King's Research Portal

DOI:

[10.1080/09515089.2017.1328491](https://doi.org/10.1080/09515089.2017.1328491)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Doody, P. (2017). Is there evidence of robust, unconscious self-deception? A reply to Funkhouser and Barrett. *PHILOSOPHICAL PSYCHOLOGY*, 1-20. <https://doi.org/10.1080/09515089.2017.1328491>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Is there evidence of robust, unconscious self-deception? A reply to Funkhouser and Barrett

Author: Paul Doody

Abstract

Robust self-deception, in Funkhouser and Barrett's sense, consists in the strategic pursuit of the goal of misleading oneself with respect to some proposition. Funkhouser and Barrett's thesis is that an evaluation of the relevant empirical literatures reveals that the unconscious mind engages in robust self-deception. If Funkhouser and Barrett are correct, the psychological evidence vindicates an account of self-deception that challenges the orthodox motivationalist approach, and makes clear the distinction between self-deception and other forms of motivated belief formation such as wishful thinking. The aim of this paper is to show that the evidence adduced by Funkhouser and Barrett fails to establish the existence of robust, unconscious self-deception.

Acknowledgements

I would like to thank Julien Dutant, Kevin Lynch and an anonymous referee for *Philosophical Psychology* for their invaluable feedback on earlier versions of this paper. Many thanks also to the National University of Ireland for its support under the Travelling Studentship Scheme, and to the London Arts and Humanities Partnership.

1. Introduction

Suppose that I aim to make you believe something that I know, or have good reason to believe, is false — say, that your partner is cheating on you. Suppose also that I execute this project by using resources of the kind employed in ordinary cases of deception.¹ Perhaps I draw your attention to the fact that your partner's affection towards you has been uncharacteristically sporadic of late, or maybe I comment on their close relationship with an attractive colleague. Finally, suppose that my project is successful: After repeated instances of this manipulation of your evidence, you come to believe that your partner is having an affair. Call this kind of case *standard deception*.

Does self-deception simply consist in an agent's perpetrating standard deception upon herself? An affirmative answer gives rise to a host of well-known puzzles concerning both the doxastic *state* of the self-deceived subject — the so-called “static puzzle” (see Mele, 2001, p. 7-8) — and the *strategy* required for successful self-deception — the “dynamic puzzle” (*ibid.*, p. 7-8). Early accounts of self-deception sought to accommodate these puzzles. The most prominent version of such accounts was the *mental partitioning account*: The self-deceiver is divided into two parts (deceiver and deceived); the deceiver is understood as a quasi-agent or ‘sub-system’ whose goal is to install in the person's ‘main system’ the relevant false or unwarranted belief (Pears, 1984, p. 91-92).² Owing partly to a perceived lack of empirical support, the mental partitioning account has largely been abandoned in the contemporary literature. Indeed, rather than attempt to provide accounts of self-deception which meet the

aforementioned puzzles, most contemporary philosophers deny that self-deception is analogous to standard deception. The orthodox account is now *motivationalism* — according to which self-deception is simply a species of motivationally biased belief (see Lynch, 2012, 2013; Mele, 2001; Nelkin, 2002; Scott-Kakures, 2002).

Still, motivationalism is not without its detractors. A persistent worry is that it is unable to capture the distinctive tension in self-deception (see Audi, 1997; Losonsky, 1997; Porcher, 2012) or that it collapses the distinction between self-deception and other kinds of motivated belief formation, such as wishful thinking. Relatedly, Funkhouser and Barrett maintain that motivationalism's characterisation of self-deception "fail[s] to involve deception in any robust sense" (Funkhouser and Barrett, 2016, p. 682). The mental partitioning account seems to avoid these worries. If this account is vindicated by the relevant empirical literatures, then, far from failing to be a viable account of self-deception, it might be *preferable* to motivationalism. Funkhouser and Barrett attempt to provide such vindication in their paper. Specifically, they claim that evidence from social and cognitive psychology confirms the reality of *robust, unconscious self-deception*: The unconscious mind engages in tactics and strategies the aim of which is to make the subject in question (or the part of her mind that is distinct from the deceptive part) believe some falsehood. Unfortunately, as I will argue in what follows, the evidence presented by Funkhouser and Barrett fails to establish the reality of robust, unconscious self-deception.

2. What is Robust, Unconscious Self-Deception?

In order to assess the evidence that Funkhouser and Barrett adduce in support of the reality of robust unconscious self-deception, it is first necessary to make explicit what the authors mean

by this term. Let's start with the notion of unconsciousness in play here, as this is relatively straightforward. For Funkhouser and Barrett, unconscious mental states or processes are mental states or processes which are not verbally reportable, in the following sense. The subject either simply cannot report the state or process in question or she can only report it with the assistance of some special aid — such as, for example, the employment of excessive self-scrutiny or therapy (Funkhouser & Barrett, 2016, endnote 1 p. 693). Admittedly, this characterization is fairly rough. But it will suffice for our discussion.

Funkhouser and Barrett identify four elements which they claim are jointly sufficient for robust self-deception. First, the subject must have the goal to mislead herself regarding some state of affairs. The subject must then meet a further three conditions. She must engage in:

1. the *strategic* pursuit of that goal [the goal to mislead herself],
2. in a way that is *flexible* to the nuances of possibly changing situations, and
3. which involves some *retention of the truth* (or at least a non-trivial doubt).

(Funkhouser & Barrett, 2016, p. 683).

Call conditions 1 to 3 *the strategy condition*, *the flexibility condition*, and *the truth-retention condition*, respectively. Each of these conditions requires some unpacking.

Let's start with the strategy condition. Funkhouser and Barrett expand on the notion of strategic goal pursuit as follows: "Strategy requires a *goal* and *rationality* — strategizing is the rational pursuit of a goal. But we need not think of such rationality as conscious, deliberate, or ideal" (Funkhouser & Barrett, 2016, p. 683). The concept of strategizing is very closely connected to the more familiar concept of deliberation. Perhaps, then, the strategizing that

Funkhouser and Barrett refer to consists in deliberation on, and selection of, what are perceived by the putative deceiver as the best means towards her end, namely, the deception of herself (or, perhaps more precisely, the deception of the *part of herself* that has been targeted for deception — viz., the part of her mind that is consciously accessible). This reading is supported by Funkhouser and Barrett's remarks on automatic processes.

[U]nconscious responses are often described as automatic reactions that are too simplistic to generate the rational pursuit of a goal. It is certainly true that *some* unconscious knowledge is like this. For example, the galvanic skin responses supposedly indicative of unconscious recognition of familiar voices found in the Gur and Sackeim study are simple and automatic responses that are not incorporated into any kind of rational goal pursuit. It is not strategic because the information is not used toward any end[.]

(Funkhouser & Barrett, 2016, p. 685-686)

As we can see, Funkhouser and Barrett wish to distinguish the kind of unconscious response that they are interested in — the strategic, goal-oriented response — from the kind of response demonstrated by Gur and Sackeim (1979). Gur and Sackeim present evidence to the effect that subjects tend to exhibit increased galvanic skin response (GSR) when presented with audio recordings of their own, and familiar, voices. This kind of response — whether it evinces the belief that one is hearing one's voice or merely indicates *sub-doxastic* recognition of that proposition (for the latter interpretation see Mele, 2001, p. 82) — is mechanical or dumb, insofar as the information in question is not poised for further activity in the subject's cognition or behaviour. This seems to be why, according to Funkhouser and Barrett, it fails to be strategic: "It is not strategic because the information is *not used toward any end*" (emphasis added) (Funkhouser & Barrett, 2016, p. 686).

To count as strategic (and not merely automatic), then, an unconscious response must be more than a mere signalling of the presence or absence of a particular kind of stimulus. The response must be one that uses information in a process which aims at some end. This sounds just like deliberation — though, as Funkhouser and Barrett emphasise, such deliberation need not be ideal or conscious.³ With an understanding of the strategy condition in place, let's turn our attention to the flexibility and truth-retention conditions.

The flexibility condition maintains that the strategy must be pursued in a way that is flexible to changes and nuances in the deceiver's situation. This is explained by Funkhouser and Barrett as follows:

[Robust] Deceivers actively adapt to situations so as to better avoid the truth or perpetuate a falsehood. Such flexibility further supports the claim that the behavior truly is rational...and goal-seeking (e.g., purposefully aimed at deception, the avoidance of a truth).

(Funkhouser & Barrett, 2016, p. 683)

The truth retention condition requires that the deceiver retains some knowledge of the true proposition. This knowledge, Funkhouser and Barrett claim, is typically expressed in the specific way in which the deceiver is flexible to her environment:

In many cases of self-deception, the particular form this flexibility takes also reveals that the subject retains some of the truth or at least has a non-trivial doubt. For example,

strategically avoiding evidence in some settings but not in others, so as to better self-deceive, can reveal unconscious knowledge of that very evidence.

(Funkhouser & Barrett, 2016, p. 683)

Let's summarize the discussion so far with the following schematized characterisation of robust, unconscious self-deception.

A subject *S* engages in robust, unconscious self-deception with respect to a proposition *p* iff:

- (1) *S* strategically pursues the goal of misleading herself regarding whether *p*.
- (2) *S*'s pursuit of this goal is sensitive to possible and actual changes in her situation.
- (3) In pursuing this goal, *S* retains knowledge of the truth regarding whether *p*.
- (4) The fact that *S* is engaging in this strategy is either: (i) completely unreportable by *S*; or (ii) reportable by *S* only with the assistance of some special aid.

3. Assessing the Evidence

Funkhouser and Barrett's strategy is to present "experimental psychology literature that favors the idea that unconscious deception is real and ubiquitous in human (and animal) behavior" (Funkhouser & Barrett, 2016, p. 687). Specifically, they offer studies which, they argue, confirm that the unconscious engages in, or has the capacity to engage in, the following tactics of deception:

- (1) The avoidance of threatening information.
- (2) The production of misleading information.
- (3) Selective attention to favourable information.
- (4) Biased interpretation of threatening information.
- (5) Rationalization.
- (6) Memory distortion.
- (7) Acting as-if.

Recall that, for the execution of a given tactic to instantiate robust deception, it must meet the strategy, flexibility, and truth-retention conditions. Focussing specifically on the strategy condition, I argue that the studies presented by Funkhouser and Barrett fail to support the claim that the unconscious engages in, or has the capacity to engage in, these tactics of deception.⁴

3.1 Avoidance of threatening information

People often avoid information, or sources of information, which may confirm or underline an uncomfortable truth. For instance, I might avoid opening a reply to my application for a research grant, for fear that it might be a rejection letter. Here the avoidance is *conscious*. And it is based on either the suspicion that the letter is a rejection letter or the belief or knowledge that there is a *significant chance* that it is a rejection letter.⁵ It is not based on knowledge that it is, in fact, a rejection letter. Funkhouser and Barrett acknowledge that many cases of evidence avoidance are amenable to this kind of explanation (Funkhouser & Barrett, 2016, p. 687). However, they insist that “at other times it seems almost undeniable that the person avoids evidence or situations because the person *unconsciously knows* what it would reveal”

(emphasis added) (*ibid.*, p. 687). In support of this claim, Funkhouser and Barrett first appeal to a study by Simpson, Ickes and Blackstone (1995) on motivated empathic inaccuracy. I will briefly explain Simpson *et al.*'s study and then discuss Funkhouser and Barrett's interpretation of the results.

Couples in romantic, heterosexual relationships were jointly presented with a slide show featuring images of males and females from what they were told was a "local 'dating pool'" (Simpson et al., 1995, p. 631). The couples were first asked to rate, and discuss together, the attractiveness of the males and females presented in the slide show. They were then separated and taken to rooms in which they independently viewed video recordings of their discussion. Each subject watched the tape twice. During the first viewing, the subject's task was to stop the tape when "he or she remembered having had a specific thought or feeling at that moment in the interaction" (Simpson et al., 1995, p. 633). They had to note down that thought or feeling as well as the exact time that it occurred. They also had to specify whether the "entry was positive, neutral, or negative in its overall emotional tone" (*ibid.*, p. 633). During the second viewing of the tape, each subject was presented with a form containing the information of the times at which *their partner* had a thought or feeling. On the basis of the behaviour exhibited by their partners in the recording, the subjects were asked to infer, and rate the emotional tone of, their partners' thoughts or feelings at those times. In general, the results suggested that subjects in the high threat condition (i.e. the condition in which the images presented to their partners in the slide show were of highly attractive individuals) tended to perform poorly on this mind reading task compared to those in the low threat situation (in which the pictures were of less attractive persons). According to Simpson *et al.*, this study evinces motivated empathic inaccuracy: Subjects inaccurately infer their partners' thoughts in

situations in which they perceive themselves, or their relationship, to be under threat. Funkhouser and Barrett offer the following interpretation:

This suggests that there is a specific, strategic avoidance of threatening information. When the stakes are low, the appropriate information is gathered and used to make a reliable judgment about one's partner's thoughts and feelings. When the stakes are higher, however, the information is simply ignored and their judgment is compromised. (Funkhouser & Barrett, 2016, p. 688)

According to Funkhouser and Barrett, then, the unconscious mind evaluates a piece of information in the subject's environment as threatening — namely, the attractiveness of the persons presented in the slide show (Funkhouser & Barrett, 2016, p. 688) — and directs the subject (or the conscious part of her mind) away from that information. This results in the subject failing to take that information into account when drawing inferences about her partner's mind. A review of Simpson *et al.*'s study reveals that this interpretation is contentious. Simpson *et al.*'s study is silent on the mechanisms responsible for motivated empathetic inaccuracy and their relation to *consciousness*. In the discussion section of their paper, the authors remark that “[a]nother major question for future research involves whether, or the extent to which, psychological mechanisms that produce motivated inaccuracy operate at the level of conscious awareness” (Simpson *et al.* p. 639). Simpson *et al.* do go on to say that “[i]t is conceivable that the selective information processing which takes place in conditions of high threat may occur outside of conscious awareness” (*ibid.*, p. 639). But the mere *conceivability* of unconscious, robust self-deception is not enough for Funkhouser and Barrett. Funkhouser and Barrett aim to present clear evidence that the unconscious mind *in fact* has the capacity to execute tactics of robust deception.

Funkhouser and Barrett might reply that their hypothesis offers the most *plausible* explanation of the data.⁶ Unfortunately, their interpretation is not significantly more plausible than alternatives. Consider, for instance, the explanation offered by the model of hypothesis testing which Mele appeals to — the Trope-Liberman Model (TLM) (Trope & Liberman, 1996).⁷ The basic idea of the TLM is this: The higher the perceived cost of a false positive (accepting a hypothesis when that hypothesis is false), the greater the quantity/quality of evidence required to accept the hypothesis; the higher the perceived cost of a false negative (rejecting a hypothesis when that hypothesis is true), the greater the quantity/quality of evidence required to reject that hypothesis. Importantly, since the perceived cost of a false positive can differ from the perceived cost of the corresponding false negative, the quantity/quality of evidence that the subject requires to accept a hypothesis might differ significantly from the quantity/quality of evidence that she requires to reject that hypothesis (Mele, 2001, p. 34). Now, for someone in a close monogamous relationship, it is plausible that the perceived cost of falsely accepting the hypothesis that their partner finds another person attractive is quite high (plausibly there will be immediate costs related to self-esteem), whereas the perceived cost of falsely rejecting that hypothesis is relatively low. If this is correct, then, compared to participants in the low threat condition, those in the high threat condition will require more evidence to attribute to their partners the thoughts or feelings ostensibly suggested by their behaviour. This would explain the between group asymmetry in empathic accuracy exhibited in Simpson *et al.*'s study.

It seems, then, that Simpson *et al.*'s results are amenable to motivationalist explanations that do not impute a motive or strategy to avoid information to those in the high threat condition. But there is also a plausible *non-motivational* explanation of the results — one that

Simpson *et al.* did not control for. It is plausible that persons in close, monogamous relationships have idealized beliefs about their partners or the factors that will contribute to the posterity of their relationship; they might believe, for instance, that their partner is unlikely to find other people attractive or that s/he is not attracted to 'conventionally' handsome or beautiful people (note that the pictures presented in the high threat condition were of college age models). If this is correct, then, for those in the high threat situation, there is conflict between their idealized beliefs and the thoughts or feelings that are ostensibly indicated or expressed by their partners' verbal or nonverbal behaviour. In the low threat condition, there is no such conflict. It might be, then, that the diminished empathic accuracy exhibited by those in the high threat condition is the result of an attempt to reconcile (false) idealized beliefs with recalcitrant evidence, rather than by motivated mishandling, or avoidance of, the relevant information.

The other study that Funkhouser and Barrett present in support of the reality of unconscious evidence avoidance is a study on the inhibitory effects of attention on eating (Polivy, Herman, Hackett, & Kuleshnyk, 1986). Participants in this study were left alone to taste, and record their ratings of, two kinds of candy. After completing their ratings, they were told that they could eat as many candies as they wished. Polivy *et al.* tested the differences in candy consumption between those in the following different conditions: The "public attention condition," the "self-attention condition," and the "no self-attention/no public attention condition" (Polivy *et al.*, 1986, p. 1255). Those in the public attention condition were instructed to leave the wrappers of consumed candies on the table; the accumulation of candy wrappers would make salient to the participant herself, and to the experimenter upon return, the amount of candies that she had eaten. In the self-attention condition, participants were instructed to allow the candy wrappers to accumulate on the table while eating, but to discard the wrappers

in a wastebasket, located a few yards away, when they had finished. These subjects would thus be attentive to the amount of candy that they had consumed but they would have no reason to think that the experimenter would be aware of this. Finally, in the no self-attention/no public attention condition, subjects were instructed to discard, in a wastebasket beside the table, each wrapper upon consumption of the candy it contained. Discarding the wrappers on a “candy-by-candy basis” ensured that “[n]o residual evidence regarding [the] exact amount consumed was available either to the subject or to the experimenter” (Polivy *et al.*, 1986, p. 1255). As predicted, those in the no self-attention/no public attention condition ate a greater number of candies than those in the other conditions.

Funkhouser and Barrett advance the following interpretation of these results:

[D]iscarding the evidence (wrappers) encouraged more eating by allowing the dieters to have the impression that they have thereby hidden the truth from others (e.g., the experimenter).... But we need not think that the dieters are, or at least must be, *consciously* sensitive to the evidence or how hiding it has misled their audience. This ability to unconsciously appreciate manipulations of the evidence in the service of deception can be turned inward. It is not hard to recognize how we can then be incentivized to unconsciously destroy or hide from evidence—emptying ashtrays or wearing baggy clothing—in order to better deceive ourselves.

(Funkhouser and Barrett, 2016, p. 688)

This requires some unpacking. First, notice that Polivy *et al.*’s study is not presented by Funkhouser and Barrett as *direct* evidence of unconscious information avoidance. Rather, they claim that the study supports the existence of an ability, or tendency, to unconsciously

appreciate and manipulate evidence, and that such an ability could be employed in the service of unconscious information avoidance. Second, Funkhouser and Barrett claim that, in the study in question, this ability is employed in the service of *interpersonal* deception: The subjects in the no self-attention/no public attention condition are presented as unconsciously manipulating evidence in order to deceive the experimenter.⁸ Specifically, Funkhouser and Barrett claim that these subjects consumed more candies than those in the other conditions because they unconsciously believed (or ‘had the impression’) that, by discarding the candies in the way they were instructed to, they could hide the extent of their candy consumption from the experimenter.

The above interpretation seems quite convoluted. Is it warranted? Closer scrutiny of the study suggests that it is not. To see this, consider Polivy *et al.*’s remarks in the discussion section of their paper:

It has generally been assumed that...monitoring or self-attention reduces the consumption of overeaters by informing them of the full extent of their eating, of which they are presumably otherwise unaware. Although the accuracy of self-reported consumption data from Study 2 could be said to support this interpretation, similar data from Study 1 suggest that something more is provided by self-attention. Postexperimental questioning in Study 1 revealed that 98% of the subjects, regardless of condition, could accurately report their consumption. The discarded candy wrappers in Study 1 were thus serving as something beyond a simple counting aid; they must have functioned prescriptively to induce adherence to regulatory norms.

(Polivy *et al.*, 1986, p. 1259)

As we can see, according to Polivy *et al.*, it is in virtue of failing to attend to regulatory norms that subjects in the no self-attention/no public-attention condition ate more candy than those in the other conditions. This failure to attend to regulatory norms is explained, they suggest, by the absence of any *cues* for such norms in this condition. In particular, the self-attention cue was absent: The extent of their candy consumption was not made salient to these subjects. And the public-attention cue was absent: These subjects had no reason to believe, and thus likely did *not* believe, that the experimenter would know how many candies they had eaten.⁹ The difference in restraint between subjects in the no self-attention/no public attention condition and subjects in the other two conditions is explicable in terms of the presence or absence of cues for regulatory norms. It is gratuitous, then, to maintain that those in the no self-attention/no public-attention condition were (consciously or unconsciously) under the impression that they could hide the evidence of how many candies they had eaten from the experimenter. No such assumption is necessary for the purposes of explanation.

Let's take stock of the discussion so far. In an effort to show that the unconscious is capable of strategically avoiding and hiding evidence in the service of deception, Funkhouser and Barrett present us with two studies: One on motivated empathic inaccuracy and another on the effects of self and public attention on eating. They claim that the former study demonstrates unconscious evidence avoidance, and that the latter demonstrates a "strategically sensitive [response] to evidence even when its significance is not consciously recognized" (Funkhouser & Barrett, 2016, p. 688). A review of these studies, however, shows Funkhouser and Barrett's claims are tenuous. Simpson *et al.*'s study is silent on the mechanisms that give rise to empathic inaccuracy. Moreover, it admits of plausible explanations which do not attribute strategic information avoidance to those in the high threat condition. Polivy *et al.*'s study on inhibited eating does not demonstrate unconscious sensitivity to evidence. Rather, it only shows that in

the absence of cues that make salient norms related to restraint and temperance, subjects are less likely to inhibit their eating behaviour.

3.2 The production of misleading evidence

Suppose that in a last ditch effort to make you believe that your partner is unfaithful, I spray one of her shirts with the kind of perfume used by her attractive colleague, and leave it in a place where you are likely to find it. Funkhouser and Barrett claim that the unconscious is similarly capable of producing misleading evidence. In support of this claim, they present research on vocal pitch adjustment:

Fraccaro et al. (2011); Hughes, Farley, and Rhodes (2010); and Leongomez et al. (2014) each show that people alter the pitch of their voice when conversing with others they find attractive. Men speak with a lower pitch when speaking to females they find attractive, while women speak with a higher pitch when speaking to males they find attractive. This counts as misleading evidence because women and men both have (different) pitch-level preferences for their partners.... And, of course, there is no reason for thinking that on all or even most occasions men and women *consciously* alter their voices.

Again, take note of the selectivity of the phenomenon. People do not indiscriminately alter their voice pitch when around members of the opposite sex; rather, they alter their voices only in those circumstances where the other is an attractive member of that sex. Thus the unconscious processes underlying this [sic] phenomena exhibit flexibility to a changing environment.

(Funkhouser & Barrett, 2016, p. 688-689)

The phenomenon described above is clearly flexible and situation-sensitive; subjects alter their vocal pitch in a specific way in specific situations. Moreover, perhaps it is implausible to think that this is something that people do consciously. But is this response *strategic*? Note that the mere flexibility of a response to specific situational factors is not sufficient for its being strategic. For the automatic response described by Gur and Sackeim (1979) is flexible in just this way: Subjects exhibit increased GSR in response to video and audio recordings of *familiar* voices (even when they do not consciously recognise those voices). Just as the pitch altering response is flexible to the *attractiveness* of the stimulus, so the response demonstrated by Gur and Sackeim is flexible to the *familiarity* of the stimulus.

The pitch altering response must meet other conditions if it is to count as strategic. In particular, it must be shown that the response is “incorporated into...rational goal pursuit” (Funkhouser & Barrett, 2016, p. 685-686) — the goal here presumably being to mislead persons one finds attractive as to the typical pitch of one’s voice. But Funkhouser and Barrett do not show this. What they offer is evidence of an unconscious, discriminative response which tends to mislead listeners as to the typical pitch of the subject’s voice. The fact that this response tends to do this (combined with the fact that listeners have pitch-preferences) might well explain why this response came about or was selected for. That is, its tendency to mislead prospective mates in this way might constitute a plausible *ultimate explanation* of this phenomenon — i.e., an explanation that refers to environmental factors “the causal influence of which is spread over generations” (Fedyk and Kushnir, 2014, p. 143). Funkhouser and Barrett, however, are making a claim about the *proximate explanation* of vocal pitch alteration — in particular, they claim that instances of this phenomenon result from the unconscious goal

to mislead. However, as Fedyk and Kushnir explain, it is fallacious to infer a proximate explanation of a trait or behaviour from an ultimate explanation of that trait or behaviour. This is because for any ultimate explanation, there is a set of proximate explanations, any one of which might be true. The point is made vivid by the following example:

[I]magine a pattern of behaviour that is adaptive for some species in some environment, like cooperating with conspecifics. Then list several different proximate causes of the pattern of behaviour, like inherent prosociality or reputation management. Natural selection will not care which of these causes is the right one, and so generally speaking there is no inference from the fact that a pattern of behaviour is adaptive to any particular proximate conclusion.

(Fedyk and Kushnir, 2014, p. 143)

In sum, the vocal pitch alteration response does not seem to be any more strategic than the response presented by Gur and Sackeim. And an attempt to defend the claim that it is strategic, or that it aims at deception, on the basis of the claim that it was selected for because of its propensity to deceive prospective mates is to fallaciously infer a proximate explanation from an ultimate explanation.

3.3 Selective attention to favourable information

Another common tactic of deception involves directing the victim's attention towards information which supports the relevant belief and away from information which conflicts with it. Funkhouser and Barrett claim that the unconscious is capable of selectively screening

information in this way. In this connection, they cite a paper by Jiang, Costello, Fang, Huang and He (2006) on the effects of subliminally presented sexual images on attention.

Jiang *et al.*'s study can be summarized as follows. Invisible images were presented to each side of the subject's visual fixation point: An erotic image on one side and a scrambled, meaningless image on the other.¹⁰ A Gabor patch was then briefly presented to either the left or right side of fixation. Jiang *et al.* showed that subjects performed better at judging the orientation of the Gabor patch (clockwise or anti-clockwise) when it was presented on the side of fixation in which the suppressed erotic image appeared, compared to when it was presented on the side in which the control appeared. This suggests that unconscious perception of an arousing stimulus in a particular part of the visual field can result in the direction of one's attention to that part of the visual field. Funkhouser and Barrett extrapolate on these findings as follows:

The subjects flexibly respond to their environment (male or female nudes) based on their sexual goals (viewing male or female nudes). Attention is not captured simply by anything that appears, but is preferentially attracted to specific stimuli and selectively ignores others.

(Funkhouser & Barrett, 2016, p. 689)

Two features of Funkhouser and Barrett's remarks are striking. First, notice that they attribute to the subjects in these studies an unconscious sexual *goal* — namely, the goal of viewing male or female nudes. Second, they describe the subject as unconsciously *ignoring* certain stimuli (the scrambled control image). The unconscious is described as having the goal of viewing (or, perhaps more precisely, the goal of making the *conscious part* of the subject

view) male or female nudes. In pursuit of this goal, it directs the subject's attention to the part of their visual field where it (the unconscious) last saw the arousing stimulus, and it also directs attention away from, or 'selectively ignores,' those parts of her visual field in which less interesting stimuli were perceived. Funkhouser and Barrett claim that the applicability of this response, so described, to self-deception "should be obvious" (Funkhouser & Barrett, 2016, p. 689). Presumably the idea is that just as the unconscious can direct the subject's attention to certain aspects of her visual field in the service of the goal of viewing erotic images, so can it direct the subject's attention to certain aspects of her visual field to facilitate the goal of making her believe some false proposition.

This goal-oriented characterisation of unconscious selective attention is unnecessary for the purposes of explanation.¹¹ To see this, consider an analogy with conscious attention. Suppose that, upon entering his kitchen, Jones (consciously) sees a mouse scuttle across the kitchen floor and into a hole in the wall. This arousing stimulus — the scuttling of the mouse across the floor — draws Jones's attention to the floor and the hole the mouse hurried into. His attention to that area of the floor is explicable in terms of his perception of the arousing stimulus that just occurred and its location. It does not require, in addition to this, that Jones has the *goal* of seeing the mouse. Analogously, the direction of a subject's attention to a part of her visual field in which a suppressed image appeared, does not require that she have the unconscious goal of seeing that arousing image (or images of that kind) in that area. It is plausible that the subject's attention is drawn to that area in the same way that Jones's attention is drawn to the kitchen floor. Funkhouser and Barrett's characterisation of the response as goal-oriented, then, lacks motivation.

The description of the response as one that *selectively ignores* certain stimuli also lacks motivation. Consider again the example of the kitchen floor. However, suppose that this time there is no mouse. Jones enters the kitchen and makes his way towards the kettle. In this scenario he does not attend to the part of the floor or the hole in the wall that he attended to in the previous scenario. Jones does not attend to that location because there is nothing there to arouse his attention. But he does not selectively *ignore* that part of the floor or the mouse hole. Rather, he just fails to give them any consideration at all. There is a difference between selectively ignoring something and failing to attend to it. Selectively ignoring something implies some kind of activity, whereas failing to attend to something implies inactivity. Funkhouser and Barrett's description of the unconscious as selectively ignoring certain stimuli (such as the scrambled control image), then, seems to rest on a confusion between the failure to attend to some phenomenon and attending, but refusing to give further consideration, to that phenomenon.

In sum: The evidence presented by Funkhouser and Barrett shows that unconsciously perceived cues can result in selective attention. However, it does not show that such cues trigger goal-oriented responses of the kind required for robust deception.

3.4 Biased interpretation of information

The presentation of evidence that provides an incomplete, and thereby misleading or distorted, picture of reality is a familiar strategy of deception. According to Funkhouser and Barrett, "this kind of deceptive tactic is also found in unconscious processes" (2016, p. 689). They shore up this claim with a study by Aarts, Custers, and Veltkamp (2008) on unconscious goal activation. Aarts *et al.*'s study purports to shed light on "how the mental apparatus directs and motivates

people to either pursue a primed goal or not in the absence of conscious will” (Aarts *et al.*, 2008, p. 556). Their experiments indicate that the pursuit of unconsciously activated goals is facilitated by effects on *perception*. As Funkhouser and Barrett explain,

[S]ubjects were presented subliminally with goal-related words (those having to do with the completion of a puzzle) and with visibly presented positive or neutral words. Then they were asked to judge the height of a crossword puzzle or puzzle book. Remarkably, the participants given the achievement words judged the objects to be taller than participants who were given neutral words.

(Funkhouser & Barrett, 2016, p. 689-690)

We can see how this kind of phenomenon might be useful in less artificial settings. In particular, plausible speculations as to its evolutionary advantage are readily available: If the goal to evade predators causes potential prey to overestimate the size of a predator in their environment, this might allow them to better ready themselves for evasion. But what has this to do with biased interpretation and self-deception? To demonstrate its relevance, Funkhouser and Barrett invoke the example of a man who attempts to impress his date by presenting an incomplete, rose-tinted picture of his character. They claim that “[o]ur perceptual faculties...mislead us by making certain features more salient than others, just as the single man misleads by making his positive traits more salient to his date” (Funkhouser & Barrett, 2016, p. 690).

Let’s make explicit the analogy that Funkhouser and Barrett are drawing here. The single man’s primary goal is to impress his date. He does this by presenting her with evidence that paints a distorted picture of his character. By the same token, in cases of unconscious goal activation, subjects have a certain unconscious primary goal (to complete a puzzle) and they

pursue this goal by providing themselves with a distorted picture of their environment. Just as the single man deceives his date about the nature of his character in order to achieve some goal (namely, the goal of impressing his date), the unconscious mind deceives the conscious mind regarding the magnitudes of objects in its environment in order to achieve some goal (the completion of a puzzle).

Closer scrutiny suggests that the analogy is not as clean as Funkhouser and Barrett would like us to suppose. Recall that the single man's primary or ultimate goal is to present a good first impression. And that he attempts to achieve this goal by achieving another, instrumental goal — the goal of presenting to her a curated, rose-tinted picture of his character. It is in virtue of pursuing this instrumental goal that he qualifies as engaging in deception. The problem is that it is not clear that the unconscious can be attributed the analogous instrumental goal. That is, it is not clear that the process by which unconsciously primed goal-concepts result in misperceptions (perceiving objects as larger than they are) should *itself* be understood in goal-directed terms. Rather, it seems that the primed goal-concept simply affects the salience of aspects of the subject's environment in a way that facilitates the achievement of goals associated with the relevant concept (e.g. completing puzzles). Again, the evidence falls short of demonstrating strategic, unconscious self-deception.

3.5 Rationalization

Rationalization consists in explaining away threatening information. As Funkhouser and Barrett explain, “[i]t is usually taken that rationalization is *conscious* activity that happens in response to unconscious motivation unexpectedly manifesting itself in the subject's behavior” (Funkhouser & Barrett, 2016, p. 690).¹² Still, they maintain that “there is some evidence that

unconscious rationalization occurs” (*ibid.*, p. 690). In support of this claim, Funkhouser and Barrett present a study on hidden motives by Snyder, Kleck, Strenta, & Mentzer (1979).

Participants in Snyder *et al.*’s study were given the choice to view one of two televisions, each of which was showing a short movie. The televisions were in identically designed rooms which were divided by a partition. In one room was an able-bodied confederate, posing as another participant, and in the other was another confederate, posing as a disabled participant. There were two conditions: The *same-movie condition* (each TV played the same movie) and the *different-movie condition* (each played a different movie). The results showed a marked difference between the choices of the subjects in these conditions: The majority of those in the different-movie condition (83%) chose to view the television that was also being viewed by the able-bodied person. That is, they avoided the person who appeared to be disabled. In the same-movie condition, however, there was a slight preference (58%) for viewing the movie with the disabled-appearing person (Snyder *et al.*, 1979, p. 2300). Funkhouser and Barrett advance the following interpretation:

[S]ubjects were more likely to choose to watch along with the able-bodied person in the real choice case because they had an excellent fall-back plan in case of any doubts about the subjects’ seat preference: simply rationalize that they were more interested in the movie shown on that television.

(Funkhouser & Barrett, 2016, p. 690)

Funkhouser and Barrett then make the further claim such rationalization, or the plan to rationalize, is *unconscious*: “The unconscious must rationalize on the basis of plausible features of the environment or subject” (Funkhouser & Barrett, 2016, p. 690). This is highly

contentious. To see this, consider first the goal of Snyder *et al.*'s study. The purpose of the study was to "illustrate a general strategy for detecting motives that people wish to conceal" (Snyder *et al.*, 1979, p. 2303). In the case of their study in particular, the authors assumed (1) that people are motivated to avoid disabled people, and (2) that they wish to conceal this motive. Snyder *et al.*'s goal was to create conditions in which the concealed motive could be plausibly inferred from subjects' choices. And they did this by arranging things, in the different-movie condition, so that the expression of that motive (avoiding the disabled person) was co-extensive with the expression of a more benign motive (watching one movie as opposed to another). The prediction was that subjects would take advantage of the resulting ambiguity, and act on their concealed motive. Whether these subjects were, in fact, acting on the concealed motive could then be determined by the overall pattern of behaviour of the participants in this condition — and with 83% of subjects in the first experiment choosing the movie that the able-bodied person was watching, we should accept that these subjects acted on the unsavoury motive.

Now, when asked to explain their actions, these subjects did appear to engage in rationalization — specifically, as predicted, the majority of subjects in the different movie condition cited a preference for the particular movie as the reason for their choice (Snyder *et al.*, 1979, p. 2302). But this might just show that these subjects are unwilling to admit to the experimenter their true motives — which is just what we would expect when those motives are objectionable. Note, also, that the subject had to *talk* with the experimenter about the reasons for her choice (*ibid.*, p. 2302). It is plausible that this face to face encounter compounded these subjects' desires to conceal and rationalize their true motives. We are presented with no reason to think that these subjects were unconscious of their concealed motives. Nor are we given reason to think that their rationalizations of their choices were unconscious.

In sum, Snyder *et al.*'s experiments do not show, nor were they designed to show, that subjects unconsciously rationalize choices when those choices are based on objectionable motives. The experiments also do not show that subjects have the unconscious *plan* to rationalize their choices. It is perfectly compatible with the results, and indeed plausible, that subjects in the different-movie condition avoided the disabled person because they *consciously* believed that the basis of their choice would be ambiguous. Again, it seems that the evidence presented by Funkhouser and Barrett provides insufficient support for the reality of the unconsciously executed tactic of deception.

3.6 Memory Biases

Clearly, memory loss can give rise to false beliefs. Funkhouser and Barrett present what they claim is evidence of “*strategic* memory loss,” wherein the subject “unconsciously knows exactly what to forget” (emphasis added) (Funkhouser & Barrett, 2016, p. 691). Specifically, they appeal to a study on mood congruent memory (MCM) biases (Watkins, Vache, Verney, Muller, & Mathews, 1996).

MCM bias is the tendency to remember events or facts that are congruent with one's mood. Watkins *et al.* tested differences in implicit recall between depressed and non-depressed subjects. Subjects were first exposed to a word set comprising even numbers of positive, negative and neutral words. These subjects were tasked with imagining themselves in a scene related to the word that appeared on-screen (Watkins *et al.*, 1996, p. 36). Later in the study, after a distraction task, subjects were presented with cue words. Their task was to produce, for each cue word, as many one-word associations as they could think of (within a time limit of 30

seconds per cue word). Researchers then measured the between group difference in recollection of the words presented in the imagination task. The production of words in the free association task is supposed to demonstrate *implicit* recollection, as subjects were led to believe that they were engaging in a free association task as opposed to a recall task. That is, they are not *asked* to recall words from the previous task.

The results showed that depressed subjects associated, and thereby remembered, a greater number of negative words from the first task and a smaller number of positive words compared to non-depressed subjects. It seems, then, that mood affects the content of implicit recollection. As Funkhouser and Barrett remark, “this sort of partiality could deceive a depressed person into, say, interpreting a situation as far worse than it actually is by remembering selectively only the sad events that surround it” (Funkhouser & Barrett, 2016, p. 691). This is eminently plausible. Indeed, Watkins *et al.*’s study is specifically concerned with the role of MCM in the maintenance of depression (see Watkins *et al.*, 1996, p. 34). But there is a difference between a bias’s being such that it *causes* (or tends to cause) false beliefs, and its being such that, when activated, the subject (or the relevant part of her) can be described as engaging in deception (recall the strategy, flexibility and truth-retention conditions). This is the difference that Funkhouser and Barrett are at pains to emphasise when distinguishing their account from motivationalist accounts such as Mele’s (see Funkhouser & Barrett, 2016, p. 684). MCM may well cause false beliefs. But does the subject (or the unconscious part of her), in succumbing to that bias, engage in deception? A modest response is that the study cited by Funkhouser and Barrett is agnostic on this question. For that study only evinces a phenomenon (in this case, implicit MCM bias) without shedding light on the mechanisms responsible for it.

3.7 Acting As-If

The final tactic that Funkhouser and Barrett offer as a putative means of self-deception is that of acting as if one's preferred proposition is true: "We can...deceive ourselves by simply acting as-if what we desire to believe is true" (Funkhouser & Barrett, 2016, p. 291). I take it that by 'acting as if what one desires to believe is true', Funkhouser and Barrett mean acting as if one *believes* some desired proposition (or believes that some desired proposition is true). Acting as if one believes a proposition, p , then, consists in engaging in the kinds of behaviour typical of someone who believes that p . How can belief imitation of this kind give rise to false belief? Funkhouser and Barrett gesture towards an answer: "we might 'infer' that there is a god from our as-if behavior to the effect that there is a god (say, by praying and participating in religious rituals)" (*ibid.*, p. 691). They then point to the well documented phenomenon of unconscious imitation (see Chartrand & Bargh 1999), in an attempt to show how one might unconsciously engage in belief imitation:

We unconsciously imitate the behavior and mannerisms of those around us.... Since we can unconsciously imitate the distinctive behaviors of a group or role model, it is reasonable that we could unconsciously imitate in the pursuit of an ideal (e.g., attractiveness) or favored situation (e.g., belief in a god).

(Funkhouser & Barrett, 2016, p. 691).

The idea, then, is that one can come to falsely believe that p on the basis of an inference from certain (unconsciously executed) imitation behaviour.

There are a number of problems with this proposal. First, consider the inference that such subjects are supposed to make: "we might 'infer' that there is a god from our as-if behavior

to the effect that there is a god” (*ibid.*, p 691). The inference here is from $\langle S \text{ acts as if she believes that } p \rangle$ to $\langle p \rangle$. It seems that this kind of inference is only warranted in very peculiar circumstances. Suppose, for instance, you know that Jones knows whether or not some proposition, p , is the case. Suppose also that you know that Jones’s behaviour is a reliable indicator of what he believes. It seems that, in this scenario, you are warranted to infer $\langle p \rangle$ from Jones’s behaving as if he believes that p . However, it is hard to imagine how this kind of inference could be warranted when it is made about oneself, on the basis of one’s own behaviour. One might reply that the fact that an inference is unwarranted does not entail that it is psychologically impossible; we are not ideally rational agents, after all, and are thus liable to make unwarranted inferences. But irrational inferences of this kind nonetheless require explanation. The burden is on Funkhouser and Barrett to show how the subject manages to do this. No such explanation is offered in their paper.

The other problem concerns the imitation behaviour on the basis of which the subject is supposed to make this inference. Funkhouser and Barrett seem to make the following claim. When in an environment in which one’s peers act as if they believe that p , one might unconsciously imitate this behavior; and this imitation behavior might then lead one to infer that p . Although collective self-deception is certainly possible, it is not the paradigm case; what is striking about self-deception is that the subject seems to believe her preferred proposition *in spite of* the protestations of her peers. The acting-as-if tactic would thus have limited explanatory scope. Moreover, notice that Funkhouser and Barrett’s claim is, in fact, much stronger than the characterisation just given. It is not that the subject simply falsely infers that p on the basis of acting as if she believes that p . Rather, the claim is that the imitation behaviour is unconsciously executed *in pursuit of the goal of acquiring that belief*. “we could unconsciously imitate [behaviour] in the pursuit of...[a] favored situation (e.g., belief in a

god)” (*ibid.*, p. 691). This characterisation is important for the purposes of Funkhouser and Barrett’s thesis. This is because, for the deception to count as strategic, the unconscious must be attributed the goal of making the subject acquire the relevant belief. But no case is made for the claim that the propensity to mimic others’ behaviour can be co-opted in pursuit of goals in this way. Indeed, in the paper cited by Funkhouser and Barrett, this mimicry is understood as an *automatic* phenomenon — one that is caused by an automatic perception-to-behaviour mechanism (see Chartrand & Bargh, 1999, p. 893-894).

In sum, Funkhouser and Barrett’s claim that we have the capacity to deceive ourselves by unconsciously engaging in belief mimicry is dubious for the following reasons. First, it attributes to subjects a disposition to make bizarre inferences — namely, inferences from propositions of the form $\langle S \text{ acts as if she believes that } p \rangle$ to those of the form $\langle p \rangle$. Second, it requires a strategic characterisation of unconscious mimicry which is not supported by the relevant evidence.

4. Conclusion

Funkhouser and Barrett appeal to a substantial body of empirical literature in an attempt to revive what was once the dominant account of self-deception: the mental partitioning account. In contrast to the now orthodox motivationalist account of self-deception, the mental partitioning account preserves the close analogy between self and interpersonal deception, and thereby describes a phenomenon that is deceptive in an intuitively satisfying sense. Arguably, this account is too quickly dismissed in contemporary discussions of self-deception. At the very least, then, Funkhouser and Barrett should be commended for putting it back on the table.

Unfortunately, however, the account lacks the empirical support that Funkhouser and Barrett attribute to it. Specifically, as argued in the preceding sections, for each study presented by Funkhouser and Barrett, the phenomenon illustrated therein either fails to clearly qualify as unconscious or fails to satisfy the conditions for robust deception.

It would be premature to conclude that there is no such thing as robust, unconscious self-deception — or that the mental partitioning account is false. The most cautious conclusion to be drawn is that this picture of self-deception lacks the necessary support from the relevant empirical literature.

¹ I do not manipulate your beliefs by using drugs, hypnosis or brain surgery, for instance (see Lynch, 2014).

² The mental partitioning account was sketched in preliminary form by Donald Davidson (1982, 1986) and developed by David Pears (1984, 1991).

³ How are we to square the strategizing-as-deliberation reading with the claim that such strategizing need not be ideal or deliberate? If, by ‘deliberate,’ we take Funkhouser and Barrett to mean any process that involves deliberation, then the reading advanced here must be wrong.

The adjective ‘deliberate,’ however, has various senses. It can be used to describe the temporal property of a process of reasoning — as “not hasty or rash” (*OED Online* 2017). It is also synonymous with ‘careful’ and ‘considered’. When ‘deliberate’ is understood in either of the latter ways, the claim that strategizing need not be deliberate is compatible with the strategizing-as-deliberation reading.

⁴ I focus primarily on the strategy condition, as this is the condition that figures most prominently in Funkhouser and Barrett’s discussion of the evidence presented in their paper.

⁵ For a detailed discussion and analysis of this phenomenon, see Lynch (2016).

⁶ Thanks to an anonymous reviewer for pressing this point.

⁷ The TLM underpins Mele’s account of the process of self-deception. For detailed discussion, see Mele (2001, p. 31-46).

⁸ Funkhouser and Barrett seem to maintain that if the unconscious has the capacity to engage in interpersonal deception, then it is plausible that it has the capacity to engage in self-deception also. I do not question this inference here. Rather, what I want to show is that there is no reason to suppose that subjects in Polivy *et al.*’s study engaged in unconscious interpersonal deception.

⁹ Notice that the fact that one *fails to believe* some proposition does not entail that one believes the *negation* of that proposition. So we need not (and indeed should not) assume that these

subjects also (consciously or unconsciously) *believed* that the experimenter would not know how many candies they had eaten.

¹⁰ The images were rendered invisible by interocular suppression . For details, see Jiang *et al.*, (2006, p. 17048).

¹¹ Jiang *et al.* themselves do not mention any such goals.

¹² See Wilson and Nisbett's famous stockings experiment (Nisbett and Wilson 1977).

References

Aarts, H., Custers, R., & Veltkamp, M. (2008). Goal priming and the affective-motivational route to nonconscious goal pursuit. *Social Cognition*, 26, 555–577.

Audi, R. (1997). Self-deception vs. self-caused deception: A comment on professor Mele. *Behavioral and Brain Sciences*, 20, 104-104.

Chartrand, T., & Bargh, J. (1999). The chameleon effect: The perception-behavior link and social interaction. *Journal of Personality and Social Psychology*, 76, 893–910.

Davidson, D. (1982). Paradoxes of irrationality. In R. Wollheim (Ed.) *Philosophical Essays on Freud* (pp. 289-305). London: Cambridge University Press.

— (1986). Deception and division. In J. Elster (Ed.) *The Multiple Self* (pp. 70-92). Cambridge: Cambridge University Press.

deliberate, adj. (2017). *OED Online*. Oxford: Oxford University Press. Retrieved from <http://www.oed.com>.

Fraccaro, P., Jones, B., Vukovic, J., Smith, F., Watkins, C., Feinberg, D., Little, A., & DeBruine, L. (2011). Experimental evidence that women speak in a higher voice pitch to men they find attractive. *Journal of Evolutionary Psychology*, 9, 57–67.

Funkhouser, E. & Barrett, D. (2016). Robust, unconscious self-deception: Strategic and flexible. *Philosophical Psychology*, 29, 682–696.

Fedyk, M. & Kushnir, T. (2014). Development links psychological causes to evolutionary explanations. *Behavioral and Brain Sciences*, 37 (2):142-143.

Gur, R., & Sackeim, H. (1979). Self-deception: A concept in search of a phenomenon. *Journal of Personality and Social Psychology*, 37, 147–169.

Hughes, S., Farley, S., & Rhodes, B. (2010). Vocal and physiological changes in response to the physical attractiveness of conversational partners. *Journal of Nonverbal Behavior*, 34, 155–167.

Jiang, Y., Costello, P., Fang, F., Huang, M., & He, S. (2006). A gender-and sexual orientation-dependent spatial attention effect of invisible nudes. *Proceedings of the National Academy of Sciences*, 103, 17048–17052.

Leongomez, J. D., Binter, J., Kubicova, L., Stolarova, P., Klapilova, K., Havlicek, J., & Roberts, S. C. (2014). Vocal modulation during courtship increases proceptivity even in naive listeners. *Evolution and Human Behavior*, 35, 489–496.

Losonsky, M. (1997). Self-Deceivers' intentions and possessions. *Behavioral and Brain Sciences*, 20, 121–122.

Lynch, K. (2016). Willful ignorance and self-deception. *Philosophical Studies*, 173, 505-523.

— (2014). Self-deception and shifts of attention. *Philosophical Explorations*, 17, 63–75.

— (2013). Self-deception and stubborn belief. *Erkenntnis*, 78, 1337–1345.

— (2012). On the ‘tension’ inherent in self-deception. *Philosophical Psychology*, 25, 433–450.

Mele, A. (2001). *Self-deception unmasked*. Princeton, NJ: Princeton University Press.

Nelkin, D. (2002). Self-deception, motivation, and the desire to believe. *Pacific Philosophical Quarterly*, 83, 384–406.

Nisbett, R. E., & Wilson, T. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, 84, 231–259.

“deliberate, adj.”. *OED Online*. Oxford University Press. Retrieved from <http://www.oed.com>.

Pears, D. (1984). *Motivated irrationality*. Oxford: Oxford University Press.

Pears, D. (1991). Self-deceptive belief-formation. *Synthese*, 89, 393–405.

Polivy, J., Herman, P., Hackett, R., & Kuleshnyk, I. (1986). The effects of self-attention and public attention on eating in restrained and unrestrained subjects. *Journal of Personality and Social Psychology*, 50, 1253–1260.

Porcher, J. (2012). Against the deflationary account of self-deception. *Humana. Mente*, 20, 67-84.

Scott-Kakures. (2002). At 'permanent risk': reasoning and self-knowledge in self-deception. *Philosophy and Phenomenological Research*, 65, 576-603.

Simpson, J., Ickes, W., & Blackstone, T. (1995). When the head protects the heart: empathic accuracy in dating relationships. *Journal of Personality and Social Psychology*, 69, 629–641.

Snyder, M., Kleck, R., Strenta, A., & Mentzer, S. (1979). Avoidance of the handicapped: An attributional ambiguity analysis. *Journal of Personality and Social Psychology*, 37, 2297–2306.

Trope, Y. & Liberman, A. (1996). Social hypothesis testing: cognitive and motivational mechanisms.” In E. Higgins and A. Kruglanski (Eds.) *Social Psychology: Handbook of Basic Principles* (pp. 239–70). New York: Guilford Press.

Watkins, P., Vache, K., Verney, S., Muller, S., & Mathews, A. (1996). Unconscious mood-congruent memory bias in depression. *Journal of Abnormal Psychology*, 105, 34–41.