technically, clearly is. This might raise the further question, "What, then, does it take for some utterance to fail to be Communicatively Cooperative?" This might raise the further, related question, "Is every utterance that appears to violate a maxim a case of the speaker failing to be Communicatively Cooperative?" The answer to this second question is 'no'.

A maxim of conversation may appear to be violated in cases of implicature. This is central to Grice's theory of conversation. The fact that a maxim appears to have been violated at the level of what the speaker said is what tips off hearers that the speaker may be implicating something at the level of what the speaker meant⁷ (1989: 33; Neale 1992).

There are three types of implicature in Grice's theory:

- 1) Group A—those in which no maxim is violated, or at least in which it is not clear that any maxim is violated.
- 2) Group B—those in which a maxim is violated, but its violation is to be explained by the supposition of a clash with another maxim.
- 3) Group C—those in which a maxim is flouted, or exploited.

It is clear that Grice recognizes that if one fails to uphold a maxim it

⁷ For more details on what the speaker said, or made as if to say, and what the speaker meant see Grice (1989: 33–34) and Neale (1992: 13–16).

may be justified (because of a clash with another maxim as in Group B or because a maxim is flouted as in Group C).

A speaker who makes a Group B or Group C implicature does so by abiding by the Cooperative Principle. Such cases are not counterexamples to Grice but features of his theory. If the speaker appears to fail to uphold a maxim and it is not a Group B or Group C implicature, this may mean 1) the speaker is *misleading* the hearer with an unostentatious violation—which includes behavior such as lying, or 2) the speaker is *opting out* of the conversation. This can be understood according to the following chart:

Participating or Not Participating in a Cooperative Transaction	Type of Conversational Implicature
<p>A Speaker Appears to Fail to Fulfill Maxims:</p>	<p>GROUP B: Examples in which a maxim is violated, but its violation is to be explained by the supposition of a clash with another maxim</p> <p>A: Where does C live? B: Somewhere in the South of France</p> <p>B's answer is less informative than is required to meet A's needs. Infringement of the Maxim of Quantity can be explained by the supposition that B is aware that to be more informative would be to say something that infringed on the Maxim of Quality. B implicates that he does not know in which town C lives.</p>
<p>Communicatively Cooperative</p> <p>1. Clash Faced by a clash – cannot uphold one maxim without violating another</p> <p>2. Flout a maxim – blatant failure to fulfill - this situation is one that characteristically gives rise to a conversational implicature; a maxim is being exploited</p>	
<p>Not Communicatively Cooperative</p> <p>3. Violate a Maxim Quietly and unostentatiously violate a maxim – will be able to mislead</p> <p>4. Opt Out from the operation of the maxims and the CP – unwilling to cooperate</p>	<p>GROUP C: Examples that involve exploitation, that is, a procedure by which a maxim is flouted for the purpose of getting in a conversational implicature by means of something of the nature of a figure of speech</p> <p>A: How often do you color? B: I can't color enough</p> <p>The hearer is entitled to assume that the maxim, or at least the overall Cooperative Principle, is observed at the level of what is implicated.</p>
<p>If the speaker seems to have violated one of the maxims, hearer will work out that a particular conversational implicature is present by relying on: (1) conventional meaning of the words (2) the Cooperative Principle and Maxims (3) context of the utterance (4) other background knowledge (5) fact that (1-4) are available to both</p>	

A quiet and unostentatious violation of a maxim is a case where the speaker is not being Communicatively Cooperative, but hopes the hearer will proceed as though the speaker is. The speaker formulates utterances with an audience that takes the speaker to be acting in a Communicatively Cooperative way. As Grice notes this can lead to manipulation of the hearer. Lying is one such form of manipulation achieved this way.

Opting out occurs when two potential participants in a talk exchange do not have a common immediate aim, and, in fact, may have immediate aims that are in conflict. For instance, imagine the following three scenarios. You 1) receive a letter in the mail that says “Action Required: Important Survey for Residents”. You then notice the return address says it is from People for the Ethical Treatment of Animals. You’ll not open it. You see PETA is trying to get you to open their mail by deceptive means and opt out, tossing the unopened envelope in the recycling. 2) A woman is walking down the street on her way to a date. A man sitting near the sidewalk says, “God damn”. She will not reply or make eye contact. She opts out of the exchange. 3) You are walking by the Empire State Building on a visit to CUNY for a conference, trying to get to Starbucks on 33rd Street and back to 34th Street before the next talk starts. A man in a red vest carrying pamphlets about the Empire State asks you, “Going up?” You do not want to go to the top of the Empire State Building. You will not reply. Grice recognizes these uncooperative possibilities.

But, in any given case, before we can conclude that an apparently violated maxim means the speaker is being uncooperative, we must make sure the violation does not result from an implicature of the Group B or Group C kind.

With these cases in mind, we can now make a necessary final addition to our understanding of Communicative Cooperation.

Communicative Cooperation

Demanded by Grice’s Cooperative Principle

1. Speakers abide by Grice’s maxims of conversation in formulating their utterances.
2. The conversational participants have enough immediate aims in common that they will not abandon the conversation altogether.
3. If one appears to fail to uphold a maxim it is justified within Grice’s theory of implicature—e.g. because of a clash with another maxim as detailed in the chart on the previous page.

Ultimate Cooperation

Not Demanded by Grice’s Cooperative Principle

1. Shared desire for some outcome with respect to the conversation, overall interaction, or some further long-term aim.

Second Premise: Veiled Speech is Uncooperative

Equipped, now, with a clearer picture of Grice's theory, and what is demanded by his Cooperative Principle, we can return to Pinker's implicit argument. I will now argue that the second premise is false.

Pinker's Argument

1. Veiled speech is inefficient.
2. Because veiled speech is inefficient it is uncooperative.
3. Grice only considers communication that occurs in contexts of pure cooperation, as given by his "Cooperative Principle".
4. Therefore, veiled speech cannot be accounted for by Grice, and we need Pinker's account instead.

In moving toward showing Pinker's second premise to be false we can next ask, "Are there any circumstances in which a speaker can be intentionally obscure while still being Communicatively Cooperative?"

In the section where Grice details a number of Group C implicatures he considers a case of apparent obscurity that arises as a result of the presence of a third party. He writes,

Obviously if the Cooperative Principle is to operate I must intend my partner to understand what I am saying despite the obscurity I impart into my utterance. Suppose that A and B are having a conversation in the presence of a third party, for example, a child, then A might be deliberately obscure, though not too obscure, in the hope that B would understand and the third party not. (Grice 1989: 37)

For illustration of such a scenario, we can consider a scene from a fiction (King 2008), which Grice seems to have anticipated almost exactly. In this scene, four women—Miranda, Carrie, Charlotte, and Samantha—are meeting for brunch. Charlotte has brought along her young daughter, Lily. Carrie and Lily are coloring at the table.

Miranda: How often do you guys have sex?

Lily: Sex!

Charlotte: Miranda, please! (Points to her daughter, Lily)

Miranda: What? She's 3! She doesn't know what it means. I'm 41 and I still don't know what it means.

Charlotte: I know, but she is repeating everything.

Samantha: If I had known that girl talk was going to be on lockdown I wouldn't have flown 3,000 miles.

Charlotte: No, we can talk. Let's just not use the word.

Miranda: Fine. How often do you guys ...

Carrie: (looks up from coloring) ... color?

Charlotte: Thank you!

Samantha: Well, I can't color enough. I could color all day, every day, if I had my way. I would use every crayon in my box.

Carrie: We get it. You love to color. (Turns to Miranda.) Why are you asking? (King 2008)

We see in this example exactly the sort of case Grice mentions as a hypothetical. In this interaction there is intentional and careful obscurity used in the conversation in the form of the adoption of a new word in place of an ordinary word, and we see novel, metaphorical riffs off the new word by Samantha. The speakers in this conversation are being intentionally obscure in the hopes that the other women at the table will understand their meaning and the child will not.

Note that, at the same time, the substitute phrase Carrie comes up with is one that would not stand out to the child as being something that would be strange for people to discuss. If, on the other hand, Carrie had suggested that they discuss “how often they have noodles for arms” even the child would likely pick up on this as being strange and ask what they are talking about. Also, we see in Samantha’s objection, a reason why the topic is not abandoned altogether for a time when the child is absent.

Of course, there are many techniques a speaker might employ if she has an utterance she wishes only one person to hear. She may meet with some audience alone. She may attempt to make her utterance known to only one audience member by manipulating acoustic means—that is, by whispering. However, whispering is recognized as behavior one engages when excluding some potential audience member who otherwise would hear, and, thus, if done the presence of company who feels entitled to be included, will not be used for reasons of politeness. A speaker could switch to a language one hearer knows and another does not. Although this can, again, be seen as rude and may be objected to by the excluded hearer. Thus, in the presence of audience members who for social reasons feel they should not be ostentatiously excluded from the conversation by means such as whispering, switching language, or asking to speak to someone alone, obscurity remains one of the least socially costly options for excluding some audience.

For Grice, a speaker can be intentionally obscure while still being Communicatively Cooperative in cases where there are two audiences the speaker has in mind. The speaker uses veiled or obscure wording so that one audience will understand the meaning of the utterance and the other will not. In the case of the child overhearing an adult conversation it is her presence that explains the use of veiled speech. The speaker is cooperating with the adult listeners insofar as the speaker works to ensure they can still understand what she means by the veiled speech. Such use of veiled speech requires formulating an utterance that will be understood by one audience but not the other.

This gap in interpretive understanding is to be explained by differences in background knowledge and interpretive tendencies of the hearers, which was gauged by the speaker. In the “coloring” example, if the speaker incorrectly gauges the background knowledge of her audience she may: 1) produce an utterance that the child understands and/or 2) produce an utterance that the adults do not understand.

Alternative Game-Theoretic Framework

So, we see that for Grice apparent obscurity can sometimes be explained in terms of a speaker wishing to be understood by one audience in terms of the literal content of their utterance and by the second audience in terms of the implicated content of their utterance. In the police bribe scenario Pinker has it be the case that the speaker implicates a bribe so that there is less chance he will get arrested. However, if the speaker properly implicated the bribe (taking into account the interpretive capacities of the hearer), there would be no lack of clarity from the perspective of the police officer. In order to implicate some content, p , a speaker must intend that his audience will recognize p (Grice 1989: 39). So *to whom* is the content of the implicature in this example really less clear?

I would like to propose that what is going on in the bribing case is that the speaker has formulated the utterance with two audiences in mind (as in the case Grice proposed). The implicature is clear to one audience and unclear to another audience.

The second audience the speaker has in mind does not have to be present. The speaker will think this person could later hear about their utterance and judge their behavior on the basis of it. If the second audience never in fact hears of the utterance this does not change what it meant (recall Neale on ‘determination’).

With this in mind, let us revisit the game-theoretic framework Pinker presents to explain the police bribe case.

	Honest Officer	Dishonest Officer
Don't bribe	Ticket	Ticket
Bribe	Arrest for bribe	Pay bribe; go free
Implicate bribe	Ticket	Pay bribe; go free

Recall that one assumption of Pinker's framework is that in order for a speaker to implicate the bribe along the lines of the third row, it must be the case that the speaker “knows that the officer can work through the implicature and recognize it as an intended bribe, and he also knows that the officer knows that he couldn't make a bribery charge stick in court because the ambiguous wording would prevent a prosecutor from proving his guilt beyond a reasonable doubt”⁸ (Pinker 2007a: 394). There is no reason to make this assumption.

This case can instead be understood first in terms of two different types of interpreters in the officers: 1) those officers who will understand the implicature, and 2) those officers who will not. This is not equivalent to the difference Pinker details between honest officers and

⁸ Although Pinker later contradicts himself and writes that for certain bribes that are more heavily veiled there is “some risk that they might go over the head of a bribable officer” (Pinker 2007a: 395).

dishonest officers. For Pinker the difference between honest and dishonest officers is spelled out in terms of *the officer's response to the implicature, not whether or not he or she understand an implicature was made*. There is a gap between whether some person understands an implicature was made and how they act on the basis of this recognition. This is an important gap to recognize and the following framework will do so.

If we understand the officers as two different classes of interpreters that constitute two different audiences, we can see this case as a close parallel to the one Grice presents, where the officers who do not understand an implicature was made are akin to the child, in the sense that they understand the utterance in terms of its literal meaning, and those who do understand an implicature was made are akin to the adults, in the sense that they understand the utterance in terms of its implicated meaning.

This means that we must revise Pinker's game theory framework to reflect these two audiences.⁹

Audience 1. Savvy Police Interpreter—will understand the implicature

	Honest Officer	Dishonest Officer
Don't bribe	Ticket	Ticket
Bribe	Arrest for bribe	Pay bribe; go free
Implicate bribe	Arrest for bribe	Pay bribe; go free

Audience 2. Naïve Police Interpreter—will not understand the implicature

	Honest Officer	Dishonest Officer
Don't bribe	Ticket	Ticket
Bribe	Arrest for bribe	Pay bribe; go free
Implicate bribe	Ticket	Ticket

Veiled speech has the added benefit that some honest officers who might have otherwise arrested a driver for bribery will not understand that a bribe was made. At the same time, this framing highlights one of the risks of using veiled speech with an audience one is unfamiliar with: there is a possibility the intended audience will not work it out—in this case, losing out on officers who would take the bribe if they could work out one was offered.

A negotiation between these two possible audiences leads to (part of) the benefit of having some bribe be more or less explicit. A more explicit bribe increases a) the chances an officer who would take the bribe will work out one is being made and b) the chances an officer who will arrest for a bribe will work out one is being made. A more veiled bribe increases a) the chances an officer who would take the bribe will not

⁹ Matrices modeled after Farrell and Gibbons (1989).

work out one is being made and b) the chances an officer who will arrest for a bribe will not work out one is being made.¹⁰

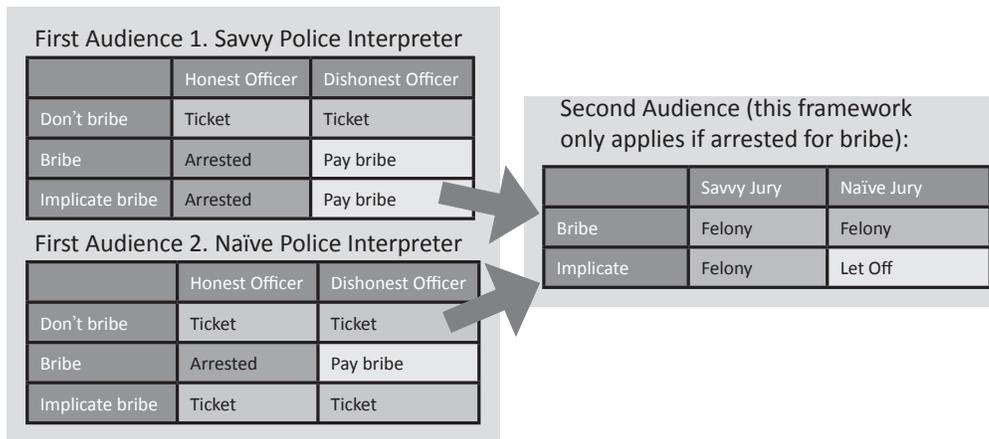
I am not arguing that this is what Pinker really meant but that what is really going on in the sorts of cases he discusses is best modeled with this framework that allows for more nuance, and makes clear the ways in which his account can be brought into conjunction with what Grice says about multiple audiences.

To most fully model the circumstances of the example we must also consider another audience. Should the driver be arrested, another audience becomes relevant: the jury. We should also consider the utterance with respect to the second set of audiences. In doing so, we separate two issues: 1) whether or not the police officer will understand the implicature, and 2) (bringing in a new potential audience) whether or not a jury will understand the implicature. If the speaker does formulate his utterance with two audiences in mind, as I suggest, this means that we must further revise Pinker's game theoretic framework to reflect these two audiences.

If it is right that the utterance was formulated with both of these pairs of audiences in mind—both 1) a savvy and naïve officer, and 2) a savvy and naïve jury¹¹—the final game-theoretic framework would be a combination of two matrices.

¹⁰ Let me provide a brief note on how this utterance could plausibly be received as something other than a bribe. When I presented this talk at the 2015 Language and Linguistics Conference in Dubrovnik, along with many of the other papers in this volume, there was an interesting discussion stemming from a professor who lives and works in Mexico struggling to understand how this utterance could be considered anything other than a bribe. In fact, in the U.S., where Pinker lives and works, the rules for how speeding tickets are paid vary by state. So, this example makes more sense in its setting in America where, because of the variance in rules by state, it is likely that there may be genuine ignorance or confusion on the part of some drivers about what the rules are in that state and where bribing is less common than in many other countries (Walton 2013).

¹¹ I realize these names “naïve” and “savvy” do not fully capture what is going on here but have decided to stick with them because they do capture some piece—although in perhaps too dramatic a fashion. “Literal” and “read-between-the-lines” could be another way of spelling it out. Also, this is not meant to be a fixed feature of the agent per se, but the agent faced with some particular utterance. Some police officer might be very good at recognizing a bribe in his native language, but more likely to interpret literally in his second language, for example.



On the left, calculation of what the best move is in this step depends first on the values given in each of the cells and the ratio of honest and dishonest police officers. This is calculated within each matrix. Then the matrices are weighted with respect to the odds of a police officer being savvy or naïve. On the right, calculation of what the best move is in this step is done within the grid and depends on the ratio of savvy to naïve jurors.

The honesty or dishonesty of the police officer comes from the decision to put the utterance up for evaluation by the second audience. This feature is independent of this audience's ability to work out implicatures. A similar thing is going on in Pinker's maître d' case. In each case, who the relevant audiences are depends on who the speaker has in mind, which is a function of who the speaker believes may respond in a way that has costs or benefits for him/her.

Takeaway

We can now, once again, return to Pinker's argument.

1. Veiled speech is inefficient.
2. Because veiled speech is inefficient it is uncooperative.
3. Grice only considers communication that occurs in contexts of pure cooperation, as given by his "Cooperative Principle".
4. Therefore, veiled speech cannot be accounted for by Grice, and we need Pinker instead.

I have argued here that premise 2 is false, and argued in previous sections that premise 3 is false. This means that Pinker's argument is not sound and that we do not reach his intended conclusion.

Despite the flaws in Pinker's argument, the first premise in his argument—which I have not argued against and which points to the inefficiency of indirect speech as an important puzzle—is true, and is an important point that has been overlooked in much literature on implicature. Implicature is costly: it is difficult for a speaker to produce, difficult for a hearer to interpret, and there is increased risk that the message will go wrong along the way (Kruger 2005). Because of these potential

costs to the use of indirect speech we ought to seek an understanding of why and when speakers choose to use implicature rather than state content literally. Pinker's idea of using a game theoretic framework to model the costs and benefits of certain utterance types—although flawed in the details of its execution and especially in his understanding of Grice—is a very helpful way to consider this question. The multiple audience framework I have presented here could be utilized to help us understand why implicature is used in a number of cases.¹²

The multiple audience framework provides one piece of the puzzle of when we use implicature. There are many other pieces of the puzzle and reasons for using implicature in addition to those discussed here.¹³ It is a puzzle that anyone who makes use of Grice's framework of conversational implicature or other similar accounts of veiled speech should be interested in solving. Even if there is a quibble about my treatment of Pinker's specific cases here I hope to have succeeding in making the more general point against the general contours of Pinker's arguments against Grice.

Up to this point I have sought to stress the ways that Pinker mischaracterizes Grice on cooperation. I also presented a new game theoretic framework that deals with apparent uncooperative obscurity in terms of two audiences, and allows us to better explain Pinker's bribing case. My ultimate goal, however, is a more general one about the aims of the Gricean project, the tools at hand to help us meet those aims, and overlooked but worthy questions relevant to that project.

a. Deniability

I will now consider some objections and addendums to the proposal as presented thus far. Someone responding to the ideas put forth in this paper may object that this analysis is unnecessary because the benefits afforded by using veiled speech can be accounted for by saying that the veiled speech is used so that the implicated content is deniable. Pinker, himself, appeals to a notion of "plausible deniability" to explain the bribing cases.¹⁴ However, if we are to say that some content is deniable, we must be clear about what is demanded by deniability. Four points of needed clarification emerge.

¹² I have in mind Grice's tea party and letter of recommendation cases.

¹³ For example, one other piece of the puzzle might be in our judgments that those who express criticism with sarcasm rather than with literal speech are seen as being more in control (Dews, Kaplan, and Winner 1995; Pinker 2007a: 379).

¹⁴ I seem to be at odds here, too, with Pinker, who writes that one who has uttered "They never seem to have enough salt shakers at this restaurant" can deny having asked for the salt, and that someone hearing such an utterance could simply ignore it without being rude. He also writes that for this reason, ostensibly following Brown and Levinson (1987), utterances such as "It's too dark to read" as a request for the lights to be turned on are more polite than direct requests, again dubious claim. It seems such a speaker, who assumes it is the duty of the hearer to quell all concerns, speaking to anyone other than a servant, would be quite rude and unpleasant.

First, in saying that some claim is deniable, what are we saying about the likelihood that this act of denying succeeds? Surely this objection rests on deniability being something stronger than just that some claim *could* be denied—for any claim whatever *could*, in principle, be denied. Consider the following utterances made by the same person:

- 1) I'm from Athens.
- 2) Oh—but I don't mean I'm from Athens.

Or, drawing from a recent political news event,

- 3) "Perhaps there are two Donald Trumps."
- 4) "I don't think there's two Donald Trumps. I think there's one Donald Trump." (Nguyen 2016)

Such a series of remarks would confuse and require clarification by the hearer. It seems not that the speaker has succeeded in denying the content of the first utterance but simply said two things that are contradictory. Deniability, then, seems to demand some reasonable chance of success of convincing an interlocutor that the speaker did not mean the denied content.

Second, we have a further point of clarification demanded of anyone who relies on a notion of deniability. That is, he or she must take a stance on whether succeeding in denying some utterance is an act of 1) getting the hearer to believe that the asserted proposition is not true, 2) getting the hearer to believe that the speaker did not mean the asserted content (although it may be true), or 3) Both 1) and 2)—that is, an act of getting a hearer to believe both that the asserted proposition is not true, and that the speaker did not mean the asserted content.

Third, at the same time, we would want to isolate deniability from clarification. Consider the following modification on two of the utterances above.

- 5) I'm from Athens.
- 6) Oh—but I don't mean I'm from Athens, Greece; I'm from Athens, Georgia.

In this case we do not have a true case of deniability because in this instance the speaker (assuming he had a genuine intention to speak about Athens, Georgia in the first utterance) is not denying the proposition he meant in the first utterance. What he was likely doing in this case is realizing after he made the utterance that he had been misunderstood, and wanted to rectify this by making a more explicit utterance. Such a case would result from a speaker not accurately assessing how he would be interpreted. Thus, we have a further distinction between deniability applying to 1) what the speaker meant or 2) what the hearer took the speaker to mean, or 3) what the speaker thought (non-factive) the hearer took her to mean. Any appeal to deniability would require resolution of what the required target of the act of denying is.

Fourth, in saying that some content is deniable the further question also arises of to whom the content is to be deniable, especially if we are considering utterances formulated with multiple authors in mind.

b. Cancellability

Such discussions of deniability can be somewhat clarified by appealing to Grice's notion of cancellability. For Grice, if some content was implicated, rather than stated literally, the speaker should be able to cancel the implicature (39). Consider the following example:

- 7) Mr. X has wonderful handwriting.
- 8) Of course, I am not saying he is no good at philosophy.

In this example there is a reasonable chance that the speaker would succeed in cancelling the content of utterance 7—or at least a better chance than if he had stated the implicated content literally in the first utterance.

Grice's notion of cancellability, can, at times appear to be at odds with his spelling out of implicatures. For, in order for some speaker *s* to have formed a genuine meaning-intention with respect to some content *p* and some audience *a* by uttering some utterance *r*, *s* must expect that *a* will work out *p* on the basis of *r*. If this is right, why would *s* then think he could subsequently get *a* to believe $\sim p$?

One way around this apparent problem with cancellability is to say that what is going on here is not that *s* get *a* to believe $\sim p$, but that *s* get a second audience, *a**, to believe $\sim p$. We could say, then, that the cancellability of some content *p* is a feature of the ability to get *a** to believe $\sim p$ after hearing *r*. We see, then, that deniability—as understood in terms of Grice's notion of cancellability—is not a true alternative to the multiple audience framework but a concept that is already present in the Gricean framework and could be brought together with the multiple audience framework to be fully fleshed out.

c. The Demands of the Multiple Audience Framework

Some object to Grice because his account of speaker meaning is thought to be too psychologically demanding. Jennifer Hornsby, for example, writes the following of Grice's theory: "I think that this ought to seem ludicrous. Real people regularly get things across with their utterances; but real people do not regularly possess, still less act upon, intentions of this sort...notice that an enormous amount would be demanded of hearers, as well as speakers, if such complex intentions really were needed to say things" (Hornsby 2000: 95). The framework I present here is an especially complex application of the Gricean framework. What, then, of this objection to Grice, which seems doubled here?

There are many behaviors we engage in that seem dizzyingly complex when spelled out, such as athletes catching a baseball, musicians playing a saxophone solo, or firefighters deciding when to run out of a burning building (Kahnemann 2013). The complexity of these behav-

iors themselves is not enough to insist that we do not really engage in them, but further reasons to be in awe of what the human mind can do.

With regard to the multiple audience framework, there are undoubtedly many cases where, because of the potential costs and benefits, speakers do successfully navigate the waters of multiple audiences. The “coloring” example above is one such instance. A starting place of the objection—the belief that formulating some meaning intention with respect to multiple audiences is highly demanding—is right, and this is why it routinely goes wrong. However, the psychological demandingness appealed to in this proposal is not enough on its own to reject it.

Applying the Multiple Audience Framework to Political Speech

Failure to recognize that some utterance may be put up for interpretation by a second audience can lead to disastrous consequences. For my concluding example, I will draw on a real-world case that is an instance of such a disaster. This case makes it clear that my suggestion that a speaker could or should formulate an utterance while keeping in mind the possible interpretations of a second audience is not something cooked up for hypothetical examples, but something we routinely engage in.

Mitt Romney lost to Barack Obama in the 2012 presidential election. The blow that may have cost him the presidency resulted from failing to take into account the interpretation a second audience would make of one of his utterances. The gaffe occurred at a Boca Raton fundraiser, which would have been filled with very right-leaning donors to Romney from whom he would like to solicit as much money as possible.¹⁵ With this audience it is most advantageous to use strong rhetoric. Romney told the crowd he was unconcerned with the 47 percent of the electorate who supported Obama (that is, “those people”, “victims” who take no “personal responsibility”) (Leibovich 2014).

However, unbeknownst to Romney, someone in the audience filmed his utterance and posted it online. Thus, this utterance that first took place “behind closed doors” became available for interpretation by a second audience, the general public, who were outraged by what Romney had said. A headline in *Mother Jones*, which broke the story, read, “Secret Video: Romney Tells Millionaire Donors What He REALLY Thinks of Obama. When he doesn’t know a camera’s rolling, the GOP candidate shows his disdain for half of America” (Corn 2012). After this footage came out, Romney’s poll numbers slipped irreparably.

¹⁵ Note here the two audiences are made up of two types of interpreters, not two individuals. The multiple audience framework can be applied either to individual audiences, or audiences grouped according to some similar interpretive tendencies.

On September 30, 2014, nearly two years after he lost the election, Romney did an interview with the *New York Times* in which he reflected on why he lost.

‘I was talking to one of my political advisers,’ Romney continued, ‘and I said: ‘If I had to do this again, I’d insist that you literally had a camera on me at all times’ — essentially employing his own tracker, as opposition researchers call them. ‘I want to be reminded that this is not off the cuff.’ This, as he saw it, was what got him in trouble at that Boca Raton fund-raiser. (Leibovich 2014)

He continued, “My mistake was that I was speaking in a way that reflected back to the man,” Romney said. “If I had been able to see the camera, I would have remembered that I was talking to the whole world, not just the man” (Leibovich 2014). As he later recognized, Romney would be better off if he had formulated his utterance by weighing the costs and benefits with respect to his multiple audiences. He felt he had erred in not doing so in the first place.

This tension is common in political discourse. Politicians are constantly negotiating the line between energizing supporters and making comments that will appeal to—or at least not greatly offend—the general public (Economist 2005; Fear 2007; Haney López 2014; Nguyen 2016). The proposed game theoretic framework of cooperation with multiple audiences could be used to model such negotiations. The framework presented here has wide applicability to a number of types of discourse.

Conclusion

In this paper I have argued against some of the claims made by Steven Pinker in his book *The Stuff of Thought* and in subsequent papers. In particular, I argued that Pinker is misguided in a number of objections he makes to the theories of philosopher H. P. Grice. In arguing for Grice I developed positive notions of the sort of cooperation Grice demands and explained how cases of obscure language can be cooperative within Grice’s framework. Attention to the role multiple audiences play in leading to intentional, cooperative obscurity led me to return to some cases Pinker presented against Grice within Grice’s treatment of multiple audiences. I concluded by addressing some possible objections to this proposal and by suggesting how the multiple audiences framework I presented could be applied to cases of political discourse.

Acknowledgements

This paper was presented at the CUNY Graduate Center Workshop on Relevance, Pragmatics, and Game Theory in New York in 2014, the Philosophy of Language and Linguistics Conference at the Interuniversity Center in Dubrovnik in 2015, and the Thursday Seminar Series at the Australian National University in Canberra in 2015. I am grateful to the audiences at all these talks for their helpful comments and ques-

tions. In addition, I would especially like to thank Justin Bruner, Ryan DeChant, Michael Devitt, Maita Escurdia, Michael Glanzberg, Michelle Johnson, Dunja Jutronic, Karen Lewis, Eliot Michaelson, Nenad Mišcević, Stephen Neale, David Papineau, Ronald Planer, Anne Reboul and Deirdre Wilson.

References

- Brown, P. and Levinson, S. C. 1987. *Politeness: Some Universals in Language Usage*. Cambridge: Cambridge University Press.
- Dawkins, R. and Krebs, J. R. 1978. “Animal Signals: Information or Manipulation?” In Krebs, J. R. and Davies, J. B. (eds.) *Behavioral Ecology*. Oxford: Blackwell.
- Dessalles, J. L. 2007. *Why We Talk: The Evolutionary Origins of Language*. Oxford: Oxford University Press.
- Dews, S, Kaplan, J. and Winner, E. 1995. “Why Not Say It Directly? The Social Functions of Irony”. *Discourse Processes* 19: 347–367.
- Economist. 2005. “High Pitch, Low Politics”. *The Economist*. March 23. Accessed online.
- Fear, J. 2007. “Under the Radar: Dog-Whistle Politics in Australia”. *The Australia Institute*. Discussion Paper No. 96. September.
- Farrell, J. and Gibbons, R. 1989. “Cheap Talk with Two Audiences”. *The American Economic Review* 79 (5): 1214–1223.
- Godfrey-Smith, P. and Martínez, M. 2013. “Communication and Common Interest”. *PLOS Computational Biology* 9 (11): 1–6.
- Goltsman, M. and Pavlov, G. 2011. “How to Talk to Multiple Audiences”. *Games and Economic Behavior* 72: 100–122.
- Grice, H. P. 1989. *Studies in the Way of Words*. Harvard: Harvard University Press.
- Haney López, I. 2014. *Dog Whistle Politics: How Coded Racial Appeals Have Reinvented Racism and Wrecked the Middle Class*. Oxford: Oxford University Press.
- Hornsby, J. 2000. “Feminism in Philosophy of Language”. In Fricker, M and Hornsby, J. (eds.) *The Cambridge Companion to Feminism in Philosophy*. Cambridge: Cambridge University Press: 87–106.
- Kahneman, D. 2013. *Thinking Fast and Slow*. New York: Farrar, Straus and Giroux.
- King, M. P. 2008. *Sex and the City: The Movie*. (Writer, Director). Film. New Line Cinema.
- Kruger, J. et. al. 2005. “Egocentrism Over E-mail: Can We Communicate as Well as We Think?” *Journal of Personality and Social Psychology* 89 (6): 925–936.
- Lee, J. and Pinker, S. 2010. “Rationales for Indirect Speech: The Theory of the Strategic Speaker”. *Psychological Review* 117 (3): 785–807.
- Leibovich, M. 2014. “Mitt Isn’t Ready to Call It Quits”. *New York Times*. September 30. Accessed online.
- Lumsden, D. 2008. “Kinds of Conversational Cooperation”. *Journal of Pragmatics* 40: 1896–1908.
- Marmor, A. 2011. “Can the Law Imply More Than It Says? On Some Pragmatic Aspects of Strategic Speech”. In Marmor, A. and Soames, S. (eds.)

- Philosophical Foundations of Language in the Law*. Oxford: Oxford University Press: 83–104.
- Neale, S. 1992. “Paul Grice and the Philosophy of Language”. *Linguistics and Philosophy* 15 (5): 509–559.
- Neale, S. 1999. “Colouring and Composition”. In Stainton, R. (ed.) *Philosophy and Linguistics*. Boulder: Westview Press: 35–82.
- Neale, S. Forthcoming. *Linguistic Pragmatism*. Oxford: Oxford University Press.
- Nguyen, T. 2016. “The Two Trumps: 5 Signs Donald Trump Is Having a Full-Blown Identity Crisis”. *Vanity Fair*. April 22. Accessed online.
- Papineau, D. 2005. “The Cultural Origins of Cognitive Adaptations”. In O’Hear, A. (ed.) *Biology, Evolution, and Culture*. Royal Institute of Philosophy Supplement No. 56. Cambridge: Cambridge University Press: 291–324.
- Pinker, S. 2007a. *The Stuff of Thought: Language as a Window into Human Nature*. London: Penguin.
- Pinker, S. 2007b. *Games People Play: Indirect Speech as a Window into Social Relationships*. Lecture. Presented at the Institute of Philosophy, London, UK. October 17, 2007. Video accessed online at: http://library.fora.tv/2007/10/15/Steven_Pinker_Games_People_Play.
- Pinker 2011 “Indirect speech, politeness, deniability, and relationship negotiation Comment on Marina Terkourafi’s ‘The Puzzle of Indirect Speech’”. *Journal of Pragmatics* 43: 2866–2868.
- Pinker, S., Nowak, M. and Lee, J. 2008. “The Logic of Indirect Speech”. *PNAS* 105 (3): 833–838.
- Reboul, A. Forthcoming a. “Is Implicit Communication a Way to Escape Epistemic Vigilance?” In Avrassimopoulos, S. (ed.) *Pragmatics at Its Interface*. Berlin: Mouton.
- Reboul, A. Forthcoming b. *Communication and Cognition in Language Evolution*. Oxford: Oxford University Press.
- Sterelny, K. 2003. *Thought in a Hostile World: The Evolution of Human Cognition*. Oxford: Wiley-Blackwell.
- Sperber, D. and Wilson, D. 1989. *Relevance: Communication and Cognition*. Oxford: Blackwell.
- Terkourafi, M. 2011a. “The Puzzle of Indirect Speech”. *Journal of Pragmatics* 43: 2861–2865.
- Terkourafi, M. 2011b. “Why Direct Speech Is Not a Natural Default: Rejoinder to Steven Pinker’s ‘Indirect Speech, Politeness, Deniability, and Relationship Negotiation’”. *Journal of Pragmatics* 43: 2869–2871.
- Tomasello, M. 2010. *Origins of Human Communication*. Cambridge: Bradford.
- Walton, J. et al. 2013. “Map: Which Country Pays the Most Bribes?”. *BBC News*. Accessed online. July 9.