


Happy Egrets Strike Back?

Francesco Orsi

In: Garcia, A., Gunnemyr, M. & Werkmäster, J. (2023) *Value, Morality & Social Reality: Essays dedicated to Dan Egonsson, Björn Petersson & Toni Rønnow-Rasmussen*. Lund: Department of Philosophy, Lund University. DOI: <https://doi.org/10.37852/oblu.189>

ISBN: 978-91-89415-65-2 (print), 978-91-89415-66-9 (digital)



Published by the Department of Philosophy, Lund University.
Edited by: Andrés Garcia, Mattias Gunnemyr, and Jakob Werkmäster
Cover image by Fabian Jones. Cover layout by Gunilla Albertén.

DOI: <https://doi.org/10.37852/oblu.189.c529>



This text is licensed under a Creative Commons Attribution-NonCommercial license. This license allows reusers to distribute, remix, adapt, and build upon the material in any medium or format, so long as attribution is given to the creator. The license does not allow for commercial use.

(License: <http://creativecommons.org/licenses/by-nc/4.0/>)

Text © Andrés Garcia, Mattias Gunnemyr, and Jakob Werkmäster 2023.
Copyright of individual chapters is maintained by the chapters' authors.

Happy Egrets Strike Back?

Francesco Orsi¹

1. Introduction

The fitting attitude account of value (FA) claims that to be good or bad is to be a fitting target of a pro-attitude or a fitting target of a con-attitude. Toni Rønnow-Rasmussen has recently defended FA from a new version of what is variously called the solitary goods objection (Bykvist 2009), the wrong kind of value problem (Reisner 2015), or the too little value problem (Rowland 2019, chapter 7):² there seem to be objects or states of affairs which are good (or bad), but it is not the case that it is fitting for anyone to favour (or disfavour) them. If the objection is correct, then, contrary to what FA holds, for x to be good or bad cannot be for x to be a fitting target of a pro-attitude or a fitting target of a con-attitude. In this contribution I argue that advocates of FA have a better reply to give to the new version of the solitary goods objection than Rønnow-Rasmussen's somewhat defeatist defence. (For the record, I say this as someone who has often been on the side of those who are sceptical about FA, see Orsi & Garcia 2021, and explored alternatives to it, see Orsi 2013b. But I do find the solitary goods objection to FA unconvincing, as I did in Orsi 2013a.)

A typical example of the solitary goods objection asks us to consider a state of affairs that seems good, but which by its very nature implies that no one is in a position to have a fitting attitude towards it:

Happy Egrets: there being happy egrets but no past, present or future agents (Bykvist 2009: 5).

¹ This is a nod to “The Strike of the Demon” (Rabinowicz and Rønnow-Rasmussen 2004), which contains a section titled “The Demon Strikes Back” (pp. 419 ff.).

² Dancy (2000) should be credited for first stating the problem. Bykvist (2015) replies to Orsi (2013a).

No one *in the same world* where *Happy Egrets* obtains was, is, or will be in a position to favour *Happy Egrets*. One might conclude that FA is false, because *Happy Egrets* is (or at least can be) clearly good even if it is not fitting for anyone to favour it.

A natural reaction is to say that at least it is fitting for us, contemplating *Happy Egrets* from the actual world, to favour it, for example by taking contemplative pleasure in it (Orsi 2013a). However, Kent Hurtig (2019) has recently argued that this kind of reply may not always work. In particular, when a state of affairs akin to *Happy Egrets* is indexed to the actual world, it cannot be the case that it is fitting for subjects in a *non-actual world* to favour or disfavour that actual state of affairs, with the result that such states of affairs may be good (or bad), without anyone's attitudes being fitting towards them.

In what follows I first articulate Hurtig's argument—making it more precise, if possible, than Hurtig himself does. Then I discuss Rønnow-Rasmussen's response to it and show why it is somewhat defeatist. Finally, I provide a response to Hurtig that illustrates a broader point about why arguments from solitary goods against FA are doomed to fail: if—due to their location in modal space—the relevant states of affairs cannot even be evaluated as good or bad (and a fortiori favoured or disfavoured) by readers, then such cases are dialectically powerless against FA; if, on the other hand, they can be evaluated as good or bad—despite their location in modal space—then they can also be favoured or disfavoured, and it will be fitting for us (if for no one else) to favour or disfavour them, thus defusing the challenge.

2. World-specific Values

Hurtig argues that FA fails to account for the value of a state of affairs such as this:

S: [*p* is a significant true proposition, and no one in the actual world at any time has any attitude toward *p*] (3245)³

Hurtig suggests that *S*'s actually obtaining is bad for its own sake, presumably because if *p* is a significant true proposition, it would be good to know that *p*, and a fortiori it would be good that someone in the actual world had some attitude toward *p*.⁴ In order to account for the value of *S*, FA must find a suitable truth or fact about fitting attitudes towards *S*, for example, the fact that it is fitting to disfavour *S* for its own sake. What are the available candidates?

³ All page-only references are to Hurtig (2019). I have explicitly included 'is a significant true proposition' to Hurtig's own formulation, because he himself describes *p* as a significant true proposition (3245).

⁴ I say "presumably" because Hurtig himself appears to just stipulate the badness of *S*.

Happy Egrets Strike Back

Since it is a feature of *S* that no one in the actual world has any attitude towards *p*, and *S* includes *p*, it follows that no one in the actual world can have an attitude of disfavour towards *S*, because if they did, then they would have an attitude towards *p* as well, be that as non-committal an attitude as merely entertaining *p*. In other words: there is no coherent scenario where a subject is both part of the actual world *and* disfavours *S*. And if there is no such coherent scenario, then it is not possible to actually disfavour *S*. On the assumption that fittingness implies can, it follows that *actually* disfavours *S* cannot be the fitting attitude towards *S*.⁵

The natural alternative is to say that it is fitting for a *non-actual* subject, i.e. for a subject existing elsewhere than in the actual world, to disfavour the actual state of affairs *S*. Given the content of *S*, only a non-actual subject could have some attitude towards *p*, and thus towards *S*. This is analogous to the move I suggested (Orsi 2013a) in response to *Happy Egrets*. Since in the world of *Happy Egrets* it is not fitting for anyone to favour *Happy Egrets*, FA can only locate fitting responses to *Happy Egrets* in a world where *Happy Egrets* does not obtain, for example, in our own world. Hurtig's case would seem to be the reverse of that: as he writes, "the evaluating—if there is to be any at all—*has* to take place from a non-actual world" (3247, his italics). In both cases, it seems that FA will resort to what has been called trans-world fittingness (Reisner 2015).

Hurtig, however, rejects the idea that it can be fitting for a *non-actual* (i.e. counterfactual) subject to disfavour the *actual* state of affairs *S*. Here is a reconstruction of his argument (3247):

- P1. If it is fitting for a non-actual subject N to disfavour the actual state of affairs *S*, then N must be able to disfavour the actual state of affairs *S*.
- P2. In order for N to be able to disfavour the actual state of affairs *S*, N's evaluation must be able to uniquely be about *S*'s obtaining in the actual world.
- P3. There is no causal link between the actual world and N's world.
- P4. If there is no causal link between the actual world and N's world, then N's evaluation cannot uniquely be about *S*'s obtaining in the actual world.
- C1. Therefore, N's evaluation cannot uniquely be about *S*'s obtaining in the actual world.
- C2. Therefore, N is not able to disfavour the actual state of affairs *S*.
- C3. Therefore, it is not fitting for a non-actual subject N to disfavour the actual state of affairs *S*.

⁵ A reviewer pointed out that, e.g., A. C. Ewing did not accept "fittingness implies can". But the "can" in this case is one of logical possibility.

And of course, if *S* is bad for its own sake, but it is not fitting for *N* (or any other actual or non-actual subject) to disfavour it, then FA is false.

P1 is an application of the idea that “normativity implies can” (3243, 3248). Premises P2 to P4 are Hurtig’s paraphrases of passages in Brogaard and Salerno (2019), where the latter cast doubt on the possibility of counterfactual knowledge of actual truths.⁶ It seems that their doubts, if sound, would carry over to the case of counterfactual evaluation and favouring of actual states of affairs that are impossible actually to evaluate, such as *S*. Due to lack of a causal link, counterfactual evaluators have no way of “latching onto” the actual world as opposed to any other world in which *S* obtains, and thus have no way of latching onto the value of *S* as it obtains in the actual world (3247). I will now discuss Rønnow-Rasmussen’s response.

3. Rønnow-Rasmussen’s Response

Rønnow-Rasmussen responds to Hurtig by essentially conceding to the challenge: “Like Hurtig, I believe it is impossible to have an attitude in the evaluating world that latches onto the actual world” (2022: 116). However, he goes on to explain why “the problem is not quite as serious as it appears to be” (ibid.). He distinguishes two scenarios: (a) non-universalizable features of a state of affairs are not value-makers; (b) non-universalizable features are or can be value-makers.

In the first scenario, then, the value of a state of affairs like *S* depends only on its universalizable value-making features, and therefore not on non-universalizable features like the identity of the individuals involved (say, “Charlie”) or—as I understand Rønnow-Rasmussen—even the particular modal location of the state of affairs. If so, then “a proponent of FA analysis ought to be quite satisfied with the counterfactual evaluator evaluating a class of Charlies (i.e. those individuals in possible worlds that share Charlie’s universal value-making features) even if his attitude does not ‘latch on’ to the Charlie in the actual world” (ibid.: 116-117). In the case of state *S*, then, the counterfactual evaluator can still evaluate a class of states of affairs that share *S*’s universal value-making features, even if her evaluating does not latch onto *S* as belonging to the actual world. The assumption made by this reply to Hurtig is that evaluating the class of such states of affairs does not itself require latching onto any particular world or individuals, or (practically equivalently) that the ‘latching onto’ required in this case is possible for any

⁶ “If there is such non-actual knowledge, there is non-actual thought about an actual situation. So the non-actual thinker somehow has a concept of an actual situation. But how is it possible for a non-actual thinker to have a concept that is specifically about situations in this the actual world. It will not do for the thinker to express the thought ‘actually *p*’, since ‘actually’ will designate rigidly only situations in her own world. Moreover, since there is no causal link between the actual world *w*₁ and the relevant non-actual world *w*₂, it is unclear how non-actual thought in *w*₂ can be uniquely about *w*₁” (Brogaard & Salerno 2019, in turn referring to Williamson (1987)).

evaluator whatever their location in modal space. (I'll grant this assumption in what follows.)

I believe that this response concedes a much larger defeat for FA than Rønnow-Rasmussen supposes. If Hurtig is right, then *whenever* the allegedly required latching onto a certain world or individual does not or cannot take place, it will not be fitting to favour exactly *that* valuable state of affairs (or the individuals therein), but only the class of states of affairs sharing the universalizable features. But the required latching onto can fail in a myriad of cases, even when the evaluating world and the evaluated world coincide. It is a highly contingent matter whether anyone's attitudes do or do not latch onto a given states of affairs. In turn, it will be a highly contingent matter whether anyone can favour that particular state of affairs, and thus whether it is fitting to favour that particular states of affairs.

This predicament puts FA before two unpleasant alternatives. The first alternative is to hold a disjunctive account of the objects of fitting attitudes: a state of affairs P is good if and only if either it is fitting to favour P or, failing that, it is fitting to favour something like P's better relative P*: that is, P minus any non-universalizable feature which would require the evaluator's latching onto exactly P's world or P's individuals. But this account is not a great solution: ideally, we would like the objects of fitting attitudes (*what* it is fitting to favour) to be, always, exactly the same—i.e. the same tokens under the same description—as the objects bearing the value property (*what* is good). If Charlie's being happy is good, then it is Charlie's being happy that should be favoured, and not, even as a second best, simply the state [someone's being happy]. Even if one agrees that the value of a state of affairs depends only on a state of affairs' universalizable features, one may still require that the object of the fitting attitude be the state of affairs including its non-universalizable features, because that is after all how the relevant value bearer is presented in this case (e.g. as Charlie's being happy). One thing is the question about legitimate value-makers, another thing is the question about what are the legitimate targets of fitting attitudes.

The second unpleasant alternative is to stipulate that FA only account for the value of value bearers stripped of any non-universalizable feature. On this view, FA should account neither for the value of Charlie's being happy, nor for the value of someone's being happy in the actual world, but only for the value of someone's being happy. Of course FA advocates are free to select the subject matter of FA as they please. But it seems to me FA would lose some of its appeal. After all, we ordinarily ascribe value to states of affairs that include non-universalizable features such as the identity of individuals, times, places. Moreover, we may be tempted to hold the view that, for example, "agent *a*'s pleasure is valuable, but [...] no other agent's pleasure is valuable, however similar it is in terms of its universalizable features" (Rønnow-Rasmussen 2022: 117). FA, then, had better find a way to also cover the value of value bearers with non-universalizable aspects rather than ignore them.

This point is particularly pressing when we consider Rønnow-Rasmussen's second scenario: non-universalizable features are or can be value-makers of states like Hurtig's *S* or Charlie's being happy in the actual world. As in the example above, one could hold that only a certain agent's pleasure is valuable, regardless of similarities with other agents, thus making the value dependent, in part, on who the agent is. In fact, one could hold that God's pleasure is the highest good, or that beatific vision of God is the highest good, as distinct from, say, "the pleasure of perfect beings is the highest good" and "beatific vision of perfect beings is the highest good". In these cases it is God's identity that matters, over and above his perfection or other universalizable features he (and maybe only he) possesses. Whether these substantive axiologies are plausible or not, it would be a significant cost for FA to decide to leave them outside of its sphere of analysis.

In this connection, Rønnow-Rasmussen notes that "*any* view suggesting that non-universalizable features can be value-making features will owe us an explanation" (ibid.). This is true. But the special problem for FA is that FA faces the extra burden of making sense of the object of fitting attitudes in these cases, *if*, as Rønnow-Rasmussen appears to concede to Hurtig, having fitting attitudes towards these particular states of affairs requires the possibility of referring or latching onto particular worlds or individuals, and this possibility may not always be given to the relevant evaluator. In this sense, *if Hurtig is right*, then it is probably better for FA to altogether give up on accounting for the value of value bearers with non-universalizable features (despite the cost of this move), and focus on finessing the first alternative above to make it more digestible.⁷ However, it will be clearly even better for FA if one can reject Hurtig's argument in the first place. This is what I do in the next section.

4. A Different Response: Modal Relocation

The first thing to note is that Hurtig's argument, as it stands, would generalize to cases that, for all Hurtig says, FA *can* account for. As reminded above, FA already needs trans-world fittingness in order to explain the value of states of affairs that are non-actual, like *Happy Egrets*. In this case, all we know is that such a state obtains in some possible, non-actual, world. We do not seem to need any causal link with

⁷ I find Rønnow-Rasmussen's first response (a counterfactual evaluator can still have fitting attitudes towards a class of states of affairs, if not towards *S* itself) similar to the response I gave to the 'distance problem' in Orsi 2013a. I argued that 'x is good to degree n' can be defined as 'it is fitting to favour x to degree n from behind a veil of ignorance regarding the evaluator's distance (personal, temporal, modal, even epistemic) from x'. However, I am not sure that all factors to be 'veiled' necessarily match non-universalizable features of the state of affairs—they are rather facts about me than facts about the state of affairs. The question also remains whether, after veiling all these factors, we still need to be able to latch onto x in order to evaluate *it* and not something else.

Happy Egrets Strike Back

any of the possible worlds where *Happy Egrets* obtains in order to contemplate and favour it. But if lack of a causal link is a challenge for FA in the case of *S*, then it should be so also in the case of *Happy Egrets*. Since it doesn't seem to be a challenge in the latter case—and Hurtig appears to agree (3245)—then it is down to Hurtig to explain why it is a challenge in the case of *S*.

Hurtig is likely to answer that *S* is different from other solitary goods or evils in that *S* is world-specific (3246). In fact, in the article there is a crucial shift in the content of the relevant state of affairs from simply

S: [*p* is a significant true proposition, and no one in the actual world at any time has any attitude toward *p*]

to what I will call

Actual S: [*S* obtains in @]—where '@' designates the actual world (3246-7).

Hurtig does not seem to appreciate that *Actual S* is different from *S*. *Actual S* is a state of affairs *indexed* to a particular location, namely the actual world, while *S* is not. *S* only says something *about* the actual world, namely that in this world no agents have attitudes at any time towards *p*. But *S* itself could in principle be indexed to a different world. So, to make Hurtig's argument work, we should now insert '*Actual S*' in place of 'the actual state of affairs *S*'.⁸

So *Actual S* (rather than just *S*) is supposed to be importantly different from *Happy Egrets*. Since the latter is not indexed to any specific modal location, an evaluator's latching onto the possible worlds where it obtains comes on the cheap, i.e. without the need for any causal link. (Or perhaps there is no need for latching onto them in the first place in order to have the relevant fitting attitude.) But when a state of affairs is indexed to the actual world, like *Actual S*, the thought must be that a causal link between evaluating world and evaluated world is required, so that the evaluator is able to pick exactly the right location of the valuable state of affairs. Were she to pick a state exactly like *Actual S*, however located in a world that is not the actual world, she would make a mistake and end up evaluating not *Actual S* but a different state of affairs. In the absence of a causal link, then, a counterfactual evaluator is not able to evaluate *Actual S*, hence it cannot be fitting for her to disfavour *Actual S*.

I will now provide a response to this modified version of Hurtig's argument. The starting point is that Hurtig does not sufficiently explain whether there is supposed to be a special problem for non-actual evaluators to uniquely pick the actual world among other similar worlds, or whether there is also a parallel problem for actual evaluators to uniquely pick a specific non-actual world *W*₁, when a valuable state

⁸ In fairness to Hurtig, by the time he presents the argument reconstructed above, the shift to *Actual S* has already occurred. But he doesn't register the shift.

of affairs is indexed to *W1* (think *Happy Egrets in W1*). But the latter claim seems more plausible.⁹ In other words, it seems that if a causal link is required for a counterfactual evaluator to have an attitude towards a world-indexed actual state of affairs, then by parity of reasoning a causal link must also be required for an actual evaluator to have an attitude towards a world-indexed non-actual state of affairs. And if this is true, then it follows that in the absence of a causal link with the relevant non-actual world we, as actual evaluators, would not be able to favour or disfavour a world-indexed non-actual state of affairs, because we would not be able to uniquely pick this one among other, similar, states of affairs occurring in other worlds.

However, the latter implication should give us pause. As I noted in (Orsi 2013a), whenever a seemingly good or bad state of affairs is put up for consideration as a counterexample to FA, it must at least be fitting *for the readers* to regard it as good or bad. If it is not even fitting for readers to evaluate it, then it can hardly be eligible as a counterexample to FA. But if it is fitting for readers to evaluate it, then it must be possible to evaluate it. And if we, as readers, can evaluate it, then this means that our ability to evaluate such states of affairs holds regardless of our location in modal space vis-à-vis the location of the state of affairs. Where we stand with relation to the state of affairs doesn't seem to matter. A fortiori, the absence of a relevant causal link between the reader's world and the world where the state of affairs obtains is neither here nor there. Since there does not seem to be any additional challenge in going from evaluating a state as good or bad to favouring or disfavouring it, it also follows that we, as readers, can have the relevant fitting attitude. In other words, it is tempting to suggest that the whole literature on solitary goods must be premised on the assumption that it is at least possible, and fitting for someone, namely the reader, to favour or disfavour the putative solitary good or evil in some way.¹⁰

World-specific or indexed solitary goods and evils are no exception. In fact, Hurtig concurs, as he writes that "it is coherent to think that *S*'s obtaining *in the actual world* [i.e. *Actual S*] is bad for its own sake" (3245, his italics). Now, this needs to be refined, since it is not coherent for us to *both* regard ourselves as actual evaluators of *Actual S* and think that *Actual S* is bad. By hypothesis, *Actual S* cannot have *actual* evaluators. What is coherent, instead, is for the reader to think that *Actual S* is bad while regarding herself as non-actual—placing herself in a non-actual world. This sort of modal relocation must be possible for us, or else it is not clear *for whom* it is coherent to think that *Actual S* is bad for its own sake. And here is the catch: if it is possible for us, readers, to place ourselves in a non-actual world and evaluate *Actual S* *from there*, without there being any apparent causal link between the non-actual world we would inhabit and the actual world where *Actual*

⁹ It seems that Hurtig would agree, as he writes that the challenge for FA is, in general, to show "how it is possible to favour *specific* worlds, situations, or states of affairs" (3248, his italics).

¹⁰ This point applies also to Reisner's "causal entanglement" case (2015). See Rowland (2019, ch. 7) for a detailed response to Reisner's arguments.

S obtains, then it must be possible for *any* counterfactual evaluator to evaluate and disfavour *Actual S*, without the need for any causal link between *their* non-actual world and the actual world. *Qua* placed in a non-actual world, we are in a no more privileged position with respect to *Actual S* than any other non-actual evaluator of *Actual S*. (It's not as if by virtue of, in fact, inhabiting the actual world, we can somehow smuggle our way to *Actual S*. When we evaluate *Actual S*, we stand firmly in a non-actual world.) Hence either one should reject premise P4 in Hurtig's argument (if there is no causal link between the actual world and N's world, then N's evaluation cannot uniquely be about *S*'s obtaining in the actual world), or Hurtig must accept that *Actual S* is a state of affairs not even his readers can coherently evaluate.

Hurtig may want to buy into the second horn of the dilemma. He might suggest that putting up *Actual S* for consideration as a legitimate counterexample to FA does not require the readers' ability to evaluate *Actual S*. It only requires the ability to contemplate its general features, namely *Actual S* minus the world-specific index. In other words, Hurtig may claim to be entitled to present *Actual S* as a counterexample, even if *Actual S* is a state of affairs that (by Hurtig's own lights) *as such* we cannot properly grasp and evaluate, since we—forced by the nature of *Actual S* to take the position of non-actual evaluators—cannot uniquely pick *Actual S* from other, similar states of affairs. It is not fitting for us to think that *Actual S* is bad, yet we are to somehow take it that *Actual S* is bad.

It is not clear whether Hurtig would endorse this reply line, as this involves the same mismatch between value bearer (here, *Actual S*) and object of the fitting attitude (*Actual S* minus the world-specific index) pointed out in my reply to Rønnow-Rasmussen above. Of course, since Hurtig is arguing *against* FA, such a mismatch need not be a problem for *him*. But this reply still involves two hefty commitments that plausibly *everyone* should steer clear of.

First, if, despite our inability to uniquely pick it from similar states of affairs, *Actual S* works as a counterexample to FA, then it follows that there are, or could be, good or bad states of affairs that, by their very nature, are not *as such* graspable by any subject in any world, and a fortiori it cannot be fitting for anyone to evaluate them, let alone favour or disfavour them. The best we can do is relate to similar states of affairs that do not include a problematic index to a world we have no causal link with. I will not discuss whether such a view is coherent. Arguably it is a view that those who think that normativity is optional to value might be happy to endorse: there are valuable states of affairs that it is fitting for no one to even evaluate—because no one could evaluate them. But at this point we might wonder whether, overall, FA offers a better package than any such view. It is worth remarking that extant theories on the relation between value and normativity do not go so far as to *deny* the rather trivial claim that if something is good, then it is fitting to regard it as good. If a theory can account for the value of *Actual S* only by denying this trivial claim, then such a theory earns a benefit at a very significant cost.

The only way out of this problematic commitment is to deny premise P4 in Hurlig's argument: despite the absence of an appropriate causal link between *Actual S* and the non-actual world into which we "relocate" when contemplating *Actual S*, we are able to evaluate *Actual S* as such, and so are other counterfactual evaluators. That is why it is coherent and fitting for us (more precisely, for the counterfactual "us") to think that *Actual S* is bad. And if it is coherent and fitting for "us" to think that *Actual S* is bad, then it is a short step to it being fitting for "us" to disfavour *Actual S*, as FA has it. Note: on this view, the object of fitting disfavour is indeed *Actual S*, not just the class of states of affairs sharing universalizable features with *Actual S*. Indexing states of affairs to particular modal locations cannot make them completely inaccessible to evaluation and other fitting attitudes. It thus seems that, whatever may be true regarding belief or knowledge, evaluation is a kind of attitude that can tolerate lack of appropriate causal links between evaluating world and evaluated world. The broader implications of this point will need to be explored elsewhere.¹¹

The second hefty commitment is not so much one in value theory, but rather in the ethics of argumentation. Suppose Hurlig does endorse the idea that *Actual S* is a valid counterexample to FA, even though we cannot really judge *it* bad. Then he would need to defend the fairness of objecting to a view on the basis of a counterexample that, by its very nature, readers (and author) are unable to properly grasp, but only able to "get somewhere near". Whether such argumentative moves are ever legitimate is a complicated question I cannot address here, but it seems Hurlig has taken upon himself the burden to address it.¹²

References

- Brogaard, B. & Salerno, J. (2019). "Fitch's Paradox of Knowability", The Stanford Encyclopedia of Philosophy (Fall 2019 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/fall2019/entries/fitch-paradox/>.
- Bykvist, Krister (2009). 'No Good Fit: Why the Fitting Attitude Analysis of Value Fails'. *Mind* 118, 469: pp. 1–30.

¹¹ A reviewer helpfully pointed out that states like *S* and *Actual S* may be 'axiological blindspots', analogous to epistemic blindspots like 'It is now raining but no one believes that it is now raining'. Epistemic blindspots are true propositions that cannot be truthfully or rationally believed. Axiological blindspots would be states of affairs that are good (bad) but cannot be fittingly (dis)favourable or even evaluated. One question, here, is whether epistemic blindspots can at least be coherently *entertained* by someone—not much more than this would seem to be needed for evaluation.

¹² This research has been supported by the European Union through the European Regional Development Fund (the Centre of Excellence in Estonian Studies, TK 145), and the University of Tartu, grant PHVF121914. I thank the editors for their invitation to contribute and a reviewer for their insightful comments.

Happy Egrets Strike Back

- Bykvist, Krister (2015). 'Reply to Orsi'. *Mind* 124, 496: pp. 1201–5.
- Dancy, Jonathan (2000). 'Should We Pass the Buck?' In A. O'Hear (ed.), *Philosophy, The Good, the True, and the Beautiful*. Cambridge: The Press Syndicate of the University of Cambridge: pp. 159–73.
- Hurtig, Kent (2019). 'The Fitting Attitude Analysis of Value: An Explanatory Challenge'. *Philosophical Studies*, 176, 12: pp. 3241–49.
- Orsi, Francesco (2013a). 'Fitting Attitudes and Solitary Goods'. *Mind* 122 (487): 687–698.
- Orsi, Francesco (2013b). 'What's Wrong with Moorean Buck-Passing?'. *Philosophical Studies* 164, 3:727–746.
- Orsi, F. & Garcia, A. G. (2021). 'The Explanatory Objection to the Fitting Attitude Analysis of Value'. *Philosophical Studies* 178 (4):1207–1221.
- Rabinowicz, W. and Rønnow-Rasmussen, T. (2004). 'The Strike of the Demon: On Fitting Attitudes and Value'. *Ethics* 114, 3: pp. 391–423.
- Reisner, Andrew (2015). 'Fittingness, Value, and Trans-World Attitudes'. *Philosophical Quarterly* 65, 260: pp. 464–85.
- Rønnow-Rasmussen, Toni (2022). *The Value Gap*. Oxford: Oxford University Press.
- Rowland, Richard (2019). *The Normative and the Evaluative: The Buck-Passing Account of Value*. Oxford: Oxford University Press.
- Williamson, Timothy (1987). 'On the Paradox of Knowability'. *Mind* 96, 382: pp. 256–61.