

Two Dilemmas for Value-Sensitive Technological Design

Mona Simion

Cogito Epistemology Research Centre, University of Glasgow

For *Feminist Philosophy and Emerging Technologies*, Eds. Mary Edwards and Orestis Palermos, Routledge.

1. Introduction

Value Sensitive Design (VSD) has been developed to explicitly address the ethical nature of design. The thought is that values underlying the design process should be made explicit in order to be the transparent and proper subject of ethical scrutiny. The design process, thus, should not only be conceptualised as aimed towards generating product P to fulfil F function; rather, the value-theoretic credentials of the function F themselves should be scrutinised: is this a good aim to have/a good function to fulfil to begin with? In this, VSD aims at integrating values of ethical importance in engineering design in a systematic way, and thereby generating human wellbeing.

This paper has three aims: first, it puts forth a dilemma for VSD having to do with value conflicts; to this effect, I put forth a case in which VSD either promotes gender-normative structures, or fails as user-sensitive VSD (Section §2). Second, it looks into the semantics of taste predicates and argues that a realist account serves us better for escaping the dilemma (§3). Third, the paper puts forth an etiological functionalist framework for understanding design function. Finally, the paper argues that, under this framework, VSD is faced by a new dilemma, between design functions and etiological functions fulfilment (§4).¹

2. Value Conflicts and VSD

The classical example of Value Sensitive Design is Privacy by Design: the latter is concerned with respecting the privacy of personally identifiable information in systems and processes. There are two reasons why this is an excellent case for understanding the nature of VSD: first, it highlights the departure from merely designing systems and processes with prudential functions in mind (efficiency, cost-efficiency) towards incorporating moral values and human rights in the design process. Second, and even more importantly for our purposes today, Privacy by Design is an excellent example of

¹The dilemmas discussed in this paper generalise to design more generally insofar as the value-theoretic VSD ambitions are shared.

VSD handling value conflicts: incorporating attention to moral values is expensive, and thus comes into conflict with prudential values. It is easy to see that permissible VSD will consist in VSD that aligns with the all-things-considered normative picture at the context: since moral values override (to some extent) prudential constraints concerned with cost-efficiency, permissible VSD will be privacy-informed VSD.

Relatedly, it is also interesting to see how theorists of VSD conceive of the incorporation of relevant values in the design process. According to the tradition (e.g. Friedman & Hendry 2019, Friedman et al 2006), the design process comprises three types of investigations: conceptual, empirical and technical. Conceptual investigations aim at understanding and articulating the various stakeholders of the technology, as well as their values and any values conflicts that might arise for these stakeholders through the use of the technology. Empirical investigations are qualitative or quantitative design research studies used to inform the designers' understanding of the users' values, needs, and practices. Finally, technical investigations can involve either analysis of how people use related technologies, or the design of systems to support values identified in the conceptual and empirical investigations.

Note that this traditional way to conceive of the design process features two important value-theoretic foci: first, a focus on users' perceived values, and a focus on intra-user value conflicts – e.g. cost vs. sustainability, cost vs. privacy etc.

What remains less explored on this traditional model are inter-user value conflicts and conflicts between value *simpliciter* and perceived value. To see what I have in mind, consider the following case from Sally Haslanger:

The Crop-Top

Daughter: “Can I have some money to buy a crop-top like Ashley’s to wear to school?”

Parents: “You can have a new top, but not a crop-top. Crop-tops are too revealing.”

Daughter: “But Mom[Dad], you’re just wrong. Everyone knows that crop-tops are cute; and I don’t want to be a dork.”

Parents: “I’m sorry, sweetie, crop-tops *are not* cute, and you *won’t* be a dork if you wear your track suit.”

(adapted from Haslanger 2007)

In The Crop-Top, Daughter and Parents are disagreeing over whether crop-tops are cute: Parents think they are not – because too revealing – , while Daughter thinks they are cute because everyone at school thinks so. Under the circumstances it seems that there is something right about Daughter’s reply to Parents, and their reply is not enough. And yet, aren’t the parents right? Are crop-tops not a paradigmatic instantiation of design incorporating problematic gender-normative values?

A parallel interesting example of technology design featuring the same value problem has to do with computer programmes and games aimed specifically to be enjoyed by women and girls: Girl-tech: computer programmes, games aimed to be enjoyed by women and girls (often referred to as

‘Girl-Tech’). These are commonly built on gender-normative assumptions: rewards are lipsticks, avatars are dressed in pink etc. The initial goal of creating such tech was arguably a morally admirable one: to attract more women and girls to technology. The problem, however, is that in order to achieve this goal the relevant technology has incorporated problematic gender-normative assumptions, and is thereby liable for furthering such problematic values.

The question that the next section will address is: are crop-tops and Girl-Tech successful value-sensitive design? If yes, how do we explain away the problematic gender-normative components (henceforth, The Problematic Intuition)? If no, how do we explain the intuition that since, as a matter of fact, users would fail to be attracted to these products unless they displayed these problematic properties, these are instances of user-sensitive VSD (henceforth also The User-Friendly Intuition)?

3. VSD and the Semantics of Taste Predicates

This variety of conflicts are familiar from the literature on the semantics of aesthetic and taste predicates. In what follows, I survey the main views and what they predict about value conflicts in VSD.

Contextualism

Very roughly, a contextualist semantics for a certain term holds that the term can have different contents in different contexts. As a result, sentences featuring context sensitive terms can have different contents, i.e. they express different propositions, in different contexts also. To take just one example of a context sensitive term, consider ‘here’. ‘Here’ can have different contents in different contexts. Suppose that Zac is currently in Zermatt in the middle of the winter. It does not require a stretch of the imagination to see how ‘It is snowing here’ may come out true in Zac’s context. In contrast, suppose that Barbara is in Bermuda where the weather is fair and it’s 25C. Again, it’s easy enough to see how ‘It’s not snowing here’ may come out true in Barbara’s context. A contextualist semantics for ‘here’ helps us explain how it can be that both Zac’s and Barbara’s assertions can come out true. ‘It is snowing here’ expresses different propositions in the two contexts. This, in turn, is because ‘here’ has different contents in the two contexts. In Zac’s context, the semantic value of ‘here’ is Zermatt. The proposition expressed by ‘It is snowing here’ is that it is snowing in Zermatt. In contrast, in Barbara’s context, the semantic value of ‘here’ is Bermuda. The proposition expressed by ‘It is not snowing here’ is that it is not snowing in Bermuda.

Contextualist approaches to aesthetic disagreement² venture to explain the intuitions about aesthetic disagreement by appealing to a contextualist semantics for aesthetic terms. It is easy enough to see that contextualist approaches to aesthetic

²For recent defences see e.g. Sundell 2011, Plunkett & Sundell 2013.

disagreement excel when it comes to handling the faultlessness intuition. Suppose that aesthetic terms are context sensitive. If so, they can have different contents in different contexts. What's more, sentences featuring them may express different propositions in different contexts.

For instance, in the Crop-top case, if 'cute' affords a contextualist semantics, the sentence 'The crop-top is cute' may express different propositions depending on whether it is uttered in at school or at home. In this way, contextualism can allow that Daughter's assertion that the top is cute and Parents' assertion that it is not cute are both true, in the relevant contexts. In this way, contextualism promises to accommodate both the User-Friendly Intuition and the Problematic Intuition: this instance of VSD is successful in the school context (where 'Crop-tops are cute' is true) and unsuccessful at home (where 'Crop-tops are cute' is false).

There is a serious problem with this result, however: contextualism predicts that this instance of VSD is successful in the school context, in spite of it being problematically incorporating and reinforcing gender-normative values. In this, it would seem, the Problematic Intuition is, after all, not fully accommodated: what we want is a result that holds context-independently.

Note, also, that contextualism will predict - contra intuition - that there is no disagreement between Daughter when she thinks to herself that 'Crop-tops are cute' at school and Parents when they deny this at home, since these contents are not incompatible across contexts. One popular way to go about accommodating the disagreement intuition within a contextualist framework is to appeal to metalinguistic disagreements: according to one prominent view, due to David Plunkett and Tim Sundell (e.g. 2013), for instance, what is going on in The Crop-top case is a disagreement about whether crop-tops *should* be included in the extension of 'cute' at that context or not. Now the main reason this is relevant for present purposes is that if contextualists can successfully accommodate the disagreement intuition, then they might also be able to explain the Problematic Intuition: maybe what is wrong in this case has to do with the fact that, even though crop-tops are in the extension of 'cute', they *shouldn't* be. On this way to look at things, Parents and Daughter disagree, about both contexts at stake, about whether crop-tops *should* fall in the extension of 'cute' or not: Parents think 'no', Daughter thinks 'yes'.

The question that arises concerns the type of should at issue in the metalinguistic disagreement in question. One way to go would be to take the should in question to be an all-things-considered should: Parents think moral considerations having to do with not reinforcing gender norms prevail, while Daughter disagrees. However, it seems as though we can easily imagine cases in which Daughter and Parents agree that crop-tops should - all things considered - be included in the extension of 'cute' at a particular context - say, because they both have a gun to their head. However, even in a case like this, it seems like they would still be in disagreement as to whether crop-tops actually

are cute or not, and, at the same time, the Problematic Intuition with regard to this instance of VSD survives.³

What all these considerations suggest is that we want to work with a suitable, stronger version of the Problematic Intuition henceforth, that is one that holds across contexts.

Relativism

Relativist approaches to the semantics of aesthetic and personal taste predicates (e.g. Kölbel 2002, MacFarlane 2014) aim to improve on contextualism by avoiding the idea that aesthetic terms are context sensitive. As a result, sentences featuring them will express the same proposition across all contexts. For instance, ‘Crop-tops are cute’ expresses the same proposition no matter whether Daughter or Parents utter it, i.e. the proposition that crop-tops are cute. Relativists distinguish between context and circumstance of evaluation. The circumstance of evaluation includes a range of parameters that determine the truth value of the proposition determined at the context. Standard views include worlds and times in the circumstance of evaluation. Suppose I utter ‘Brasilia is the capital of Brazil.’ The context of use determines a proposition that is the content of my utterance, here the proposition that Brasilia is the capital of Brazil. In order to evaluate whether the proposition is true, we need to have a world and time. For instance, the proposition is true at the actual world and at present. However, it is not true at the actual world in 1959 (when it was Rio de Janeiro) or at some possible world (at which Brasilia was never built). The key idea of relativist approaches to aesthetic disagreement is to add a judge parameter in the circumstance of evaluation. On this view, then whether a given proposition is true will turn, in addition, on who is judging it. Relativists accommodate the disagreement intuition in this way: While the sentence ‘Crop-tops are cute’ expresses the same proposition no matter whether Mary or Ann utters it, whether it is true or not will turn in addition on who is the judge in question. From the point of view of Daughter, crop-tops are cute. From the point of view of Parents, they are not. And the thought is that when Daughter is the judge, the proposition comes out true, while when Parents are the judge, it comes out false. In this way, one would think, at first glance, relativism could accommodate both the User-Friendliness (true for Daughter!) and the Problematic (False for Parents!) intuitions.

Unfortunately, things are not as easy as they first seem after all: to see this, note that VSD will have to predict that this is an instance of successful VSD: after all, crop-tops are for

³ According to a more plausible alternative, the ‘should’ at stake is an aesthetic should. Unfortunately, this suggestion does not fare much better, for at least two reasons. To see the first, consider the question as to what, if anything, grounds this aesthetic should. The intuitively most plausible answer is that it is the aesthetic truths on the ground. For instance, what explains why crop-tops should (aesthetically) fall under ‘cute’ is that they are cute. However, this answer isn’t available for champions of the present proposal. After all, this leads us right back into the territory of first order aesthetic disagreement.

Daughter. Since it is true (for Daughter) that they are cute, and since the Problematic Intuition is only explained with regard to Parents' view point, it seems to follow, at a minimum, that at least when considered relative to the consumer 'Daughter', this is an instance of fully successful VSD. However, this seems clearly false: cute-tops are bad for Daughter, in virtue of their reinforcing gender-normative values. As such, it seems, a successful account of what is going on needs to be able to not only accommodate the Problematic Intuition more generally, concerning what is going on in this case, but more narrowly, with respect to Daughter as the target consumer.

Realism

We have seen so far that Contextualism and Relativism about the semantics of taste predicates have failed to deliver the result we wanted about the target instance of Value Sensitive Design. We have also, along the way, figured out that what we want is for our view of the relevant semantics to be able to explain the Problematic Intuition across contexts of assessment, and when considered for Daughter as target consumer.

According to realist views, there is no semantic mystery about aesthetic disagreements such as the one in the Crop-Top case: one of the two parties is right, and the other one is wrong. On this view, short of occasional cases of indeterminacy, disagreements about aesthetic matters are garden variety factual disagreements: there is a fact of the matter as to whether crop tops are cute or not, and thereby one of the parties to the disagreement is making a mistake. Unsurprisingly, realists will have no problem at all accommodating the Problematic Intuition: Daughter is just wrong, crop-tops aren't cute. Kids like them, but they shouldn't. And that's because they reinforce problematic gender norms. This is an instance of failed Value Sensitive Design.

It becomes clear, however, that Realism will struggle with accounting for the intuition of User-Friendliness: after all, Realism predicts outright that 'Crop-tops are cute' is false, and, at first glance, lacks the resources to predict that there is anything good about crop-tops at all: sure, Daughter likes them, but she's wrong. Why is this a problem? After all, we are often wrong about matters of fact – and this is just one such case: this case of design is sensitive to pseudo-values – i.e. to what Daughter thinks is good for her – rather than values proper – what actually is good for her.

The problem with a reply along these lines is that, as a matter of fact, if Daughter goes to school wearing a track suit rather than a crop-top, her social life will very plausibly be negatively affected, along with her general wellbeing. Even if gender-normative social structures are wrong, they are real, and they affect our wellbeing. From this perspective, there is a clear sense in which it is good for Daughter to wear a crop-top at school rather than a track-suit.

Can realism recover this result? I think it can, when considered alongside a correct theory of wellbeing. The next section will develop this idea.

4. VSD, Realism, and Wellbeing

It has become standard in wellbeing research to distinguish theories of the nature of well-being as either hedonist theories, desire theories, or objective theories.⁴ Alas, all these theories face problems hard to surmount. As a result, positive psychology (the branch of psychology concerned with human happiness) has been known to import and work with different theories of wellbeing. This, of course, resulted in psychologists often talking past each other and delivering incommensurable research results, which got in the way of progress in the field.

In answer to the question, ‘What does well-being consist in?’, Hedonism answers, roughly, ‘The greatest balance of pleasure over pain’. This view was first expressed by Socrates in Plato’s dialogue, *Protagoras*. Hedonism—as is demonstrated by its ancient roots—has long seemed an obviously plausible view. However, it has mostly been abandoned in contemporary research due to one major objection, due to Robert Nozick (1971): Imagine that I have a machine that I could plug you into for the rest of your life. This machine would give you experiences of whatever kind you thought most enjoyable—writing a great novel, bringing about world peace, attending an early Rolling Stones’ concert. You would not know you were on the machine. Would you plug in? Most importantly, compare this life with one in which all these things actually happen: you bring about world peace etc. – but with only one exception: the coffee at Starbucks tastes better in the experience machine life than in the real life. Would you choose to plug in? The correct answer is clearly ‘no’. Hedonism, however, will have to answer ‘yes’: this theory mistakenly predicts that life in the experience machine is better, since it’s (slightly) more pleasant. This result is strongly counterintuitive. Furthermore, we don’t need to appeal to philosophical thought experiments to see that this is so: most of us think that, for instance, a life on hallucinogenic drugs would be inferior to a reasonably fulfilled, non-chemically enhanced life. All this seems to suggest that wellbeing is about more than just experiencing pleasure.

The experience machine thought experiment is one motivation for the adoption of a desire theory of wellbeing (e.g. Heathwood 2019). When you are on the machine, many of your central desires are likely to remain unfulfilled. Take your desire to write a great novel. You may believe that this is what you are doing, but in fact it is just a hallucination. And what you want, the argument goes, is to write a great novel, not the experience of writing a great novel. According to desire theories, roughly, wellbeing consists in desire satisfaction.

The classic objection the desire theories comes from Derek Parfit (1984): Imagine that you can start taking a highly addictive drug, which will cause a very strong desire in you for the drug every morning. Taking the drug will give you no pleasure; but not taking it will cause you quite severe suffering. There will be no problem with the availability of the drug, and it

⁴See (Crisp 2021) for an excellent overview.

will cost you nothing. Is your wellbeing increased by getting addicted to this drug? Intuitively, the answer is ‘no’. However, the desire-theorist will have to say that it is, since you get one extra desire easily satisfied.

These problems with hedonism and desire-based theories of wellbeing led researchers to propose objective theories of wellbeing. These theories are sourced in Aristotle’s theory of *eudaimonia* (human flourishing) – and are objective in that they claim that an objective list of things are constitutive of our wellbeing, independently of whether they produce pleasure in us or not, or whether we desire them or not: knowledge, love, physical health, good social interactions, moral thoughts and actions are things that often show up on such objective wellbeing lists. The main worry for these views, however, is a theoretical one: how to justify including one item or another on such an objective list? What grounds this take on the nature of wellbeing? On my preferred, broadly Aristotelian, objective-list theory of wellbeing, the latter is grounded in the nature and normativity of biological functions. Human systems serve etiological cognitive, emotive, moral, social, and physical functions, which, in turn, generate biological norms for proper functioning. Consider the heart. It is a paradigm case of a functional trait. Its function is to pump blood. What are functions? According to the leading account in the philosophy of science, the etiological theory of functions,⁵ functions turn on histories that explain why the item exists or operates the way it does. In the case of the heart, tokens of the type pumped blood in our ancestors. This was beneficial for our ancestors’ survival, which explains why the heart has continued to exist in creatures like us. As a result, the heart acquired the etiological function (henceforth also ‘e-function’)⁶ of pumping blood.

The heart acquires its e-function via natural selection over generations. Not all functional traits follow the model of the heart: there will be cases where a requirement of selection over generations for function acquisition will seem implausibly strong (Sosa 1993). The paradigmatic case is that of beneficial macro-mutations, so-called ‘hopeful monsters’ (Graham 2014, 30). Most mutations are harmful (think of extreme birth defects); once in a while, though, a happy accident happens: someone is born with an almost entirely new trait or organ, very different in kind from its ancestral trait, which actually benefits the recipient. Since they are mutations, they don’t have an evolutionary history; they are ‘first generation’ traits. Still, they can have functions. What matters is that the

⁵ Defended by people like e.g. David J. Buller (1998), Ruth Millikan (1984), Karen Neander (1991), and Larry Wright (1973). The etiological theory of functions is, by far, the most widely endorsed view in the literature, due to its normative import. That being said, the view defended here does not rest on the etiological view: any account of functions that delivers the highly desirable result that functions generate norms (i.e., that there is such a thing as a malfunctioning/properly functioning heart) will do for our purposes.

⁶ For applications of the etiological account of functions to epistemology, language, and feminist philosophy, see e.g. (Kelp 2018, 2021) (Simion 2016, 2018a, b, 2019a, b, c, d, 2021), (Graham 2012), (Millikan 1984).

existence/continuous existence of a trait is explained via a history of positive feedback:

[Etiological] functions arise from consequence etiologies, etiologies that explain why something exists or continues to exist in terms of its consequences, because of a feedback mechanism that takes consequences as input and causes or sustains the item as output (Graham 2014, 35).

Your heart has the etiological function to pump blood in your circulatory system. On the etiological theory of functions, that is because hearts pumped blood successfully in your ancestors, which lead to their survival, and which, in turn, explains the continuous existence of hearts. Because this function contributes to the explanation of why hearts continue to exist, it generates norms for what it is for a heart to be a properly functioning heart: it is for it to work in the way in which, in normal environmental conditions, it reliably enough fulfils its function (in the case of the heart, beat a particular rate). All human capacities afford a similar functionalist unpacking. Since the human system has several capacities, it will only be properly functioning as a whole – and thereby instantiate wellbeing – insofar as all these capacities are properly functioning – i.e. working in the way that reliably leads to function fulfilment in normal environmental conditions. Full wellbeing arises when these norms are met across all dimensions to a contextually determined sufficiency threshold.

5. A Functionalist Account of Successful Value Sensitive Design

In previous work (Simion & Kelp 2020), I have defended a functionalist account of design success. Not all things have etiological functions. Some things have *design functions* (d-functions). The dishwasher is a paradigm example of an item with a d-function. Its d-function is to clean dishes. Items with d-functions have their d-functions in virtue of the intentions of the designer. The reason why the dishwasher has the d-function of cleaning dishes is that the inventor of the dishwasher intended it to clean dishes. Crucially, e-functions are different from d-functions. For starters, there are things that have e-functions but not d-functions. The heart is a clear example here. At the same time, there are things that have d-functions but not e-functions. The reason for this is that e-functions require a history of success. The heart could only acquire the e-function of pumping blood because token hearts successfully pumped blood in the past. Exhibits in the Museum of Failure have d-functions but no e-functions: they just didn't work.

At the same time, many items with d-functions also acquire e-functions. Consider, in particular, new products which are launched on a competitive market. These products have d-functions. They are meant to do something. If they are successful, they will in addition acquire a certain kind of e-

function. If all goes well, these products will initially be bought to do what they are designed to do and subsequently will continue to be in demand (if all goes well) in virtue of the fact that they did the things they were designed to do and that this was beneficial to consumers – it generated wellbeing in them. But given that such products continue to be in demand, they will continue to be produced and bought. It is now easy to see that we have exactly the kind of feedback loop going that is characteristic of e-functions. In this way, d-functional items may acquire e-functions. On my view, the design of a particular product is successful only the d-function/s of the product turn/s into e-function/s of the sort just mentioned. By the same token, one important success condition for a d-functional item is for its d-function to be converted into the relevant kind of e-function.

Going back to The Crop-top case: recall that the problem for Realism was that, as a matter of fact, if Daughter goes to school wearing a track suit rather than a crop-top, her social life will very plausibly be negatively affected, along with her wellbeing. Even if gender-normative social structures are wrong, they are real, and they affect our wellbeing via the social effects of non-conformity. From this perspective, there is a clear sense in which it is good for Daughter to wear a crop-top at school rather than a tracksuit (the User-Friendliness Intuition).

Can realism, in conjunction with my wellbeing functionalism, recover this result, while at the same time predict crop-tops are an instance of failed VSD (The Problematic Intuition)? I think it can.

Note first that, of course, on my preferred account of wellbeing, experienced happiness need not imply wellbeing – pleasure-based and desire-based theories are wrong. The realist prediction that just because Daughter takes pleasure in wearing crop-tops it does not follow that they are not problematic is vindicated: crop-tops are bad for Daughter in several ways – socially, emotionally - due to their gender-normative design. Compatibly, of course, m functionalism does not predict that wellbeing will be fully independent of pleasure and desire-satisfaction: after all, emotional wellbeing is highly influenced by these factors. What realism predicts, then, for this case, is a familiar intra-user value conflict – to be solved via all-things-considered permissibility considerations: if negative considerations having to do with the effect of incorporating and reinforcing gender-normative values prevail – as I think they do – Daughter’s general wellbeing will be better served via stopping the production of crop-tops. If not, the reverse will hold.

At the same time, it is interesting to note that, if this picture is correct, designers will be faced with the following dilemma sourced in conflicts between real and perceived values: e-functions are not always good functions. Sometimes, they are generated by perceived rather than actual values. Fizzy drinks are bad for us, but we like them: in this way, they acquire an etiological function that is grounded in a pleasure component of wellbeing – they keep being consumed, and thus keep being

produced – while at the same time being bad for our general wellbeing.

In contrast, there will be cases of successful VSD – i.e. VSD that tracks our general wellbeing – that is not perceived as such, and thus the design function fails to be accompanied by an etiological function in virtue of the products being discontinued: healthy but tasteless drinks will likely have this faith.

6. Conclusion

VSD is hard. In order to instantiate successful VSD, the artifact needs to be designed around a ‘real’ value, sourced in our general wellbeing. But if this real value is not a perceived value, the artifact will likely be discontinued. On the other hand, observing perceived values alone runs the risk of VSD failure, since often perceived values are not furthering our general wellbeing. The case of VSD furthering gender normative structures is a case in point.

What is the solution to this new dilemma for VSD? An answer to this question falls outside of the scope of this paper. On a first approximation, however, it would seem as though one way forward would be one whereby VSD ensures continuous existence via serving enough perceived values as well, alongside ‘real’ values; ideally, the former should be at least neutral with regard to our general wellbeing.

References

- Buller, D. 1998. Etiological theories of function: a geographical survey. *Biology and Philosophy*, 13:505–527.
- Friedman, B., Hendry, & D. G. (2019). *Value Sensitive Design: Shaping Technology with Moral Imagination*. MIT Press.
- Friedman, B., Kahn Jr, P.H., Borning, A., & Kahn, P.H. (2006) Value Sensitive Design and Information Systems. *Human-Computer Interaction and Management Information Systems: Foundations*. ME Sharpe, New York, 384-372.
- Graham, P. 2012. Epistemic entitlement. *Noûs*, 46:449–82.
- Graham, P. 2014. Functions, warrant, history. In Fairweather, A. and Flanagan, O., editors, *Naturalizing Epistemic Virtue*. Cambridge University Press, Cambridge.
- Haslanger, S. (2007). “But Mom, Crop Yops Are Cute! Social Knowledge, Social Structure and Ideology Critique” *Philosophical Issues* Vol. 17, *The Metaphysics of Epistemology* (2007), pp. 70-91
- Heathwood, C. (2019), “Which Desires are Relevant to Well-being?”, *Nous*, 53: 664-88.
- Hopkins, R. (2000). Beauty and Testimony. In O’Hear, A. ed. *Philosophy, the Good, the True and the Beautiful*. Cambridge: Cambridge University Press, 209–236.
- Hopkins, R. (2011). How to Be a Pessimist about Aesthetic Testimony. *Journal of Philosophy* 108, 138–157.

- Kelp, C. 2018. 'Assertion: A Function First Account.' *Noûs* 52, 411-42
- Kelp, C. & Simion, M. (2017). *Commodious Knowledge. Synthese*, vol 194(5): 1487-1502.
- Kelp, C. and Simion, M. *Sharing Knowledge: A Functionalist Account of Assertion*. Cambridge: Cambridge University Press.
- Kölbel, M. (2002). *Truth Without Objectivity*. London: Routledge.
- MacFarlane, J. (2014). *Assessment Sensitivity: Relative Truth and its Applications*. Oxford: Oxford University Press.
- Millikan, R. 1984. *Language, Thought, and Other Biological Categories*. MIT Press, Cambridge/MA.
- Neander, K. 1991b. The teleological notion of 'function'. *Australasian Journal of Philosophy*, 69:454–68.
- Nozick, R., 1974, *Anarchy, State, and Utopia*, Oxford: Basil Blackwell.
- Parfit, D., 1984, *Reasons and Persons*, Oxford: Clarendon Press.
- Plunkett, D., and T. Sundell. 2013. disagreement and the Semantics of normative and evaluative Terms. *Philosophers' Imprint* 13/23.
- Simion, M. & Kelp, C. (2020). Conceptual Innovation, Function First. *Nous*. 54/4: 985-1002.
- Smion, M. (2016). Perception, History and Benefit. *Episteme*, 13(1): 61-76.
- Simion, M. & Kelp, C. (2018). Assertion: The Constitutive Norms View. In Goldberg, S. (ed), *Oxford Handbook of Assertion*. Oxford University Press.
- Simion, M. (2019a). Knowledge-First Functionalism. *Philosophical Issues*. Ed: Lisa Miracchi. Online First.
- Simion, M. (2019b). Saying and Believing: The Norm Commonality Assumption. *Philosophical Studies* Volume 176, Issue 8: 1951–1966.
- Simion, M. (2021). *Shifty Speech and Independent Thought*. Oxford: Oxford University Press.
- Simion, M. (2018a). The Explanation Proffering Norm of Moral Assertion. *Ethical Theory and Moral Practice*. 21(3), pp. 477-488
- Simion, M. (2019c). Hermeneutical Injustice as Basing Failure. In Bondy, P. and Carter, J.A. (eds.), *Well Founded Belief: New Essays on the Epistemic Basing Relation*, Routledge.
- Simion, M. (2019d). Assertion: The Context Shiftiness Dilemma. *Mind & Language*. 34: 503–517
- Simion, M. (2018b). Epistemic Trouble for Engineering 'Woman'. *Logos and Episteme*, vol. 9(1): 91-98.
- Sundell, T. (2011). Disagreements about Taste. *Philosophical Studies* 155 (2): 267–288.

Wright, L. (1976). *Teleological Explanations: An Etiological Analysis of Goals and Functions*. Berkeley, CA: University of California Press.