

Luca Barlassina

University of Sheffield

Beyond good and bad: Reflexive imperativism, not evaluativism, explains valence

Forthcoming in *Thought*

Evaluativism (Carruthers 2018) and reflexive imperativism (Barlassina and Hayward 2019) agree that valence—the (un)pleasantness of experiences—is a natural kind shared across all affective states. But they disagree about what valence is. For evaluativism, an experience is pleasant/unpleasant in virtue of representing its worldly object as good/bad; for reflexive imperativism, an experience is pleasant/unpleasant in virtue of commanding its subject to get more/less of itself. I argue that reflexive imperativism is superior to evaluativism according to Carruthers’s own standards. He maintains that a theory of valence should account for its phenomenology and role in imagination-based decision-making. I show that it is reflexive imperativism, rather than evaluativism, that fits this explanatory bill.

“One loves ultimately one’s desires, not the thing desired.”

—Friedrich Nietzsche, *Beyond Good and Evil*

Introduction

There is a growing consensus in cognitive science that valence—the (un)pleasantness of experiences—is a natural kind shared by affective states as diverse as pleasures, pains, emotions, and desires (Levy and Glimcher 2012). If this consensus is on the right track, the recent philosophical debate on pain’s unpleasantness (Bain 2013; Jacobson forthcoming; Martínez 2011) may be too narrow in focus. We should rather aim at developing theories of (un)pleasantness across the affective spectrum.

In this article, I contrast two philosophical theories that have the required generality: the version of *evaluativism* put forward by Carruthers (2018) and *reflexive imperativism*, the

form of imperativism I developed with Max Khan Hayward (Barlassina and Hayward 2019). According to the former, an experience is pleasant/unpleasant in virtue of representing its worldly object as good/bad; according to the latter, it is pleasant/unpleasant in virtue of commanding its subject *Get more of me!/Get less of me!*

I argue that reflexive imperativism is superior to evaluativism according to Carruthers's own standards. He maintains that a theory of valence should account for its phenomenology and role in imagination-based decision-making. I show that it is reflexive imperativism that fits this explanatory bill.

1. A kind of (un)pleasantness

Affective states come in many forms. There are sensory pleasures (e.g., orgasms) and pains (e.g., backaches); propositional (dis)pleasures (e.g., being (dis)pleased that the sun is shining); basic emotions (e.g., fear) and higher-cognitive emotions (e.g., guilt); positive and negative moods (e.g., elation and depression). Still, Carruthers (2018) argues that one can find in the cognitive science literature three reasons to think that valence, i.e., (un)pleasantness, is a natural kind shared by all affective states.

First, valence is underpinned by a common neurobiological network comprising the anterior insula, striatum, dorsomedial and ventromedial cortices (Bartra et al. 2013; Leknes and Tracey 2008). This network is activated across multiple stimulus types (e.g., food and money) and modalities (e.g., vision and touch). In fact, it is recruited both while having an (un)pleasant experience and while imagining a future (dis)pleasure.

Second, Carruthers notices that the same interventions influence the valence of different affective states. For example, acetaminophen—the active principle in paracetamol—

blunts the unpleasantness of both physical and emotional pain (Dewall et al. 2010). More than this, it dulls both pleasantness and unpleasantness (Durso et al. 2015).

Third, Carruthers maintains that valence occupies a distinctive functional role: it allows you to make decisions about initially incommensurable outcomes by generating a common currency for outcome comparison (Levy and Glimcher 2012). To decide whether to go to the movies or to a restaurant, you can imagine either activity and use the elicited valences ('this seems more pleasant than this') to guide your choice. This comparison, however, is possible only if valence is a property shared by different types of experience—otherwise, you couldn't compare the aesthetic pleasure of watching a movie with the gustatory pleasure of eating food.

Carruthers's arguments might not provide conclusive evidence for the claim that valence is a natural kind. But this is irrelevant. In fact, Hayward and I accept this claim too. Our disagreement with Carruthers has to do with how to explain valence.

2. A tale of two theories

Both evaluativism and reflexive imperativism are intentionalist theories of valence—they explain the (un)pleasantness of an experience in terms of the experience's intentional content. However, the intentional contents they resort to are pretty different.

2.1 Evaluativism

You are having a toothache. Your experience has a certain sensory component: you feel that there is something throbbing in your mouth. According to Carruthers, your experience has this feature in virtue of representing that such-and-such a bodily condition obtains. But your

experience has valence as well—it is unpleasant. This is so—Carruthers says—because your experience also has *evaluative content*: it represents your bodily condition *as bad*.

Carruthers is not the first to defend evaluativism about pain's unpleasantness (Bain 2017). What's distinctive about his proposal is that this explanatory strategy gets applied to valence across the board. You are afraid of a bear; you are happy that your team won; you feel an itch in your wrist. For Carruthers, these experiences are pleasant/unpleasant in virtue of representing their *worldly objects*—the bear, the victory or your team, the physiological condition of your wrist—as good/bad.

The content of a mental state is first-order if and only if it does not feature any mental state. Moreover, Carruthers says that the relevant representations of badness and goodness are nonconceptual. Therefore, another way to articulate evaluativism is this: an experience has valence in virtue of possessing first-order, evaluative, nonconceptual content.

2.2 Reflexive imperativism

Hayward and I explain valence in terms of *imperative content*—a kind of intentional content that directs rather than evaluates. Our theory is a form of imperativism (Klein 2015). What distinguishes ours from other forms of imperativism is that the imperative content in question is reflexive: an experience is pleasant/unpleasant in virtue of commanding its subject *Get more/less of me!* Here 'me' picks out the *type* of experience, rather than the token experience. So, the content should be read as follows: *Get more/less of the type of experience of which I am a token!*

An example should get the idea across. You are having a valence-neutral experience S—say, a sensation of warmth in your neck. Hayward and I agree with Carruthers that S has sensory phenomenology in virtue of representing that your neck is warm. But while Carruthers says that in order to become pleasant, S needs to come to represent the state of

your neck as good, we propose that S has to acquire reflexive imperative content *Get more of me!*

Reflexive imperative contents aren't first-order. Are they higher-order? If 'higher-order' simply means 'about a mental state', they are. However, if a mental state has higher-order content in virtue of being about *some other mental state*, then reflexive imperative contents aren't higher-order: the mental state targeted by the imperative content is *the very self-same state* that possesses this content. This is why Hayward and I follow Kriegel (2006) in characterising reflexive imperative contents as *same-order*.¹ Still, these contents differ from first-order ones in being experience-directed, rather than world-directed.

Now that we know how evaluativism and reflexive imperativism work, we can establish which works better. Carruthers maintains that a theory of valence should explain its phenomenology and role in imagination-based decision-making. I agree. *Pace Carruthers*, however, it is reflexive imperativism that offers the best account of these *explananda*.

3. The phenomenology of valence

3.1 Felt (un)pleasantness

Evaluativism and reflexive imperativism make different predictions about the phenomenology of valence, i.e., about felt (un)pleasantness. For the former, such a phenomenology is *evaluative* and *world-directed*: given that valence reduces to a representation of the goodness/badness of worldly objects, the phenomenology of valence consists in worldly objects appearing as good/bad. When you have an orgasm, the felt

¹ Hereafter, I use 'higher-order' only to refer to cases in which the content of a mental state M is about *a distinct mental state M**.

pleasantness of your experience is such that the condition of your genitals appears good to you. When you are afraid of a bear, the felt unpleasantness of your experience makes the bear seem bad to you.

Reflexive imperativism proposes that an experience has valence in virtue of possessing imperative content *Get more/less of me!* It thus predicts that the phenomenology of valence is *experience-directed* and *conative*: the felt (un)pleasantness of an affective state is experienced as an inclination to get less/more of that very state. When feeling the unpleasantness of a toothache, you feel the inclination to get rid of that very experience.

I now present four cases in which the phenomenology of (un)pleasantness comes apart from the phenomenology of worldly objects seeming (bad) good. Since evaluativism identifies these two phenomenologies, these cases pose a serious threat to this theory. More precisely, my contention is that even *if* Carruthers found a way to deal with each of these cases, he could do that in a piecemeal fashion only—by changing some aspects of evaluativism to respond to one case, some other aspects to accommodate another one, and so on. *At a minimum*, these cases thus force Carruthers to sacrifice explanatory unification for explanatory adequacy. In contrast, reflexive imperativism can account for these cases *at once* in a principled way; in doing so, it also explains the connection between valence and experience-directed inclination.

3.2 Four stories of pleasure and pain

The cake. Olivia is hungry and passes by a patisserie. She is admiring a cake on display, when she realises that she has left her wallet at home. She desires that cake so badly: it looks so *good* to her. But her desire feels *unpleasant*, and she would like to get rid of it.

The general and the saint. General Bouba and Saint Kiki have done so much to prevent the bandits from destroying the village. The villagers are now celebrating them in the main square. They both feel so proud. Their pride represents themselves and their actions as *good*. Bouba's pride feels pleasant—if only he could feel like that every day. Kiki's pride, however, feels unpleasant—she knows that a servant of God has to remain humble.

The dentist. My dentist is a butcher—he has really smashed up my mouth. At least, he gave me tons of anaesthetic. Now I am looking at my mouth in the mirror. It looks so messed up, meaning: my visual experience represents the condition of my mouth as *bad*. But I am so numbed by the drugs that my visual experience is entirely valence-neutral: it does not feel either pleasant or unpleasant. I really don't care about it.

Bonjour tristesse. Jean-Paul has just woken up and feels miserable. He doesn't know why he feels that way. He cannot even detect what his misery is directed at. There is nothing in his life that appears bad to him. Still, Jean-Paul's misery feels *unpleasant* and he would like it to go away.

Let's start with *The cake*. Carruthers maintains that desires have valence. Therefore, he is forced to say that they have first-order evaluative content. Since what we desire typically appears *good* to us, Carruthers proposes that desires represent their objects *as good*. But this combination of evaluativism and the good-based theory of desire (Oddie 2005)

entails that *all* desires feel *pleasant*, which is clearly false. As we have seen, Olivia's desire feels the opposite.

To solve this problem, Carruthers might suggest that (unfulfilled) desires represent the non-obtaining of a certain state of affairs *as bad*. Accordingly, Olivia's desire feels unpleasant in virtue of representing Olivia's not eating the cake as bad. But this won't do either, since it wrongly predicts that *all* (unfulfilled) desires feel *unpleasant*. A guitarist is driving towards the concert hall where she will play in a few hours. Her desire to play in front of the audience is currently unfulfilled, but it feels extremely pleasant.

Reflexive imperativism is not saddled with finding a first-order evaluative content that could explain why some desires feel pleasant while others don't. Even though desires typically have first-order content (evaluative or otherwise), this content has nothing to do with their valence. Rather, desires feel (un)pleasant in virtue of having reflexive imperative content. The guitarist's desire feels pleasant in virtue of commanding her *Get more of me!*, and Olivia's desire feels unpleasant in virtue of commanding her *Get less of me!* This also explains why the phenomenology of desire is conative and experience-directed, i.e., why Olivia feels the inclination to get rid of her desire, while the guitarist experiences exactly the opposite inclination.

Let's turn to General Bouba and Saint Kiki. Their pride represents their actions and themselves as good. As predicted by evaluativism, Bouba's pride feels pleasant. However, Kiki's pride feels unpleasant. Carruthers might attempt to accommodate this case by saying that: (i) even though *in the great majority of cases* valence reduces to a first-order evaluative content, it *sometimes* reduces to a higher-order one; and (ii) Kiki's pride feels unpleasant exactly because it is targeted by a higher-order state of shame with content *The pride that I am experiencing is bad*.

I wouldn't recommend Carruthers taking this route. First, the idea that valence reduces to a first-order evaluative representation *sometimes-but-not-always* makes evaluativism patchy. Second, and more importantly, this manoeuvre won't solve the initial problem anyway. Assume for argument's sake that Kiki's pride is in effect represented as *bad* by Kiki's shame. Still, Kiki's pride represents Kiki and her actions as *good*. Evaluativism should then predict that Kiki is undergoing an *ambivalent* affective state, i.e., she is experiencing *both* pleasantness and unpleasantness. But she is not feeling any pleasure at all. *This* is what evaluativism cannot explain.

Reflexive imperativism has a simple explanation of the case of Kiki and Bouba: Bouba's pride feels pleasant because it has reflexive imperative content *Get more of me!*; Kiki's pride feels unpleasant because it has reflexive imperative content *Get less of me!* This also accounts for the fact that Bouba is experiencing the inclination to feel more pride, while Kiki is experiencing exactly the opposite inclination.

In *The dentist*, my visual experience is such that the condition of my mouth appears bad to me. This should *not* be read as follows: my visual experience represents that my mouth is in such-and-such a condition, and this leads me to *judge* that such a condition is bad. Rather, it is the very visual experience that nonconceptually represents the condition of my mouth as bad. However, *contra* evaluativism, my visual experience does not feel unpleasant—nor does it feel pleasant. For reflexive imperativism, it is easy to capture this phenomenon: the drugs are preventing my visual experience from acquiring reflexive imperative content. This being the case, my experience is devoid of any experience-directed, conative phenomenology: I don't feel any inclination for or against my visual experience.

What about Jean-Paul? Some moods have first-order intentional content. Depression can be about the world as a whole, or about the state of one's body (Kind 2014). But this is not what is happening to Jean-Paul. His misery isn't directed at any worldly object. Thus,

pace Carruthers, his mood cannot feel unpleasant in virtue of evaluating a worldly object as bad. But unpleasant it is, so an alternative explanation is needed. Reflexive imperativism provides one. On many occasions, reflexive imperative content is just one of the contents possessed by an affective experience. For example, Kiki's pride has reflexive imperative content *Get less of me!* as well as evaluative content *Kiki and her actions are good*. Sometimes, however, affective experiences only have reflexive imperative content. If the content is *Get more of me!*, one experiences "pure elation". If it is *Get less of me!*, one experiences pure misery—like Jean-Paul. This explains why he is bothered by his misery: he experiences it as a "free-floating" inclination to get rid of itself.

4. Imagination-based decision-making

Reflexive imperativism gives a better account of the phenomenology of valence than evaluativism. But this is just half of the story. Carruthers maintains that one should also capture the role played by valence in imagination-based decision-making. In this section, I show that reflexive imperativism is superior to evaluativism in this regard too.

4.1 The architecture of imagination-based decision-making

For both theories, imagination-based decision-making is a three-step process. You don't know whether to eat pizza or pasta. To make up your mind, you first imagine the outcome of each decision. This is followed by *affective forecasting* (Gilbert and Wilson 2005), i.e., each imagining elicits a certain degree of (un)pleasantness. Finally, these elicited valences are compared and the option associated with the greater amount of pleasantness is selected.

The agreement between the two theories ends here. In line with the view of affective forecasting endorsed by the great majority of psychologists (Gilbert and Wilson 2007;

Miloyan and Suddendorf 2015), Hayward and I propose that imagination-based decision-making is *experience-directed*. You first imagine the *experiential outcomes* of your decisions. You imagine, say, *tasting pizza*. This is an instance of *experiential imagination*: you mentally simulate an experience (Barlassina and Gordon 2017). Accordingly, it is simulated experiences that feel (un)pleasant during affective forecasting. For reflexive imperativism, this depends on their reflexive imperative content—your imagining tasting pizza feels pleasant because it commands you *Get more of me!* Your decision-making mechanism then outputs a decision by comparing the strength of the experiences' imperative contents.

For Carruthers, the cognitive architecture of imagination-based decision-making is instead *world-directed*. You first imagine the *worldly outcomes* of your decisions. That is, you don't imagine tasting pizza. Rather, you imagine *pizza*. During affective forecasting, the expected values of these worldly outcomes get estimated. For example, the pleasantness elicited by imagining pizza represents it as good. The decision-making mechanism then produces an output by comparing the expected values of such worldly outcomes.

4.2 Three cheers for experiential imagination

A key difference between these two architectures concerns the type of imaginings upon which decision-making operates: reflexive imperativism endorses the mainstream account according to which affective forecasting takes experiential imaginings as input; evaluativism recommends ditching the received view in favour of the idea that it is world-directed imaginings that are at work here. Here are three lines of evidence in support of the received view.

First, if affective forecasting takes experiential imaginings as input, it should activate brain areas that underpin experiential projections. This is the case. Affective forecasting engages the default network (Gilbert and Wilson 2007; Spreng et al. 2009), a brain system

that subserves many forms of experiential projections, including episodic memory, perspective-taking, and, most importantly, experiential imagination itself (Buckner and Carroll 2007).

Second, the main region responsible for comparing valences and making decisions on this basis is the ventromedial prefrontal cortex (Bartra et al. 2013). Crucially, this region is known to be recruited during first-person thinking and thus hypothesised to deal with experiential imaginings (Benoit et al. 2014; D'Argembeau, 2013).

Third, experiential imagination is an instance of *cognitive re-use* (Hurley 2008). In order to simulate experience E (say, fear), one takes the system responsible for E (the fear system) “off-line” and re-uses it to generate a “pretend experience”. Reflexive imperativism therefore predicts that the current state of the recruited system influences valence-based decision-making. This expectation is so widely confirmed that it has earned its own name, i.e., *projection bias* (Lowenstein et al. 2000). Just to give an example, many studies have found that one’s *current* hunger influences one’s decisions about what food to eat *later* (Gilbert et al. 2002; Read and van Leeuwen 1998).

4.3 Self-sacrifice and evolution

Can we conclude that imagination-based decision-making is experience-directed? Not yet. Carruthers has in fact two arguments for his world-oriented model. The first starts from the observation that people sometimes sacrifice their lives for a cause. This is understandable given evaluativism: “somebody imagines the state of the world [after the revolution] and it seems ... better than does own’s own continued existence in the absence of the revolution” (Carruthers, 2018, p. 667). How can reflexive imperativism make sense of this? I surely cannot imagine experiencing pleasure *after* the revolution if I imagine dying *during* the revolution. Or can I?

Here is a way. If I identify not with my physical body, but with another entity—e.g., my social group—I can imagine experiencing pleasure after the destruction of my body, provided that I imagine my social group surviving. Accordingly, reflexive imperativism predicts that those who decide to die for a cause need to conceive of themselves not simply in terms of their bodily existence. Strikingly, this prediction is dead right. A number of studies (see Whitehouse 2018 for a review) have found that extreme forms of self-sacrifice among people as different as religious fundamentalists, armed insurgents, tribal warriors, etc. are *all* preceded by intense collective experiences resulting in *identify fusion*—a sense of oneness with the group (Swann et al. 2012). In this way, one can imagine existing, and thus experiencing pleasure, after one’s “bodily death”. Whitehouse (2018, p. 3) puts it nicely: “Fused individuals see their personal selves as encompassed ... by the group. For them, the survival of the group constitutes a form of immortality”.

The first argument given by Carruthers against an experience-directed architecture has turned out to be an argument in its favour. What about the second one?

Evolution couldn’t care less about how one feels. ... It would therefore be surprising that people should have evolved a decision-making architecture that pays attention ... to the prospect of good and bad experiences. ... One might have expected that evolution would have utilised representational resources ... about worldly options.

(Carruthers, 2018, p. 670)

We can unpack this idea as follows. The ability to predict, and make decisions on the basis of, the value of future worldly outcomes increases the fitness of an organism. Evaluativism hence offers a plausible account of the selection history of imagination-based decision-making. In contrast, the expected (un)pleasantness of an experience is an imperfect proxy for the fitness value of a decision—after all, a decision might result in a *bad* outcome

accompanied by a *good* experience, and vice versa. It is thus unclear why evolution should have selected an experience-directed architecture over a world-directed one.

Here is why. Evolution is a tinkerer: it builds new structures by introducing small modifications to pre-existing ones. Accordingly, the already existing structures impose architectural constraints on the evolutionary process. Future thinking in human beings is typically episodic in nature: we tend to simulate future *experiences* (Boyer 2008; Schacter et al. 2017). This explains why an experience-directed architecture for imagination-based decision-making was selected: because it had the right representational format to interface with experiential future thinking.

5. Better be good

Carruthers thinks that valence is a natural kind and that an adequate theory should explain its phenomenology and role in imagination-based decision-making. In this article, I argued that reflexive imperativism accounts for these explananda *better than* Carruthers's own evaluativism. But being better than somebody doesn't mean being *good*. If reflexive imperativism faces major difficulties of its own, it may have to be jettisoned as well. In this final section, I consider two issues concerning the picture of decision-making emerging from reflexive imperativism.

5.1 Unpleasant desires

The first issue concerns the role of *unpleasant desires* in decision-making. You have an unpleasant desire for water. According to reflexive imperativism, the unpleasantness of your desire depends on the desire's reflexive imperative content *Get less of me!* But then why is

your decision-making system (DMS) likely to respond to such a desire with the decision to *satisfy* it (i.e., with the decision to drink water), rather than with the decision to *suppress* it?

The key to answering this question is to remember that reflexive imperative content is typically *just one of the contents* possessed by an affective experience (see section 3.2). Your unpleasant desire, being a desire *for water*, also has a “water-directed” first-order content—let’s suppose it is the first-order imperative content *Drink water!* Accordingly, upon getting an unpleasant desire for water as input, your DMS receives not one, but two “requests”: drinking water *and* not desiring water any longer. Since the job of the DMS is to maximise request satisfaction (in the light of the available information), it is easy to see why your DMS is *ceteris paribus* going to output the decision to drink water, rather than the decision to suppress the desire: only the former decision is capable to satisfy *both requests at once*—drinking water is an excellent way to stop desiring water.

One might respond that this generates an *exclusion problem* for reflexive imperativism. Given that the first-order content *Drink water!* is sufficient to elicit the decision to drink water, what’s the point of having a reflexive imperative content on the top of it? It appears that reflexive imperativism predicts that, at least in the case of unpleasant desires, valence is *motivationally idle*.

Appearances are misleading though. Suppose that while watching a *boring* movie at home, you form a *valence-neutral* desire for water. In such a case, this desire is likely to cause the decision to go to the kitchen and drink some water. But if you *enjoy* the movie, your desire to keep on watching it is probably going to *override* your valence-neutral desire for water. What if your desire for water were *unpleasant*? In that case, it would have a *greater motivational force* and thus be more likely to bring about the decision to go to the kitchen.

5.2 Simulated experiences are not on the menu

Another issue emerges in the context of *imagination-based* decision-making. On the one hand, reflexive imperativism says that this type of decision-making involves *experiential imagination*—like when your DMS outputs the decision to order pizza because you have imagined tasting pizza and that simulated gustatory experience felt pleasant. On the other hand, reflexive imperativism states that an experience is pleasant in virtue of having content *Get more of me!*, i.e., *Get more of the type of experience of which I am a token!* An inconsistency appears to emerge here. Given that (i) your imagining tasting pizza belongs to the type *simulated gustatory experience* and (ii) the only way to produce a simulated experience is to imagine it, it seems that reflexive imperativism should instead predict that your DMS is going to output the decision to keep on imagining tasting pizza. How can this tension be solved?

Here's how. In section 4.2, I said that mental simulation is a form of *cognitive re-use*: a simulated experience is produced by *the very same cognitive system* responsible for the generation of its genuine counterpart. This explains why a genuine experience (say, the experience of tasting pizza) and its simulated counterpart (i.e., imagining tasting pizza) are *similar* in many respects—why they have overlapping functional roles, similar phenomenology, and the same neural realisers (Barlassina and Gordon 2017; Goldman 2006). Given this, it would be unsurprising if there were cognitive systems that treated genuine and simulated experiences as belonging to *the same type*—this, after all, is the very point of having simulated experiences: you can use the simulation *in lieu of* the real thing. The DMS is one such system. Accordingly, for the DMS, both the decision to order pizza and the decision to keep on imagining tasting pizza count as ways of satisfying the request *Get more of the type of experience of which I am a token!*—they both are routes to the experience of tasting pizza, be it genuine or simulated.

But why then does your DMS opt for the decision to order pizza? The answer is simple. You are imagining tasting pizza in the context of a *specific decision-making problem*, namely: ‘What food shall I order?’ And while the decision to order pizza satisfies the request *Get more of me!* in a way that solves the decision-making problem at hand, the decision to keep on imagining tasting pizza doesn’t. Imagining tasting pizza isn’t on the menu.²

References

- Bain, D. (2013). “What makes pains unpleasant?” *Philosophical Studies*, 166: 69-89.
- Bain, D. (2017). “Evaluativist accounts of pain’s unpleasantness.” In J. Corns (ed.) *Routledge Handbook of the Philosophy of Pain*.
- Barlassina, L., and Gordon, R.M. (2017). “Folk psychology as mental simulation.” In E. Zalta, (ed.), *Stanford Encyclopedia of Philosophy*.
- Barlassina, L., and Hayward, M.K. (2019). “More of me! Less of me! Reflexive imperativism about affective phenomenal character.” *Mind*, 128, 512: 1013-1044.

² Earlier versions of this article were presented at the universities of Bern, Glasgow, and Sheffield. In addition to an anonymous referee for, and the Editors of, *Thought*, I would like to thank the following people for their criticisms and suggestions: David Bain, Avner Baz, Andrea Blomkvist, Davide Bordini, Emma Borg, Jennifer Corns, Ryan Doran, Alex Duval, Max Hayward, Antti Kauppinen, Viktoriia Kononova, Steve Laurence, Michelle Liu, Olivier Massin, Neil McDonnell, Michael Messerli, Kevin Reuter, Mauro Rossi, Abraham Sapien-Cordoba, Johanna Schnurr, Mona Simion, Yonatan Shemmer, David Stromahier, and James Turner.

- Barlassina, L., and Hayward, M.K. (forthcoming). “Loopy regulations: The motivational profile of affective phenomenology.” *Philosophical Topics*.
- Bartra, O., et al. (2013). “The valuation system: A coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value.” *NeuroImage*, 76: 412-427.
- Benoit, R.G., et al. (2014). “Ventromedial prefrontal cortex supports affective future simulation by integrating distributed knowledge.” *PNAS*, 111(46): 16550-16555.
- Boyer, P. (2008). “Evolutionary economics of mental time travel”, *Trends in Cognitive Science*, 12(6): 219-224.
- Buckner, R.L., and Carroll, C.D., (2007). “Self-projection and the brain.” *Trends in Cognitive Science*, 11(2): 49–57.
- Carruthers, P. (2018). “Valence and value.” *Philosophy and Phenomenological Research*, 97(3): 658-680.
- D’Argembeau, A. (2013). “On the role of the ventromedial prefrontal cortex in self-processing: The valuation hypothesis.” *Frontiers in Human Neuroscience*, 7(372): 1-13.
- Dewall, C.N., et al. (2010). “Acetaminophen reduces social pain: behavioral and neural evidence.” *Psychological Science*, 21(7): 931-937.
- Durso, G.R., et al. (2015). “Over-the-counter relief from pains and pleasures alike: Acetaminophen blunts evaluation sensitivity to both negative and positive stimuli.” *Psychological Science*, 26(6): 750-758.
- Gilbert, D. et al. (2003). “The future is now: Temporal correction in affective forecasting.” *Organizational Behavior and Human Decision Processes*, 88(1): 430-444
- Gilbert, D., and Wilson, T. (2005). “Affective forecasting: Knowing what to want.” *Current Directions in Psychological Science*, 14: 131-134.

- Gilbert, D., and Wilson, T. (2007). "Prospection: Experiencing the future." *Science*, 317: 1351-1354.
- Goldman, A. (2006). *Simulating Minds*. Oxford University Press.
- Hurley, S. (2008). "Understanding simulation", *Philosophy and Phenomenological Research*, 77(3): 755–774.
- Jacobson, H. (forthcoming). "Not only a messenger: Towards an attitudinal-representational theory of pain." *Philosophy and Phenomenological Research*.
- Kind, A. (2014). "The case against representationalism about moods." In U. Kriegel (ed.), *Current Controversies in Philosophy of Mind*. Routledge.
- Klein, C. (2015). *What the Body Commands: The Imperative Theory of Pain*. MIT Press.
- Kriegel, U. (2006). "The same-order monitoring theory of consciousness." In U. Kriegel and K.W. Williford (eds.), *Self-Representational Approaches to Consciousness*. MIT Press.
- Leknes, S, and Tracey, I. (2008). "A common neurobiology for pain and pleasure." *Nature Review Neuroscience*, 9(4): 314-320.
- Levy, D.J., and Glimcher, P.W. (2012) "The root of all value: a neural common currency for choice." *Current Opinion in Neurobiology*, 22(6): 1027-1038.
- Loewenstein G.F., et al. (2000). "Projection bias in predicting future utility." Unpublished manuscript. Carnegie-Mellon University.
- Martínez, M. (2011). "Imperative content and the painfulness of pain." *Phenomenology and the Cognitive Sciences*, 10(1): 67-90.
- Miloyan, B., and Suddendorf, T. (2015). "Feelings of the future." *Trends in Cognitive Science*, 19(4): 196-200.
- Oddie, G. (2005). *Value, Reality, and Desire*. Oxford University Press.

- Read, D., and van Leeuwen, B. (1998). "Predicting hunger: The effects of appetite and delay on choice." *Organizational Behavior and Human Decision Processes*, 76(2): 189-205.
- Schacter, D.L., et al. (2017). "Episodic future thinking: Mechanisms and functions." *Current Opinion in Behavioral Sciences*, 17: 41-50.
- Spreng, R.N., et al. (2009). "The common neural basis of autobiographical memory, prospection, navigation, theory of mind, and the default mode: a quantitative meta-analysis." *Journal of Cognitive Neuroscience*, 21(3):489-510.
- Swann W. B., et al. (2012). "When group membership gets personal: A theory of identity fusion." *Psychological Review* 119 (3):441–56.
- Whitehouse, H. (2018). "Dying for the group: Towards a general theory of extreme self-sacrifice." *Behavioral and Brain Sciences*, 41(192): 1–62.