Research Article

Charles Forceville*

Relevance theory as the foundation for an inclusive theory of communication

https://doi.org/10.1515/mc-2025-0036
Received August 12, 2025; accepted September 10, 2025; published online September 23, 2025

Abstract: Sperber and Wilson's relevance theory (RT) is well-equipped to develop into an inclusive theory of multimodal – and indeed all other – communication. However, in its "classic" variety, it typically focuses on spoken communication in a face-to-face situation. To fulfil its promise, it needs to be adapted and expanded to be able to accommodate *all* media, modes, and genres. Forceville (2020. *Visual and multimodal communication: Applying the relevance principle.* Oxford: Oxford University Press) launches proposals that enable Sperber and Wilson's classic RT to account for visual and multimodal mass-communication. This paper presents key concepts of RT, responds to criticisms of Forceville's version of the theory by Ntouvlis (2021. Review of Forceville (2020). *Multimodal Communication*), and discusses six visual/multimodal case studies to demonstrate RT's value for further developing multimodality as a robust scholarly discipline.

Keywords: relevance theory; visual and multimodal communication; semiotics; systemic functional linguistics

1 Introduction

Sperber and Wilson (Sperber and Wilson 1995; Wilson and Sperber 2012; see also Clark 2013) did groundbreaking work to show how communication is ultimately rooted in the relevance principle. Relevance theory (henceforth: RT) is indebted to Grice's (1975) Cooperative Principle, but departs from Grice's model in two crucial ways. In the first place, Sperber and Wilson promote one of Grice's four maxims of conversation (namely, the maxim of relation/relevance) to the all-encompassing principle to which the other three maxims (of quantity, quality, and manner) are subservient. In the second place, whereas Grice formulated his four maxims as recommendations that a good communicator should adhere to, RT's relevance principle is deeply entrenched in communication partners' cognition, and thus does not need to be in any sense negotiated between them: the relevance principle is automatically activated.

RT has the potential to develop into an inclusive theory of communication. However, for this to materialise the theory needs to be adapted and extended in various ways, since "classic" RT focuses almost exclusively on face-to-face communication. This is thus monomodal communication of the spoken-verbal variety – although it may also draw on the "bodily action" mode in the form of gestures and facial expressions (for more discussion of what is, or could qualify as, a mode, see Forceville 2021). Forceville (2020) shows how relevance theory can, and must, be adjusted to be able to account for visual and multimodal communication (see also Forceville 1996, 2005, 2014, 2022; Forceville and Clark 2014; Forceville and Sánchez Querubín 2022). Since such communication is typically mass-communicative, these studies also address the issue how exchanges are affected when they are of the one-to-many variety instead of the one-to-one type characteristic of classic RT.

For a long time, approaches and theories focusing on visual and multimodal communication primarily built on (social) semiotics and Systemic Functional Linguistics/SFL (e.g., Adami 2017; Archer and Breuer 2015; Baldry and Thibault 2006; Djonov and Zhao 2013; Jewitt 2014; Jewitt et al. 2016; Kress 2010; Kress and Van Leeuwen 2001,

***Corresponding author: Charles Forceville**, Department of Media Studies/ACLC, University of Amsterdam, Amsterdam, Netherlands, E-mail: c.j.forceville@uva.nl. https://orcid.org/0000-0002-6365-500X

2021; Norris 2012; Royce and Bowcher 2007; Serafini 2022; Van Leeuwen and Jewitt 2001 – for critical reviews of several of these studies, see my Google Scholar profile; see also scholars drawing on different models, e.g., Bateman et al. 2017; Klug and Stöckl 2016; Wildfeuer et al. 2019). Sperber and Wilson (1995: 2) rather summarily reject semiotics as a candidate for an inclusive theory of communication because it is rooted in a code-model of communication. By contrast, RT holds that communication combines decoding with inferring. I agree that RT is a better candidate for providing an inclusive model of communication than social semiotics/SFL, but I also think that these paradigms' best insights can, and should, be integrated in that inclusive RT model. For instance, an inclusive theory can benefit from semiotics' broader interpretation of the concept of "encoding/decoding" (Forceville 2020: 74–80).

In this paper I will first explain the basics of RT (there is no space to discuss all its aspects). Then I will respond to Ntouvlis' (2021) SFL-oriented criticisms of Forceville (2020). After that six case studies are discussed to show how RT can steer the analysis of multimodal discourses. I end with some summarizing remarks. Please note that while I think my proposals are completely true to the spirit of classic RT, I sometimes take liberties with it so as to do justice to how visual and multimodal mass-communication differs from the face-to-face variety. My own work can thereby be considered as a "post-classic" version of RT.

2 A brief overview of RT's key concepts as adopted in Forceville (2020)

2.1 Relevance is relevance to the envisaged audience

Human beings – and probably *all* living creatures – constantly monitor their environment for information that is important to them. Crudely simplified, information is of two kinds: either it bodes good for an individual, because it (potentially) serves the interests of that individual, or it bodes bad for him, because it (potentially) harms his interests. This idea is captured in Sperber and Wilson's Cognitive Principle of Relevance: people are constantly alert to anything happening around them that may benefit or endanger their chances of survival and well-being. But while nature is indifferent to the interests of humans, fellow humans, particularly those who are members of the same group, are naturally disposed, other things being equal, to help each other survive in daily life literally and metaphorically. This is probably less a matter of benign altruism than an awareness that helping others within one's community ultimately means helping oneself, too. A crucial way to ensure survival is constantly engaging in communication with conspecifics in a relevant way. Thus, Sperber and Wilson claim, the Cognitive Principle of Relevance gives rise to the Communicative Principle of Relevance, which they define as follows: "Every act of ostensive communication communicates a presumption of its own optimal relevance" (1995: 260). It is the Communicative Principle of Relevance that constitutes the heart of RT.

The central claim of RT can be summed up in a single sentence: every ostensive act of communication comes with the presumption of optimal relevance to its envisaged audience. From the envisaged audience's point of view this can be phrased as: the audience at which an act of communication is directed expects this act of communication to be optimally relevant to it. Let me zoom in on each of the key terms of this pithy claim.

An *ostensive act of communication* is any sign or signal in any mode or combination of modes which carries a more or less specific content that its sender – the communicator – wants to convey to one or more addressees. Aiming this message, discourse, text, picture, tune, gesture, sound, bodily movement, etc. emphatically at the envisaged audience, the communicator tacitly promises to this audience: you are invited to pay attention, since what I am conveying to you is relevant to you. The form that I, the communicator, have chosen for this "ostensive stimulus" (Sperber and Wilson 1995: 153) is optimal in the sense that it is the best one I can and wish, in the current situation, to provide for you.

Relevance. "Relevance" itself is a technical term in RT, although it is rooted in its informal, everyday meaning. RT stipulates that a message's relevance has two components: its *effect* on the envisaged audience, and the *effort* it takes the envisaged audience to retrieve this effect. We can understand these as the benefit and the cost side of

processing the message, respectively. When cognitive effect increases, relevance increases; when effort increases, relevance decreases. Relevance thus grows proportionally with the extent to which the message conveyed is important for the envisaged addressee. Thus, in most circumstances, "I want to marry you" has many more effects for – and is thus more relevant to – the addressee, than "there is a hole in your sock". By contrast, given the greater amount of mental effort needed, "Looking downwards to the part of the world in which your shoes touch the ground, it has come to my attention that the textile fabric covering your feet and disappearing in one of your shoes has an opening showing part of the lower part of your left leg" is by that criterion less relevant than "there's a hole in your sock". It is to be noted that processing cognitive effects in, for instance, hermetic poems, abstract paintings, and philosophical treatises requires considerable mental effort – and their typical addressees are aware of this. It is thus up to them to decide, for any such ostensive stimulus, whether they want to consider themselves part of the envisaged audience. If so, they have to pay the price of investing substantial mental effort in order to be able to reap the cognitive effects.

The presumption of optimal relevance. The word "presumption" is crucial, since a presumption is by no means a guarantee. The idea is that whenever somebody begins to address you, your "gut reaction" is that whatever it is this person is going to convey to you is, in one way or another, relevant to you. But you may be disappointed in that expectation. Indeed, often what people convey to you turns out to be irrelevant. But this leaves intact the expectation that every act of communication directed at you comes with the promise that it is optimally relevant to you.

The envisaged audience. Classic RT states that relevance is always relevance to an individual (Sperber and Wilson 1995: 142–151), whereby it is important to realize that for the sake of clarity RT has adopted the convention that the communicator is always female and the addressee is always male. In the many examples in Sperber and Wilson (1995), "Mary" needs to be optimally relevant to only one person, "Peter" – and to nobody else. This means that she needs to take into account only the sum total of things she knows or believes Peter knows, believes, hopes, wants etc. that constitute his "cognitive environment" (1995: 38–45). Since one gets the impression from the many examples in classic RT that Mary and Peter know each other very well (they are probably married), Mary's acts of communication can usually leave out or keep unspecified numerous things, because she knows that Peter already knows them, and to save him from having to spend unnecessary mental effort in processing her utterance, she will typically favour briefness wherever possible. Thus, when she sees Peter looking at her askance at the breakfast table, her ostensive stimulus "on the top shelf" is supposed to be optimally relevant for him: Peter will mentally complete that lapidary phrase into, say, "the marmalade, which Mary knows I like putting on my morning toast, is on the top shelf in the left cupboard in the corner of the kitchen". But if the asking look to Mary would come from Alwin, Peter's brother, who is staying with the couple for a few days, her optimally relevant ostensive stimulus might be, "if you should be looking for anything to put on your toast, you will find marmalade on the top shelf in the left cupboard in the corner of the kitchen" – which would inevitably require more mental effort on Alwin's part.

Now while all this is relatively simple in Mary-to-Peter or Mary-to-Alwin communication, things become considerably more complicated in mass-communication. The moment a communicator addresses more than a single person, she needs to take into account the cognitive environment of at least two (but more typically, hundreds, thousands, or millions of) individuals. Since each individual has a unique cognitive environment, the communicator has to finetune her message in such a way that it will be optimally relevant to all individuals she addresses. For this reason I adapted, in Forceville (2020), "relevance to an individual" to "relevance to the envisaged audience". Given the size of most mass-communicative audiences, each individual in it with his/her/+'s unique cognitive environment, it is possible that the individuals in the envisaged audience will not interpret the message in precisely the same manner. I will come back to this in Sections 2.4 and 2.5.

Moreover, whereas in classic RT the communicator is always Mary, in mass-communication it is frequently impossible to attribute the communicator role to a single individual. Often a message comes from an organization, say "the tax office" or "the government" or "Amnesty International". That said, the presumption/expectation of relevance remains in operation: the communicator trusts that the envisaged audience will find the message useful, interesting, insightful etc. In order to achieve this, the communicator thinks of the best way to formulate the message, so as not to force the envisaged audience to expend unnecessary mental effort to

understand it. Of course the larger and more varied the envisaged audience, the less opportunity the communicator has to ensure that the message "lands well" for each individual member of it. While Mary only has to think of how to optimize the message so that Peter understands (and hopefully accepts) its contents, in mass-communication, one size has to fit all. Moreover, while in face-to-face exchanges the roles of communicator and addressee typically alternate all the time, enabling constant fine-tuning and calibrating to ensure understanding, this is not the case in most varieties of mass-communication. Messages (which in this paper comprise anything ranging from single word utterances to entire discourses, the latter often involving information in different modes) are, as we saw in the marmalade example above, usually incomplete in the sense that the communicator counts on it that she does not need to spell out every detail, since she expects the envisaged audience to be able to retrieve pertinent details from its general or specific background knowledge. It is thanks to this that messages can often be short, and thereby spare the envisaged audience unnecessary mental effort in processing them. In specific situations, truncated messages such as "There!" "Three days ago", and "Alwin" are perfectly adequate. Messages, moreover, can consist of neutral, factual information but often carry emotional and/or attitudinal overtones that the envisaged audience must register for communication to be successful: it matters whether a "yes" is spoken in a happy, irritated, surprised, or hesitant way. Utterances such as "What a fantastic risotto this is!" and "You are a complete asshole!" do not convey information so much as emotions and attitudes – which is why Yus has proposed to relabel the notion of "cognitive effect" as "cognitive reward" (2011: 65), which includes "*affective effects*" (Yus 2023: 19, italics in original). And just to be sure, while communication usually works amazingly well, it can also go slightly or dramatically wrong: "Communication is governed by a less-than-perfect heuristic. On this approach, failures in communication are to be expected: what is mysterious and requires explanation is not failure but success" (Sperber and Wilson 1995: 45).

2.2 Attracting the envisaged audience's attention and conveying information to that audience

The key issue in communication is, of course, whether the communicator has succeeded in conveying the intended message to her envisaged audience. But actually this issue by definition becomes only pertinent after another condition has been satisfied: the awareness of the envisaged audience *that* the communicator wants to engage in communication. That is, the communicator first of all must, in one way or another, attract the attention of the envisaged audience and make clear that her message is aimed at this envisaged audience – whether consisting of one individual (say, "Peter") or of dozens, or millions of individuals. Two substages in the "attracting attention" part of the communicative process can be distinguished. First, the (aspiring) communicator needs to ensure that the envisaged audience becomes aware of the fact that she wants to engage in communication. This stage is taken for granted in classic RT's Mary-to-Peter utterances, but is of course a major concern in mass-communication. An advertiser, a TV presenter, a cartoonist, the writer of a scholarly article – all of them must first of all catch the attention of the envisaged audience. If the latter is not aware of the communicator's desire to convey something, communication breaks down before it has even properly started. Technically, the communicator needs the envisaged audience to *recognize* her intention to communicate. Only if the envisaged audience accepts the invitation to be addressed (technically: *fulfils* the communicative intention), the next stage comes into play: ensuring that the envisaged audience *understands* what the communicator wants to convey. If so, the communicator's intention to inform results in comprehension of the message by the envisaged audience (technically: the informative intention has been *recognized*). But understanding a message is not the same as accepting it as (probably) true or appropriate. After all, we may perfectly well understand something that we consider complete nonsense. Only when the informative intention is not just recognized but also *fulfilled*, we can say that communication has been successful. It is worth emphasizing that communication may break down at each of these stages, and for various reasons. Here is an example to illustrate the various stages. Imagine you are travelling by train, find a left-behind newspaper on one of the seats, and start flipping through the pages. Scenario 1: you accidentally turn two pages instead of one, and thereby miss the cartoon on page 3 (hence you do not recognize the cartoonist's intention to communicate something to you, as a prospective member of the envisaged

audience). Scenario 2: you access page 3, seeing from the corner of your eye there is a cartoon, but you don't bother to look at it (that is, you recognize, but do not fulfil the communicative intention of the cartoonist). If, by contrast, you look at it, you not only recognize but also fulfil the communicative intention (scenario 3). Now, you look at it, but unfortunately you don't understand the cartoon (scenario 4: communicative intention recognized and fulfilled; informative intention not recognized). If you understand the cartoon you recognize the informative intention (scenario 5). However, you judge that the cartoonist unfairly ridicules your favourite politician (scenario 6: informative intention recognized but not fulfilled). Only if you both understand and accept the critical and (hopefully) witty perspective on the political state of affairs depicted, the communicative process is successful (scenario 7: both the communicative and the informative intentions have been recognized and fulfilled). To be sure, while this completes the range of theoretical possibilities, it remains an issue for discussion what would be the precise criteria for deciding whether a given informative intention is fulfilled (for more discussion and examples of the recognition and fulfilment of the communicative and informative intentions, see Forceville 2020: 35–40).

2.3 The first interpretation of a message by the addressee is typically the right one

One of RT's most daring and insightful claims is that communication can be so fast and, usually, successful because the envisaged audience of a message stops interpreting it as soon as it has arrived at a satisfactory interpretation. So, the envisaged audience typically takes the first interpretation that comes to mind as the one that the communicator presumably intended. The rationale behind this is that if the envisaged addressee would have to consider a second, or third interpretation, that would decrease relevance because it requires extra mental effort. Of course, just as the presumption of optimal relevance may in reality disappointingly fail to be satisfied, the first interpretation of any message may actually *not* be the one that the communicator expected the addressee to derive. In such a case, communication went awry – and we would have to say retrospectively that the communicator did not choose an adequate way to formulate the message. But in everyday exchanges, the mechanism of the "first interpretation is the intended one" works well in by far most cases. That said, as we saw above, the creators of messages in certain genres (such as hermetic poems), expect their envisaged audience to be prepared to think beyond their first interpretation, spending extra mental effort to gain extra cognitive rewards.

2.4 Relevance comes with degrees of commitment to the communicated content

RT distinguishes between explicit and implicit meaning ("implicatures" and "explicatures" in Sperber and Wilson 1995: 182). For explicatures holds that the addressee fully understands the message, since all the pertinent information is "objectively" part of the message. "You will find the wine bottles in the cellar" is explicit – provided, of course, that the addressee knows English, and is familiar with the concepts "wine bottle" and "cellar" (and "finding" and "in" …). Even explicit messages, that is, can yield explicatures only on the basis of background knowledge that is shared between the communicator and the addressee.

But often, as we have seen in 2.1, information in a message is not in itself complete enough to enable the recipient to derive the "cognitive reward" that makes it relevant to him. Typically, that is, for the message to achieve relevance the addressee needs to combine it with additional information that is available to him – and that the communicator *expects* to be available to him. Such additional information is often part of the background knowledge that the communicator trusts the addressee to be able to recruit from his cognitive environment. Combining information in the specific spatio-temporal situation with presumably available background knowledge leads the envisaged addressee to derive one or more implicatures, which are to be inferred. For instance, someone may say to a talking person in a Dutch train, "Sorry, but this is the Silence compartment", assuming that the addressee knows that one is supposed to be silent in the train's Silence compartment – and thus infer the implicature, "I should not talk here". But the pertinent information that needs to be inferred can also be part of specific background knowledge shared between the communication partners. When my wife and I arrive

home from work, late and tired, she may look at me, and say, "I still owe you a dinner". In the given situation, I may infer from this the implicature that she suggests we go out for dinner in a restaurant rather than buy and cook food ourselves, since this is the first interpretation of "I still owe you a dinner" that makes the message relevant to me.

Now whereas in this example I will probably have little difficulty deriving the implicature that my wife intended me to, in other situations it may be less clear for me to decide on what, precisely, she wanted to convey. Imagine it is a Saturday morning in July, and my wife looks at me, saying, "The sun is shining". How am I to make this piece of explicit information relevant by inferring the implicature that she intends me to? A good guess would be that she suggests we go out to benefit from the good weather. But does she suggest we go walk in the park? Cycle to the beach for a swim? Put out on our terrace, at long last, the garden table and chairs? Go to the chemist's and buy sun-block? Or is this a hint that I go up to the attic and pull down the shutters to ensure the house will not become too hot? Her rather general utterance leaves room for a lot of interpretation on my part. Not each of these interpretations, to be sure, is equally plausible. If we had, in the preceding days, repeatedly said to each other that if it would be sunny on Saturday we would go and walk in the park/cycle to the beach/put out the garden set/buy sun-block/close the shutters in the attic, then this conversation on this particular topic would clearly be part of our jointly accessed, "mutually manifest" (Sperber and Wilson 1995: 42) background information – and therefore be a highly probable candidate to constitute the intended implicature. In such a situation, the implicature would be a "strong" one. But if we had not discussed any of these possible activities for a sunny Saturday before, it would be unclear which of these implicatures (or any of a host of others) my wife intended me to derive. Any of them would, thereby, be rather weakly implicated: what she wanted to convey is quite vague. If in the preceding days we repeatedly discussed the possibility of cycling to the beach, but none of the other actions or activities, my wife saying "it is sunny" strongly implicated that we should cycle to the beach. An appropriate response to that would be, for instance, "OK, I will unlock the bikes". But if in that situation I were to respond, somewhat angrily, "You always expect *me* to run up all the stairs to close the shutters in the attic!" my wife would presumably be highly surprised, and rightly so, and she might retort, "That's not at all what I meant!" RT would say that the "please close the shutters in the attic" implicature would be a very weak one in the situation at stake – if it should count as an implicature at all – and that I misinterpreted her.

These examples show that there is a continuum from highly explicit to highly implicit meaning, which teaches us a very important lesson: the more a message verges toward the explicit pole of the continuum, the more the responsibility for its derivation resides with the communicator. By contrast, the more it verges toward the implicit pole, the more that responsibility shifts to the envisaged audience. Of course, in everyday life, the question whether a given message contains strongly, weakly, or very weakly implicated meanings is often a bone of contention, and in view of who is responsible for their derivation the judgment which type of meaning is at stake can have serious consequences.

2.5 Intentional versus unintentional meaning

As we have seen in Section 2.4, meaning can range from more or less fully explicit to very weakly implicit. In the case of a weak implicature the addressee infers a meaning for which he holds the communicator, rightly or wrongly, at least somewhat responsible. But, we can ask, is there any kind of meaning that can be said to be not even very weakly implicated? That is, can there be completely *un*intended meaning? Surely the answer has to be affirmative. Something can be meaningful, even extremely meaningful, whereas it is crystal clear that the communicator did not want to share this meaning with the audience that interpreted it. Let us say you received a mail from a trusted colleague, telling about some measures the boss is going to take that will negatively affect the team you are a member of. You are not all amused and vent your discontent, vilifying the boss, in your response to the trusted colleague. Unfortunately, you failed to notice that your colleague sent her mail to the entire team, and instead of replying to her only, you accidentally ticked the "reply all" option, the "all" including the boss herself … The envisaged audience of your message consisted only of the trusted colleague who sent the original mail – and no one else. To all others, your message may be very insightful, and thus relevant, but you did not intend to be

relevant *to them* (particularly not the boss!). Another example of such unintended meaning would be a lawfully wedded woman, whispering to her husband, in the throes of love-making, not his name, but that of her secret lover.

In both these cases, any meaning inferred by the audience may be highly relevant, but it was not intended for them. In Forceville (2020) I propose the adjective "symptomatic" (borrowed from Bordwell 1989: 9) for meaning that is unintentionally conveyed. Possibly running counter to classic RT, I moreover propose that "unintentional/symptomatic communication" is a contradiction in terms, and that it therefore falls outside the provenance of what should count as communication *per se*. But of course, it is ultimately a moot point whether this is accepted or not, as long as "symptomatic meaning" is understood to be fundamentally different from meaning in the explicit-implicit spectrum. Two issues are crucial here. First, it is in practice often difficult or impossible to decide whether a given chunk of meaning should be considered weakly implicated or symptomatic. Indeed, communicators can try to fob off responsibility by dishonestly claiming that the addressee derived meaning that was wholly unintended by her. (Perhaps your "reply all" action was not so accidental, after all …?) Second, the study of inadvertent, unintended meaning is a highly pertinent scholarly pursuit. Researchers critically evaluating ideological issues are often keen to expose symptomatic meaning: they are interested in meanings that the communicator inadvertently "exuded" (Yus 2022: 23) or "leaked" (Yus 2022: 29).

3 Intermezzo: a response to Ntouvlis (2021)

Among the criticisms raised in reviews of Forceville (2020), those by Ntouvlis (2021) are specifically worth discussing here, because this helps clarify some crucial differences between RT and social semiotics approaches to multimodality. Ntouvlis' main objection pertains to what he sees as my commitment to "evolutionary reductionism":

> The cognitivist evolutionary tenets presented in Chapter 1 suggest an approach to human communication focusing on nature rather than nurture which is based on a highly selective review of interdisciplinary literature. […] the studies discussed stretch over an array of disciplines so wide that this often results in considerable loss of nuance. […] The discussion of communication in humans and primates highlighting some commonalities in the species' behavior (e.g., cooperation) and the mention of explanatory notions like "human universals" (p. 27) point to the theory's tendency towards ontological reductionism in Lachapelle's (2000) sense, whereby lower-level explanations (say, on the basis of posited evolutionary drives) are seen as preferable to higher-level ones (causation being viewed as a bottom-up process). […] This suggests an epistemological bias that comes at the expense of what Lachapelle (2000) calls "explanatory pluralism". In contrast with the book's approach, explanatory pluralism is agnostic as to whether higher- or lower-level explanations are generally preferable, accepting that different phenomena may necessitate different sorts of explanations. (Ntouvlis 2021: 330)

The list of authors discussed in my Chapter 1, including psycholinguists, primatologists, film scholars, evolution-oriented humanities researchers, and art theorists, was deliberately as wide-ranging as possible. The point I wanted to make was that, irrespective of the specific kind of ostensive behaviour people (as well as chimpanzees, and presumably other species, such as dolphins, Forceville 2025) display, they attempt to be optimally relevant to their envisaged audience in the sense explained in Section 2. If this is correct, this is hardly surprising, since the communicative principle of relevance is rooted in the cognitive principle of relevance – the latter the principle stating that all human (and presumably other) beings are predisposed to constantly monitor, and evaluate, their environment to ensure they can benefit from potentially good things in that environment and avoid potentially bad things in it. This is a key survival mechanism, and can therefore be explained by evolution principles. That said, I completely agree with Lachapelle's (2000) claim that bottom-up, biological evolution cannot suffice to explain how humans attribute meaning to acts of communication. Indeed, if "for explanatory pluralists, biological explanations and cultural explanations are seen as completing one another, not as conflicting" (Lachapelle 2000: 352), I am definitely an explanatory pluralist.

Ntouvlis further objects to my supposed "equating professional film critics to 'ordinary viewers' in [the] discussion of Bordwell's (1989) work on film interpretation (p. 14), thereby largely ignoring the wide divergence between these two posited categories of individuals in terms of their social learning histories underlying their

film-viewing practices" (Ntouvlis 2021: 330). This is not a fair criticism. Actually, I claim that Bordwell's (1989) central point is that "most of film viewing does not require any special skills" (Forceville 2020: 13), supporting this by the following quotation: "Filmic storytelling usually relies on our everyday assumptions about why people act as they do, how they will respond to others, and how they come to decisions" (Bordwell 2013: 30). I explicitly point out that "I will by and large bypass Bordwell's (1989) main concern, which is charting the processes and routines that film *critics* recruit in their interpretation of films, and instead focus on those adopted by ordinary viewers" (Forceville 2020: 14, italics in original) considering this justifiable inasmuch as,

> in Bordwell's view, the main difference between critics and ordinary viewers is that the former earn their bread interpreting films, which means that they have to work under certain institutional constraints, while the latter presumably watch films for pleasure. *For my purposes* [emphasis added, ChF], this difference is not very important. Bordwell would probably agree: "the critic takes as given, at least initially, a posited 'ordinary' viewing that makes referential sense of the film (identifies agents and settings, follows a story or argument). … The 'output' of these processes of comprehension creates, more or less tacitly, points of entry for interpretation' (1989: 132–133)". (Forceville 2020: 14)

In short, given my argument about how people make sense of films, and Bordwell's views on this, I considered it fully legitimate to focus on what film critics and ordinary film viewers have in common rather than on what sets them apart.

Here is a third objection Ntouvlis raises: he sees merit in the case study chapters in Forceville (2020), and acknowledges the analyses are sometimes informed by frameworks like Critical Discourse Analysis. However, "unfortunately, this more socially-oriented perspective is only brought up here and not incorporated later in the author's analyses. For example, despite the dynamic nature of genre being acknowledged, the case studies omit considerations of, say, how the processing of particular messages relies on genre-based cognitive schemata but also actively informs and 'updates' them" (Ntouvlis 2021: 330). My response would be that the brief reference to Critical Discourse Analysis is meant to remind the reader that, as I emphasize throughout the book, RT provides no less, but also no more than a model, offering a value-free template for analysing any given message. This template offers the starting point for further analysis and (aesthetically or ideologically) coloured interpretations. My reference to CDA is thus a demonstration of how scholars with specific research goals or agendas need to draw on other theories to complement RT. Elsewhere in my monograph I make the same point by referring in passing to Conceptual Metaphor theory and Blending Theory.

Ntouvlis' objections have something in common: he would have liked me to focus on differences (namely, between various disciplines, between film critics and regular film viewers, and between the highly specific socio-cultural contexts that need to be taken into account to do justice to each of the case studies analysed), whereas I draw attention to commonalities. Unlike much social semiotics-oriented research, my cognitivist-oriented interest is in charting opportunities to generalize, not to particularize. Inasmuch as the search for relevance has strong biological roots, it is universal – and therefore is a strong candidate for an inclusive theory of communication. That said, in my take on RT each act of communication is unique and unrepeatable. This uniqueness, captured in Sperber and Wilson's "relevance is always relevance to an individual" and my "relevance is always relevance to the envisaged audience", means that by definition numerous cultural factors need to be taken into account. "Nature" is almost always to be complemented by "nurture".

One issue Ntouvlis raises is of another kind. He rightly points out that "intentionality […] forms the basis for the book's very definition of communication, which excludes natural/'symptomatic' meaning. […] RT's ambitious scope might be seen as clashing with this intentionality-oriented account of communication, most notably in the case of art" (Ntouvlis 2021: 330–331). I realize my insistence on intentionality as essential in communication is controversial. So let me reiterate that while in my view, "unintentional communication" is a contradiction in terms, I hasten to repeat that studying communication in a broader social/cultural/political context (in the spirit of Van Dijk 2014) of course requires taking into account "symptomatic meaning" – a type of meaning that is central in much aesthetic and ideological criticism. But while I propose to keep "symptomatic meaning" outside of the provenance of communication-proper, I am not dogmatic about my position, and could accept that "symptomatic meaning" be considered part and parcel of "communication" – and then insist that the difference between

"weakly implicit" and "symptomatic" meaning is of a more fundamental nature than that between explicit/strongly implicit/weakly implicit meaning – and therefore deserves special attention.

On a final note in this Section: precisely since much art verges toward the "weak" pole of the spectrum, it allows more freedom of interpretation than other genres. Art, after all, is often creative in breaking the conventions of its genre, and may give rise to a host of idiosyncratic inferences by different individuals – some of them even beyond what the artist may have intended. Because of this I don't think that RT has much of interest to say about artistic genres (but see Pilkington 2000). In Forceville (2020) I have restricted myself to the "ninth art" of comics – which in its dominant variety is governed by narrativity. Since narrative is enormously steered and constrained by specific expectations of relevance, RT has something to say about it. That said, I think that in the end RT discussions of narrative can benefit much more from narratology (see e.g., Bal 2017; Forceville 2023) than vice versa.

4 Six case studies

In this Section I discuss six visual or multimodal mass-communicative messages from an RT perspective. My conviction that RT can develop into an inclusive theory of communication is rooted in the claim that RT is better equipped than any other communication model (including social semiotics, *pace* Ntouvlis 2021) I am aware of to do justice to the central importance of the following interrelated questions. Their general applicability demonstrates their pertinence to *all* communication, and provides insight into on the one hand the ideologically neutral aspects of a given message, and on the other hand into why and how a message can lead to different interpretations (or: misinterpretations) by different recipients:

(1) To what genre does the message belong, and how does this knowledge steer interpretation?
(2) Who/which agent communicates (and how do we know)?
(3) Who/which is the envisaged addressee/audience (and how do we know)?
(4) What background knowledge/values/spatio-temporal context does the communicator expect the envisaged audience to have access to?
(5) Is the message explicit? Implicit? Or should we consider the meaning to be symptomatic?

4.1 A pictogram

The five key RT questions can for Figure 1 be answered as follows:

(1) The various forms of this visual belong to the genre of signs that aid navigation in public buildings. The location of this (as of any) pictogram is crucial for its interpretation.
(2) The communicating agent is the authority in charge of the public building at stake, which needs to ensure optimal clarity to and safety of the people moving within it.
(3) All people moving in a building looking for the "exit" belong to the envisaged audience. In practice, these will mainly be people who are not (very) familiar with the building's ground plan.
(4) The communicator expects the envisaged audience to previously have learned the meaning of this coded sign.
(5) The sign presents fully explicit meaning.

Figure 1: Four pictograms signalling "here is the exit" (Google Images).

4.2 A shop name

The five key RT questions can for Figure 2 be answered as follows:

(1) The idiosyncratically stylized "Syriously" appears above a shop window, which is typically the place where names of shops appear. So, this word belongs to the genre: "shop name".

(2) The owners of the shop are the communicators: it is them that have an interest in, and the right to, placing the name of their shop above its front window.

(3) The name of a shop functions to attract attention to the wares sold in that shop, so the envisaged audience consists of (prospective) customers of Syriously.

(4) Regular customers would of course know what kinds of goods would be sold in the shop. Others would be expected to have enough knowledge of English to understand that the shop name is a pun merging "Syria" and "seriously". The audience would ideally also recognize that the style adopted for the shop name aims to resemble Arabic letters.

(5) The punning name of the shop can be decoded and is therefore explicit. The "Syrian/Arabic" connotations are strongly implicit – but only viewers recognizing the Arabic style will find it relevant. They will infer that the wares sold in the shop are somehow "Syrian" – and that the owners adopt a somehow "serious" attitude to their products. Note that one would either have to be already familiar with the shop or walk to the shop (or find it online) to realize that it sells Syrian take-away food.

4.3 A political cartoon

The five key RT questions can for Figure 3 be answered as follows:

(1) This four-part (narrative) series of pictures belongs to the genre of the political cartoon. The genre is governed by the audience's expectation that it offers a critical comment on a political state of affairs, and does so in a purportedly humorous or witty manner.

(2) The (Slovakian) Marian Kamensky is the cartoon's creator. We know this because his signature appears in the bottom right corner. Of course, to assess this we would have to be familiar with Kamensky's signature.

(3) Typically, regular readers of media in which Kamensky's cartoons appear constitute the envisaged audience. They would thus be familiar with this particular cartoonist's style. However, in this case the cartoon was never published, as newspapers found it too controversial (e-mail from Marian Kamensky to the author, 19 April 2025), but it circulated on internet, downloaded from the cartoonist's website. To be sure, any recontextualization (in a blog, or newsletter, or via X …) alters the identity of the cartoon's audience.

Figure 2: Name of shop, above shop window, at Botermarkt, Haarlem, The Netherlands (28 December 2023, photographed by the author).

Figure 3: Unpublished cartoon *Mutti* (by © Marian Kamensky) reproduced with permission of Marian Kamensky. Thanks to Martin Poduska and Maciej Grzenkowicz for alerting me to the cartoon.

(4) Kamensky expects the envisaged audience to recognize the (then) chancellor of Germany, Angela Merkel, and to be aware that Merkel said "Wir schaffen das" ("We can manage that") during a speech on 31 August 2015 about refugees wanting to come to Germany. That is, the envisaged audience is supposed to construe the metaphor "admitting (too) many refugees into Germany is feeding (too) many pigeons".

(5) The message is not explicit, as no code exists that specifies that "pigeon" is synonymous with "refugee". It is strongly implicit, as the metaphor must be construed to make the cartoon relevant. Minimally, the envisaged audience needs to infer the strong implicature that, supposedly, Merkel has underestimated the difficulty of accommodating large numbers of refugees in Germany. More weakly implicated are inferences such as "Merkel was unwise/foolish/over-optimistic …", "Merkel wants Germany to accommodate too many refugees", "If a country tries to accommodate too many refugees, they end up chasing the native inhabitants out of their own country". Inasmuch as these are weak implicatures, these inferences cannot be attributed to Kamensky; the responsibility for these inferences shifts to the addressee who derives them. Significantly, the ideological signature (left-wing, right-wing, sympathetic to immigration, opposed to immigration) of any re-publishing medium is bound to encourage some implicatures rather than others.

4.4 A comic strip

The five key RT questions can for Figure 4 be answered as follows:

(1) Figure 4 is an instalment of the genre "daily newspaper strip", specifically "The Wizard of ID", first published in 1964. By 2002 it appeared in more than 1,000 newspapers across the world (https://en.wikipedia.org/wiki/The_Wizard_of_Id, accessed 5 May 2025). Such strips are supposed to be light-heartedly humorous – so ideally the aimed-for "cognitive reward" is a smile.

(2) The communicators are the creators of the series, Brant Parker and Johnny Hart, which we know thanks to their names above the strip. Habitual reader-viewers would also recognize the distinctive style of the strip, and thus immediately know who are the communicators.

(3) The envisaged audience would typically consist of subscribers of the newspaper in which the strip appears – and more specifically of fans of the strip. It is such viewer-readers who, in the pre-digital era of 1977, would come across the strip.

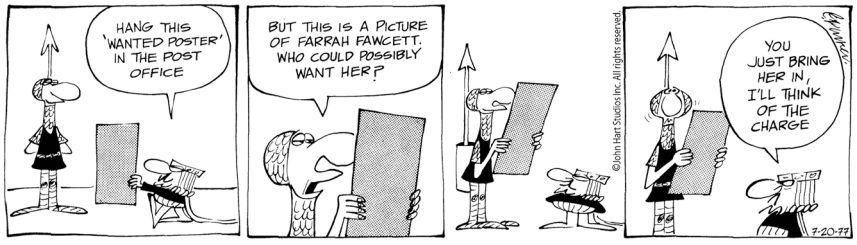

Figure 4: Instalment of *The Wizard of ID* (by Bryant Parker and Johnny Hart, first published 20 July 1977). Reprinted by permission of © John Hart Studios Inc.

(4) Regular readers of the strip would recognize the two characters, the King and Sir Rodney the Chicken-Hearted, and, having seen many previous instalments, would know that the King is a tyrannical man, and Sir Rodney cowardly, not overly intelligent, and eager to please his monarch. For this particular instalment, the envisaged audience would be expected to know about the convention to stick up posters with pictures of "wanted" criminals in public space. Moreover, the envisaged audience is supposed to be familiar with the actress Farrah Fawcett, who at the time was considered one of the sexiest women on earth.

(5) The language part can be decoded, and yields explicatures such as "the King wants Sir Rodney to stick up Wanted posters with Sara Fawcett's photograph so that she can be arrested". The strip achieves relevance only if the audience recruits the awareness that the verb "want" is used in two of its meanings: "sought-to-be-arrested" and "sexually desired" – the latter meaning explaining why the King has not yet considered the charges that should normally, of course, be clear before anybody can be legally "wanted". The double meaning is a strong implicature: without inferring both meanings, the strip is not funny, hence does not yield cognitive rewards, and is thus not relevant to a given addressee.

4.5 A print advertisement

The five key RT questions can for Figure 5 be answered as follows:

(1) The picture-plus-text is a commercial advertisement. Reader-viewers thus know that the communicator wants to promote something, in this case a conference, and realize that participating in conferences usually costs money.

(2) The name of the conference *ECOM21* and the logos of sponsors, Decta, KPMG, and Mastercard make clear who is/are the communicator(s).

(3) The envisaged audience consists of potential attendees of this "Baltics' Top Fintech Gathering". The organizers think that advertising a business conference in the capital of Latvia in an inflight magazine of Baltic airlines will reach many people of the target audience.

(4) Background knowledge and contextual circumstances the communicator expects the envisaged audience to be able to recruit include that business people may be interested in finance-related conferences, that they typically fly, that they may have time during the flight to leisurely consider the ad, and that these are people who can afford the fee (which turns out to start at €199,-).

(5) The verbal information in the ad is largely coded, and yields explicit information concerning details of the conference, its venue, participants, countries represented, partners and speakers, a QR code, and a discount code. The logos of sponsors also constitute explicatures. The photograph shows musicians performing on a stage during *ECOM21* in 2023, and may yield implicatures such as: "ECOM21 mixes business with pleasure" and "ECOM21 finds artistic performances important". These are relatively weak implicatures, and their derivation, unlike that of the written information and the logos, is thus largely the responsibility of the audience. The typo of "countries" in "30+ countires" [sic] is clearly unintended, and thereby constitutes symptomatic meaning. If found relevant it is so thanks to the Cognitive, not the Communicative Principle of Relevance (leading, possibly, to the inference, "If the organizers are so careless as to miss a typo in the ad they commissioned, their conference may be similarly sloppy – so I won't go there").

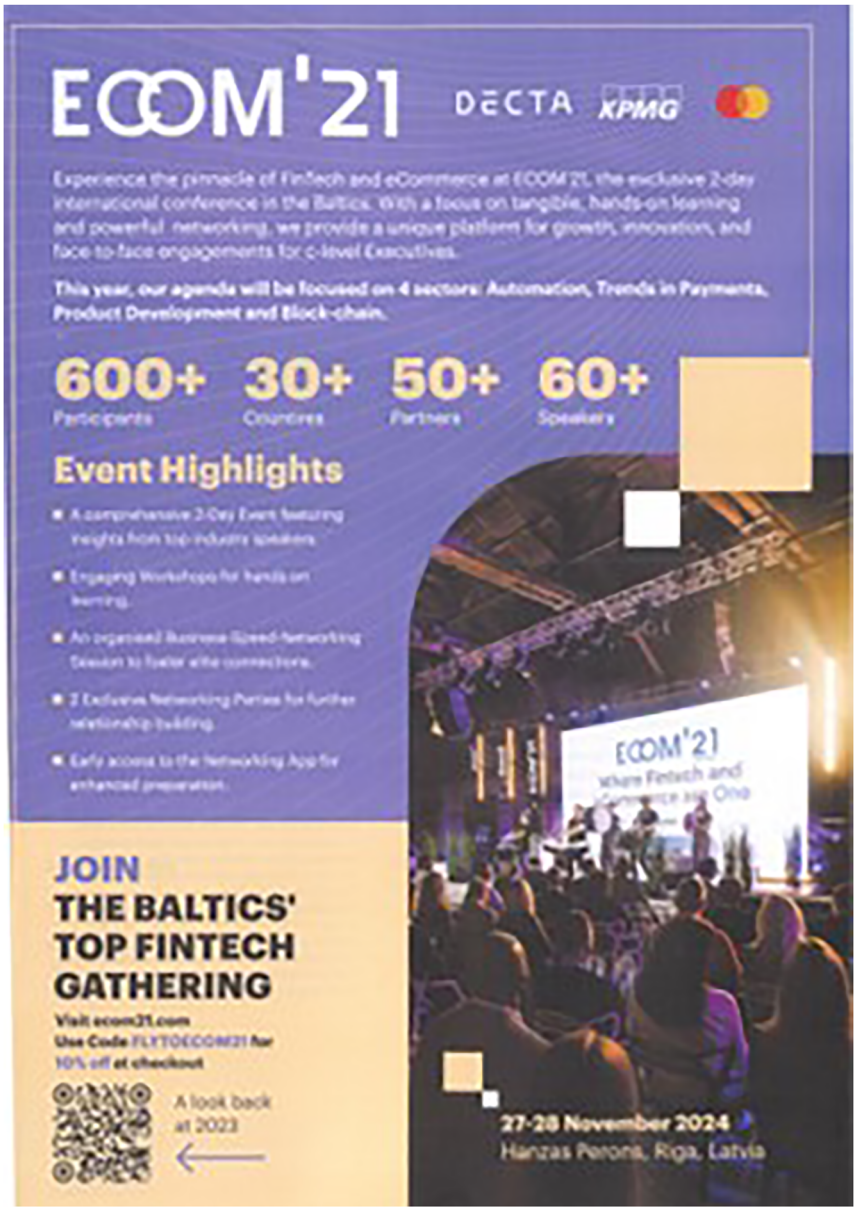

Figure 5: Advertisement *ECOM21 Conference*, 27–28 November 2024, Riga, Latvia (October 2024, Baltic Outlook, Inflight Magazine).

4.6 A Zoom meeting

The five key RT questions can for Figure 6 be answered as follows:

(1) Figure 6 is a screenshot that is part of the genre "recorded Zoom-meeting video".

(2) In Zoom meetings, which became hugely popular during COVID, participants typically take turns being communicators and addressees.

(3) See (2).

(4) Shared background knowledge importantly pertains to awareness how this medium works, such as that participants sit in front of their screens, mute themselves when not talking, and switch off their cameras when they leave the meeting.

(5) As Yus (2011) proposes, cognitive effects may also be non-propositional, and for instance include (non-verbally conveyed) emotions, evaluations, and attitudes. Dressing and behaving professionally is thus arguably part of the "ostensive stimulus" of communication in a Zoom meeting. Clearly the man in the top

Figure 6: Screenshot of just-finished Zoom meeting (https://www.youtube.com/watch?v=rOWGe7uOuPU, accessed 5 May 2025).

left screen forgot to switch off his camera immediately after the meeting ended, and now reveals he was wearing shorts, and indulges in uninhibited bodily behaviour. Obviously, he did by no means intend to communicate any of this – and thereby this is an instance of symptomatic meaning.

Of course, much more can be said about each of the case studies. SFL and semiotics-oriented scholars might want to engage critically with issues such as that the stylized human person in Figure 1 is a man and not a woman. Perhaps they would want to explore the ideological implications of the metaphor "refugees are pigeons" in Figure 3 in more detail (is this bird's inclination to shit all over the place one of the features supposedly to be metaphorically mapped onto "refugees" …?). Or maybe they would like to discuss the sexism of the King in Figure 4. All these are legitimate, pertinent questions. But they all arise *after* assessing the communicative situation in the cool, "objective" manner that RT does, and I hope to have convinced the envisaged audience of this paper that, and how, RT offers the template that scholars interested in addressing ideological and ethical dimensions of multimodal discourse can (and should) build on.

5 Concluding remarks

RT's key qualities as candidate for an inclusive theory of communication comprise the following:

(1) RT acknowledges that "meaning" can never be discussed without taking into account its origin (= identity of the communicator) and its destination (= identity of the envisaged audience): meaning is always meaning for someone, or for a more or less specific group of people.

(2) RT distinguishes between coded meaning and non-coded meaning. Coded meaning can be interpreted by anybody in possession of the code; non-coded meaning can be inferred by anybody able to combine the ostensive stimulus with appropriate background knowledge and/or information in the situational, spatio-temporal context.

(3) RT distinguishes between explicatures, strong implicatures, weak implicatures, and symptomatic meaning, and thereby is able (a) to show when the responsibility for deriving interpretations of specific messages is primarily to be attributed to the communicator and when to the envisaged addressee/audience; (b) to account for diverging interpretations by different (groups of) individuals; (c) to help understand why and how miscommunication arises. The topic of symptomatic meaning – whether or not categorized as part of communication-in-the-narrow-sense – and how it can be distinguished from weakly implicated meaning, is one that deserves extensive further discussion.

(4) RT does not replace or cannibalize other approaches to meaning-making. Indeed, on the one hand it can accommodate the best insights from other approaches (social semiotics, SFL, Conceptual Metaphor Theory, CDA, Blending Theory, rhetoric, stylistics, narratology …) On the other hand, it *needs* these very insights to

> put flesh on the bones of what is in the last resort no less, but also no more, than an excellent communication *model.*

One important consequence of accepting RT is that the recontextualization (reprinting, forwarding, re-X-ing …) of messages often seriously affects whether, and how, the intended meaning is understood by the audience, since often that audience is no longer the original "envisaged audience". Clearly recontextualizing any of the varieties in Figure 1 by positioning it above the door to the toilets makes it misleading/wrong/false, while the same holds for repositioning the name in Figure 2 over another shop's window. The weak implicatures derived from the cartoon in Figure 3 will undoubtedly differ when read by subscribers of a right-wing versus a left-wing publication, respectively. *The Wizard of ID*'s humour in Figure 4 is presumably general enough to be widely understandable in the same way, as transpires from the fact that by 2002 the strip appeared worldwide – but being familiar with Farah Fawcett is indispensable background knowledge. The fact that the ad in Figure 5 is typically accessed by people flying with Baltic Air favours its envisaged audience, although the ad is sufficiently context-independent to communicate its meaning to everybody who knows English. The audience of the unintentionally conveyed information in Figure 6 consisted originally of the participants of the Zoom meeting: by being recontextualized as part of a YouTube video consisting of blunders during Zoom meetings, the genre changes, and thereby the envisaged audience – and the cognitive reward shifts from "wow, how stupid/naïve/ridiculous our colleague is!" to "what a funny video this is!"

Clearly, developing RT into an inclusive theory of communication is still in its early stages. Further work requires analyses focusing on different media (e.g. film, games, digital platforms …), modes (at present the most often researched multimodality ensemble is, as in this paper, the combination of static visuals with short written texts), and genres (where the sky is the limit).

Acknowledgments: I am indebted to two anonymous reviewers for their insightful comments, many of which I have tried to accommodate. I thank Sabine Heinemann for editorial suggestions. Of course, I alone remain responsible for all contents and interpretations.

References

Adami, Elisabetta. 2017. Multimodality. In Ofelia García, Nelson Flores & Massimiliano Spotti (eds.), *The Oxford handbook of language and society*, 451–472. Oxford: Oxford University Press.

Archer, Arlene & Esther Breuer (eds.). 2015. *Multimodality in writing: The state of the art in theory, methodology and pedagogy*. Leiden: Brill.

Bal, Mieke. 2017. *Narratology*, 4th edn. Toronto: University of Toronto Press.

Baldry, Anthony & Paul J. Thibault. 2006. *Multimodal transcription and text analysis: A multimedia toolkit and coursebook*. London: Equinox.

Bateman, John, Janina Wildfeuer & Tuomo Hiippala. 2017. *Multimodality: Foundations, research and analysis – a problem-oriented introduction*. Berlin: De Gruyter.

Bordwell, David. 1989. *Making meaning: Inference and rhetoric in the interpretation of cinema*. Cambridge: Harvard University Press.

Bordwell, David. 2013. The viewer's share: Models of mind in explaining film. In Arthur P. Shimamura (ed.), *Psychocinematics: Exploring cognition at the movies*, 29–52. Oxford: Oxford University Press.

Clark, Billy. 2013. *Relevance theory*. Cambridge: Cambridge University Press.

Djonov, Emilia & Sumin Zhao (eds.). 2013. *Critical multimodal studies of popular culture*. New York: Routledge.

Forceville, Charles. 1996. *Pictorial metaphor in advertising*. London: Routledge.

Forceville, Charles. 2005. Addressing an audience: Time, place, and genre in Peter van Straaten's calendar cartoons. *Humor: International Journal of Humor Research* 18(3). 247–278.

Forceville, Charles. 2014. Relevance theory as model for analysing multimodal communication. In David Machin (ed.), *Visual communication*, 51–70. Berlin: De Gruyter.

Forceville, Charles. 2020. *Visual and multimodal communication: Applying the relevance principle*. Oxford: Oxford University Press.

Forceville, Charles. 2021. Multimodality. In Xu Wen & John R. Taylor (eds.), *The Routledge handbook of cognitive linguistics*, 676–687. London: Routledge.

Forceville, Charles. 2022. Visual and multimodal communication across cultures. In Istvan Kecskés (ed.), *The Cambridge handbook of intercultural pragmatics*, 527–555. Cambridge: Cambridge University Press.

Forceville, Charles. 2023. Narrating and focalizing visually and verbal-visually in comics and graphic novels. *Pragmatics & Cognition* 30(1). 183–211.
Forceville, Charles. 2025. Book review of Denise L. Herzing, *Is anyone listening? What animals are saying to each other and to us* (The University of Chicago Press 2024). *Leonardo* 58(3). https://leonardo.info/review/2025/03/is-anyone-listening-what-animals-are-saying-to-each-other-and-to-us (accessed 17 July 2025).
Forceville, Charles & Billy Clark. 2014. Can pictures have explicatures? *Linguagem em (Dis)curso* 14(3). 451–472.
Forceville, Charles & Natalia Sánchez Querubín. 2022. Relevance theory perspectives on web-based communication. In Herbert L. Colston, Teenie Matlock & Gerard J. Steen (eds.), *Dynamism in metaphor and beyond*, 325–340. Amsterdam: Benjamins.
Grice, H. Paul. 1975. Logic and conversation. In Peter Cole & Jerry L. Morgan (eds.), *Syntax and semantics.* Vol. 3: *Speech acts*, 41–58. New York: Academic Press.
Jewitt, Carey (ed.). 2014. *The Routledge handbook of multimodal analysis*, 2nd edn. London: Routledge.
Jewitt, Carey, Jeff Bezemer & Kay O' Halloran. 2016. *Introducing multimodality*. London: Routledge.
Klug, Nina-Maria & Hartmut Stöckl (eds.). 2016. *Handbuch Sprache im multimodalen Kontext*. Berlin: De Gruyter.
Kress, Gunther. 2010. *Multimodality: A social semiotic approach to contemporary communication*. London: Routledge.
Kress, Gunther & Theo Van Leeuwen. 2001. *Multimodal discourse*. London: Arnold.
Kress, Gunther & Theo Van Leeuwen. 2021. *Reading images: The grammar of visual design*, 3rd edn. London: Routledge.
Lachapelle, Jean. 2000. Cultural evolution, reductionism in the social sciences, and explanatory pluralism. *Philosophy of the Social Sciences* 30(3). 331–361.
Norris, Sigrid (ed.). 2012. *Multimodality in practice: Investigating theory-in-practice through methodology*. London: Routledge.
Ntouvlis, Vinicio. 2021. Review of Forceville (2020). *Multimodal Communication* 10(3). 329–331.
Pilkington, Adrian. 2000. *Poetic effects*. Amsterdam: Benjamins.
Royce, Terry & Wendy L. Bowcher (eds.). 2007. *New directions in the analysis of multimodal discourse*. Mahwah: Lawrence Erlbaum.
Serafini, Frank. 2022. *Beyond the visual: An introduction to researching multimodal phenomena*. New York: Teachers College Press.
Sperber, Dan & Deirdre Wilson. 1995. *Relevance: Communication and cognition*, 2nd edn. Oxford: Blackwell.
Van Dijk, Teun. 2014. *Discourse and knowledge: A sociocognitive approach*. Cambridge: Cambridge University Press.
Van Leeuwen, Theo & Carey Jewitt (eds.). 2001. *Handbook of visual analysis*. London: Sage.
Wildfeuer, Janina, Jana Pflaeging, John A. Bateman, Chiao-I Tseng & Ognyan Seizov (eds.). 2019. *Multimodality: Disciplinary thoughts and the challenge of diversity*. Berlin: De Gruyter.
Wilson, Deirdre & Dan Sperber. 2012. *Meaning and relevance*. Cambridge: Cambridge University Press.
Yus, Francisco. 2011. *Cyberpragmatics: Internet-mediated communication in context*. Amsterdam: Benjamins.
Yus, Francisco. 2022. *Smartphone communication: Interactions in the app ecosystem*. London: Routledge.
Yus, Francisco. 2023. Finding relevance in smartphone advertising. *Internet Pragmatics* 6(1). 1–41.