


ARTICLE

Deception as Mimicry

Marc Artiga¹  and Cédric Paternotte²

¹Universitat de València, Department Filosofia, València, Spain and ²SND Research Team, Sorbonne Université, Paris, France

Corresponding author: Marc Artiga; Email: marc.artiga@uv.es

(Received 20 February 2023; revised 19 October 2023; accepted 26 October 2023; first published online 15 November 2023)

Abstract

What is the nature of deception? What does it take for an organism to deceive another one? In this article we address these questions by appealing to the concept of mimicry. More precisely, we argue that a fruitful perspective regards deception as an instance of mimicry rather than the reverse. Conceiving of deception as an instance of mimicry has a number of interesting consequences: It draws connections between different areas of research, vindicates a functional approach to deception by providing a satisfactory answer to some recent objections, and suggests some worries for game-theoretic approaches to deception.

1. Introduction

Philosophical attempts to define deception can be generally classified into two groups. Some accounts focus on human deception and tend to provide an intellectualized analysis that cannot be straightforwardly applied to other living organisms (Carson 2010; Mahon 2007). In contrast, others seek a more general approach that could also accommodate nonhuman deception (Birch 2019; Fallis and Lewis 2019; Krstić and Saville 2019; Martínez 2019; McWhirter 2016; Shea et al. 2018; Skyrms 2010). The latter tend to develop game-theoretic models, which focus on deceptive signals and have fueled lively recent debates regarding the correct definition of biological deception. Following this tradition, most of the present article focuses on the phenomenon of deception as it appears across the natural world—in humans and nonhuman animals, even perhaps in plants or bacteria.¹

Our main goal is to argue that deception can be regarded as an instance of mimicry, which does not necessarily involve signals. In the natural world mimicry and deception are often intermingled, although their exact relation is a matter of dispute (Dalziell and Welbergen 2016; Quicke 2017; Font 2019). Indeed, both phenomena are difficult to capture conceptually. In this article, we provide a broad

¹ We will concentrate on the latter game-theoretic tradition, especially in section 4, even if it will gradually become clear that our account is intended to span all varieties of deception.

analysis of mimicry, not to deliver the best definition of this concept for biological practice on mimicry, but to highlight important similarities between these two phenomena. In particular, we suggest deception can be understood as an instance of mimicry rather than the reverse. We will argue that the explanatory scope of such an account goes beyond that of extant theories of deception—it encompasses a wider array of deceptive phenomena.

We will show that conceiving of deception as an instance of mimicry has significant consequences. It allows one to draw fruitful connections between different areas and suggests new questions for future research. In particular, our approach vindicates a relatively broad analysis of mimicry and connects with recent functionalist theories of deception and representation. Furthermore, it provides a response to recent objections raised by Birch (2019) and Fallis and Lewis (2021) against a functional theory of deception (Artiga and Paternotte 2018; Krstić 2021) and thus refines and extends this theory. Additionally, it suggests that extant game-theoretic approaches fail to provide an adequate analysis of deception and explains why. These works have proven extremely fruitful over the years, with a wealth of theoretical examples that have provided heuristic guidelines or constraints for possible definitions of deception. However, the standard way of framing the problem using game theory also tends to suggest the wrong kind of answers regarding the definition of deception, partly because of their reliance on signaling games and of their emphasis on the concept of deceptive signal—or so we argue.

In section 2, we motivate and present our account of mimicry. Section 3 argues that deception can be conceived as an instance of mimicry and identifies this approach as a version of a functional theory of deception, which as a result also includes the case of intentional deception. Section 4 considers game-theoretic approaches to deception and, in the light of our account, explains why they tend to provide the wrong kind of answers. Section 5 discusses some objections to, and consequences of, our analysis.

2. Mimicry

The use of “mimicry” to refer to a specific kind of functional resemblance between organisms dates from the middle of the nineteenth century. Henry Walter Bates suggested that individuals of one species may gain an advantage by mimicking members of another species only three years after the publication of Darwin’s *On the Origin of Species* (Forbes 2011). Nowadays, different proposals exist as to how to define this phenomenon. Current disputes on mimicry concern not only how to define it but also which distinctions should be made between them or whether phenomena such as camouflage or masquerade should be included in the analysis (Ruxton et al. 2018; Quicke 2017). To avoid being overwhelmed by these debates, let us begin by considering uncontroversial cases before we move on to more difficult ones.

2.1. Types of mimicry

We are all familiar with common cases of mimicry, such as stick insects or inoffensive droneflies copying the red and yellow stripes of the honeybee to avoid being preyed upon (Brower and Brower 1962). A general definition of mimicry, however, is hard to come by. As a first approximation, mimicry usually (perhaps always) involves a model

and a mimic. The model is the entity being mimicked, whereas the mimic is the organism that changes so as to develop some sort of resemblance with the model. For this resemblance to be classified as an instance of mimicry, however, it has to serve certain purposes. Let us consider some illustrative examples.

A first paradigmatic category is “Batesian mimicry,” in which a harmless organism deters predation by imitating a harmful model. A classic example is provided by nonvenomous snakes that have evolved the characteristic tri-color ringed patterns of the venomous coral snake (Greene and McDiarmid 1981). Other related phenomena involve rewardless flowers imitating the color or shape of rewarding flowers to attract pollinators (Dafni 1984)—reproductive mimicry—and the *Photuris* female firefly imitating the light pattern of *Photinus* female fireflies to lure males and devour them (Lloyd 1975)—aggressive mimicry. More generally, a key aspect of these kinds of mimicry is that the mimic somehow imitates the model and enjoys a benefit, whether it concerns protection, predation, or reproduction, without paying the cost.²

By contrast, as traditionally understood, in Müllerian mimicry both mimic and model benefit from the resemblance. Monarch and Viceroy butterflies provide the best-known example: Their mutual resemblance facilitates predator learning. As both butterflies taste unpleasant, any predator that eats one of them will quickly learn and come to avoid both species (by failing to distinguish them). Müllerian mimicry is often presented as involving only two species, but it often includes a larger number of species that form “mimicry rings” (Mallet and Gilbert 1995).

Batesian, Müllerian, reproductive and aggressive mimicry represent some paradigmatic categories, but many other phenomena are sometimes classified as mimetic: background matching, countershading, dazzling, and so forth (Merilaita et al. 2017; Ruxton et al. 2018). To simplify our discussion, we will focus on only two of these. On the one hand, background matching³ involves a set of strategies that avoid detection by resembling the background against which the organism is observed. On the other hand, masquerade consists in resembling an inanimate object, typically to avoid being preyed on. Peppered moths (*Biston betularia*) and flat-tailed geckos (*Uroplatus fimbriatus*) are well-known examples of the former case and stick insects (order *Phasmatodea*) and the common sea dragon (*Phyllopteryx taeniolatus*) (which resemble branches and seaweed, respectively) of the latter.

A crucial distinction between these different phenomena is that they tend to involve different evolutionary dynamics. Batesian and aggressive mimicry, for example, typically exemplify negative frequency-dependent selection: Mimics benefit from resembling the model without paying a cost, whereas the existence of parasitic mimics has some cost for the models (their mimicked features risk becoming less efficient, e.g., in deterring predators). If mimics are frequent enough, models might thus tend to change their appearance, while the mimics’ interest is to keep this similarity. By contrast, Müllerian mimicry involves positive frequency-dependent selection, as both species usually benefit from sharing an advertisement display (Roy

² Various classifications of mimicry exist. From the Batesian/Müllerian distinction, many theorists moved to cost/benefit criteria (e.g., Vane-Wright 1976). For a comparative discussion, see e.g., Quicke (2017, ch. 1).

³ What we call “background matching” is often labeled as “camouflage” or “crypsis” in the specialized literature.

and Widmer 1999). Other phenomena like background matching and masquerade typically involve resemblance to nonliving objects, so neither of these frequency-dependence evolutionary models apply.

Still, not all strategies for avoiding attack and damage fall under the category of mimicry. For instance, the purpose of disruptive coloration is to avoid or minimize damage by presenting some color or shape that hides or hinders the correct perception of the organism's shape or edges. Similarly, it is controversial whether organisms that avoid detection by being transparent (e.g., the brush-footed butterfly *Greta oto*) should qualify as mimics.

2.2. The definition of mimicry

Developing a satisfactory concept of mimicry is a complex enterprise. As suggested in the preceding text, most definitions are controversial for various reasons, including the fact that theorists disagree as to whether masquerade or camouflage should be included. The goal of this article, however, is not to provide a general definition of mimicry that is useful for biologists across the board. Rather, we need a concept of mimicry that may play a certain theoretical role, namely, help illuminate the nature of deception.

Which analysis of mimicry can best serve our theoretical purposes? Batesian, Müllerian, aggressive, and reproductive mimicry are central cases of deception, which the analysis should certainly include. However, because we aim for a broad definition of mimicry that enables us to draw fruitful connections with deception, and because background matching and masquerade include paradigmatic instances of deception such as geckos, stick insects, and the common sea dragon, we consider these should be included as well. As a first approximation our analysis of mimicry should encompass all these phenomena.

Having settled on this minimal set of desiderata, we can draw several lessons from some accounts of mimicry found in the biological literature. Consider, for instance, one of the classic definitions:

Mimicry is defined in terms of a system of three *living* organisms, model, mimic and operator (*signal-receiver*), in which the mimic gains in fitness by the operator *identifying* it with the model. (Vane-Wright 1980, 1; emphasis added)

This analysis faces various difficulties. First, it assumes mimicry involves signals, which is not necessarily so. For example, orchids in the genus *Ophrys* display a resemblance with female bees that attract males that attempt to copulate with them and end up spreading the plant's pollen (Schiestl 2005; Anders Nilsson 1992). Consequently, mimicry need not involve signals, unless one assumes an extremely liberal notion of signal, according to which almost anything counts as such (and which would confuse cues with signals). Likewise, some organisms mimic nonliving entities (e.g., rocks, sand, twigs).

Secondly, requiring the *identification* of mimic and model implies a relatively sophisticated capacity on the part of the receiver. Bolas spiders attract moths by mimicking female pheromones, but whether something like identification takes place in the moth's brain is unclear. Moreover, the focus on identification could unduly exclude possibly deceptive phenomena, for instance background matching, which

involve a “failure of awareness” rather than “mistaken identification” (Font 2019; the terms are taken from Vane-Wright 1980).

More recent accounts are also problematic for these or other reasons. Consider, for instance, Dalziell and Welbergen’s proposal (2016, 612; emphasis added):

[In mimicry] a receiver *perceives the similarity* between a mimic and a model and as a result changes its behaviour in a manner that provides a selective advantage to the mimic.

Here, mimicry is analysed in terms of perceiving similarity, which is a relatively sophisticated capacity: to perceive a similarity between A and B implies more than merely perceiving A and B. It requires representing that A and B are similar in certain respects and might even require the possession of some concepts. We do not think pollinators that are misled by rewardless mimics, or moths by pheromones released by bolas spiders, need to possess the capacity to distinguish between similar and different objects. Similarity can lead a receiver to error without being represented as such by the receiver.⁴

Still, mimicry does affect some of the receiver’s cognitive states (broadly understood), as it works by altering the receiver’s capacity for discriminating, identifying, or recognizing organisms (Ruxton et al. 2018). Orchids mislead pollinators such as bees and stick insects mislead predators. Thus we would at least say that some minimal form of perceptual, cognitive, or representational capacity is required on the part of the receiver.

Let us now present a working definition of mimicry that will allow for an account of deception that does not necessitate signaling or some cognitively demanding processing involving conceptual representations. As we have seen, mimicry involves a partial copy of a model (typically, an organism or an object) that need not be understood in terms of signaling. But if mimicry is not a matter of communicating a particular content, what function does it perform? How does an organism benefit from resembling another entity? As seen in the preceding text, it seems hard to deny that the function of mimicry is to influence the receiver in certain ways. More precisely, its purpose is to produce a certain representational state in the receiver⁵; this is why current definitions use expressions like “discriminating,” “identifying,” or “recognizing.” But which representational state is supposed to elicit mimicry? We suggest that resemblance is useful because it tends to produce in the receiver *the same representational state that the model would have generated*. Accordingly, we suggest this is the function of mimicry.

Putting all this together, we may define mimicry as follows:

(MIM) An organism’s trait X mimics a model’s trait Y iff X has the function of causing in the receiver the same representational state as Y by means of being a partial copy of the model.

⁴ Of course, more than two definitions exist; for a review, see Quicke (2017, ch. 1).

⁵ Even if mimicry and deception need not involve signals, for simplicity we sometimes use the standard expression “receiver” of mimesis. Similarly, we will assume a broad understanding of the notion of “representational state.” For a discussion, see section 5.

We think this analysis has all the properties a satisfactory definition should include⁶ and can accommodate the cases mentioned already:

The dronefly's color pattern mimics the honeybee's color pattern because the former has the function of causing in the receiver the same representational state as the latter by means of being a partial copy of honeybees (Batesian mimicry).

The viceroy's wing pigmentation mimics the monarch's wing pigmentation because the former has the function of causing in the receiver the same representational state as the latter by means of being a partial copy of monarchs (Müllerian mimicry).

Certainly, among all the suggested definitions of mimicry, this is a relatively broad one because it includes background matching and masquerade as well as Batesian and Müllerian mimicry. Nonetheless, it is not too broad: It still excludes some mechanisms for avoiding predators, such as disruptive coloration, which do not mimic any model.

The definition relies on the notion of a *partial copy*. This is a purposely vague concept; whereas by definition mimics need to resemble the model in various ways, this copy will always be partial in an objective sense: An identical copy of a model—which has all and only the model's properties, including its DNA, and so forth—would just be another instance of the model (e.g. another bee, rather than a mimic of a bee). The mimic is always an imperfect copy. Of course, how much it needs to resemble the model and which aspects are worth copying will depend on each case (Ruxton et al. 2004, 159–61). This is influenced, for instance, by limited receiver discrimination, evolutionary lags, or cost-benefit constraints (Kikuchi and Pfennig 2013): It may be enough that the predator hesitates, identify the target a bit more slowly, or is deceived because it cursorily checks its environment—at high speed and/or with little attention to detail. For instance, the tail of the spider-tailed viper (*Pseudocerastes urarachnoides*) is only broadly similar to a spider: It would not bear close scrutiny but suffices to lure birds. It may also stem from the simultaneous mimicking of multiple models, or from cases of exaggerated or “supernormal” traits, such as for the *Ophrys* orchid (Dalziell and Wenbergen 2016, box 2). Relatedly, note that the interesting sense of a copy is observer-relative: What matters is that the copy appears similar to the model relative to the observer's perceptual abilities. As Font (2019, 13) suggests, mimicry is “in the eye of the beholder.”

Finally, some cases, such as transparency, are not clearly included or excluded (Gomez et al. 2021; Barnett et al. 2020). Transparency does not produce a copy of the background; but the transparent producer does create the same representational state as its surroundings would, by virtue of a partial resemblance (often, only part of the body is transparent). The issue hinges on the weight one puts on the distinction between being misinterpreted and not being seen (as hiding is not a case of mimicry), which we need not decide here.

Having provided a preliminary definition of mimicry, our next question is: What distinguishes *deceptive* mimicry from other forms of mimicry?

⁶ Our definition does not explicitly refer to fitness, let alone to fitness gains or benefit gains or benefits, which also contrasts with some of the classic definitions (see section 5).

3. Deception

3.1. A mimicry-based definition

What is the relationship between mimicry and deception? The original proposal we would like to explore is that deception is a specific kind of mimicry. This may sound surprising and even counterintuitive, as mimicry is often considered as a kind of deception rather than the reverse, although seldom for conceptual reasons. However, even theorists of mimicry do not agree on this point. For instance, whereas Font (2019: 11) argues that “deception (understood as perceptual error) is a hallmark of mimetic relationships,” Jamie (2017) emphasizes the existence of “non-deceptive mimicry,” or “rewarding mimicry,” which benefits the models as well as the mimics (and so includes some cases of Müllerian mimicry). Moreover, insofar as mimicry, which may target signals as well as cues, seems more widespread than signal-based deception, a unified analysis that includes the latter within the former rather than the reverse has better chances of success.

If deception is a form of mimicry, what makes it distinctive from other forms? We suggest that deception is supposed to cause a representational state *with false content*. As a result, we can modify the previous description of mimicry as follows:

(DEC) An organism’s trait X deceptively mimics a model’s trait Y iff X has the function of causing in the receiver the same representational state as Y, but with false content, by means of being a partial copy of the model.

For example:

The stick insect’s color and shape deceptively mimic a branch color and shape because the former have the function of causing in the receiver the same representational state as the latter, but with false content, by means of being a partial copy of a branch.

First, as stated in the definition, deceptive mimicry only necessitates a *partial* copy of the model. As we argued, mimicry is always imperfect. The definition also covers cases in which deceptive mimicry has a double function, for example both hiding from prey and escaping predators (Aaronson 1983).⁷

Secondly, note that this definition of deception does not require deceptive traits to be signals (for more on that, see the text that follows). Nevertheless, because signals are a specific kind of trait, this approach can easily include them.⁸ Consider the following example:

Photuris light patterns deceptively mimic *Photinus* firefly light pattern because the former has the function of causing in the receiver the same representational state as the latter, but with false content, by means of being a partial copy of a *Photinus* firefly. (See Artiga 2014.)

⁷ Also note that being a partial copy is necessary, but not sufficient for mimicking or deceiving.

⁸ Note that this goes against proposals such as Dalziell and Wenbergen’s (2016), which aim to bring back signaling theory as fundamental for mimicry. By contrast, we see nonsignal-based mimicry as the basis for deception.

Interestingly, this approach not only correctly predicts which traits should be classified as deceptive but it also rightly excludes some mimics as not engaging in deception. For instance, Müllerian mimicry is typically not an instance of deception because its function is not to produce a false representational content, but more something like hiding the truth: Given the similarity between monarch and viceroy butterflies, an organism that perceives one of them is not in position to distinguish whether it is in front of a monarch or a viceroy. This is not revealing the whole truth, rather than producing a false representation and that seems to be the right result.

Background matching is a more difficult case, but the explanation for this requires some stage setting. First of all, note that in DEC the expression “causing a state with false representational content” can be understood in various ways (Chisholm and Feehan 1977). For instance, one could interpret it as causing the receiver to *acquire* a false representational state or as causing it to *maintain* a false representational state. Some favor acquisition, for instance when choosing to define mimicry on the basis of behavior alteration, and thus exclude masquerade and camouflage (Font 2019). However, one might suggest that causing an organism *not to acquire* a *true* representational state should also be included in the definition. Although a full discussion of these options would be complex and beyond the scope of this article, we think the perspective offered here suggests that the first two options (acquiring or maintaining state) should be ruled in, whereas the latter (not acquiring a state) should be ruled out. In short, if hiding truth sufficed for deception, then probably all cases of mimicry would be classified as deceptive, which would blur some distinctions that we consider worth keeping. Therefore, our tentative suggestion is that deception should be interpreted as involving producing or maintaining a false representational state, whereas causing an organism not to acquire a true representation should be excluded.⁹

If this reasoning is on the right track, then whether background matching should be classified as deception depends on whether it causes the receiver to produce or maintain a false representational state or not. This might happen in some cases, but not in others. For instance, in some situations the receiver might possess a mental state that we can approximately capture with the expression “there is no prey around,” in which case background matching would be maintaining a false representational state. In other scenarios, however, the receiver might lack any contentful mental state about prey, so background matching only causes it to fail to acquire a true representational state, which we have argued does not suffice for deception. Thus, some cases of background matching probably count as deceptive, whereas others do not. Not much seems to hinge on this issue.

So far we have discussed cases of mimicry and deception that involve relatively stable features (shapes, colors) and recurring signals (fireflies’ color patterns). Yet, because one of the main goals of this article is to illuminate the nature of deception in

⁹ Indeed, “producing or maintaining” is also likely to be a simplification. For instance, a false forecast might not cause our belief that it will rain tomorrow, but merely increase our confidence in this belief (Fallis and Lewis 2021). So perhaps we should only require increasing the probability of a false representational state. These questions would need to be addressed in a full account of deception.

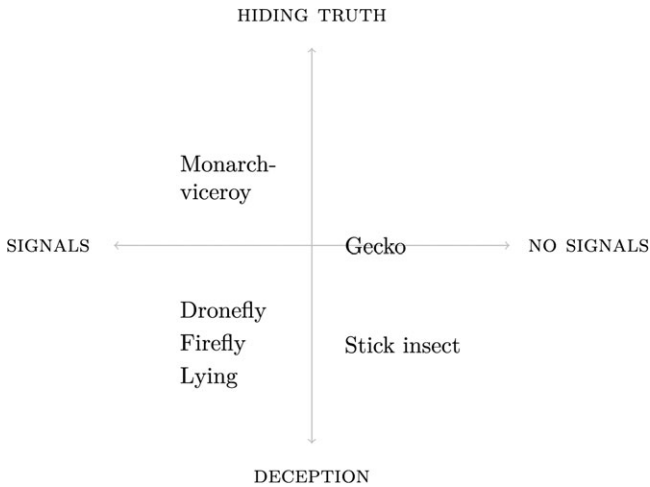


Figure 1. Examples of mimicry classified by two axes: signal/no signal and deceptive/nondeceptive.

general, our approach should also encompass one-shot deceptive signals. Fortunately, we think they can also be accommodated by DEC:

Bill Clinton’s utterance “I did not have sexual relations with that woman” deceptively mimics an utterance of a true sentence by an honest speaker because the former has the function of causing in the receiver the same representational state as the latter, but with false content, by means of being a partial copy of it.

Note that utterances can also have functions. The vast literature analysing the concept of “function” has established that functions can be grounded in different processes (Millikan 1984; Neander 1991). In orchids and fireflies, functions are probably grounded in natural selection, whereas in the case of human utterances, they might derive from human intentions. Hence, we think our account of deception in terms of mimicry can account for human and nonhuman forms of deception.

In short, we think that according to this definition of mimicry, deception can be regarded as an instance of mimicry. The relationship between some of the phenomena discussed so far is illustrated by Figure 1.

3.2. *The functional theory of deception*

Despite its original inclusion of deception within mimicry rather than the reverse, our approach is coherent with several functional approaches found in the literature. It is in line with most definitions of mimicry, which are also functional (e.g., Dalziell and Wenbergen 2016). Furthermore, it has strong affinities with Artiga and Paternotte’s (2018) proposal of a functional definition of deception. Indeed, we will argue that the differences between DEC and their functional theory of deception are minimal. Nonetheless, we think our perspective makes some important contributions. First, we arrive at a similar conclusion from a substantively different point of departure: from an original analysis of mimicry and the claim that all deception is an instance of

mimicry. Second, in addition to its other virtues, highlighted in the previous sections, the perspective offered here is also slightly superior, for two reasons: It shows how their account can deal with a recent objection raised by some philosophers and enhances its universality. The goal of this section is to describe and develop these aspects.

According to Artiga and Paternotte's proposal (2018):

(FT1) X is a deceptive state iff X has the function of causing a misinformative state (or failing to acquire a particular piece of information).

One reason for the convergence between FT1 and our account is that both avoid characterizing deception in terms of fitness benefits or costs—in particular costs that incur to the receiver (Christy 1997). As Font (2019) rightly states: “Mimetic signals are deceptive in the mechanistic sense because they elicit receiver errors, i.e. misperceptions or categorical mistakes, not because they are costly.” A functional definition allows a deceptive trait to be favored by natural selection without having to specify in virtue of what fitness combinations render it advantageous.

Now, a central aspect of this original analysis remains unexplained: How should we understand the receiver's “misinformative state”? What is it for a receiver to possess a misinformative state in the relevant sense? Although Artiga and Paternotte's (2018) avoid commitment to any specific interpretation, Birch (2019) and Fallis and Lewis (2021) have recently argued that some straightforward ways of understanding this notion do not work. On the one hand, if “misinformative state” is interpreted as false belief, then organisms such as insects or reptiles, insofar as they do not have beliefs (understood as mental states), cannot be deceived. That would be an unwelcome result because a central motivation for the functional theory of deception is to accommodate human and nonhuman deception. As an alternative, one could interpret “misinformative state” as a state that increases the probability of a nonactual state or decreases the probability of an actual state (along the lines of Skyrms 2010). However, on that understanding, many cases of hiding truth would be wrongly classified as deceptive (Birch 2019). For instance, the appearance of the monarch butterfly increases the probability of a nonactual state (a viceroy). More generally, nonfully informative signals would be classified as deceptive, insofar as they increase the probability of a nonactual state. That result puts some pressure on Artiga and Paternotte: Unless they can find some interpretation of “misinformative state” that yields the right results, their account is in jeopardy.

However, we propose that considering deception as an instance of mimicry provides an answer to this worry: “misinformation state” in FT1 should be understood as a state with false representational content. Note that we are using “representational content,” very broadly, to refer to a state that represents or is *about* certain aspect of the world and that has truth/accuracy conditions. We certainly don't presuppose that having representational states requires possessing any conceptual or sophisticated cognitive ability. If one has qualms with the notion of “representation,” we would be happy to use some other expression like “intentional state,” “intentional icon” or the like.

In any case, it is here that our account of deception as mimicry can be understood as a modification of Artiga and Paternotte’s proposal (2018). The analysis of deception we offered already was the following:

(DEC) An organism’s trait X deceptively mimics a model’s trait Y iff X has the function of causing in the receiver the same representational state as Y, but with false content, by means of being a partial copy of the model.

By relegating the identity of the model to the background, we can derive a definition of deceptive state that looks very similar to Artiga and Paternotte’s (2018) proposal:

(FT2) X is a deceptive state iff X has the function of causing in the receiver a false representational state by means of being a partial copy of a model (based on Artiga and Paternotte 2018).

Apart from the specification that “misinformative state” now refers to “false representational state,” there are only two differences between FT2 and Artiga and Paternotte’s FT1. First, FT2 definition requires deception to piggyback on a model. Note, however, that we are assuming a very liberal notion of model: A model need not be a living being or be present in the environment of the actors. Indeed, a model need not be an actual organism, object, or state of affairs because this role can also be fulfilled by a noninstantiated type. Although classical theories of mimicry predicted that this phenomenon only exists in the presence of an actual model (reviewed in Ruxton et al. 2018), more recent research puts this assumption into question. There are a number of mechanisms that can explain the maintenance of mimicry in allopatry, such as gene flow, range contraction/expansion, and different forms of selection (Pfennig and Mullen 2010). Some mimics even lack a specific model. For instance, the Australian orchid genus *Diuris* engages in “guild mimicry”: They evolved a color and shape similar to a guild of yellow and brown pea plants (Fabaceae) to deceive pollinators. The caterpillar of the moth *Hemeroplanes triptolemus* mimics a snake to deter predators, but it doesn’t resemble any specific species of snake. Likewise, in crypsis the heterogeneity of background type might lead to a mimetic compromise between them (Quicke 2017, 31). All these cases can be thought of as involving a copy of a noninstantiated type (which, of course, bears interesting relations with instantiated types).

This point is important because it enables us to generalise the notion of mimicry to other domains. In particular, some human forms of deception also mimic noninstantiated types—for example, a new string of words that no one has uttered before.¹⁰ If I lie to you by uttering the sentence “Bob moved from Bristol to Canberra to meet Cailin,” this is deceptive in virtue of being a copy of the sentence an honest person would utter, but that has the function of causing a false representational state.¹¹ If this sentence has never been uttered before, it makes my utterance a copy of

¹⁰ In section 4, we briefly discuss whether game-theoretic models could cover such cases.

¹¹ To be clearer: We are assuming here that the model is an honest speaker, the trait being mimicked is the honest speaker uttering a true sentence and the deceptive state is the actual sentence being uttered by the liar. Nonetheless, there might be other ways to apply our framework to this situation.

an uninstantiated type (which, again, bears certain interesting relations with instantiated types). Our previous reasoning suggests that this is not an *ad hoc* move to account for human deception; mimicking noninstantiated types is a common phenomenon in the natural world. Thus, this point highlights the universality of such a functional account, which, as Artiga and Paternotte (2018) noted, covers both nonhuman and human deception.

Secondly, in contrast to FT1, FT2 does not include failing to acquire a piece of information as being sufficient for deception, for the reasons already outlined: Hiding truth should be distinguished from deceiving. In any case, Artiga and Paternotte were not holding that failing to acquire a piece of information suffices for deception; rather, they suggested that some people might consider it sufficient. Thus, we do not think this is a real disagreement between FT1 and FT2.

Why then does interpreting “misinformative state” as false representational state help to assuage Birch (2019) and Fallis and Lewis (2021) worries? It is because the previous counterexamples do not apply any more. The color pattern of a monarch butterfly is not supposed to produce a false representational state, but at most to cause a failure to acquire a true representational state. More generally, not fully informative signals typically aim at hiding truth, rather than causing a false representational content. Thus, this response grants critics of the functional theory that the various interpretations of “misinformative state” provided are inadequate but replies that there is a more promising alternative that yields the correct results, as well as being motivated by an independent theory of mimicry. Thus, we cannot agree with Skyrms (2010, 76) when he writes:

What is important is that our definition [of deception] is information-based, rather than depending on imputed propositional content that is false.

On the contrary, as we have argued, it is important that a definition of deception depends on imputed propositional content that is false, rather than being information based. In section 5 we will argue that one way of understanding this concept of representational content resorts to a teleological theory. But another way to supplement our account with an explicit definition of false representational content would be to use the game-theoretic literature on signaling and deception. Several recent works of this kind have suggested possible definitions of the propositional content of a signal; we further discuss these in section 4. For now, it suffices to note that for a definition of propositional content to fit our account, it would only need to entail that in cases of deception, states involve a false propositional content. That is, it should not hold that a deceptive state has true content because it is understood disjunctively, for example, as “I am either a sexual partner or a predator” (see section 4 for more details).

Overall, regarding deception as mimicry not only provides an original and potentially illuminating perspective on deception in general but also adequately refines and extends Artiga and Paternotte’s proposal.¹² Furthermore, crucially, it

¹² Krstic (Forthcoming) also argues that analyses of deception should not be based on signals; although he thinks that mimicry, e.g. insects, may involve signals in a sense (more on this case in section 4.4). Still, he notes the advantages of a functional approach to deception (as well as of his own manipulation-based account).

addresses a recent objection suggested by Birch, and Fallis and Lewis. Therefore “misinformative state” should be interpreted as “false representational state.”

4. Game-theoretic approaches to deception

Now that we have laid down and defended our approach to deception, we would like to raise some concerns with game-theoretic analyses of deception, which further justify the need for our alternative account. The reason for addressing this point here is that some of the reasons that speak in favor of conceiving deception as a form of mimicry also suggest some worries for these accounts.

Game-theoretic accounts of deception are typically motivated by some paradigmatic examples of deception that involve signaling; in some scenarios (e.g., *Photuris* fireflies) sending certain signals counts as deception whereas in other cases it does not. Thus, the goal is to provide an analysis that delivers this result and that can account for the explanatory role associated with deceptive states. To this end, theorists usually develop models and seek to identify which features of the model are distinctive of deceptive *signals* (Birch 2019; Fallis and Lewis 2019; Martínez 2019; McWhirter 2016; Shea et al. 2018; Skyrms 2010). Game-theoretic models are extremely fruitful and useful, partly because they provide an elegant and precise way of describing paradigmatic cases of deception and of distinguishing them from nondeceptive interactions.

Yet we also think focusing on game-theoretic analyses of deceptive signals might systematically lead researchers into wrong theories of the nature of deception. The first worry, of course, is that as we saw in the preceding text, some deceptive states are not signals (see also Artiga and Paternotte 2018). Stick insects are not sending signals (in the same way branches are not). Thus, the existence of deception without signaling represents a first difficulty for these approaches.¹³

This concern, however, is probably not decisive. After all, one could grant that definitions of deception based on deceptive signals cannot be the full story, but still hold that paradigmatic instances of deception involve signals, and that other cases might be accommodated as marginal instances. Moreover, the evolutionary dynamics of deceptive traits may be the same whether signals are involved or not, which would render game-theoretic examples sufficient for our understanding of deception. Accordingly, one may argue that by developing game-theoretic analyses, one does capture the central aspects of this phenomenon, which should eventually be extended to nondeceptive signals. Even if more work needs to be done to make this view plausible (why are stick insects or orchids less paradigmatic examples of deception than wasp beetles or *Photuris* fireflies?), this is nonetheless a viable option.

However, we think there is a more pressing worry for these approaches, which does not derive from any explicit requirement of the models, but from the solutions they suggest. This is that game-theoretic approaches tend to suggest the wrong kind

¹³ An exception is Fallis and Lewis (2021, 122), who do consider twigs as “degenerate senders” with “constant interests” and argue that mimicry can thus be understood as involving signaling and modelled in game theory. On the contrary, we have argued that mimicry is the general category under which deception falls, and that its definition does not necessitate payoff-based considerations. (This does not entail that their central point, the definition of the propositional content of a deceptive signal, is incorrect.)

of definitions of deception, and the reason for this derives from some features of game-theoretic models. Let us elaborate.

A game is a mathematical object defined by appealing to certain features (payoffs, number of players, set of states, signals and acts, etc.). A player is any entity that can be associated with a certain payoff and that can adopt different strategies (or have different types in a population, in an evolutionary setting) that can be understood as a function from states to signals (in senders) or as a function from signals to actions (in receivers). Crucially, note that these models do not assume that players possess representational states. Senders and receivers can be as simple-minded as you like; they need not represent the world in certain ways. Players only need to have certain interests (payoff) and give a response to certain situations (in terms of sending a signal or acting). Indeed, this makes a lot of sense if one is interested, for instance, in the evolution of communication because communication can take place between very simple organisms.¹⁴ There is also a theoretical motivation for this assumption: The notion of representational content is a hotly disputed concept, and it is unclear how to include an assumption that certain players have states with certain representational content within the model.

Nonetheless, the fact that in these models players usually lack states with representational content has a remarkable consequence: Because to define deception one can only appeal to features of the model, and as the representational content of the player's mental life is absent, one cannot include it in the analysis of deception. Thus, if one begins by focusing on *signals* and then seeks to specify what is special about *deceptive* signals, deception should be analyzed in terms of those features included in a game. In particular, theorists might appeal to sender/receiver costs and benefits, false/misinformative signals, and population strategies or counterfactual strategies, for instance. However, partly for the reasons suggested already, these strategies seem limited if one aims to provide a general account of deception. For instance, deception cannot be defined in terms of false/misinformative signals because many deceptive states are not signals and, indeed, it is possible to deceive by using a *true* signal, as in double bluffing: For instance if, knowing that you don't trust me, I tell you that a road is very dangerous (which I think is true) so that you believe that it is not. Similarly, the existence of altruistic and spiteful deception suggests that simple analyses in terms of sender benefit or receiver cost are inadequate (see Christy 1997). An appeal to receiver behavior has to accommodate the fact that successful deception might not alter behavior in a given situation. Of course, one can develop more sophisticated analyses than the simple accounts suggested here, but our point is that they are unlikely to succeed because deception has something to do with causing the receiver to misrepresent the world, which usually is not a parameter of game-theoretic models.¹⁵

Even if these models cannot by themselves fit the purpose at hand, that is, provide a definition of deception, this does not detract from their value for other theoretical purposes; in fact we consider these models as complementary to our approach, for three reasons. First they illuminate the nature and dynamics of signal-based

¹⁴ We come back to this point in section 5.

¹⁵ In principle, one could develop game-theoretic models that included the receiver's representational content as one of the parameters. However, including this aspect in models might be far from easy.

deception, which, as already stated, is far from being negligible. Second and more importantly, these works suggest fruitful ways to define the propositional content of a signal, and so allow one to make progress on the conditions under which one can talk of a false representational content. For instance, several recent works (Birch 2014; Skyrms and Barrett 2019; Shea et al. 2018; but see Artiga 2020) have suggested that games and game-theoretic equilibria can be used to define propositional content so that it doesn't suffer from the limits of Skyrms's (2010) account: They do not attribute to deceptive signals a disjunctive content (either a sex partner or a predator is there) that would render them always true. Although such accounts are not equivalent, all respect the intuition that deceptive signals can have false content, and thus can be used to fill in the details of what counts as false representational content (at least in the family of cases involving signal-based deception). Finally, these recent game-theoretic accounts and ours can be seen as sharing the view that no account of deception can be purely informational. The former typically go beyond *actual* information by saying that the propositional content of a deceptive signal depends on what would be the case in a similar nondeceptive situation.¹⁶ Our account goes beyond information by requiring agents with minimal representational capacities.

The game theoretic literature may offer routes for future work on deception that fits some of this article's suggestions. We already pointed out that our account fits the novel lies or deception involving noninstantiated signals. Game-theoretic models may be helpful for understanding under which conditions a nondeceptive signal may be mimicked (more or less faithfully) to deceive a receiver, by using and further developing original insights regarding the invention of new signals (McKenzie Alexander et al. 2012), the gradual complexification of games (Barrett and Skyrms 2017; Barrett 2020), or the dynamics of how agents may learn to lie (Skyrms and Barrett 2019, sec. 6). In other words, despite their limits regarding a general definition of deception, game theoretic models have been and should remain an illuminating source of its evolution and dynamics.¹⁷

5. Discussion

In this final section, we would like to address several issues with or objections to the view defended in this article.

First, one goal of Artiga and Paternotte's (2018) functional account was to encompass cases of altruistic or spiteful deception, in which the producer does not benefit (and that typically involve kin selection). However, it isn't clear whether altruistic or spiteful mimicry exists, or is even possible. We have said that it need not involve a cost to the receiver; but all known examples include a fitness benefit for the producer. If this is correct, then as a logical matter of fact, deception could not possibly be a kind of mimicry.

¹⁶ For Birch (2014), a similar situation is given by the closest separating equilibrium, or the equilibrium in the closest game; for Skyrms and Barrett (2019), it corresponds to the part of the game (the subgame) in which agents have common interests. These could still be called informational insofar as they involve modal or counterfactual information; what matters is that they are not informational in the more straightforward sense of Skyrms's (2010) account.

¹⁷ We thank two anonymous referees for helping us clarify the link between our account and game-theoretic approaches.

However, this conclusion would be misguided. To argue for our mimicry-based account of deception, we considered nonsignal-based accounts of mimicry. These may necessitate fitness benefits while signal-based cases do not. In any case, our definition makes no mention of such benefits or costs, and so escapes the risk of logical contradiction.

Still, why would deception differ depending on whether it involves signals or not? This difference may be a matter of degrees rather than categorical. Altruistic deception is already rare when signal based, probably because it is evolutionarily unstable. Thus if signals are more flexible and so evolve faster than physiological traits, and if even this barely allows deceptive signals to be maintained by evolution when they are altruistic, then we should expect nonsignal-based altruistic deception to be even less common, that is, virtually nonexistent (even if conceptually possible).

Second, as we have recalled several times, we put forward a concept of mimicry that is intended to encompass human and nonhuman cases, which includes, among other things, human lying. Even granting that such a broad concept is useful for certain purposes, one might be concerned that it is certainly not the kind of concept in which biologists are interested when considering mimicry.

We think this complaint is justified. As a matter of fact, biologists might prefer to further restrict the notion of mimicry. For instance, one might require the mimic to resemble a specific model (Schaefer and Ruxton 2009) or restrict the notion of mimicry to the resemblance with a living being (Vane-Wright 1980). Regarding the first approach, for example, a worm that resembles a snake, but not a specific species of snake would not count as mimicry. Regarding the second, stick insects would not count as mimics. Note, however, that our point was not primarily to devise a concept of mimicry that fits the interests of biologists working with this concept, but to develop a concept that could be used to better understand the phenomenon of deception by highlighting certain continuities between these phenomena. Biologists working on mimicry tend to have other goals in mind, so they might be interested in a different concept of mimicry. If our use of the term “mimicry” is in question, we could happily employ a different concept such as “mimesis.”

A third concern has to do with our appeal to “representational states.” One might complain that we provided an interpretation of “misinformative state” in terms of representational state without explaining in detail what makes it the case that certain organisms have states with representational content. Nevertheless, there are different approaches and our proposal is compatible with many of them. For instance, one could endorse Millikan’s (1984), Neander’s (2017), or Shea’s (2018) teleological views. Note that with these views, the definition of a representational state in turn involves an attribution of functions grounded in some sort of selection or stabilizing process, such as learning or natural selection. For instance, on a simple teleological theory, a receiver has as a representation *R* with content *C* iff it possesses an internal mechanism whose function is to token *R* whenever *C* is the case. Interestingly, once applied to our definition of deception, such a theory would involve a two-step account: first, representations evolve and, second, deception evolves by causing false representations. The dynamics of deception may even get particularly complex when representations and deception coevolve. Certainly, our account might not be compatible with some naturalistic proposals, such as some informational theories, precisely for the reasons provided by Birch (2019) and Fallis and Lewis (2021)

developed in section 4.1.¹⁸ Nonetheless, because our approach is consistent with many accounts (in particular game-theoretic ones, as argued in the previous section), we prefer to not to take sides in the debate between them here.

The fourth, final worry we consider has to do with the relation between communication and deception. If even simple organisms such as bacteria can communicate while lacking representational states, why would deception (or something close enough) be impossible for them as well?¹⁹ If genuine deception among bacteria is possible, one might worry that our definition only provides sufficient conditions for deception, just like the cost/benefit accounts. At least two different but compatible replies are possible though. First, one may hold that communication tends to go hand in hand with some minimal form of intentional content on the part of the sender and/or receiver. In particular, if one holds that bacteria possess some minimal form of cognition (Shapiro 2007; Figdor 2018), then our framework is fully compatible with them engaging in genuine deception. Second, whatever one thinks of the prospects of this first reply, we can explain why some cases intuitively *seem* to involve deception although they do not. Although according to our account deception is defined in terms of mimicry, representational states and so on, we of course accept that common (even paradigmatic) instances of deception possess additional features. Some of them have been highlighted in previous theories of deception: Deception tends to involve a receiver cost and a sender benefit, it often involves signals, and so forth. Indeed, it might even be useful to use some of these features in scientific practice to study or identify potential cases of deception. Now, our point is that whenever we find these properties in nature, we might be tempted to classify them as a genuine case of deception. Thus, although we contend that the deceptive state of bacteria utterly depends on whether they satisfy our analysis, we can explain why they *seem* to be an instance of genuine deception.

6. Conclusions

In this article we have shown that there are important continuities between biological phenomena typically classified under the headings “mimicry” and “deception,” and that exploring this perspective sheds some light on the nature of both phenomena. To highlight these continuities, we developed a broad concept of “mimicry” and argued that perhaps surprisingly, deception can be regarded as a specific kind of mimicry, rather than the reverse.

We also derived certain important consequences of this approach. This perspective reveals that an analysis of deception has to appeal to the receiver’s representational content and that deception need not involve signals. We also showed this analysis offers a superior version of the functional theory of deception to the previous version because it provides an answer to recent objections (Birch 2019; Fallis and Lewis 2021). Furthermore, it explains why game-theoretic analyses of deception tend to go astray. Certainly, some questions still remain open, as discussed in sections 4 and 6. For instance, one question to be explored is whether this approach can accommodate all cases of deception. Another concern is whether this broad definition of mimicry can

¹⁸ We would like to thank Nicholas Shea for helping us develop this point.

¹⁹ We thank an anonymous referee for raising this point.

replace extant definitions in the context of biological research. All in all, we believe considering deception as an instance of mimicry provides an illuminating perspective that opens up new avenues for future research.

Acknowledgments. We would like to thank Marta Campdelacreu, Josep Corbí, Don Fallis, Manolo Martínez, Sergo Oms, Vladimir Krstic, Hannah Rubin, Carlota Serrahima, Nicholas Shea, Chon Tejedor, Víctor Verdejo, and the audiences at the Valencia Colloquium in Philosophy (Valencia, February 2022), the Workshop on Deception (Barcelona, June 2021), and the VI Congrés Català de Filosofia (Manresa, June 2023) for their helpful suggestions and criticisms. Thanks to John Wright for proofreading the article. Part of the research conducive to the ideas developed in this paper was funded by the projects PID2021-127046NA-I00 and PID2019-106420GA-I00 funded by the Ministerio de Ciencia, Innovación y Universidades and the projects CIGE/2021/160 and CISEJI/2023/51 funded by the Conselleria d'Educació, Universitats i Ocupació.

References

- Aaronson Richard B. 1983. "Foraging Behavior of the West Atlantic Trumpetfish, *Aulostomus Maculatus*: Use of Large, Herbivorous Reef Fishes as Camouflage." *Bulletin of Marine Science* 33 (1):166–71.
- Anders Nilsson, L. 1992. "Orchid pollination biology." *Trends in Ecology & Evolution* 7 (8):255–59.
- Artiga, Marc. 2014. "Signaling without Cooperation." *Biology and Philosophy* 29 (3): 357–78. <https://doi.org/10.1007/s10539-014-9436-0>
- Artiga, Marc. 2020. Models, information and meaning. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 82:101284.
- Artiga, Marc, and Cédric Paternotte. 2018. "Deception: A Functional Account." *Philosophical Studies* 175 (3):579–600. <https://doi.org/10.1007/s11098-017-0883-8>
- Barnett, James B., Constantine Michalis, Hannah Anderson, and Innes C. Cuthill. 2020. "Imperfect Transparency and Camouflage in Glass Frogs." *PNAS* 117 (23):12809–85. <https://doi.org/10.1073/pnas.1919417111>
- Barrett, Jeffrey. 2020. "Self-Assembling Games and the Evolution of Salience." *British Journal for the Philosophy of Science* 74 (1):75–89. <https://doi.org/10.1093/bjps/axv043>
- Barrett, Jeffrey, and Brian Skyrms. 2017. "Self-Assembling Games." *British Journal for the Philosophy of Science* 68 (2):329–53. <https://doi.org/10.1093/bjps/axv043>
- Birch, Jonathan. 2014. "Propositional Content in Signaling Systems." *Philosophical Studies* 171:493–512. <https://doi.org/10.1007/s11098-014-0280-5>
- Birch, Jonathan. 2019. "Altruistic Deception." *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 74:27–33. <https://doi.org/10.1016/j.shpsc.2019.01.004>
- Brower, Jane V. Z., and Lincoln P. Brower. 1962. "Experimental Studies of Mimicry. 6. The Reactions of Toads (*Bufo terrestris*) to Honeybees (*Apis mellifera*) and Their Dronefly Mimics (*Eristalis vinetorum*)." *The American Naturalist* 96:297–307. <https://doi.org/10.1086/282237>
- Carson, Thomas L. 2010. *Lying and Deception: Theory and Practise*. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199577415.001.0001>
- Chisholm, Roderick M., and Thomas D. Feehan. 1977. "The Intent to Deceive." *Journal of Philosophy* 74 (3):143–59. <https://doi.org/10.2307/2025605>
- Christy, John. 1997. "Deception: The Correct Path to Enlightenment?" *Trends in Ecology & Evolution* 12 (4):160. [https://doi.org/10.1016/s0169-5347\(97\)89785-6](https://doi.org/10.1016/s0169-5347(97)89785-6)
- Dafni, Amots. 1984. "Mimicry and Deception in Pollination." *Annual Review of Ecology and Systematics* 15 (1):259–78. <https://doi.org/10.1016/bs.abr.2016.10.005>
- Dalziel, Anastasia H., and Justin A. Welbergen. 2016. "Mimicry for All Modalities." *Ecology Letters* 19 (6):609–19. <https://doi.org/10.1111/ele.12602>
- Fallis, Don, and Peter J. Lewis. 2019. "Toward a Formal Analysis of Deceptive Signaling." *Synthese* 6:2279–2303. <https://doi.org/10.1007/s11229-017-1536-3>
- Fallis, Don, and Peter J. Lewis. 2021. "Animal Deception and the Content of Signals." *Studies in History and Philosophy of Science Part A* 87:114–24. <https://doi.org/10.1016/j.shpsa.2021.03.004>

- Figdor, Carrie. 2018. *Pieces of Mind: The Proper Domain of Psychological Predicates*. Oxford: Oxford University Press. <https://doi.org/10.1093/oso/9780198809524.001.0001>
- Font, Enrique. 2019. "Mimicry, Camouflage and Perceptual Exploitation: The Evolution of Deception in Nature." *Biosemiotics* 12 (1):7–24. <https://doi.org/10.1007/s12304-018-9339-6>
- Forbes, Peter. 2011. *Dazzled and Deceived: Mimicry and Camouflage*. New Haven, CT: Yale University Press.
- Gomez, Doris, Chariline Pinna, Jonathan Paireira, Mónica Arias, Jérôme Barbut, Aaron Pomerantz, Willy Daney de Marcillac, Serge Berthier, Nipam Patel, Christine Andraud, and Marianne Elias. 2021. "Wing Transparency in Butterflies and Moths: Structural Diversity, Optical Properties, and Ecological Relevance." *Ecological Monographs* 91 (4):e01475. <https://doi.org/10.1002/ecm.1475>
- Greene, Harry W., and Roy W. McDiarmid. 1981. "Coral Snake Mimicry: Does It Occur?" *Science* 213 (4513):1207–12. <https://doi.org/10.1126/science.213.4513.1207>
- Jamie, Gabriel A. 2017. "Signals, Cues and the Nature of Mimicry." *Proceedings of the Royal Society B: Biological Sciences* 284:20162080. <https://doi.org/10.1098/rspb.2016.2080>
- Kikuchi, David W., and D. W. Pfennig. 2013. "Imperfect Mimicry and the Limits of Natural Selection." *The Quarterly Review of Biology* 88 (4):297–315. <https://doi.org/10.1086/673758>. PMID: 24552099
- Krstić, Vladimir. 2021. "On the Function of Self-Deception." *European Journal of Philosophy* 29 (4):846–63. <https://doi.org/10.1111/ejop.12608>
- Krstić, Vladimir. Forthcoming. "We Should Move on from Signalling-Based Analyses of Biological Deception." *Erkenntnis* <https://doi.org/10.1007/s10670-023-00719-x>
- Krstić, Vladimir, and Chantelle Saville. 2019. "Deception (Under Uncertainty) as a Kind of Manipulation." *Australasian Journal of Philosophy* 97 (4):830–35. <https://doi.org/10.1080/00048402.2019.1604777>
- Lloyd, James E. 1975. "Aggressive Mimicry in Photuris Fireflies: Signal Repertoires by Femmes Fatales." *Science* 187:452–53. <https://doi.org/10.1126/science.187.4175.452>
- Mahon, James E. 2007. "A Definition of Deceiving." *International Journal of Applied Philosophy* 21 (2):181–94. <https://doi.org/10.5840/ijap20072124>
- Mallet, James, and Lawrence E. Gilbert. 1995. "Why Are There So Many Mimicry Rings? Correlations between Habitat, Behaviour and Mimicry in Heliconius Butterflies." *Biological Journal of the Linnean Society* 55:159–80. <https://doi.org/10.1111/j.1095-8312.1995.tb01057.x>
- Martin, Schaefer, H., and Graeme D. Ruxton. 2009. "Deception in Plants: Mimicry or Perceptual Exploitation?" *Trends in Ecology & Evolution* 24 (12):676–85. <https://doi.org/10.1016/j.tree.2009.06.006>
- Martínez, Manolo. 2019. "Deception as Cooperation." *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 77:101184. <https://doi.org/10.1016/j.shpsc.2019.101184>
- McKenzie Alexander, Jason, Brian Skyrms, and Sandy L. Zabell. 2012. "Inventing New Signals." *Dynamic Games and Applications* 2:129–45. <https://doi.org/10.1007/s13235-011-0027-2>
- McWhirter, Gregory. 2016. "Behavioural Deception and Formal Models of Communication." *British Journal for the Philosophy of Science* 67 (3):757–80. <https://doi.org/10.1093/bjps/axv001>
- Merilaita, Sami, Nicholas E. Scott-Samuel, and Innes C. Cuthill. 2017. "How Camouflage Works." *Philosophical Transactions of the Royal Society B Biological Sciences* 372 (1724):20160341. <https://doi.org/10.1098/rstb.2016.0341>
- Millikan, Ruth G. 1984. *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.
- Neander, Karen. 1991. "Functions as Selected Effects: The Conceptual Analyst's Defense." *Philosophy of Science* 58 (2):168–84. <https://doi.org/10.1086/289610>
- Neander, Karen. 2017. *A Mark of the Mental: A Defence of Informational Teleosemantics*. Cambridge, MA: MIT Press. <https://doi.org/10.7551/mitpress/7387.001.0001>
- Pfennig, David W., and Sean P. Mullen. 2010. "Mimics without Models: Causes and Consequences of Allopatry in Batesian Mimicry Complexes." *Proceedings of the Royal Society B* 277:2577–85. <https://doi.org/10.1098/rspb.2010.0586>
- Quicke, Donald. 2017. *Mimicry, Crypsis, Masquerade and other Adaptive Resemblances*. Hoboken, NJ: John Wiley & Sons.
- Roy, Bitty A., and Alex Widmer. 1999. "Floral Mimicry: A Fascinating Yet Poorly Understood Phenomenon." *Trends in Plant Science* 4 (8):325–30. [https://doi.org/10.1016/s1360-1385\(99\)01445-4](https://doi.org/10.1016/s1360-1385(99)01445-4)
- Ruxton, Graeme D., William L. Allen, Thomas N. Sherratt, and Michael P. Speed. 2018. *Avoiding Attack: The Evolutionary Ecology of Crypsis, Aposematism and Mimicry* (2nd edition). Oxford: Oxford University Press. <https://doi.org/10.1093/oso/978>

- Schaefer, Martin and Graeme D. Ruxton. 2009. "Deception in Plants: Mimicry or Perceptual Exploitation?" *Trends in Ecology & Evolution* 24 (12):676–85.
- Schiestl, Florian P. 2005. "On the Success of a Swindle: Pollination by Deception in Orchids." *Die Naturwissenschaften* 92 (6):255–64.
- Shapiro, James A. 2007 "Bacteria Are Small But Not Stupid: Cognition, Natural Genetic Engineering and Socio-bacteriology." *Studies in History and Philosophy of Biological and Biomedical Sciences* 38 (4):807–19. <https://doi.org/10.1016/j.shpsc.2007.09.010>
- Shea, Nicholas. 2018. *Representation in Cognitive Science*. Oxford: Oxford University Press. <https://doi.org/10.1093/oso/9780198812883.001.0001>
- Shea, Nicholas, Godfrey-Smith, Peter, and Cao, Rosa. 2018. "Content in Simple Signalling Systems." *British Journal for the Philosophy of Science* 69 (4):1009–35. <https://doi.org/10.1093/bjps/axw036>
- Skyrms, Brian. 2010. *Signals: Evolution, Learning, and Information*. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199580828.001.0001>
- Skyrms, Brian, and Jeffrey Barrett. 2019. "Propositional Content in Signals." *Studies in History and Philosophy of Science Part C* 74:34–39. <https://doi.org/10.1016/j.shpsc.2019.01.005>
- Vane-Wright, Richard I. 1976. "A Unified Classification of Mimetic Resemblances." *Biological Journal of the Linnean Society* 8 (1):25–56. <https://doi.org/10.1111/j.1095-8312.1976.tb00240.x>
- Vane-Wright, Richard I. 1980. "On the Definition of Mimicry." *Biological Journal of the Linnean Society* 13 (1):1–6. <https://doi.org/10.1111/j.1095-8312.1980.tb00066.x>

Cite this article: Artiga, Marc and Cédric Paternotte. 2024. "Deception as Mimicry." *Philosophy of Science* 91 (2):370–389. <https://doi.org/10.1017/psa.2023.156>